UNIVERSITY COLLEGE LONDON

Faculty of Mathematics and Physical Sciences

Department of Physics & Astronomy

TOWARDS PRECISION MEASUREMENTS OF LARGE-SCALE STRUCTURE WITH NEXT-GENERATION SPECTROSCOPIC SURVEYS

JAMES ALEXANDER FARR

Thesis submitted for the Degree of Doctor of Philosophy of University College London

SUPERVISORS:

DR. ANDREU FONT-RIBERA

PROF. ANDREW PONTZEN

EXAMINERS:

PROF. MARTIN HAEHNELT

DR. BENJAMIN JOACHIMI

Friday 22nd January, 2021

I, James Farr, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

JAMES FARR

ABSTRACT

Over recent decades, spectroscopic surveys have yielded exceptional measurements of the large-scale structure of the Universe. Notably, they have measured the baryon acoustic oscillation (BAO) scale at late times, helping to develop a tightly-constrained cosmological model by complementing measurements from the cosmic microwave background. In this thesis, we focus on late-time BAO measurements from quasar (QSO) spectra. Such measurements can be made at $1 \lesssim z \lesssim 2$ via direct QSO clustering, and at $2 \lesssim z \lesssim 4$ via the Lyman-α (Lyα) forest extracted from high-z QSO spectra.In the near future, the Dark Energy Spectroscopic Instrument (DESI) will continue to advance this field, increasing the quantity and quality of QSO spectra available via a host of technological improvements. In order to maximise the impact of its data, however, DESI will require major advances in analysis methods to be made.

This thesis describes work to develop such methods for use in two areas of DESI's QSO survey. First, we address the construction of optimal strategies for classifying QSO target spectra. We use data from existing surveys to demonstrate the performance of potential strategies, finding that high performance levels can be achieved using existing classification tools. Next, we present `LyaCoLoRe`, a package developed to produce mock Lyα forest datasets from simple simulations, to be used in Lyα BAO analyses. We describe the methods employed by `LyaCoLoRe`, and demonstrate that our mocks are suitable to be used in Lyα BAO studies present and future. We then discuss applications of the classification strategies and mock datasets presented previously, as well as a method of using BAO measurements to constrain the local cosmic expansion rate, showing results from current datasets and providing forecasts for DESI. We conclude by highlighting a number of future paths which our work could follow, with particular focus on the opportunities that will emerge from DESI.

Impact Statement

This thesis presents original work designed to improve the efficiency and reliability of data analysis in current and future spectroscopic quasar (QSO) surveys. The methods it presents will enable greater automation and rigour in such analyses, both of which are significant as the size and quality of datasets increases.

This work has had a tangible impact on the recently completed extended Baryon Oscillation Spectroscopic Survey (eBOSS): the mock datasets presented in Chapter 6 were used to validate the analysis pipeline for measuring the baryon acoustic oscillation (BAO) scale from Lyman-α forest data. As described in § 7.2, the mocks allowed the analysis pipeline to be validated in a more stringent manner, thus yielding more reliable measurements of the BAO scale. The mocks have also been used in the context of the next generation spectroscopic survey — the Dark Energy Spectroscopic Instrument (DESI) — to help test blinding methods, assess survey strategies and quantify systematic effects. Equally, the QSO selection strategies presented in Chapter 5 have been applied to initial DESI data (see § 7.1), with the results now being used to develop a tool to assist the visual inspection of future DESI spectra. They have also motivated efforts to include the classifier `QuasarNET` within the DESI pipeline, aiming to improve the automatic selection of high-z QSOs for reobservation using strategies such as those introduced in § 5.3.1.

The work presented in this thesis has reached a number of researchers from around the world through a variety of channels. It has been presented at nine collaboration meetings relating to eBOSS and DESI, including as part of a "hack session" to identify unusual QSO target spectra. Additionally, three collaborators within DESI have visited UCL to work on projects relating to the mock datasets of Chapter 6. This yielded some results presented in this thesis, as well as initialising further developments of the mocks in an ongoing project. Furthermore, this thesis' work has directly led to two first-author publications, as well as three further authorial contributions. Finally, the work on QSO classification has been presented in three talks to machine learning specialists,

once within DESI and twice to external, international groups.

Beyond academia, the skills and techniques developed as a result of the computational nature of this thesis' work are eminently applicable to the wider world. The development of realistic mock datasets is vital in ensuring the reliability of automated systems throughout modern life, including the flow of traffic in cities and consumer demand for groceries. While the context of such datasets is somewhat different to that presented in Chapter 6, the dual requirements of computational efficiency and data realism are universal. Equally, the QSO target classification tasks addressed in Chapter 5 are examples of the broader class of image recognition problems. Effective solutions to these are central to a host of modern technologies, including facial recognition and self-driving cars. Developing optimal solutions that are free from bias is a challenge that will become ever more important in the coming years.

ACKNOWLEDGEMENTS

My thanks, first and foremost, must go to my supervisors, Andreu Font-Ribera and Andrew Pontzen. The time and thought that you have put into guiding me has made the journey of my PhD substantially more straightforward, and has no doubt made shaped me as a researcher. Your dedication in both helping to diagnose the most granular of problems, as well as encouraging me to see the bigger picture has been immensely valuable and is greatly appreciated. Most importantly, your mentorship has helped to make this journey fun; I have enjoyed every step of the way and will be forever grateful to you both.

I must also thank all of those collaborators who have made invaluable contributions in supporting my research, and with whom I consider myself very fortunate to have worked. Alma González-Morales, Andrea Muñoz-Gutiérrez, Anže Slosar, David Alonso, Hélion du Mas des Bourboux, Ignasi Pérez-Ràfols, Javier Sánchez, Jean-Marc Le Goff, Jim Rich, Julien Guy, Mat Pieri, Nicolas Busca, Stephen Bailey, Suk Sien Tie, Thomas Etourneau; thank you for your feedback, your suggestions, and your time spent answering my questions. You have provided me with wonderfully varied points of view, and have been instrumental both in helping to form the work presented in this thesis, and in making the many collaboration meetings we have shared highlights of my PhD.

Beyond research, I owe a great deal to the support of the Group A community at UCL. Thanks to Constance, Harry, Luisa and Martin for showing me the ropes and guiding me to seminars through the labyrinthine passages of UCL in my first months; to Andrei and Krishna for making our walking tour of San Francisco so much fun; to Benjamin for demonstrating what a true love of coffee looks like; to Catarina for your endlessly appreciated positivity; to Chris for being a great sounding board for thoughts on both the Lyα forest and the Premier League; to Francesca for your many offers of assorted biscuits and cakes; to Max for your exceptionally valuable enthusiasm and laughter; to Roger (a de facto member of UCL now) for great times in Berkeley and Ohio; and to Xie for providing wonderfully varied lunchtime discussions. Together, you have helped to make

my time at UCL unforgettable, and I only wish that it was not over so soon.

The intensity of a PhD necessitates the occasional break from the world of academia, and I have been fortunate to have a wonderfully supportive network around me for those moments. Thank you first to all those at Hampstead and Shamley Green cricket clubs for providing ample opportunities to stand in fields around London and Surrey, and to reward our "efforts" in the bar afterwards. And of course, thank you to Catherine, Chris, Fay, Iwan, Ninghui and Simi for our friendship through all these years. Our adventures in Cambridge, London, Wales and beyond have made so many special memories, and been integral to keeping me happy and healthy.

A list of thanks would not nearly be complete without mentioning those closest to me. To Jane and Ian Price, thank you for welcoming me into your family so wholeheartedly, for providing me with an exceptional office space to write parts of this thesis through lockdown, and for your incredible support over many years. To my Mum, Dad, Rosie and Katharine, thank you for creating such a wonderful home environment, and for 26 years (and counting) of joy. The curiosity you have instilled in me has no doubt driven me through this PhD, and your love and support have been essential in allowing me to explore the world of cosmology with so much freedom. And finally, thank you to Amy, for indulging my occasional cosmological diatribes, for providing welcome reality checks where necessary and, most importantly, for being alongside me at every moment. This journey would not have been possible without you, and it has been a true pleasure to share it with you.

Contents

List of Figures

List of Tables

PREFACE

Curiosity about the origin and evolution of the Universe has inspired human thought for millennia, resulting in countless ideas both scientific and otherwise. For much of history, interpretations of the objects seen in the night sky drew as much on philosophy and religion as they did on physics, yielding a host of speculative deductions about the Universe beyond Earth (for example, the left panel of Figure 0.1). Indeed, it is only in the 20$^{\text{th}}$ Century that our curiosity grew into the subject of cosmology that we recognise today.

At the start of the Century, many astronomers believed that the Universe was "static" in its evolution, and that observing galaxies beyond our own was not possible (forming one side of the "Great Debate": Shapley and Curtis, 1921). 100 years later, however, and the combined efforts of theoretical and observational cosmologists has yielded a radically different picture. The introduction of General Relativity (GR, Einstein, 1915) provided a coherent theoretical framework with which to model our Universe, and the rapid development of astronomical instrumentation enabled the cosmological models permitted by GR to be constrained. Over the course of the century, the visible Universe was found to include galaxies beyond the Milky Way (Hubble, 1926), to be expanding (Hubble, 1929), and at an accelerating rate (Riess et al., 1998; Perlmutter et al., 1999). Thus by the end of the 20$^{\text{th}}$ Century, the observed Universe appeared much larger and more dynamic than it had done at the start.

Over the first two decades of the 21$^{\text{st}}$ century, cosmology has undergone another major transformation. As the digital age has unfolded, it has brought with it datasets of unprecedented size and quality from a range of observational fields, marking the advent of the "precision cosmology" era. Perhaps most famously, the cosmic microwave background (CMB) has provided a window onto the Universe in its youth, while large-scale structure (LSS) measurements have tracked the statistical distribution of matter through its late-time evolution. Of particular note within the study of LSS has been the use of baryon acoustic oscillations (BAO). Acting as a "standard ruler", the BAO scale has

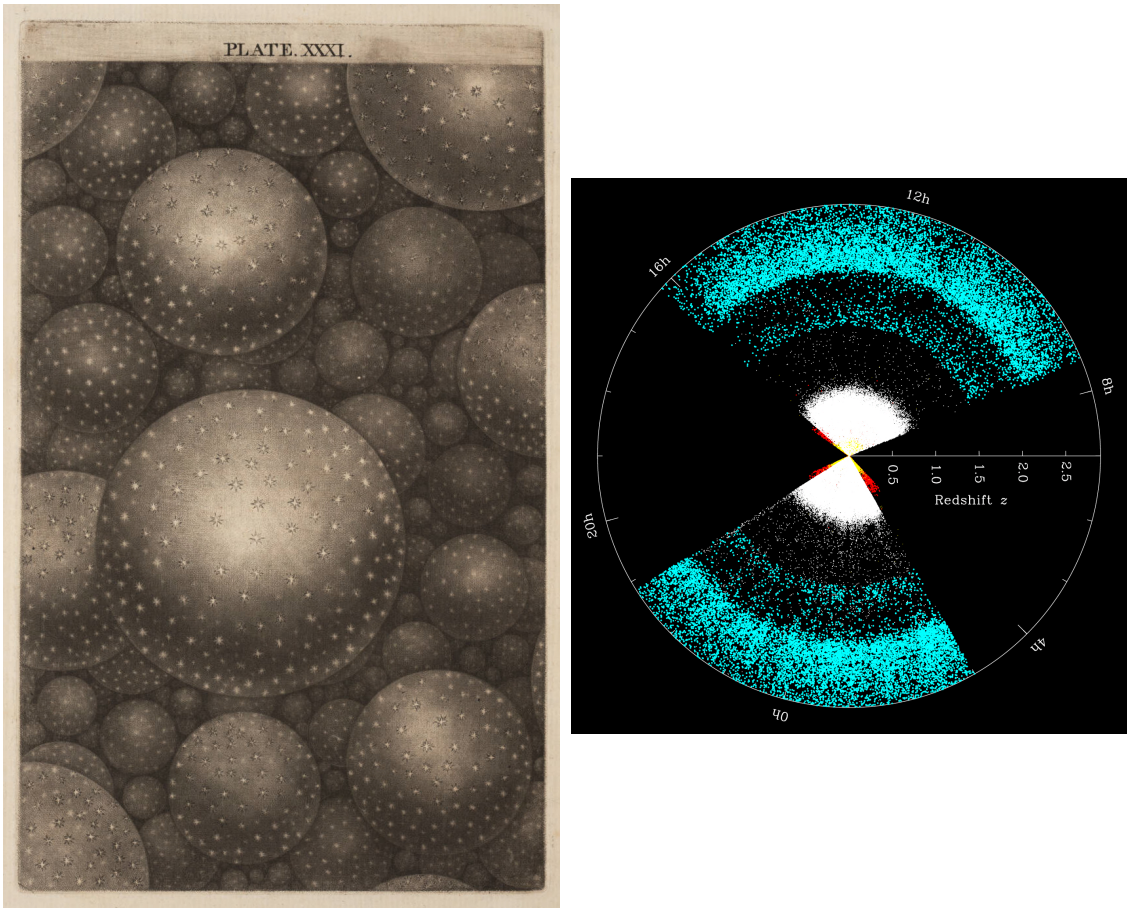

Figure 0.1. *Left:* the first depiction of multiple galaxies to be printed in a book. Each galaxy is represented by a spherical shell with inward-pointing volcanoes distributed over its surface, which are observed from the inside as stars; *Right:* a modern representation of galaxies beyond our own, with each coloured point representing a galaxy or quasar observed by the Baryon Oscillation Spectroscopic Survey (Dawson et al., 2013) during the first two years of its five-year survey. *Figure credit: Wright (1750) via Linda Hall Library; Michael Blanton and the Sloan Digital Sky Survey (SDSS) via the SDSS blog.*

enabled high-precision measurements of late-time cosmic expansion, providing complementary information to the CMB and helping to constrain a broader class of potential cosmologies.

The development of these late-time BAO measurements has been driven by the increasing power of spectroscopic surveys, which can distinguish the BAO feature in three dimensions thanks to precise redshift determinations of the objects they observe. At low redshifts, spectroscopic surveys may target a range of different galaxy types, but above $z \simeq 1$ the observed number density of galaxies becomes low. Instead, quasars (QSOs) — typically very bright due to their active galactic nucluei — may be targeted to extend the boundaries of spectroscopic observations, and thus probe the Universe in different phases of its evolution. Recent instruments have conducted extensive QSO surveys, and in the coming years, the Dark Energy Spectroscopic Instrument (DESI)

will take up this baton. DESI's main survey will begin in 2021 and will last for five years, over which time it will observe approximately 2.4 million QSOs (DESI Collaboration et al., 2016a) across a range of redshifts, allowing DESI to produce exceptional measurements of LSS. This will expand upon the work of DESI's predecessors by pushing cosmology to higher levels of precision, allowing finer distinctions to be made between cosmological models, and thus holding the potential to uncover new physics. In order to make best use of this wave of new data, the advancement of observational and analytical methods is crucial, and it is these methods which we seek to develop through this thesis.

THESIS OUTLINE

This thesis begins in Part I with a summary of the background and context of modern cosmology. Chapter 1 addresses the Universe homogeneously, deriving its behaviour on large scales assuming that the "Cosmological Principle" holds throughout, and building a basic cosmological model. We then introduce inhomogeneities in Chapter 2, deriving key results on the evolution of structure in a perturbative framework. In Chapter 3, we address modern methods of measuring our Universe, including via the CMB and LSS. We describe the measurements which have been made, and their ability to constrain cosmology. Finally, we conclude our introductory discussion in Chapter 4, where we focus on spectroscopic QSO surveys past, present and future.

Part II and Part III constitute the bulk of the work of this thesis. Both are constructed with DESI in mind, and address problems relevant to DESI's science goals. In Part II, we discuss optimal strategies for providing classifications and redshifts of QSO target spectra in spectroscopic surveys. We describe the landscape of existing classifiers and construct high-performing classification strategies using their outputs. Then, in Part III, we address the production of mock datasets for use in the Lyα BAO analyses of spectroscopic surveys. We present a new tool — `LyaCoLoRe` — for generating synthetic data, and verify that its outputs are suitable for a DESI-like survey. Further details for Part II and Part III are addressed in Appendix A and Appendix B respectively.

The thesis draws to a close with Part IV. In Chapter 7 we present three short discussions, addressing the application of our QSO classification strategies to initial DESI data; the use of the `LyaCoLoRe` mocks in the final Lyα BAO analysis of the extended Baryon Oscillation Spectroscopic Survey; and the use of BAO results more generally in constraining the Hubble constant. Finally, in Chapter 8, we summarise the findings of this thesis, discuss the onward journey of this thesis' work and conclude.

Part I

Background & Context

CHAPTER 1

THE HOMOGENEOUS UNIVERSE

We begin Part I by constructing a simplified model of our Universe. We do so within the central underlying theory of modern cosmology: General Relativity (GR, Einstein, 1915). The GR framework provides a coherent description of gravity in our Universe, linking the Universe's geometry to its contents and describing how they interact. We use GR to model the behaviour of our Universe on large scales, treating it as homogeneous to understand its origins and expansion. In this Chapter, we summarise the basic principles of GR and its application to our Universe, developing a simple model to explain large-scale phenomena that we observe today. From this baseline model, we then consider the formation of structure on smaller scales in Chapter 2 and their observational signatures in Chapter 3, before focusing on the details of spectroscopic surveys in Chapter 4.

The structure of this chapter is as follows. In § 1.1 we outline selected elements of this theory, presenting key equations and their interpretations. In § 1.2, we examine some geometrical properties that can be derived from GR, while in § 1.3 we explore the dynamical effects of different matter sources on our Universe's behaviour. We then discuss the need for an inflationary period in § 1.4, before finally summarising the thermal history of the Universe in § 1.5. Throughout, we use units with $c = 1$. The content in this chapter has been reviewed extensively in the literature, and we follow the summaries provided by Dodelson, 2003 and Weinberg, 2008.

1.1 General relativity

1.1.1 The Cosmological Principle

At the centre of General Relativity is a key guiding idea, known as the *Cosmological Principle*. First formally put forward by Newton in the 17th Century, this proposes that we as observers do not sit at a unique, or "special" point in the Universe, and thus that the Universe appears the same to all other observers on sufficiently large scales. This imposes two key properties on the large-scale Universe:

Isotropy: The Universe appears the same in every direction

Homogeneity: The Universe appears the same at every point in space

Mathematically, these properties are equivalent to rotational and translational symmetries in space; we will take advantage of these symmetries shortly.

1.1.2 The Friedmann-Lemaitre-Robertson-Walker metric

In General Relativity, the three dimensions of space and single dimension of time are fused into a four-dimensional manifold known as *spacetime*. On this manifold, we define a metric $g_{\mu\nu}$ with signature $(-,+,+,+)$, allowing us to measure distances in spacetime by converting coordinates $X^\mu = (t, x_i)$ into the line element

$$\mathrm{d}s^2 = g_{\mu\nu} X^\mu X^\nu = -\mathrm{d}t^2 + \mathrm{d}l^2. \tag{1.1}$$

Here $\mathrm{d}l^2$ is the spatial line element, and we use the Einstein summation convention to sum over all repeated indices. In general, $g = g(t, x_i)$ is a function of spacetime, with the precise nature of this dependence determined by the distribution of matter in the Universe. This dependence can be complex, and so we rely on the symmetries implied by the Cosmological Principle to derive a tractable metric for our Universe.

In particular, these symmetries allow us to decompose spacetime into a sequence of time-ordered three-dimensional spatial "slices", each of which is homogeneous and isotropic. These slices must then have uniform curvature given by a parameter K. By rescaling our coordinates with a positive "scale factor" a, we can impose that K takes values from the set $\{-1, 0, +1\}$, corresponding to Universes with negative, flat and positive curvatures respectively. These three options describe qualitatively different Universes, which we will explore in later stages of this

chapter. We can then write the spatial line element $\mathrm{d}l^2$ as

$$\mathrm{d}l^2 = a^2\gamma_{ij}\mathrm{d}x^i\mathrm{d}x^j, \text{ where } \gamma_{ij} \equiv \delta_{ij} + K\frac{x_i x_j}{1 - Kx_k x^k}. \tag{1.2}$$

Here, γ_{ij} is the spatial part of our spacetime metric $g_{\mu\nu}$, and may take various forms dependent on choice of coordinates. In order to derive our $g_{\mu\nu}$, we may substitute eq. (1.2) into eq. (1.1), allowing our scale factor a to vary with time as $a(t)$. Using spherical polar coordinates, this yields

$$\mathrm{d}s^2 = -dt^2 + a^2(t)\left[\frac{\mathrm{d}r^2}{1 - Kr^2} + r^2\mathrm{d}\Omega^2\right], \tag{1.3}$$

where $\mathrm{d}\Omega^2 = \mathrm{d}\theta^2 + \sin^2\theta\mathrm{d}\phi^2$ is the solid angle element. This defines the Friedmann-Lemaitre-Robertson-Walker metric, which, through comparison with eq. (1.1), can be seen to be diagonal in these coordinates. It is sometimes helpful to simplify eq. (1.3) via two coordinate transformations. We define the comoving distance χ via $\mathrm{d}\chi = \mathrm{d}r/\sqrt{1-kr^2}$ and the conformal time τ via $\mathrm{d}\tau = \mathrm{d}t/a(t)$. Using these two quantities, the line element becomes:

$$\mathrm{d}s^2 = a^2(\tau)\left[-\mathrm{d}\tau^2 + \mathrm{d}\chi^2 + S_K^2(\chi)\mathrm{d}\Omega^2\right], \tag{1.4}$$

where

$$S_K(\chi) = \begin{cases} \sinh\chi, & K = -1 \\ \chi, & K = 0 \\ \sin\chi, & K = +1 \end{cases}. \tag{1.5}$$

This form of the metric is especially useful when considering the movement of light in our Universe, as light travels along paths with $\mathrm{d}s^2 = 0$, and so we may simply ignore the $a^2(\tau)$ prefactor in our calculations.

1.2 Redshift & distances

Particles falling freely through spacetime move along paths called geodesics, which can be calculated by solving the geodesic equation

$$P^\nu\nabla_\nu P^\mu = 0, \tag{1.6}$$

where ∇_ν is the covariant derivative. Evaluating the $\mu = 0$ component of this equation and applying our assumption of homogeneity, we may derive that the three-momentum of any particle p must scale according to the inverse of the scale factor of the Universe: $p \propto 1/a(t)$. For photons, quantum mechanics tells us that $p = E = h/\lambda$, and thus $\lambda \propto a(t)$. As such, a photon emitted at time t_e with wavelength λ_e, and observed at time t_0 has wavelength

$$\lambda_0 = \frac{a(t_0)}{a(t_e)}\lambda_e. \tag{1.7}$$

We use this relationship to define a key quantity in observational cosmology: redshift z. This describes the proportional shift in wavelength of a photon:

$$z \equiv \frac{\lambda_0 - \lambda_e}{\lambda_e} = \frac{a(t_0)}{a(t_e)} - 1 = \frac{1}{a(t_e)} - 1, \tag{1.8}$$

where in the final equality we have made use of the conventional rescaling of $a(t)$ to set its current value $a(t_0)$ to 1[1].

We can expand $a(t_e)$ in a power series by writing $a(t) = a(t_0) + \dot{a}(t_0)(t - t_0) + \ldots$. For nearby objects and with $c = 1$, the physical distance d is simply equal to $t_0 - t_e$. We thus use eq. (1.8) to write $z = H_0 d$, where $H_0 \equiv \dot{a}(t_0)/a(t_0)$ is the local Hubble constant, describing the current rate of cosmic expansion. As shown in Figure 1.1, this measurement was first made in Hubble (1929), and there has been substantial debate about its precise value in recent years (see § 3.3).

At greater distances, this approximation does not hold. Indeed, in an expanding universe, the "distance" to an astrophysical object is not a trivial quantity, and can be defined in a number of ways. In eq. (1.5), we defined $S_K(\chi)$, which is known as the *transverse comoving distance*, also written d_m. Within this equation, we have also made use of the *comoving distance* χ, defined by

$$\chi(t_e) = \int_{t_e}^{t_0} \frac{1}{a(t)} \mathrm{d}t = \int_0^{z_e} \frac{1}{H(z)} \mathrm{d}z. \tag{1.9}$$

These two distances follow naturally from our metric definitions but, crucially, are not observable.

An alternative, observable distance can be derived by considering a distant luminous object, with absolute luminosity L. In a static, Euclidean space, the flux received by an observer F at a distance d from a source is defined straightforwardly by $F = L/4\pi d^2$, as the total flux through a spherical shell around the source must be constant. However, in our FLRW spacetime, we must modify this to account for potentially non-zero curvature and the expansion of the Universe

[1] If we do so, then the curvature K no longer takes values in $\{-1, 0, +1\}$, as that requires a different rescaling of a.

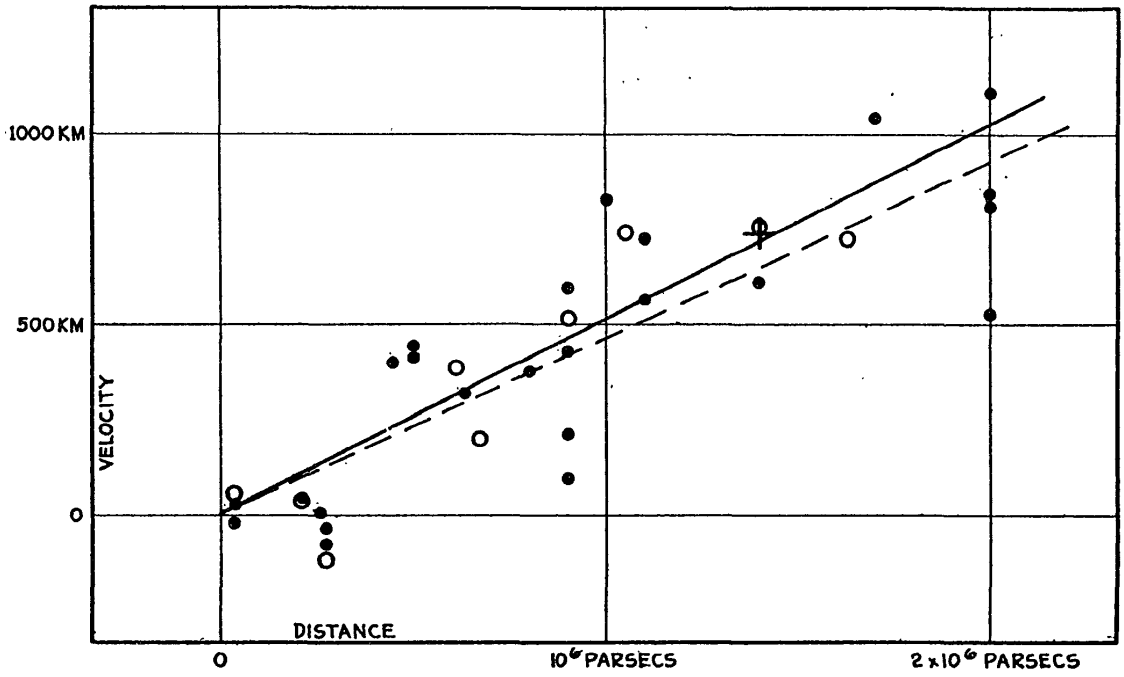

Figure 1.1. The original plot of Hubble's data from his 1929 paper (Hubble, 1929). On the horizontal axis is the distance of the galaxies d, and on the vertical axis is the recession velocity v (derived from each object's redshift). The Hubble parameter H_0 is taken as the gradient of Hubble's line of best fit, and gives a value of $\sim$500 km s^{-1} Mpc^{-1}. Modern measurements, however, tend to lie in the region of 70 km s^{-1} Mpc^{-1}. Despite the numerical inaccuracy, the data shown in this plot still represents a crucial first step towards our current understanding of the universe. *Figure credit: Hubble (1929).*

during the photons' journeys. In such a spacetime, the area of the spherical shell is given by $4\pi d_m^2$, and so we exchange d for d_m in our previous equation for F. Equally, in an expanding Universe, both the rate of arrival of photons and the energy of each photon are reduced by a factor $a(t_1)/a(t_0) = (1+z)^{-1}$, where t_1 and t_0 denote the time at which photons leave their source and arrive at the observer respectively, and z is the redshift of the source relative to the observer. As luminosity represents the energy radiated by a source per unit time, we must multiply L in our previous relation by a factor of $(1+z)^{-2}$, obtaining

$$F = \frac{L}{4\pi d_m^2 (1+z)^2} \equiv \frac{L}{4\pi d_L^2}, \tag{1.10}$$

where we have defined $d_L = d_m(1+z)$ as the *luminosity distance* to mimic the form of our initial equation.

In subsequent chapters, we will consider observations of structures of fixed physical size (which we label by d for now), which appear to observers as occupying an angular interval $\Delta\theta$ on the sky.

Again, we consider first a static, Euclidean universe, which would imply a distance $d_A = d/\Delta\theta$. Here, d_A is known as the *angular diameter distance*. Transitioning to our FLRW spacetime once again, we can deduce that

$$d = a(t_e)S_K(\chi)\Delta\theta = \frac{d_m\Delta\theta}{1+z}, \tag{1.11}$$

which yields $d_A = d_m/(1+z)$, and thus $d_L = d_A(1+z)^2$. Defining these two distances d_L and d_A provides us with observable quantities to link our measurements back to quantities more natural to the underlying theory.

1.3 Dynamics

1.3.1 Governing equations

Having considered the geometrical properties of our Universe, we now consider how it evolves dynamically. This behaviour is determined by the Einstein Field Equations, a set of partial differential equations defined by:

$$G_{\mu\nu} = 8\pi G T_{\mu\nu}. \tag{1.12}$$

This relates the geometry and curvature of the Universe — described by the Einstein tensor $G_{\mu\nu}$ — to the matter distribution of the Universe — described by the energy-momentum tensor $T_{\mu\nu}$. $G_{\mu\nu}$ can be computed from the FLRW metric $g_{\mu\nu}$ via a series of protracted calculations. Ultimately, these yield non-zero components:

$$G^0{}_0 = 3\left[\left(\frac{\dot{a}}{a}\right)^2 + \frac{K}{a^2}\right], \text{ and } G^i{}_j = 3\left[2\frac{\ddot{a}}{a} + \left(\frac{\dot{a}}{a}\right)^2 + \frac{K}{a^2}\right]\delta^i{}_j. \tag{1.13}$$

We now turn our attention to $T_{\mu\nu}$. By imposing the twin requirements of homogeneity and isotropy, we may constrain the form of $T_{\mu\nu}$ to that of a perfect fluid as seen by a comoving observer:

$$T^\mu_\nu = (\rho + P)U^\mu U_\nu - P g^\mu_\nu = \begin{pmatrix} -\rho & 0 & 0 & 0 \\ 0 & P & 0 & 0 \\ 0 & 0 & P & 0 \\ 0 & 0 & 0 & P \end{pmatrix}. \tag{1.14}$$

Here, ρ and P are the density and pressure of the fluid in its rest frame respectively, and $U_\mu = \mathrm{d}X_\mu/\mathrm{d}\tau$ is its four-velocity relative to the observer. Subsequently, by imposing conservation of

energy and momentum densities on $T_{\mu\nu}$, we may derive the continuity equation:

$$\dot{\rho} + 3\frac{\dot{a}}{a}(\rho + P) = 0. \tag{1.15}$$

Then we may combine our forms of G and T in eq. (1.13) and eq. (1.14) respectively via eq. (1.12). With some manipulation, this yields the *Friedmann equations*:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G\rho}{3} - \frac{K}{a^2} \tag{1.16}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3P), \tag{1.17}$$

where ρ and P here correspond to the total energy density and pressure of the Universe, each of which receives contributions from a variety of different component fluids.

1.3.2 Cosmic ingredients

There are numerous important components to the energy density of our Universe that behave in qualitatively different ways. Most importantly, each fluid may have a different equation of state $w = P/\rho$, which we assume to be uniform. Here, we describe the key properties of several such "cosmic ingredients", as well as highlighting major open questions.

- *Photons:* As purely relativistic particles, the energy density of photons is entirely kinetic, yielding an equation of state $P = \rho/3$. Thus they provide a source of pressure, which may oppose gravity and yield oscillatory behaviour in primordial fluids thanks to their coupling to baryons (see § 2.2.2). At a particular point during cosmic expansion, this coupling drops away, with remnant early-Universe photons then "free-streaming" across the Universe and forming the *cosmic microwave background* (CMB), the foremost early-Universe probe used today (see § 3.1).

- *Baryons:* Of all cosmological components, baryons are perhaps the most familiar, including in their ranks the protons, neutrons and electrons[2] that make up our immediate surroundings. Baryons interact gravitationally, collapsing to form large-scale cosmological structures as well as galaxies, stars and planets. They also interact electromagnetically, and are tightly coupled to photons during the early Universe. As such, oscillations are also present in

[2]Despite electrons being leptons rather than baryons, cosmologically speaking they are included under the term "baryons" for simplicity. They provide a less dominant contribution to the energy density of "baryons" than protons or neutrons, but their interaction with photons is of vital importance to cosmology.

their energy density due to pressure from photons, ultimately resulting in the formation of baryon acoustic oscillations (BAO; see § 2.2.2). Later in cosmic history, baryons and photons "decouple" (see § 1.5), and so the baryons instead follow dark matter to form cosmic structure. While stars and galaxies represent clearly visible baryonic presences in our Universe, the majority of baryonic matter is not found in such objects. Instead, it is distributed through the Universe as diffuse gas of varying properties, which may be described by terms such as the inter-stellar medium (ISM), circum-galactic medium (CGM) and the inter-galactic medium (IGM, see Meiksin, 2009; McQuinn, 2016, for recent reviews).

– *Dark matter:* Assigned the descriptor "dark" due to its lack of interaction with photons,dark matter behaves as a pressureless fluid ($w = 0$), the nature of which is not currently understood beyond its gravitational influence. Rather than directly observing dark matter — as we may do baryonic matter — we must infer its presence via a range of methods, including the rotation curves of galaxies, gravitational lensing and cosmological measurements from the CMB or large-scale structure (see Roos, 2010, for a review). Efforts are ongoing to further our understanding of dark matter, including experiments to detect it directly. A number of theories have been proposed to provide possible dark matter candidates, including weakly interacting massive particles (WIMPs), massive astrophysical compact halo object (MACHOs) and axions (see Feng, 2010, for a review). Candidates are generally classed as "hot" or "cold" — referring to whether the particles are relativistic or not when they decouple during the early Universe — while "warm" options have also recently entered consideration (e.g. Menci, Fiore, and Lamastra, 2012). This dark matter "temperature" is defined by the particles' typical velocities, and thus by the particles' mass. It has notable observational effects as "warmer" dark matter particles will typically travel faster and thus are better able to escape gravitational potential wells. This defines a mass-dependent *free-streaming scale* below which dark matter does not cluster, observable as a cut-off in the power spectrum of matter. Measurements of small-scale clustering from probes such as the Lyman-α forest may constrain the scale of this cut-off and thus probe the dark matter particle mass (e.g Viel et al., 2005; Iršič et al., 2017; Palanque-Delabrouille et al., 2020).

– *Neutrinos:* Determining the mass of neutrinos is an unsolved problem in science, which can be attacked from both a particle physics and, more recently, a cosmological perspective (see Lesgourgues et al., 2013, for a review). Cosmologically, neutrinos affect the growth of structure on small scales via the sum of their masses, though current measurements are

precise enough only to impose an upper bound on this quantity. While their exact mass is unknown, current results point to a small but non-zero value (Planck Collaboration et al., 2020b). As such, neutrinos behave relativistically in the early Universe before becoming non-relativistic at later times. Thus their equation of state follows that of photons at early times as $P = \rho/3$, before becoming pressureless $P = 0$ once they become non-relativistic. Neutrinos interact via the weak nuclear force and via gravity, and hence have very low interaction rates with other Standard Model particles. Initially, this made them candidates for dark matter particles, though their early-Universe relativistic behaviour classifies them as "hot" dark matter, a form strongly disfavoured by modern observations.

– *Dark energy:* The most mysterious of the cosmological ingredients, dark energy is an inferred component dictated by the observed late-time acceleration of cosmic expansion. The first evidence for dark energy came in 1990, when measurements of galaxy clustering were found to be poorly described by the then-favoured matter-only cosmological model (Efstathiou, Sutherland, and Maddox, 1990). The most famous evidence for dark energy, however, arrived almost a decade later, with the use of type-Ia supernovae as "standard candles" uncovering a redshift-distance relation consistent with a late-time acceleration in cosmic expansion (Riess et al., 1998; Perlmutter et al., 1999). In order to achieve this accelerating expansion ($\ddot{a} > 0$), eq. (1.17) tells us that $\rho < -3P$, with current evidence pointing towards a "cosmological constant" equation of state $P = -\rho$ (see § 3.3). While the behaviour of dark energy has been (somewhat) constrained, its nature is a complete unknown; famously, quantum field theory predicts a "vacuum energy density" approximately 120 orders of magnitude different from that implied by cosmological observations (Weinberg, 1989).

– *Curvature:* Non-zero cosmic curvature clearly affects geometrical properties of the Universe such as distances (see eq. (1.1) and § 1.2). Equally, eq. (1.16) demonstrates that curvature can also be considered as an extra cosmic ingredient with equation of state $w = -1/3$. Current evidence points towards a flat Universe ($K \simeq 0$) (e.g. eBOSS Collaboration et al., 2020), and is discussed further in § 3.3.

From the continuity and Friedmann equations, it is clear that the equations of state of the various contributors to cosmic energy density are directly tied to the behaviour of the scale factor. Considering cosmological fluids with constant equation of state $w = P/\rho$ in eq. (1.15), we may solve for $\rho(a)$ to yield $\rho \propto a^{-3(1+w)}$. We now define three different fluid types to describe the behaviour of the three qualitatively different equations of state from our list of cosmic ingredients (considering

curvature separately): pressureless matter ($w = 0$), radiation ($w = 1/3$), and a cosmological constant model of dark energy ($w = -1$). Substituting these equations of state into our general solution, we obtain

$$\rho(a) \propto \begin{cases} a^{-4}, & \text{radiation} \\ a^{-3}, & \text{matter} \\ a^{0}, & \text{dark energy } (\Lambda) \end{cases} . \tag{1.18}$$

Of course, alternative models of dark energy are also possible, with different values of w yielding different forms of $\rho(a)$ and thus different cosmological behaviours. For now, though, we focus on a cosmological constant model as preferred by recent constraints (eBOSS Collaboration et al., 2020), and we seek only to highlight important qualitative behaviours.

In reality, our Universe consists of a mixture of different components, and the total energy density has contributions from fluids with all three equations of state. As such, for a full treatment of our cosmological inventory we split the energy density and pressure of eq. (1.16) and eq. (1.17) according to our three components, writing $\rho = \rho_r + \rho_m + \rho_\Lambda$ to describe the contributions to the cosmic energy density from radiation, matter and dark energy respectively. At this point, it is helpful to define the critical density

$$\rho_c = \frac{3H_0^2}{8\pi G}. \tag{1.19}$$

Via eq. (1.16), we see that this represents the total energy density of a flat ($K = 0$) Universe today[3]. We use this to define fractional density parameters today: $\Omega_i \equiv \rho_i/\rho_c$ where i represents a contributor to the total energy density. Thus, using eq. (1.18) and defining the Hubble parameter $H \equiv \dot{a}/a$, we may rewrite eq. (1.16) as

$$H^2 = H_0^2 \left(\frac{\Omega_r}{a^4} + \frac{\Omega_m}{a^3} + \frac{\Omega_K}{a^2} + \Omega_\Lambda \right), \tag{1.20}$$

where we have further defined $\Omega_K \equiv -K/H_0^2$.

Over cosmic history, the value of the scale factor a changes by many orders of magnitude, and so the relative importance of each of the terms on the right-hand side of eq. (1.20) changes too due to their different scalings with a. Indeed, this observation suggests that the Universe has gone through distinct phases in its evolution, with long epochs of domination by individual components. Within these epochs, we may approximate the Universe as containing a single fluid, neglecting all bar one term on the right-hand side to yield a reasonable approximation of cosmic expansion.

[3]Sometimes, ρ_c is written as $\rho_{c,0}$, with the subscript 0 denoting that it represents a value *today*. We follow the convention of omitting the 0 for clarity. This is also the case for the subsequently defined quantities Ω_i

Assuming a flat Universe, consistent with current measurements (e.g. eBOSS Collaboration et al., 2020), the Universe's energy density is dominated by contributions first from radiation, then matter, and finally dark energy. In each of these cases, we may use eq. (1.20) to derive a form of $a(t)$ for that period of cosmic history:

$$a(t) \propto \begin{cases} t^{\frac{1}{2}}, & \text{radiation domination} \\ t^{\frac{2}{3}}, & \text{matter domination} \\ e^{Ht}, & \text{dark energy domination} \end{cases} . \tag{1.21}$$

These different solutions then drive different behaviours through cosmic evolution, subsequently affecting the formation of structure, as described in Chapter 2. Between the eras of single-component domination, however, lie transition regions, where at least two components have energy densities of comparable magnitude. These mark key step-changes in the Universe's history, known as radiation-matter and matter-dark energy equality. At these times, more complex solutions are necessary to determine the form of a and to describe the transition between phases of evolution.

1.4 Inflation

While the cosmological model described thus far is able to achieve a number of notable successes, it also produces a number of notable issues. These are commonly addressed via the introduction of a period of rapid expansion known as *inflation* during the very early Universe. A selection of the problems ameliorated by inflation are summarised below, along with a brief explanation of inflation's influence in each case.

– *The horizon problem:* Different regions of the Universe are said to be "in causal contact" if their past light cones intersect, and thus if there exist points in spacetime which could have influenced both regions. The *particle horizon* χ_{ph} is defined as the greatest comoving distance from which an observer could have received information emitted at some prior time, and thus defines the size of causally connected regions. At the time of the production of the CMB ($z \sim 1100$), these regions are of a size such that one would expect to see $\sim 30,000$ of them in our modern CMB observations. However, we observe the temperature CMB to be exceptionally uniform, to approximately one part in 10^5; how could apparently disconnected regions of the sky "know" to be at the same temperature? Inflation addresses this by effectively increasing the amount of conformal time before the CMB's formation,

dramatically expanding χ_{ph} such that the entirety of the observed CMB is in causal contact.

– *The flatness problem:* We currently observe Ω_K to be very close to zero (eBOSS Collaboration et al., 2020). However, this is just the present-day value of the quantity $-K/\dot{a}^2$. During the radiation- and matter-dominated phases of the Universe's history, $a(t)$ is proportional to $t^{\frac{1}{2}}$ and $t^{\frac{2}{3}}$ respectively, and thus $-K/\dot{a}^2$ grows during these phases. Given its current value, this then implies that it must have been even closer to zero before these phases began, requiring either significant fine-tuning of the Universe's initial parameters or an alternative physical mechanism to drive $-K/\dot{a}^2$ towards zero. Inflation is one such mechanism: exponential expansion results in an exponential decay in the value of $-K/\dot{a}^2$. Thus a wide range of initial values are compatible with the current measurements of Ω_K and their corresponding errors.

– *Relics:* The majority of Grand Unified Theories predict the production of relic particles such as magnetic monopoles, which are not observed in the Universe today. Inflation reduces the expected number density of such objects significantly, making it unsurprising that they are not observed.

While inflation is able to adequately solve a number of cosmological problems, neither its detailed properties nor an inflationary mechanism have been robustly determined. In general, inflation may be induced when the Universe is effectively dominated by a fluid with $w < -1/3$. There exist many potential models which fit this requirement, and in order to ensure that a given model is able to address the issues described above, we may derive a number of constraints on its properties. As an example, we now look a little closer at the popular *slow-roll* model of inflation. In this case, combining eq. (1.16) and eq. (1.15) demonstrates that an FLRW spacetime only corresponds to an inflationary period if

$$\varepsilon \equiv -\frac{\dot{H}}{H^2} = \frac{\mathrm{d}\ln H}{\mathrm{d}N} < 1, \tag{1.22}$$

where $\mathrm{d}N = \mathrm{d}\ln a = H\mathrm{d}t$ measures the number of "e-folds" of expansion, N, during inflation, equivalent to the logarithmic increase in the scale factor during inflation. This requires that the proportional change in the Hubble parameter H per e-fold of inflation must be small, but we also require a certain number of e-folds of expansion — typically around 60 — to sufficiently address the problems associated with a non-inflationary cosmology. Thus eq. (1.22) must hold for a sufficiently long period of time, and so the value of ε must remain small for this period. To satisfy this, we also

require that

$$|\eta| \equiv \left| \frac{\mathrm{d} \ln \varepsilon}{\mathrm{d}N} \right| < 1 \tag{1.23}$$

in order for inflation to last for a sufficiently long time, where η describes the proportional change in ε with each e-fold.

In addition to these requirements, we also need there to be small quantum fluctuations produced during inflation, producing slight deviations from homogeneity and isotropy. These then translate into small perturbations in our matter density field, which grow throughout cosmic evolution to form the structures visible today (discussed in Chapter 2 and Chapter 3). These are typically referred to as the *initial conditions* for structure formation as they seed perturbations that grow in the later Universe. As such, cosmological observations may allow us to measure these initial conditions and thus constrain possible models of inflation. The baseline ΛCDM model considers only perturbations with a uniform ratio between overdensities in radiation and matter. These are known as *adiabatic* perturbations, and are the only type of perturbations that may be generated by the slow-roll, single-field models of inflation mentioned previously. Within ΛCDM, the perturbations are Gaussian distributed and thus may be fully characterised by their power spectrum. This is referred to as the primordial power spectrum, and may be approximated to take the form

$$P_{\mathcal{R}}(k) = A_s \left(\frac{k}{k_0} \right)^{n_s}, \tag{1.24}$$

where $\mathcal{R}$ denotes the comoving curvature perturbation; A_s and n_s are the amplitude and slope of the power spectrum respectively; and k_0 is a pivot scale. The parameters A_s and n_s form part of the baseline ΛCDM model, while simple extensions to this may include a "running" of the spectral index, allowing n_s to vary with k by defining a parameter α via $n_s = n_s(k_0) + \alpha/2 \ln(k/k_0)$ (Kosowsky and Turner, 1995). Alternative models of inflation or theories for the early Universe may yield different initial conditions; we discuss current constraints on ΛCDM and possible extensions to it in § 3.3.

Another fundamental outcome of any successful inflationary model is that it must also be able to produce an end to inflation. This then yields the finite number of e-folds of inflation that we require, and the Universe may transition into the "Hot Big Bang" phase of cosmology. This consists of a period of reheating, during which energy is transferred into the Standard Model fields, forming a "primordial soup" of particles, which goes on to form the Universe we see today.

1.5 THERMAL HISTORY

We follow the story of this "soup" through time, describing its evolution through a number of distinct phases. As we follow this process, a key idea is that of different particle types *decoupling* from — ceasing to interact with — each other in various ways. The point at which this occurs can be determined by comparing the rate of interaction between particles, Γ, with the rate of expansion, H. If $\Gamma \gg H$, then the particles are interacting with each other faster than the Universe expands between them, and as such the particles will sit in local thermal equilibrium with each other. On the other hand, if $\Gamma \ll H$, then the Universe is expanding too fast for interactions to keep up, and so interactions effectively cease. The transition between these two scenarios depends on the forms of Γ and H. From eq. (1.20), we may deduce that $H \sim \sqrt{\rho}/M_{\rm Pl}$ where $M_{\rm Pl}$ is the reduced Planck mass. Dimensionality arguments then yield $\rho \sim T^4$, and thus $H \sim T^2/M_{\rm Pl}$. The form of Γ depends on the interaction in question. Considering the weak interaction (once electroweak symmetry has broken), $\Gamma \sim G_F^2 T^5$ where G_F is Fermi's constant, and so $\Gamma/H \sim G_F^2 T^3 M_{\rm pl}^2 \sim (T/1\ \text{MeV})^3$. In our cosmological model, then, we expect weak interactions to sit in equilibrium at early times when $T \gg 1$ MeV. As the Universe expands and the temperature drops, however, we eventually cross the threshold $T \sim 1$ MeV. Below this point, weak interactions slow to a negligible rate, and weakly-interacting particles decouple.

Immediately after inflation, we start with all particle types coupled tightly in a hot, dense thermal bath — our "soup". During this phase of the Universe's evolution, its energy density is dominated by relativistic species, and so the scale factor evolves as $a(t) \propto t^{\frac{1}{2}}$. The Universe thus expands, and the bath begins to cool. Once the temperature T reaches ~ 1 MeV, neutrinos become the first known particle to decouple, though it is possible that dark matter particles would have done so already. Shortly afterwards, at $T \sim 500$ keV, positrons and the majority of electrons annihilate, heating the surrounding photons slightly. Thus the photons are now at a fractionally higher temperature than the neutrinos, a difference which should in theory propagate to current measurements. While the temperature of the relic photons (the CMB) has been measured with exquisite precision (Fixsen, 2009), direct detection of the cosmic neutrino background has yet to be achieved and so this prediction is yet to be verified.

As expansion continues, the temperature drops further to $T \sim 100$ keV, and the vital process of big bang nucleosynthesis (BBN) occurs. This is the process by which light elements are generated from the soup of protons, neutrons, electrons and positrons. Nuclei are formed by sequential reactions, gradually adding protons and neutrons to first form deuterium, then helium and so on.

The detailed physics of BBN can be modelled by solving coupled Boltzmann equations, yielding predictions of the mass fraction of light elements through the Universe's evolution. These have been found to be in relatively good agreement with observational measurements, providing key justification for the big bang theory as a whole.

At $T \sim 0.75$ eV, we reach matter-radiation equality, when the energy density of matter takes over from that of radiation as the dominant component. Expansion thus accelerates slightly, with the scale factor $a(t) \propto t^{\frac{2}{3}}$. As the Universe enters this phase, two key events occur in quick succession. Following the process of BBN, the "primordial soup" now consists primarily of photons, electrons and nuclei. These are tightly coupled until the temperature becomes low enough for the electrons and nuclei to combine, forming the first atoms in a process known as *recombination*. As this occurs, the number density of free electrons drops dramatically, and thus the rate of interaction between photons and electrons via Compton scattering follows suit. As such, soon after recombination, photons decouple from the remaining particles and stream freely through the Universe. This is known as *decoupling*, and marks the moment of formation of the CMB.

The next few hundred million years are largely uneventful, cosmologically speaking. The only photons are those released at decoupling, while baryonic matter exists mostly as vast clouds of (mostly) neutral hydrogen. These clouds collapse very slowly due to the relatively uniform gravitational fields that they generate, and so the Universe enters the "cosmic dark ages". Gradually, though, matter begins to cluster and structures begin to form.

Eventually, "cosmic dawn" begins to break, with the formation of stars and subsequently galaxies shining new light through the Universe. While it is not clear how to define the first galaxy precisely, the processes that form such objects and the point at which they arrive in the Universe are of great scientific interest (see Barkana and Loeb, 2001; Bromm and Yoshida, 2011, and references therein for a review). Importantly, these first galaxies begin to ionise the hydrogen around them in gradually-expanding bubbles. The hydrogen transitions from being almost entirely neutral to having a neutral fraction of approximately 10^{-5} by $z \sim 3$. Initially, the ionisation state fluctuates in space according to a number of factors including gas density and proximity to nearby ionising sources, while at later times it can be treated as uniform. Meanwhile, structures continue to develop via gravity across a range of scales. This formation of inhomogenous structure, both in radiation- and matter-dominated phases of the Universe, is discussed extensively in Chapter 2.

More recently, a new form of energy has overtaken matter as the dominant contributor to the Universe's energy density. From $z \sim 0.4$ to the present day, dark energy has driven a phase of accelerated expansion of the Universe, first observed through measurements of type-Ia supernovae

during the late 1990s (Riess et al., 1998; Perlmutter et al., 1999). While the subsequent twenty years have yielded extensive research into dark energy, we currently have little understanding of the physics behind this observed behaviour, and it remains a primary motivation for cosmological surveys today.

CHAPTER 2

THE EVOLUTION OF INHOMOGENEITIES

While the homogeneity implied by the Cosmological Principle is observationally evident on the largest scales, it is equally obvious that on smaller scales our Universe is far from homogeneous, and that there exists substantial structure. Such structure exists on a wide range of scales, and encodes valuable information about cosmic evolution. In order to understand these inhomogeneities, we may introduce perturbations to the homogeneous cosmological model treated thus far, then tracking the evolution of these perturbations through time to predict the statistical properties of the Universe that we observe today. On smaller scales, matter can collapse to form dense structures such as stars and galaxies, but such processes are beyond the reach of perturbative methods and require different approaches.

In this chapter, we first introduce perturbations to a homogeneous cosmological model in § 2.1, making some simplifying assumptions in attempting to convey qualitatively the physical processes at play. We then apply these methods, tracking the growth of linear perturbations in § 2.2 and discussing baryon acoustic oscillations. Then, in § 2.3, we descend to smaller scales where non-linear structure forms, addressing the formation of collapsed structures and the development of galaxies.

2.1 PERTURBATIONS TO HOMOGENEITY

Taking the homogeneous Universe of Chapter 1 as a basis, we would like to introduce perturbations to all quantities. We assume that perturbations are adiabatic, as predicted by simple models of inflation and as considered in the baseline ΛCDM model. We express these as deviations from a homogeneous background, writing $X = \bar{X} + \delta X$ for a quantity X (such as an energy density or a pressure) with background $\bar{X}$ and perturbation δX. For perturbations to an energy density ρ_i for a fluid labelled by i, we then typically use the *density contrast* $\delta_i = \delta\rho_i/\rho_i$ in our subsequent mathematical treatment. Introducing perturbations to all quantities this way, we may then substitute into the Einstein field equations (EFEs) of eq. (1.12), separating out background quantities to derive equations governing the evolution of our perturbations.

We first consider perturbations to the left hand side of the EFEs, perturbing our FLRW metric $g_{\mu\nu}$. However, we must take care here, as writing such an equation may implicitly choose a particular set of coordinates. This choice — known as a *gauge choice* — is non-trivial and can change the appearance of perturbations. To address this subtlety, we may either define perturbations so that they are invariant under a change of coordinates, or we may fix our gauge to a particular choice. Here, we fix to the Newtonian gauge, which yields a perturbed line element

$$\mathrm{d}s^2 = a^2(\tau)\left[-(1+2\Psi)\,\mathrm{d}\tau^2 + (1-2\Phi)\,\delta_{ij}\mathrm{d}x^i\mathrm{d}x^j\right], \tag{2.1}$$

where Φ and Ψ are now our perturbation variables. By comparison with the weak-field limit of GR, we see that Ψ plays the role of a gravitational potential. This is a common gauge choice for studying large scale structure, while those investigating perturbations during inflation may commonly choose a spatially-flat gauge, and modern simulations often use a synchronous gauge.

Introducing perturbations to the right hand side of eq. (1.12), we write the perturbed energy-momentum tensor $T^\mu_\nu = \bar{T}^\mu_\nu + \delta T^\mu_\nu$. This introduces corresponding perturbations to the energy density ρ, pressure P and four-velocity U_ν — all of which were defined in eq. (1.14) — while also introducing the possibility of anisotropic stress. Here, we take the anisotropic stress to be zero, which enforces $\Phi = \Psi$. By enforcing conservation of the perturbed energy-momentum tensor, and then substituting our perturbed tensor into the Einstein field equations, we may derive a set of equations governing the evolution and interaction of our metric and matter perturbations.

2.2 Formation of linear structure

Assuming that perturbations are small relative to their corresponding mean fields, we may reduce these equations to linear order in perturbed quantities. The resultant set of equations describe the evolution of linear perturbations as the Universe progresses, while features that are not able to be described by the approximation are known as *non-linear*. This process is governed by the two competing effects of gravity and pressure. While small initial overdensities will gravitationally attract nearby matter and thus grow, this effect is resisted by internal pressure which may prevent the fluid from reaching high densities. Fundamentally, it is the balance between these two mechanisms which determines the nature of perturbation growth, mediated by the background expanding Universe. We address this task in Fourier space, considering perturbative *modes* $\delta_i(\mathbf{k})$ obtained via a Fourier transform of the configuration space perturbations $\delta_i(\mathbf{x})$:

$$\delta_i(\mathbf{k}) = \frac{1}{(2\pi)^3}\int e^{i\mathbf{x}\cdot\mathbf{k}}\delta_i(\mathbf{x})\,\mathrm{d}^3\mathbf{x}, \tag{2.2}$$

where $\mathbf{k}$ is the 3D wavevector, and we write the amplitude of this vector as $k = |\mathbf{k}|$, also referred to as the *wavenumber*. We consider perturbations to each type of fluid separately as their different equations of state result in very different behaviours.

Early in the universe, all perturbation modes are *super-horizon*: they have $k\mathcal{H}^{-1} \ll 1$, where $\mathcal{H}^{-1} \equiv (aH)^{-1}$ is the comoving Hubble radius. As the Universe evolves, the horizon expands and modes gradually cross the horizon, then becoming *sub*-horizon ($k\mathcal{H}^{-1} \gg 1$). Meanwhile, the Universe transitions from its post-inflationary radiation-dominated epoch to the matter domination of late times. The evolutionary behaviour of an individual perturbation mode varies depending on these two qualitative distinctions at any given point in time: whether the mode is super- or sub-horizon, and whether the Universe is in its radiation- or its matter-dominated phase.

Quantitatively, the evolution of a mode through these eras is described by the *transfer function*, which depends on the wavenumber, k. Equally, we will see shortly that some perturbation modes grow in amplitude at late times, in accordance with a wavenumber-independent *growth function*. As such, it is helpful to consider the schematic equation from p. 183 of Dodelson (2003):

$$\Phi(\mathbf{k},a) = \Phi_{\text{primordial}}(\mathbf{k}) \times \{\text{transfer function}\,(k)\} \times \{\text{growth function}\,(a)\}\,. \tag{2.3}$$

Thus by careful calculation of appropriate transfer and growth functions, we may relate inflationary perturbations to those we observe today. Carrying out such a calculation is the main goal of

perturbation theory, encoding the physics of perturbation evolution to provide a mapping between initial conditions and late-time observables. A full treatment of this task involves solving numerous coupled differential equations, typically treated numerically. For the purposes of this thesis, however, a higher-level discussion is appropriate, and we simply provide a summary of noteworthy features by making simplifying approximations to study the evolution of perturbations to various quantities.

2.2.1 Gravitational collapse

Evolution of the gravitational field

The gravitational field Φ provides a suitable quantity to demonstrate the key physical ideas when considering the evolution of perturbations. Assuming adiabatic perturbations, the perturbed Einstein equations yield a single equation describing the evolution of the gravitational potential:

$$\Phi'' + 3(1+w)\mathcal{H}\Phi' + wk^2\Phi = 0, \tag{2.4}$$

where w is the equation of state of the background, and $'$ denotes a derivative with respect to conformal time τ. On super-horizon scales, the final term on the left hand side is much smaller than the second term, and so can be ignored. This yields one solution that decays rapidly with conformal time, and a second, "dominant" solution $\Phi = \text{const}$ upon which we focus. This solution is independent of w and thus applicable to both radiation- and matter-dominated epochs. Once modes enter the horizon, we can no longer ignore the term $\propto k^2$, and the value of w becomes important. In the radiation-dominated era, we have $w \approx 1/3$, and so we obtain

$$\Phi'' + \frac{4}{\tau}\Phi' + \frac{k^2}{3}\Phi = 0. \tag{2.5}$$

This is solved by spherical Bessel and Neumann functions, which, on sub-horizon scales, yields a dominant solution $\Phi \propto \cos(k\tau/\sqrt{3})/(k\tau)^2$, corresponding to decaying oscillations. Equally, in the matter-dominated era we have $w \approx 0$, and so we obtain

$$\Phi'' + \frac{6}{\tau}\Phi' = 0. \tag{2.6}$$

This equation is simpler, yielding a dominant solution $\Phi = \text{const}$. Thus in the matter-dominated era, both sub- and super-horizon modes are constant.

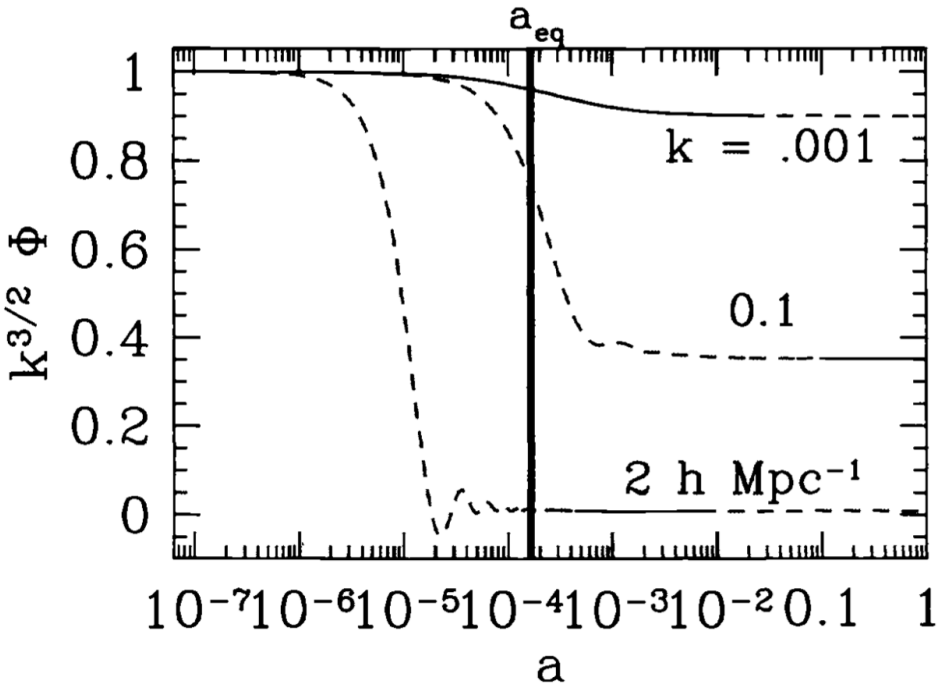

Figure 2.1. The evolution of perturbations to the gravitational field Φ at three different length scales. These scales are chosen such that the modes enter the horizon at three qualitatively different times in relation to the moment of matter-radiation equality (given by $a_{\rm eq}$): long before $a_{\rm eq}$ ($k = 2\ h\ {\rm Mpc}^{-1}$), around $a_{\rm eq}$ ($k = 0.1\ h\ {\rm Mpc}^{-1}$), and long after $a_{\rm eq}$ ($k = 0.001\ h\ {\rm Mpc}^{-1}$). As a result of these choices, the modes exhibit distinct behaviours, resulting in wide-ranging values at late-time. *Figure credit: Dodelson (2003).*

The behaviour of these perturbation modes is illustrated in Figure 2.1, originally published in Dodelson (2003). Here, the $k = 2\ h\ {\rm Mpc}^{-1}$ mode enters the horizon well before matter-radiation equality, exhibiting strongly damped oscillations and decaying (almost) to zero before equality. As a result, this mode has very low amplitude through matter domination as well. The $k = 0.1\ h\ {\rm Mpc}^{-1}$ mode, however, enters the horizon later, and thus does not decay away completely before matter domination. Finally, the $k = 0.001\ h\ {\rm Mpc}^{-1}$ mode does not enter the horizon until after matter-radiation equality, and so does not exhibit any decaying oscillatory behaviour. It does, however, reduce slightly in amplitude — by a factor of $\frac{9}{10}$ — over the epoch of equality. Deriving this effect requires a more mathematically detailed treatment than we present here, but qualitatively Figure 2.1 provides a suitable illustration of the very different evolutionary paths followed by modes of different scales.

The Jeans scale

We may now consider a perturbation to a general fluid labelled by i, with equation of state $w = \bar{P}_i/\bar{\rho}_i$. We also define the sound speed $c_s^2 = \delta P_i/\delta\rho_i$. By combining equations describing conservation of mass and momentum (continuity and Euler equations) on sub-horizon scales, we obtain a configuration-space evolution equation

$$\delta_i'' + (1-3w)\mathcal{H}\delta_i' - c_s^2\nabla^2\delta_i = (1+w)\nabla^2\Phi. \tag{2.7}$$

This is the general governing equation for fluid fluctuations, including Hubble friction and pressure in the second and third terms respectively on the left-hand side, and gravity in the right-hand side. Then applying a Fourier transform, and using the Poisson equation to remove the dependence on Φ yields

$$\delta_i'' + (1-3w)\mathcal{H}\delta_i' + c_s^2(k^2 - k_J^2)\delta_i = 0, \tag{2.8}$$

where $k_J \equiv \sqrt{4\pi G a^2\bar{\rho}(1+w)/c_s^2}$ is the *Jeans scale* (Jeans, 1902).

The solutions to this equation depend on the background cosmic behaviour via $\mathcal{H}$, a and $\bar{\rho}$, as well as the properties of the fluid via w and c_s. On small scales with $k \gg k_J$, eq. (2.8) becomes a damped oscillator equation, with pressure driving oscillations that are damped by the Hubble friction. Equally, when $k \ll k_J$, it admits growing solutions, the precise nature of which is determined by the equation of state w. This behaviour of δ prevents structure from growing on small scales, while allowing larger-scale perturbation modes to grow. Crucially, the value of the transition scale k_J depends on the behaviour of the fluid in question. For radiation, the sound speed c_s is very large, and so perturbations oscillate on all but the largest scales. For dark matter, however, c_s is very small, and so we would expect most perturbations to grow.

Clustering of dark matter

For pressureless dark matter, $w = 0$ and c_s is very small, so we may neglect the $c_s^2\nabla^2\delta$ term in eq. (2.7), yielding a Fourier-space evolution equation for sub-horizon fluctuations sourced by the gravitational potential

$$\delta_c'' + \mathcal{H}\delta_c' = -k^2\Phi. \tag{2.9}$$

where δ_c is the density contrast for cold dark matter. In the early Universe, Φ is dominated by contributions from radiation and dark matter (we ignore baryonic effects for now). However, the radiation density fluctuates very rapidly (see § 2.2.2), and thus we may neglect its contribution

in a "time-averaging" approximation (Weinberg, 2002). After some manipulation and use of the substitution $y = a/a_{\rm eq}$, we may then deduce that δ_c evolves according to the analytically tractable Mészáros equation:

$$\frac{\mathrm{d}^2\delta_c}{\mathrm{d}y^2} + \frac{2+3y}{2y(1+y)}\frac{\mathrm{d}\delta_c}{\mathrm{d}y} - \frac{3}{2y(1+y)}\delta_c = 0, \tag{2.10}$$

During radiation domination, $y \ll 1$ and eq. (2.10) has a dominant solution $\delta_c \propto \ln a$, indicating slow growth of perturbations. Physically, we may interpret this as gravitational collapse being partially thwarted by rapid oscillations in the radiation density. During matter domination, $y \gg 1$ and the dominant solution to the Mészáros equation is approximately $\delta_c \propto a \propto t^{\frac{2}{3}}$. Here, gravitational collapse causes perturbations to grow more rapidly, without being held back by radiation.

In the later Universe, the balance of energy densities changes, with dark energy taking over from matter as the dominant components. Fluctuations in pressure then become negligible, and we obtain solutions $\delta_c = \text{const}$ on all scales. This represents a suppression of growth as the Universe enters the age of dark energy-domination.

2.2.2 The effects of pressure

Evolution of radiation perturbations

We may now analogously consider how perturbations to the radiation energy density evolve, using the radiation density contrast δ_r to do so. In this case, $w = 1/3$, indicating non-negligible pressure in the fluid. Thus, for adiabatic radiation perturbations, the evolution equation becomes

$$\delta_r'' + \frac{1}{3}k^2\delta_r = -k^2\Phi. \tag{2.11}$$

During radiation domination, the potential Φ decays as $a^2 \propto \tau^2$, and so we may find approximate solutions by neglecting the right-hand side of eq. (2.11). Doing so yields a simple oscillator equation, with $\delta_r \propto (k\tau)^2\Phi \propto \cos(k\tau/\sqrt{3})$. During the matter-dominated era, $\Phi \propto \text{const}$, and so we may no longer neglect the driving term on the right-hand side. This change simply shifts the equilibrium point of our oscillations from zero to $-4\Phi(k)$. Thus in both the radiation- and matter-dominated epochs, fluctuations in the radiation density oscillate due to internal pressure.

Photon-baryon oscillations

Before decoupling, however, baryonic matter and photons are tightly coupled via Compton scattering, and can effectively be treated as a single fluid. As a result, only the combined momentum

density of this photon-baryon fluid is now conserved, and so the conservation equations which generated eq. (2.7) require modification. Writing $R = 3\bar{\rho}_b/4\bar{\rho}_\gamma$ to relate the background densities of photons and baryons, we may derive a modified version of eq. (2.11) to describe the photon and baryon density contrasts, written δ_γ and δ_b respectively (Peebles and Yu, 1970; Hu and White, 1996; Eisenstein, Seo, and White, 2007):

$$\delta_b'' + \frac{\mathcal{H}R}{1+R}\delta_b' - c_s^2\nabla^2\delta_b = \nabla^2\Phi = \frac{3}{4}\left(\delta_\gamma'' + \frac{\mathcal{H}R}{1+R}\delta_\gamma' - c_s^2\nabla^2\delta_\gamma\right). \tag{2.12}$$

Here, we have defined the sound speed c_s in the photon-baryon fluid as $c_s^2 = c^2/3(1+R)$. This equation describes driven oscillations in the fluid, with the photons providing pressure support to sustain oscillations in both baryon and photon perturbations. When decoupling occurs, however, this pressure disappears and so the oscillations freeze, leaving an imprint on the distribution of both baryons and photons known as *baryon acoustic oscillations* (BAO). Of course, perturbation modes are still able to evolve after decoupling, with baryon perturbations growing as $\delta_b \propto a$ (as was the case for dark matter).

The BAO signal imprinted upon the baryon distribution can be demonstrated more intuitively by considering the evolution of an initial, adiabatic, localised overdensity, as in the top left panel of Figure 2.2. An overdensity in the radiation sector implies an over-pressured region, and this excess pressure drives a sound wave radially outwards. As the wave travels through the tightly coupled photons and baryons, it leaves the dark matter overdensity behind, while neutrinos stream away with relative freedom (top right panel). This continues as the Universe evolves until decoupling begins, as in the middle left panel. As the photons and baryons decouple, the baryon wave decelerates rapidly while the photons stream away, yielding an overdensity of baryons at a fixed radial size in the middle right panel, with the radius determined by

$$r_d = \int_{z_d}^{\infty} \frac{c_s(z)}{H(z)}\mathrm{d}z, \tag{2.13}$$

where c_s is the speed of sound, and $z_d \simeq 1020$ is the redshift of the "drag epoch"[1]. This quantity depends on the function $c_s(z)$, and the value of z_d. The sound speed c_s is determined by balance between gravity and pressure, which depends on the photon to baryon ratio via R. Meanwhile, the drag redshift is determined by the rate of expansion in the earlier Universe, which depends on

[1] This is the redshift at which baryons decouple from photons, and differs slightly from redshift of last scattering at which the CMB forms. This difference is not important for the qualitative understanding of the BAO scale we are interested in here.

the redshift of matter-radiation equality, as defined by the energy densities of the early Universe. Current measurements of cosmological parameters yield a value $r_d = 147.57 \pm 0.22$ Mpc (Planck Collaboration et al., 2020b). At this stage, then, we are left with a shell of baryon overdensity at radius r_d, surrounding a central overdensity of dark matter. As the Universe continues to evolve through the bottom two panels of Figure 2.2, these two overdensities both grow under gravitational instability, leaving us with a central overdensity of both baryons and dark matter, surrounded by a spherical "echo", again of both baryons and dark matter.

In reality, we must consider a superposition of many such initial perturbations, and many such sound waves, all overlapping as ripples from many stones thrown into a pond. While this superposition leaves us unable to detect individual spherical shells, we may expect to observe a "preferred scale" in the late Universe: given an overdensity at one point in space, we would expect an increased likelihood of observing another overdensity at a certain "preferred" separation, defined by BAO and cosmic expansion. This ought to express itself in all quantities that trace the matter density, including galaxy counts and the IGM; we explore such methods of measuring BAO in Chapter 3. Observing this fixed-length scale enables us to use BAO as a "standard ruler"; by measuring the apparent size of the scale at a given redshift, we may use our value of its absolute size to derive the distance to that redshift. Then, by making measurements at many redshifts, we may directly probe the distance-redshift relation through the Universe's history, thus constraining cosmic expansion.

2.3 Beyond linearity

While the growth of linear structures forms a key part of cosmological theory, the differential equations yielded from the Einstein equation are inherently non-linear. This introduces coupling between different modes, and can yield complex features from even the simplest of initial conditions. Indeed, observationally, we make use of highly non-linear structures such as galaxies to look into the distant Universe, and so some understanding of their formation is important in correctly interpreting our measurements.

During the early Universe, pursuing a perturbative approach to linear order proves an accurate approximation, but at later times, non-linear features become apparent. The extent of these non-linearities depends on the scale involved. On large scales — often referred to as *quasi-linear* scales — we may continue pursuing perturbative methods, expanding to higher order terms to yield increasingly accurate solutions. In recent years, a number of more advanced techniques and

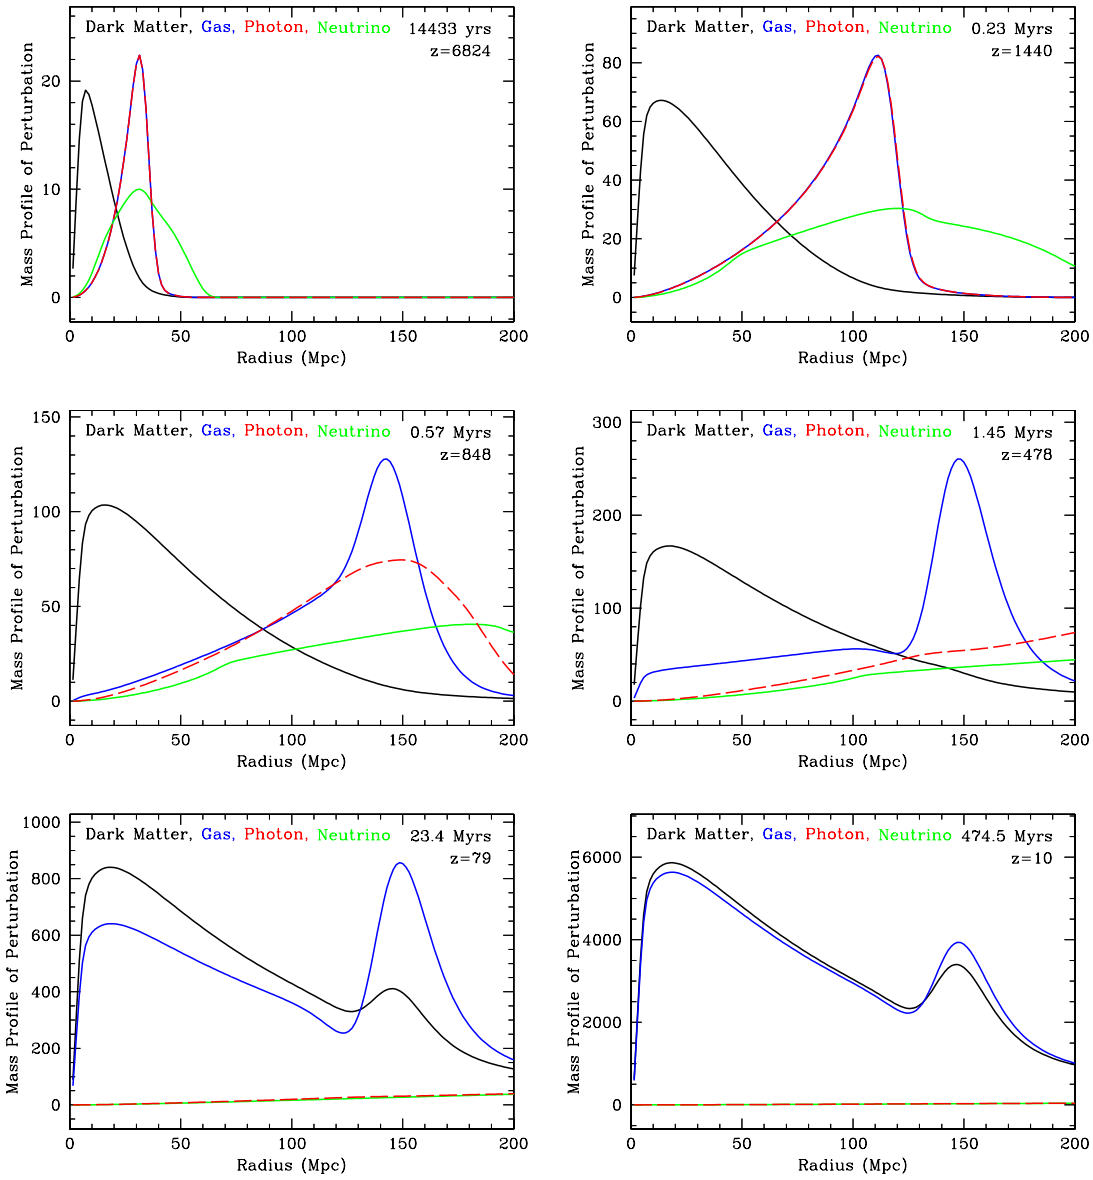

Figure 2.2. The formation of the late-time BAO feature, illustrated by tracking the 1D evolution of an initially localised, adiabatic perturbation. The process starts in the top left panel, before moving to the right and then following the same progression in the middle and bottom rows. The redshift of each snapshot is indicated in each panel. *Figure credit: Eisenstein, Seo, and White (2007).*

mathematical formalisms in this vein have been introduced, proving highly effective in conjunction with improved computational methods to quantify them. For a review of such approaches, please see Bernardeau et al. (2002).

Alternatively, analytic approaches built on the back of linear theory can provide surprisingly accurate predictions of non-linear features of our Universe. Typically, objects such as galaxies and clusters — crucial to modern cosmological observations — reside in collapsed, self-gravitating "clumps" of dark matter, known as *halos*. A number of theories predict the distribution of these halos by determining the location of *proto-halos*: areas of slight overdensity in initial fluctuations which, the theories propose, will eventually collapse to form halos. For example, the Press-Schecter formalism (Press and Schechter, 1974) defines a threshold overdensity, and proposes that halos will form at all regions where initial fluctuations meet that threshold. The breadth of this initial overdensity then determines the mass of the halo, enabling predictions to be made of the halo number density distribution (Press and Schechter, 1974) and their clustering (Mo and White, 1996). Advances have been made to extend this method (e.g. Sheth, Mo, and Tormen, 2001; Sheth and Tormen, 2002), and alternative ideas such as the peak approach have also yielded success (see Chapter 6 of Desjacques, Jeong, and Schmidt, 2018).

Ultimately, though, with the vastly increased computational power at the fingertips of modern cosmology, numerical simulations now provide the most precise predictions of structure formation on non-linear scales. Large-scale simulations may model three main processes: gravity, hydrodynamics, and astrophysical feedback. Gravity is treated by considering collisionless particles of cold dark matter, as is consistent with standard theories. A large number of these particles are placed in a simulation, each with a certain mass, and they are then allowed to evolve under gravity. This "N-body" approach (see Dehnen and Read, 2011, for a review) typically uses Newtonian approximations for gravity, though relativistic approaches do also exist. On top of dark matter particles, hydrodynamic treatments allow the simulation of baryonic gas and its inherent temperature and pressure. This behaviour is governed by the Navier-Stokes equations, substantially increasing the complexity of the simulations due to the interplay between gas dynamics and gravity. A number of methods have been developed to provide tractable yet well-motivated approximations to solutions, generally following either Lagrangian (e.g. Springel, 2010) or Eulerian (e.g. Teyssier, 2015) approaches. Finally, the addition of astrophysical feedback to simulations provides a quantification of an enormous range of astrophysical processes and their impact on the surrounding gas. This causes great difficulty due to the uncertainty in how to model each of the many forms of feedback, and, from a computational perspective, the enormous range of scales that these processes cover. Current

research in this area is extensive, with an extensive and growing literature, recently reviewed in Naab and Ostriker (2017).

CHAPTER 3

OBSERVABLES & MEASUREMENTS

In science as a whole, progress relies on the interaction between theory and measurement; either on its own may be interesting, but it is only the combination of the two which allows us to advance our understanding. The material of previous chapters provides an introduction to some of the theoretical concepts which underpin our understanding of the Universe today, and so we now turn to the other side of the coin, and consider the observations that we may make in order to test this theory. We introduce a few key measurements, and also summarise the immense observational progress that has been made in recent decades.

In § 3.1, we introduce the cosmic microwave background, the foremost early-Universe cosmological probe today. Then, in § 3.2, we focus on later stages of the Universe's evolution, looking at a selection of large-scale structure measurements at redshifts $z \sim 0.1 - 4.0$. In § 3.3, we assess the capabilities of our current observations to constrain our cosmological model, focusing on the "baseline" ΛCDM model and extensions thereof.

3.1 COSMIC MICROWAVE BACKGROUND

In § 1.5, we briefly introduced the cosmic microwave background (CMB) as the remnant photons which decoupled from the matter field at $z \sim 1100$, around 380,000 years after the Big Bang. This process of decoupling describes the formation of the "surface of last scattering": a 2D surface of fixed radius determined by the distance travelled by light since decoupling. We now observe photons coming from this surface arriving at Earth from all directions, and may measure the properties of

these photons in each direction. The CMB was first observed in 1964, as an initially unexplained residual signal at the Holmdell Horn Antenna in New Jersey (Dicke et al., 1965; Penzias and Wilson, 1965). This followed decades of speculation around its existence and predictions of its temperature. This temperature — reduced since decoupling by the expansion of the Universe — has now been measured with exquisite precision as $T = 2.726 \pm 0.001$ (Fixsen, 2009). On initial inspection, the uniformity of the CMB across the sky is striking, with fluctuations of order one part in 10^5. However, it is within these small deviations from uniformity that a great wealth of information about our Universe lies. These fluctuations were first observed by the Cosmological Background Explorer (COBE, Smoot et al., 1992; Bennett et al., 1996), with subsequent improvements in resolution made by a number of experiments, including the Wilkinson Microwave Anisotropy Probe (WMAP, Bennett et al., 2003) and *Planck* (Planck Collaboration et al., 2014). Anisotropies become apparent via several forms of measurement: temperature, polarisation and gravitational lensing, each of which is displayed in Figure 3.1.

Captured within the CMB are the imprint of oscillations in the photon-baryon fluid before decoupling — baryon acoustic oscillations (BAO) — as introduced in § 2.2.2. These oscillations cause peaks and troughs in the observed power spectrum of CMB anisotropies; a perturbation with a given wavenumber will enter the horizon and thus begin oscillating at a particular time before decoupling. Depending on whether the perturbation is reaching a maximum or a minimum in its oscillatory amplitude at the time of decoupling, we expect to see a peak or a trough respectively in the power spectrum at the scale corresponding to that particular wavenumber. These features are also affected by diffusion damping around the time of recombination and decoupling (Silk, 1968), which occurs both due to the non-zero diffusion length of photons while still coupled to baryons, as well as the finite duration of the recombination process. Ultimately, this reduces the amplitude of the peaks on small scales, leaving those at larger scales relatively unaffected (Hu and White, 1997).

Additional anisotropies may also be added to the CMB during the time since decoupling via a number of different sources. The shape of the gravitational potential at decoupling may cause photons to be red- or blue-shifted at the formation of the CMB, known as the Sachs-Wolfe effect (Sachs and Wolfe, 1967). A similar occurence — the *integrated* Sachs-Wolfe effect (Rees and Sciama, 1968) — is caused by time-varying gravitational potentials in the late Universe. Alternatively, high-energy electrons either in galaxy clusters — known as the Sunyaev-Zel'dovich effect (Sunyaev and Zeldovich, 1970; Sunyaev and Zeldovich, 1980; Ostriker and Vishniac, 1986) — or in the IGM may interact with the CMB photons via inverse Compton scattering. These two electron sources introduce a dependence of the observed CMB on astrophysical processes through

cosmic history by affecting the photons' polarisation. Finally, the integrated distribution of matter between the surface of last scattering and observation introduces lensing effects in the CMB seen today (Zaldarriaga and Seljak, 1998; Lewis and Challinor, 2006).

These sources of anisotropy add many layers to the information contained within the CMB. On the largest scales — for which perturbations remained super-horizon until after decoupling — we may use the CMB to measure the initial perturbations seeded by inflation (or an alternative theory), constraining the amplitude A_s and tilt n_s of the primordial power spectrum. At intermediate sales, we observe the peaks and troughs of BAO, the nature of which depend on combinations of energy density parameters (Hu, Sugiyama, and Silk, 1997). Different combinations determine both the acoustic properties of the primordial photon-baryon fluid, and the nature of cosmic expansion since recombination. These then define the projected size of the sound horizon at recombination, which sets the observed locations of the peaks in CMB power spectra. The locations also depend on the nature of the primordial seed perturbations, which may shift the peaks though does not affect their spacing. Beyond peak locations, the relative amplitudes of the BAO peaks are affected by the balance between baryons and radiation in the photon-baryon fluid, with baryon drag acting to enhance compression (even) peaks relative to rarefaction (odd) ones. In conjunction, measurements of the locations and amplitudes of the BAO peaks may be used as a "standardisable" ruler, breaking internal parameter degeneracies and allowing tight constraints to be made on cosmic energy densities. At smaller scales, the energy densities of baryons and matter affect the damping scale of the CMB power spectrum, which may act as a consistency check for the effects of these parameters on larger scales. For a thorough review of the cosmological dependence of CMB anisotropies, please see Hu and Dodelson (2002), or for an up-to-date summary of current constraints see Planck Collaboration et al. (2020a) and Planck Collaboration et al. (2020b).

3.2 Large-scale structure

A great deal of information is encoded within the late-time distribution of both dark and baryonic matter. A variety of observational techniques are employed to try and capture this information, from which we may apply a range of analytical techniques to extract cosmological constraints. Here, we describe three qualitatively different methods to measure large-scale structure, highlighting the uses of each in a cosmological context.

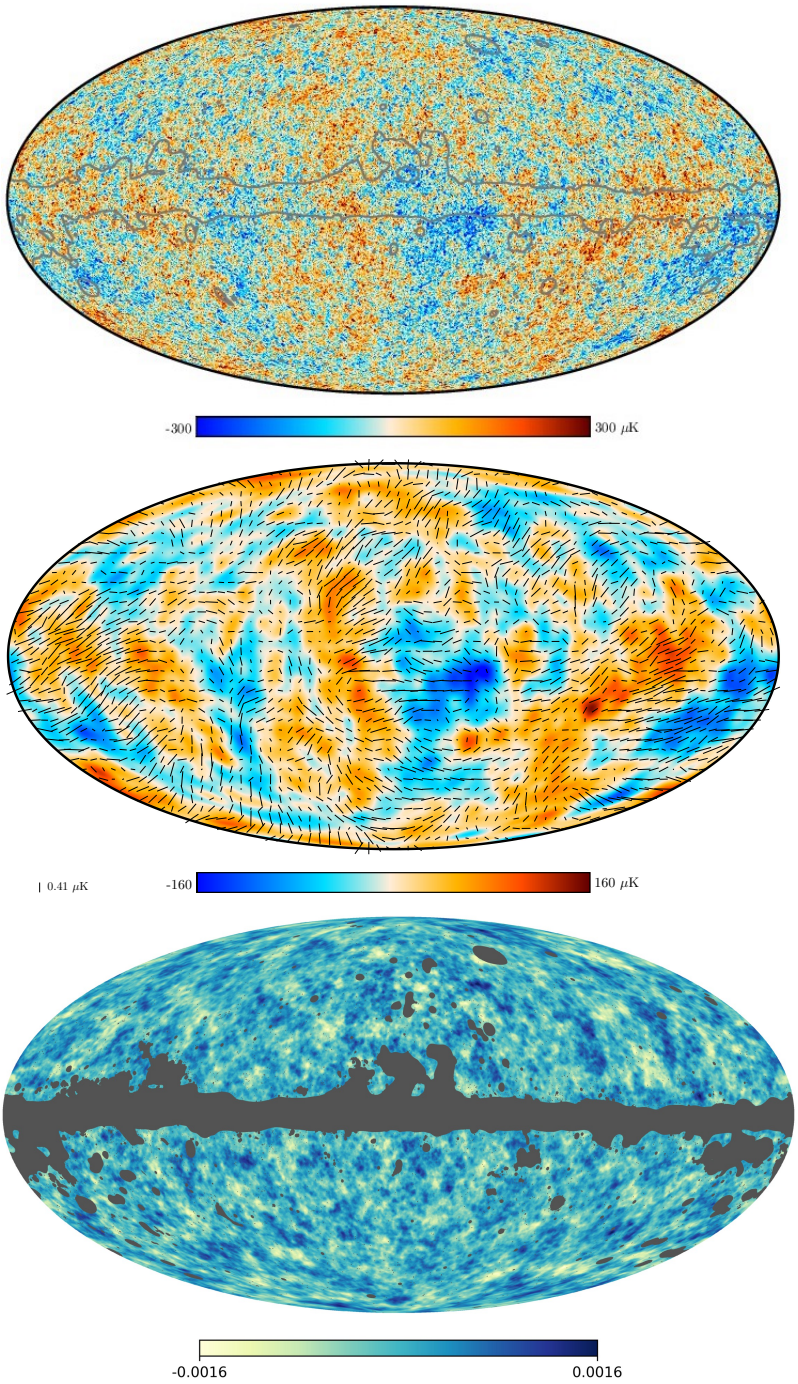

Figure 3.1. Sky maps from *Planck* in three different quantities. The top image shows the temperature map, with a masked region around the galactic plane shown by the grey line. The middle image represents the polarisation field with small black rods, superimposed on the temperature map. In this case, both polarisation and temperature are smoothed at a 5° scale for visibility. Finally, the bottom panel shows the *Planck* E-mode lensing map, which derives information from both temperature and polarisation. *Figure credit: Planck Collaboration et al. (2020a).*

3.2.1 BAO FROM SPECTROSCOPIC SURVEYS

The coupling of baryons to radiation in the early Universe dictates that the BAO signature is present in perturbations to the baryon energy density at recombination, as well as the radiation density. This signal is carried in the baryon distribution as perturbations evolve through time, and is imprinted upon all subsequent structures. It can then be measured from any matter distribution that we observe in the late Universe. We may calibrate a BAO measurement with a value of r_d, allowing us to measure the absolute distance to a given redshift. The calibration value of r_d is typically calculated using parameters derived from CMB anisotropies, though it is also possible to do so using only the temperature of the CMB along with a value of Ω_b determined from measurements of BBN abundances (Addison, Hinshaw, and Halpern, 2013; Addison et al., 2018). These two different approaches to calibration can be considered independent of each other and thus provide useful checks of consistency between data sources. Alternatively, we may treat r_d as a nuisance parameter, instead using BAO measurements at a number of redshifts to fit the relative cosmic expansion history.

We may attempt to measure the BAO signal via spectroscopic surveys, which measure the spectral distribution of light from numerous objects within a survey volume. Observing objects in this way — as opposed to doing so photometrically — allows spectral features such as emission lines to be distinguished. Distinctive lines such as the O II doublet may then be identified by their shape, while a set of broad emission lines may be used to identify a quasar. For an emission line of known rest-frame wavelength λ_{rest}, we may straightforwardly deduce the object's redshift $z = \lambda_{\text{obs}}/\lambda_{\text{rest}} - 1$. Due to the fine sampling of flux as a function of wavelength inherent to spectroscopic observation, this redshift is considerably more precise than can be obtained from photometric data (see Salvato, Ilbert, and Hoyle, 2019, for a review of methods to do so). Of course, spectroscopic observations are also more time consuming per object than their photometric counterparts, but the ability to determine precise redshifts allows us to confidently measure perturbation modes in three dimensions, which can offer increased statistical power as well as increased scientific opportunity.

Unlike in the CMB, spectroscopic surveys may make use of this third dimension to measure the BAO signal in both radial and transverse directions. In the radial direction, we may observe the BAO scale to correspond to a particular redshift separation Δz, which can be used to calculate a comoving separation using eq. (1.9). For a narrow redshift interval, this integral can be approximated to yield $\chi \simeq c\Delta z/H(z)$ (reintroducing the factor of c set to 1 previously). Thus for $\chi = r_d$, we may directly measure the Hubble parameter $H(z) = c\Delta z/r_d$ via BAO. Conversely, in the transverse direction

the BAO scale may correspond to a particular angular separation $\Delta\theta$, allowing us to measure the comoving transverse distance $d_M = r_d/\Delta\theta$, as defined in § 1.2. Typically, BAO results are presented as measurements of d_H/r_d and d_M/r_d, where $d_H = c/H(z)$ is the Hubble distance. The quantity $d_v = (z d_H d_A^2)^{1/3}$ has been used historically, and is sometimes still used when dealing with low signal-to-noise ratio measurements.

While the BAO signal certainly ought to be present in the late-time matter distribution, the gravitational coupling between baryons and dark matter — the latter of which do not carry the BAO signature immediately after recombination — dictates that the amplitude of this signal is substantially smaller than in the CMB. Thus a large survey volume is required to suitably recover the BAO signal. Provided such a volume can be achieved, a number of different structure types can be used to trace the matter density field.

Discrete tracers

Discrete objects are perhaps the easiest to observe tracer of the matter density at $z > 0.1$. As such, they provided the first large-scale structure detections of the BAO signal, using number counts of galaxies to measure the correlation function and power spectrum of galaxies from the Sloan Digital Sky Survey (SDSS, York et al., 2000) and 2-Degree Field Galaxy Redshift Survey (2dFGRS, Colless et al., 2001) respectively. In the subsequent years, a number of similar measurements have been made, primarily tracing the matter density field with galaxies (e.g. Percival et al., 2010; Beutler et al., 2011; Blake et al., 2011; Alam et al., 2017) and quasars (QSOs, e.g. Ata et al., 2018) in redshift ranges $z \sim 0.1 - 1.0$ and $z \sim 1.2 - 1.7$ respectively. Recently, the extended Baryon Oscillation Spectroscopic Survey of SDSS-IV (Dawson et al., 2016; Blanton et al., 2017) provided the most comprehensive set of BAO measurements to date. It measured the BAO signal in both Fourier- and configuration-space from two different classes of galaxies, emission line galaxies (ELGs, de Mattia et al., 2020; Raichoor et al., 2021) and luminous red galaxies (LRGs, Gil-Marín et al., 2020; Bautista et al., 2021) as well as QSOs (Neveux et al., 2020; Hou et al., 2021).

While these discrete objects make suitable tracers of the matter density field, their use in detecting the BAO signal is not without difficulty. The motion of galaxies during the late Universe due to bulk flows causes a degradation of the BAO signal, observed as a broadening of the peak in configuration space or a damping in peak amplitude in Fourier space (Seo and Eisenstein, 2005; Seo et al., 2008). This broadening does not introduce any systematic bias to the BAO scale in and of itself, though non-linear effects can introduce systematic shifts in the relative positions of galaxies. Fortunately, both simulations and theory suggest that this effect is small, shifting the peak

by $\lesssim 0.5\%$ (Crocce and Scoccimarro, 2008; Padmanabhan and White, 2009; Seo et al., 2010). Of course, we must also consider that discrete objects may trace the underlying matter density field in non-trivial ways, and using a simple linear bias $\delta_g = b\delta_m$ will likely not sufficiently capture this complexity. Introducing non-linear effects to the biasing may introduce a shift to the BAO peak (Padmanabhan and White, 2009), but such effects are small (Mehta et al., 2011) and can be accounted for by calculations of small scale clustering and use of higher order statistics.

These difficulties may be partially addressed by a process known as *reconstruction* (Eisenstein et al., 2007). This uses the observed set of galaxies to determine the gravitational potential field, from which the galaxies' linear peculiar velocities may be deduced as the gradient of the potential. We may then employ the Zel'dovich approximation (Zel'Dovich, 1970) — that objects move along straight-line paths in comoving coordinates at a rate that can be predicted by linear theory — in reverse, moving galaxies back to the positions they would have occupied if it were not for bulk flows. This allows us in some sense to "reconstruct" the BAO signal which was present in the earlier Universe, before broadening, and thus increase the significance of any detection. Having been refined over many years, reconstruction now allows for nearly all of the theoretically-available BAO information to be recovered (Seo and Eisenstein, 2007; Seo et al., 2010; Padmanabhan et al., 2012). Typically, this corresponds to an increase in precision by a factor of $\sim 1.5-2$, equivalent to expanding the volume of a survey by a factor of $\sim 2-4$. Reconstruction also helps to address the systematic shifts described previously, thus improving both the statistical and systematic errors on discrete-tracer BAO measurements (Padmanabhan et al., 2012; Anderson et al., 2012).

Beyond the challenges of non-linear effects in discrete tracers, a practical limitation is also present for realistic surveys. In order to push such measurements further back in the Universe's history by extending to higher redshifts, the volume of a survey increases dramatically but objects become fainter due to their increased distance. As such, given a certain limiting magnitude, the number density of visible objects decreases, and there are insufficient visible objects to produce a statistically significant BAO detection using discrete tracers. With current, ground-based surveys, this confines such measurements to $z \lesssim 1$ for galaxies, and $z \lesssim 1.8$ for quasars. In the near future, the Dark Energy Spectroscopic Instrument (DESI, DESI Collaboration et al., 2016a; DESI Collaboration et al., 2016b) will push these boundaries slightly higher with a large sample of emission line galaxies at $z \lesssim 1.6$ (see Table 4.1). Further, data from space-based missions such as *Euclid* (Laureijs et al., 2011) and WFIRST (Spergel et al., 2015) will enable the use of slitless spectroscopy (Glazebrook et al., 2005), yielding galaxy samples large enough to measure BAO to $\lesssim 1\%$ precision in the redshift range $1 \lesssim z \lesssim 2.2$.

The Lyman-α forest

Within the last decade, however, a qualitatively different tracer has emerged in the form of the Lyman-α (Lyα) forest, allowing the BAO feature to be detected beyond $z = 2$. The Lyα forest is a sequence of absorption features seen in the spectra of high-z QSOs, first noted in the 1960s after observations of the first such QSOs (Schmidt, 1965; Scheuer, 1965; Bahcall and Salpeter, 1965; Gunn and Peterson, 1965).

To understand the formation of the Lyα forest, we consider a photon emitted towards us from a QSO at redshift $z_{\rm QSO}$, with initial wavelength $\lambda_{\rm rest} \leq 1215.67$ Å $= \lambda_\alpha$, the wavelength of the Lyα transition of neutral hydrogen (H I). As the photon travels towards us, it is constantly redshifted and so its wavelength increases. At some point, its wavelength reaches λ_α and it may be absorbed by an H I atom in the inter-galactic medium (IGM). This absorption occurs with a certain probability, which depends on the density and kinematics of neutral hydrogen in its surrounding region at that time. This probability then determines the fraction of photons emitted at $\lambda_{\rm rest}$ from the source QSO which reach us at a wavelength $\lambda_{\rm obs} = \lambda_\alpha(1+z_{\rm HI}) = \lambda_{\rm rest}(1+z_{\rm QSO})$, where $z_{\rm HI}$ is the redshift of the neutral hydrogen where absorption occurs. Considering now the spectrum of photon wavelengths emitted by a QSO, we see that this absorption process may result in a whole sequence of absorption features at a range of different wavelengths $\lambda_{\rm obs} < \lambda_\alpha(1+z_{\rm QSO})$. These features are collectively known as the Lyα forest, and they provide a line-of-sight tracer of the density of neutral hydrogen between observer and the host QSO.

Cosmologically speaking, we are interested in using the Lyα forest to trace the distribution of matter, but we observe absorption features in QSO spectra. Translating between these two quantities is non-trivial, as the processes involved are complex. At the most basic level, we may write the optical depth as an integral over the path between source and observer:

$$\tau = \int_{\rm source}^{\rm observer} {\rm d}s \; n_{\rm HI}(z) \; \sigma_\alpha[\nu(z)], \tag{3.1}$$

where ds is a proper length interval, $n_{\rm HI}$ is the number density of neutral hydrogen atoms, σ_α is the Lyα cross section, ν is the photon frequency and $z = z(s)$ is the redshift of a photon as it moves along its path. Considering an expanding Universe with Hubble parameter $H(z)$, we may use this to derive the Gunn-Peterson relation (Gunn and Peterson, 1965):

$$\tau_{\rm GP} = \frac{\pi e^2}{m_e c} f_\alpha \lambda_\alpha \frac{n_{\rm HI}}{H(z)}, \tag{3.2}$$

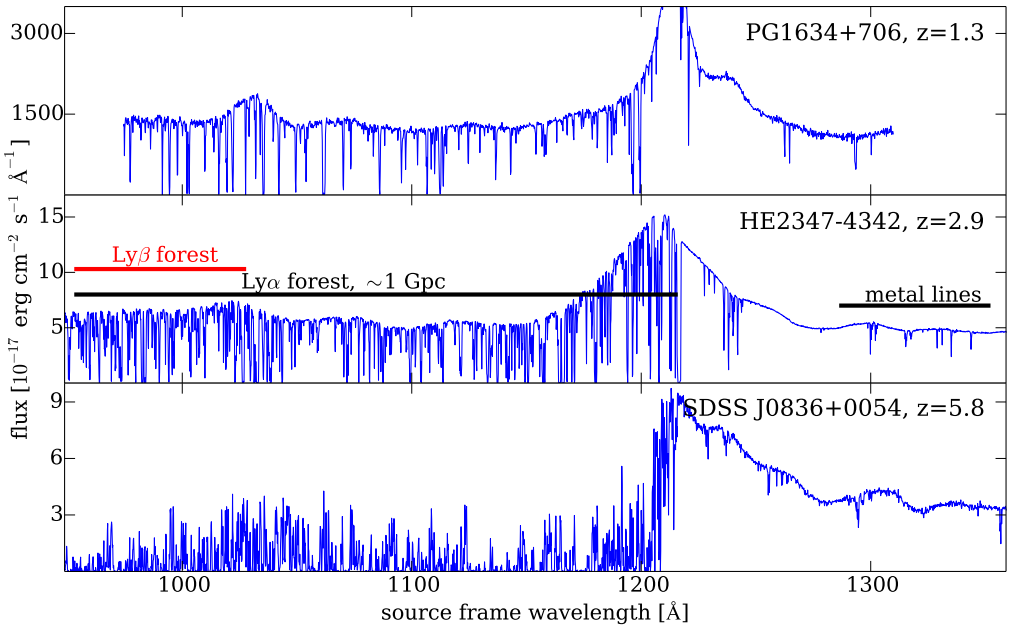

Figure 3.2. Three quasar spectra shown as a function of source (rest) frame wavelength. The three quasars are all at different redshifts, resulting in qualitatively different Lyα forests. At $z = 1.3$ (top panel), the IGM has been almost entirely reionised, and so there is relatively little H I absorption and we can make out individual features quite clearly. At $z = 2.9$, however, the level of ionisation in the IGM is somewhat lower, and we start to see a more dense forest. Also indicated here are metal absorption lines, here situated outside of the Lyα forest. Finally, at $z = 5.8$, ionisation levels are lower still, resulting in more H I absorption and a noticeable drop in the mean transmitted flux fraction. *Figure credit: McQuinn (2016).*

where f_α is the Lyα oscillator strength. This quantifies the local optical depth in terms of the local density of neutral hydrogen in a uniformly expanding Universe. We may develop this further to apply to the continuous, fluctuating IGM, and to take into account velocities in the gas beyond those from expansion alone (i.e. peculiar velocities). As in eq. (1) of McQuinn (2016), doing so yields

$$\tau_\alpha(z) = 1.3\delta_b \left(\frac{x_{\rm H\,I}}{10^{-5}}\right) \left(\frac{1+z}{4}\right)^{\frac{3}{2}} \left(\frac{H(z)/(1+z)}{{\rm d}v/{\rm d}x}\right), \tag{3.3}$$

where $x_{\rm H\,I}$ is the fraction of hydrogen that is neutral, and ${\rm d}v/{\rm d}x$ is the line-of-sight velocity gradient in the gas (including both Hubble and peculiar velocity contributions). This optical depth is then related to the observable transmitted flux fraction, F, via a non-linear mapping $F = \exp(-\tau_\alpha)$. Through eq. (3.3), it is clear that the Lyα forest we observe depends on fluctuations in three key quantities: the baryon density, the ionisation of the IGM, and the velocity gradient of the IGM. The dependence on baryon density and ionisation level is intuitive: the optical depth for photons

at the Lyα wavelength is directly proportional to the number density of neutral hydrogen atoms. The dependence on the velocity gradient is perhaps less so. Peculiar velocities in the gas result in slight deviations from the cosmic distance-redshift relation, causing absorption features to appear at slightly higher or lower wavelengths than their cosmic redshift would imply. These shifts are known as *redshift-space distortions* (RSDs). While uniform gas velocity in a region would cause all features to shift, introducing gradients to the velocity field causes the features to either spread out or cluster together. This introduces additional anisotropies to the Lyα flux distribution, but only along the line-of-sight.

RSDs are, of course, also relevant to measurements of discrete-tracer redshifts. However, the RSDs in the Lyα forest exhibit an important qualitative difference to those in discrete tracers. While RSDs may change the apparent positions of discrete tracers, they will not affect the total number count of objects. In the Lyα forest, RSDs affect the optical depth τ of the IGM, but we measure the transmitted flux fraction $F = e^{-\tau}$. The non-linear nature of this transformation means that F is not conserved by RSDs, unlike either the number counts of discrete tracers or τ (McDonald et al., 2000; McDonald, 2003; Seljak, 2012). This has important consequences for the fitting of correlation functions, and thus the constraining power of Lyα forest measurements.

When measuring the spectra of QSOs, we do not measure fluctuations in the transmitted flux fraction δ_F directly. Rather, we measure the absolute flux received by our telescope f, which we may use to calculate δ_F at a particular wavelength λ via

$$\delta_F(\lambda) = \frac{f(\lambda)}{\bar{F}(\lambda)C(\lambda)} - 1, \tag{3.4}$$

where $\bar{F}$ is the mean transmitted flux fraction a wavelength λ, and $C(\lambda)$ is the value of the unabsorbed QSO continuum. Carrying out this conversion can be tricky, particularly in low signal-to-noise ratio spectra where it is difficult to estimate the QSO continuum. Errors in continuum estimation can cause large-scale, line-of-sight fluctuations to be erased, and will introduce correlated errors in the values of δ_F from the same spectrum. Errors on values from different spectra, however, will not be correlated, and so ignoring correlations from within the same spectrum removes this systematic bias.

In addition to the relatively simple picture of the Lyα forest painted above, there are a number of processes which add layers of complexity to the problem and must be carefully modelled in order to reliably use the Lyα forest as a tracer for measuring BAO. These include contaminating absorption

from elements other than hydrogen, referred to as "metals"[1]; broadened absorption profiles from regions of high neutral hydrogen column density; and broad absorption features from fast-moving gas around the centre of the host QSO. It is also important to note that all of the physical processes described above are of scientific interest in and of themselves, and motivate significant research in the community (e.g. Pieri et al., 2014; Pérez-Ràfols et al., 2018a; Pérez-Ràfols et al., 2018b). In the context of BAO, however, they may be considered as systematics, and we seek to minimise any impact they may have on our measurement of the BAO signal.

Typically, large-scale, 3D clustering in the Lyα forest is quantified by two measurements. The first of these — the Lyα auto-correlation — measures how values of δ_F from one spectrum correlate with values from other spectra. Measuring such a signal was first discussed in (McDonald and Eisenstein, 2007), while the 3D correlation of flux transmission was first studied in (Slosar et al., 2011). The BAO signal was first detected from measurements of the Lyα auto-correlation using data from data release 9 (DR9) of BOSS (Busca et al., 2013; Slosar et al., 2013; Kirkby et al., 2013), with subsequent improvements in DR11 (Delubac et al., 2015) and DR12 (Bautista et al., 2017), as well as DR14 (de Sainte Agathe et al., 2019) and DR16 (du Mas des Bourboux et al., 2020) of eBOSS. The second measurement is known as the Lyα-QSO cross-correlation, which measures how values of δ_F correlate with the positions of high-z QSOs and was first measured in BOSS DR9 (Font-Ribera et al., 2013), with the first detection of BAO coming in DR11 (Font-Ribera et al., 2014), and improvements made in DR12 (du Mas des Bourboux et al., 2017) and eBOSS DR14 (Blomqvist et al., 2019) and DR16 (du Mas des Bourboux et al., 2020).

3.2.2 Redshift-space distortions from spectroscopic galaxy surveys

The accurate redshifts provided by spectroscopic surveys can also be used to study anisotropic clustering via RSDs (Kaiser, 1987) in a population of galaxies. As mentioned previously, this effect describes the apparent shift in galaxies' radial positions due to peculiar velocities induced by gravitational potentials, causing the apparent redshifts of objects to deviate from that which you would expect from cosmic expansion alone. As galaxies will tend to fall towards high-density regions, these RSDs lead to an apparent increase in the clustering amplitude along the radial direction relative to that along the transverse direction. The velocities probed by RSDs are the same velocities that are causing active formation of structure, and so directly measure the rate of structure formation at the redshift of the measurement. Mathematically, RSDs are used to constrain

[1] We follow the conventional approach of referring to all elements bar hydrogen and helium as "metals", despite the confusion it sometimes causes!

the product $f\sigma_8$, where the growth rate $f(z)$ is given by

$$f(z) \equiv \frac{\mathrm{d}\log \mathrm{D}}{\mathrm{d}\log a} = \frac{\mathrm{d}\log \sigma_8}{\mathrm{d}\log a}, \tag{3.5}$$

for D the linear growth function, and σ_8 the amplitude of matter fluctuations at a scale of 8 Mpc h^{-1}. The combination $f\sigma_8$ corresponds to the amplitude of the RSD power spectrum, and has been constrained by numerous RSD measurements in recent years, with the most powerful constraints coming from the recent sixteenth and final data release of eBOSS (de Mattia et al., 2020; Gil-Marín et al., 2020; Neveux et al., 2020; eBOSS Collaboration et al., 2020; Bautista et al., 2021; Hou et al., 2021).

3.2.3 Weak lensing

The bending of photon paths by gravitational potentials, known as *lensing* is a key prediction of GR, and is now observed on a regular basis. Sometimes, galaxies may be lensed strongly, producing clear visual distortions such as multiple images or visible Einstein rings. However, it is the process of *weak lensing* which is best able to probe the growth of large-scale structure. While this may not produce artefacts obvious to the human eye, weak lensing introduces subtle effects to large numbers of galaxies due to the presence of foreground matter, integrated along the line-of-sight (e.g. Blandford et al., 1991; Miralda-Escude, 1991; Kaiser, 1992). We may measure these effects in two main ways.

- *Cosmic shear:* The distorting effect of foreground matter introduces correlations in the cosmic shear measured in background galaxies, which we may attempt to quantify. The cosmic shear power spectrum has amplitude approximately $\propto \Omega_m^2\sigma_8^2$ in the linear regime, though this varies dependent on details of the cosmological model, as well as at the non-linear scales that weak lensing surveys commonly probe (Jain and Seljak, 1997). By splitting galaxies into tomographic redshift bins and measuring cosmic shear in each bin, a three-dimensional map of the matter distribution may be recovered. This enables several parameter degeneracies to be broken as well as more stringent systematic tests to be carried out.

- *Galaxy-galaxy lensing:* The cross correlation between foreground galaxies and a background shear map can be measured as an alternative to the shear power spectrum. This reduces some systematics of the shear power spectrum to noise, yet requires empirically-determined assumptions about the foreground galaxies.

Over recent decades, a number of weak lensing measurements have been made, the most recent coming from the the Dark Energy Survey (Flaugher, 2005), the Kilo-Degree Survey (de Jong et al., 2013), and the Hyper Suprime-Cam Subaru Strategic Program (Aihara et al., 2018). For a more extensive description of weak lensing — covering both theoretical and observational bases — please see, Bartelmann and Schneider (2001), or Chapter 5 of Weinberg et al. (2013).

3.3 Cosmological constraints

Modern cosmology has yielded a consistent description of our Universe under the ΛCDM model. This posits that the Universe evolves with a flat FLRW geometry, in which six cosmological parameters describe its evolution. We follow the Planck naming for these parameters (Planck Collaboration et al., 2020b): A_s and n_s, which describe the primordial power spectrum (see § 1.4); $\Omega_c h^2$, $\Omega_b h^2$, the physical energy density parameters for cold dark matter and baryons respectively; τ, the optical depth to reionisation; and θ_{MC}, the angular scale of the sound horizon. A number of other, relevant parameters may be computed from this set, including the fractional density parameters for cold dark matter, baryons and dark energy Ω_c, Ω_b and Ω_Λ; and H_0, the present day value of the Hubble rate $H(z)$.

The ΛCDM parameters are most strongly constrained by data from the CMB, namely using both temperature and polarisation results from *Planck* (Planck Collaboration et al., 2020b). Within the ΛCDM model, *Planck* data determines the energy densities of radiation, dark matter and baryons to $\sim 1\%$ precision, as well as constraining the current rate of expansion to a similar level. Additionally, when combined with B-mode polarisation data from the analysis of the BICEP2/Keck field (BICEP2 Collaboration et al., 2018), tight constraints on inflationary models can be deduced. In particular, the spectral index of primordial fluctuations is found to be scale-free at current levels of precision, and there is no evidence found for beyond-slow-roll inflation (Planck Collaboration et al., 2020c).

Results from large-scale structure measurements are also able to constrain ΛCDM parameters, but the impact of these additional probes is by no means transformational. Of course, we may naturally want to extend the ΛCDM baseline model in a variety of ways. If we allow Ω_K to vary as a free constant, CMB data alone favours a non-zero value $\Omega_K = -0.044^{+0.019}_{-0.014}$ (Planck Collaboration et al., 2020b). Equally, we may allow the equation of state of dark energy to vary from its value for a cosmological constant $w = -1$, as was presented in § 1.3. Considering this model in the context of CMB data alone results in a constraint $w = -1.58^{+0.16}_{-0.35}$. The lack of constraining power within

the CMB in these two simple extensions to ΛCDM is due to degeneracies in the cosmological dependence of CMB anisotropies, but fortunately large-scale structure allows us to break these degeneracies and thus differentiate between extended models. The extent to which large-scale structure can do so is quantified and explained in eBOSS Collaboration et al. (2020), from which visual representations are displayed in Figure 3.3. We summarise briefly the discussion of two simple extended models below:

– *Free Ω_K ($o\Lambda CDM$):* BAO measurements at different redshifts measure varying combinations of Ω_K, Ω_m and $r_d H_0$, and thus the use of multiple redshift bins (from different tracers) allows these parameters to be measured independently. As such, BAO is able to break the Ω_K degeneracies present in CMB data, prefering a flat Universe in yielding $\Omega_K = -0.0001 \pm 0.0018$ when all SDSS BAO data is used alongside data from *Planck*. Equally, measurements of the low-z growth of structure are sensitive to cosmological curvature, and thus constraints from the weak lensing analyses of DES can break the Ω_K degeneracy to yield $\Omega_K = -0.001^{+0.0043}_{-0.0038}$ when combined with *Planck*. An extensive summary of observational constraints of curvature can be found in Efstathiou and Gratton (2020).

– *Free dark energy equation of state (wCDM):* Measurements from BAO or from type Ia supernovae (SNIa) allow us to break degeneracies when considering models where the dark energy equation of state $w_{\rm DE}$ is not fixed to -1. While data from the Pantheon SNIa sample also exhibits a degeneracy when considering such models, it is in an entirely complementary direction to that of *Planck*. SDSS BAO data shows much less degeneracy, and thus combining all three data sets yields $w = -1.026 \pm 0.033$, consistent with a cosmological constant theory of dark energy. While not quite as powerful as BAO and SNIa in this respect, RSD measurements at $z \gtrsim 1$ are able to break these degeneracies as well, and combining *Planck* results with RSD measurements from SDSS yields $w = -1.09 \pm 0.11$, again consistent with a cosmological constant model.

Whilst the recent increase in precision of cosmological measurements has generally yielded excellent agreement between different probes, it has also exposed a number of tensions. The most well-known of these is the tension between early- and late-Universe measurements of the Hubble constant H_0. While local distance ladder measurements constructed using Cepheid variable stars and SNIa yield a value of 74.22 ± 1.82 km s^{-1} Mpc^{-1} (Riess et al., 2019), measurements of CMB anisotropies from *Planck* yield a value of 67.4 ± 0.5 km s^{-1} Mpc^{-1} (Planck Collaboration et al., 2020b). This tension has been known for a number of years, but now sitting at a significance of

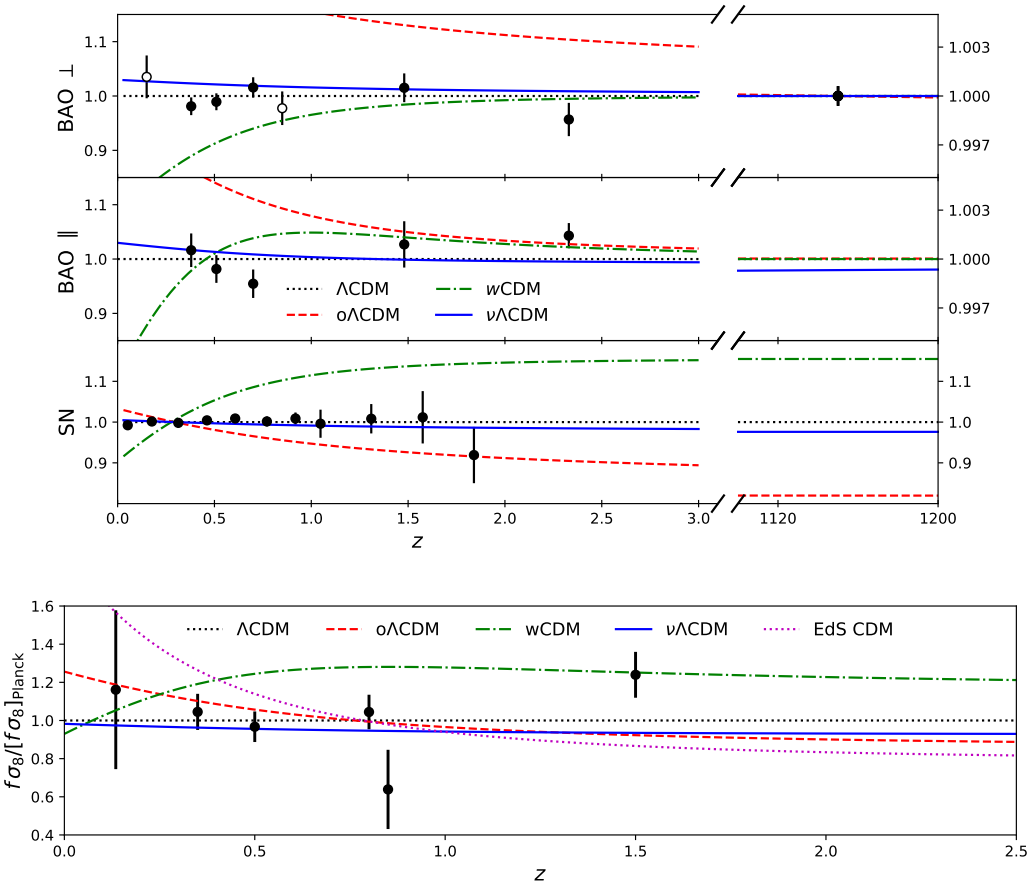

Figure 3.3. Hubble diagram residuals relative to the *Planck* best-fit ΛCDM cosmology from large-scale structure measurements. This includes BAO in both perpendicular and parallel directions from SDSS (top and middle panels, top image), type Ia supernovae from Pantheon (bottom panel, top image), and galaxy RSDs from SDSS (bottom image). The coloured lines represent best-fit *Planck* cosmologies for 1-parameter extensions to the baseline model, with Ω_K (dashed red), $w_{\rm DE}$ (dot-dashed green), and Σm_ν (solid blue) allowed to vary freely in each respective line. Evidently, while all of these models are good fits to *Planck* data, they are not consistent with the different probes of large-scale structure, demonstrating the degeneracy-breaking power of such measurements. *Figure credit: eBOSS Collaboration et al. (2020).*

4.4σ, it has become a crucial point of dispute in cosmology. The reason for the divergence in these measurements is unclear, and there are two main qualitative schools of thought. Many believe that at least one of the measurements must contain an unaccounted-for systematic error, entering either as an unknown physical contaminant or as an artefact of analysis methods. However, numerous re-analyses are yet to yield any evidence of such a systematic (e.g. Zhang et al., 2017; Feeney, Mortlock, and Dalmasso, 2018; Follin and Knox, 2018), at least not that can explain the magnitude of the current tension. Alternatively, others consider this discrepancy to be a result of the difference in cosmological eras probed by the two approaches, and perhaps thus indicative of new physics (e.g. Bernal, Verde, and Riess, 2016). A host of theoretical models have been proposed to step up to this mantle, though as yet most introduce as many problems as they solve. While by no means the only tension in modern cosmology, the H_0 discrepancy is perhaps the most widely discussed such problem, and represents an intriguing challenge as the precision of cosmological measurements continues its rapid ascent.

CHAPTER 4

SPECTROSCOPIC SURVEYS PRESENT AND FUTURE

Hereafter in this thesis, we will focus on observations of our Universe via spectroscopic surveys; sets of observations of distant objects in which the spectral distribution of light is measured as a finely-sampled function of wavelength. Broad features of an object such as spectral breaks can thus be more accurately measured, and we may also detect narrow features such as emission or absorption lines which would be lost to photometry. We focus on spectroscopic observations of the Lyα forest, which may be approached in a variety of different ways. Each of these has its own (dis)advantages and can be preferable for a given set of scientific aims. We direct particular attention to the Dark Energy Spectroscopic Instrument (DESI): the first Stage-IV[1] cosmological experiment to go on-sky. Once its 5-year main survey begins in 2021, DESI will yield high-precision measurements of large-scale structure via a range of different tracers, expanding upon the work of previous surveys and providing dramatically improved cosmological constraints.

We describe the different types of Lyα forest observations in § 4.1, introducing the different approaches to spectroscopic observations and the science needs that they address. In § 4.2, we introduce DESI in more detail, discussing its data collecting process and comparing it to previous-generation experiments in terms of instrumentation, survey strategy and (forecast) cosmological impact.

[1] A "Stage-IV" experiment is described by Albrecht et al. (2006) as one which achieves a measurement of dark energy parameters an order of magnitude more precisely than the Stage-II experiments which were available at the time. This is quantified by a "figure of merit", and is described in § 4.2.2.

4.1 Approaches to Lyα forest observations

The Lyα forest can be used for a range of different scientific purposes, both cosmological and astrophysical. Of course, each science goal will have different priorities for the properties of observational data, and thus will impose different requirements on instrumentation and observational strategy. This covers a wide range of different variables, the precise details of which we will not attempt to describe. We will, however, discuss two broad classes of Lyα forest observations in order to give context to subsequent sections: small-scale/high-SNR, and large-scale/low-SNR.

– *Large-scale/low-SNR:* Typically used for cosmological purposes, these surveys aim to measure large-scale properties of the Universe, spectroscopically observing objects to use them as tracers of the matter density field. In this case, the spectral features are needed to yield more precise redshift determination than photometric surveys, enabling the measurement of many more fluctuation modes due to the addition of a precise third dimension. As such, provided spectral features can be clearly identified and located, there is little benefit to measuring spectra at high SNR or high resolution (within reason). Meanwhile, the measurement of large-scale cosmological properties such as BAO or RSD necessitates that surveys cover a large volume of space, both to gain access to large scale modes and to drive down statistical noise. As such, collecting large numbers of low-SNR spectra provides an optimal dataset for the science goals at hand. Recent notable surveys of this kind include SDSS-I (York et al., 2000), 2dFGRS (Colless et al., 2001), 6dFGS (Jones et al., 2004), BOSS (Dawson et al., 2013) and eBOSS (Dawson et al., 2016).

– *Small-scale/high-SNR:* At the other end of the spectrum[2], these observations seek to measure spectral features with high precision, including properties such as the exact profiles of emission and absorption lines. With this information to hand, measurements can be made of numerous interesting quantities, including but not limited to: galaxies' ages and star formation rates, the dynamical properties of quasars' structure, the abundance of primordial elements, and the column density of neutral hydrogen in the IGM. In order to carry out such measurements with any degree of confidence, it is of great benefit to maximise the resolution and SNR of the input spectra (again, within reason). While increasing the number of spectra would also be beneficial in this case the effect is less substantial than increasing the spectral quality, and thus practical limitations dictate tighter limits on the number of spectra that such studies may seek to observe. Instruments used in such observations include the HIRES

[2]Pun not (entirely) intended!

Galaxy type	Redshift range	Targeting bands used	Targets per deg^2	Exposures per deg^2	Good z's per deg^2	Baseline sample
LRG	0.4–1.0	r,z,W1	350	580	285	4.0 M
ELG	0.6–1.6	g,r,z	2400	1870	1220	17.1 M
QSO (tracers)	< 2.1	g,r,z,W1,W2	170	170	120	1.7 M
QSO (Lyα)	> 2.1	g,r,z,W1,W2	90	250	50	0.7 M
Dark time total	—	—	*3010*	*2870*	*1675*	*23.6 M*
BGS	0.05–0.4	r	700	700	700	9.8 M
Bright time total	—	—	*700*	*700*	*700*	*9.8 M*

Table 4.1. The breakdown of target classes within DESI, including luminous red galaxies (LRGs), emission line galaxies (ELGs) and quasars (QSOs). The latter class is broken into two: QSOs to be used for QSO clustering ("tracers"), and QSOs to be used for measurements of the Lyα forest ("Lyα"). Also included is the bright galaxy survey (BGS), a set of nearby objects to be observed during light- and grey-sky time. The column "Exposures per deg^2" contains the total number of exposures assigned to each target class per square degree. This differs from the number of "Targets per deg^2" as some targets may not be observed at all, whereas others may be included in multiple exposures. *Table credit: DESI Collaboration et al. (2016a).*

spectrograph on the Keck telescope (Vogt et al., 1994) and the UVES spectrograph on the VLT (Dekker et al., 2000), while samples of such spectra are available as part of the KODIAQ survey (O'Meara et al., 2015; O'Meara et al., 2017).

Given the aims of this thesis, we will now focus on the former category, discussing large-scale/low-SNR spectroscopic surveys in greater detail.

4.2 The Dark Energy Spectroscopic Instrument

The Dark Energy Spectroscopic Instrument (DESI Collaboration et al., 2016a; DESI Collaboration et al., 2016b) consists of a new optical corrector and 5000-fibre spectrograph on the existing 4m Mayall telescope at Kitt Peak in Arizona. It will provide low-SNR spectra from $\mathcal{O}(34)$ million objects across a footprint of approximately 14,000 deg^2 in its main survey. These will be broken down into several target classes: luminous red galaxies (LRGs), emission-line galaxies (ELGs), and QSOs. Additionally, a bright galaxy survey (BGS) will be conducted when the moon is too bright to measure distant objects' spectra accurately. A summary of the number densities and redshift ranges of each target class is provided in Table 4.1. The DESI main survey is scheduled to start in early 2021, following delays both due to technical issues and the coronavirus pandemic, and will last for five years. Following this period, it will provide state-of-the-art measurements of BAO from clustering in galaxies, QSOs and the Lyα forest, and from RSDs in its galaxy sample.

4.2.1 Large-scale spectroscopic survey workflow

The workflow of conducting a large-scale spectroscopic survey can be divided into three main stages, corresponding to pre-, mid-, and post-observation periods. We describe the broad process of each stage in turn, focusing on the methods that will be used in DESI in the coming years, as described in DESI Collaboration et al. (2016a) and DESI Collaboration et al. (2016b).

Pre-observation: choosing targets

In order to conduct any spectroscopic survey, one first needs to specify precisely which objects to measure, rather than simply specifying an area of sky as a photometric survey would need to. The process of selecting these targets is carried out using photometric data, which can be provided by a range of instruments. In the case of DESI, three ground-based telescopes have been used to obtain *grz*-band photometry over DESI's footprint: the DECam camera (Flaugher et al., 2015) on the Blanco telescope at Cerro Tololo; the 90Prime camera (Williams et al., 2004) on the Bok telescole at Kitt Peak; and an upgraded version of the MOSAIC-3 camera (Dey et al., 2016) on the Mayall telescope, also at Kitt Peak. Meanwhile, the WISE satellite (Wright et al., 2010) has provided near-infrared photometry in two bands. These observations have now all been completed, and a final processing of photometric targeting data has been produced. A more extensive description of DESI's imaging data can be found in Dey et al. (2019).

From this data, the next stage is to select targets for each of DESI's object classes. Traditionally, this may be done using relatively simple cuts in colour space, as was originally suggested as a baseline for DESI's QSO survey (DESI Collaboration et al., 2016a). However, it is also possible to use more advanced techniques, with a variety of approaches having been explored using SDSS data (e.g. Richards et al., 2009a; Richards et al., 2009b; Yeche et al., 2009; Bovy et al., 2011; Bovy et al., 2012), some of which were used during BOSS and eBOSS (Ross et al., 2012; Myers et al., 2015). Indeed, in DESI, work is underway to use a machine-learning selection tool with a random forest architecture, as mentioned in DESI Collaboration et al. (2016a). Initial results from this approach have been successful, and it has now become DESI's baseline method (Yèche et al., 2020). In the coming months, this approach and others will be tested during DESI's survey validation phase before the final targeting scheme is confirmed.

Mid-observation: obtaining spectra

In order to obtain an object's spectrum, we must be able to align the end of a fibre-optic cable with the object's position in the sky. This requires alignment both of the telescope's field of view in the sky, and of the fibre's position in the focal plane. Determining the set of telescope positions is a process known as *field selection*, while determining exactly which fibre we will use to observe which object for each pointing of the telescope is known as *fibre assignment*. The probability that an object will be assigned a fibre in a given pointing is affected by a number of factors. For example, each instrument will have a minimum angular separation between two fibres for practical reasons, affecting the completeness of objects selected in areas of very high target number density. Equally, the number density of targets for each class will vary across the sky due to factors such as proximity to the galactic plane. At the same time, objects in different target classes will receive different observational priorities, and thus the completeness in any one class may be affected by variations in target density in other classes. Ultimately, these effects may have non-trivial impacts on DESI's clustering measurements. Thus, they must be taken into account when designing field selection and fibre assignment strategies (e.g. Blanton et al., 2017; Pinol et al., 2017; Tao et al., 2018), and clustering measurements must fully understand the strategies to ensure that no systematic bias is introduced (e.g. Burden et al., 2017; Bianchi et al., 2018; Smith et al., 2019; Sunayama et al., 2020).

Post-observation: data processing

Immediately after observations have been carried out, the data must be converted from its raw, 2D, CCD pixel form to the final, 1D, calibrated spectra. This conversion takes place in the *spectroscopic pipeline*, which for DESI will follow the five-phase method outlined in Bolton and Schlegel (2010):

1. *Pixel calibration:* First, we must calibrate the CCD pixels by measuring the properties of the detectors themselves, such as the gains, dark currents and the bias. These measurements will be made with a designated flat field illumination system (DESI Collaboration et al., 2016b), and will then be used to develop calibration routines to apply to our data in a pre-processing stage.

2. *2D point-spread function calibration:* Calibrating the optical path between photons at the focal plane to CCD images. This includes measuring the shape and centroid offset of the 2D point-spread function, using continuum lamp calibration data to locate the spectral traces on each CCD, and accounting for variation in throughput between fibres.

3. *Spectral extraction:* Using the calibrated 2D PSF, we may extract 1D spectra and associated noise vectors from the 2D CCD images. We also subtract background sky signal from the spectra. This extraction procedure follows the "spectro-perfectionisim" algorithm of Bolton and Schlegel (2010).

4. *Flux calibration and coadding:* Spectra from standard calibration stars will be compared with physical models to derive correct flux conversion vectors on an exposure-by-exposure basis. For each object, flux-calibrated spectra from the separate spectrograph arms and exposures will then be coadded to form a single output spectrum.

5. *Redshifting and classification:* Output spectra may be classified and their redshifts determined by a number of methods. In the DESI pipeline, this will be carried out by a "template-fitting" code named `redrock`, which attempts to find an optimal fit to each spectrum via linear combinations of template components. `redrock` and the classification of QSO targets is discussed more extensively in Chapter 5.

At the end of the spectroscopic pipeline, we obtain a set of coadded, calibrated spectra with well quantified noise properties, as well as classifications and redshifts with some degree of error quantification. These data products will then be processed again before conducting science analyses, for example via the construction of galaxy/QSO catalogues, or the extraction of δ_F values from Lyα QSOs.

4.2.2 DESI vs (e)BOSS

DESI is a natural successor to the highly successful BOSS and eBOSS programmes, building on their legacy in a number of ways and making use of the 10 years of experience that they provided. It marks a significant step forward in terms of its instrumentation over (e)BOSS, making both quantitative improvements in its capacity as well as introducing new, qualitatively different methods. Most obviously, DESI will use the 4m Mayall telescope compared to the 2.5m telescope of (e)BOSS, and will have 5000 fibres per exposure rather than 1000. These fibres will be positioned by robotic arms rather than by manual plugging into drilled plates, enabling the adoption of more flexible fibre assignment procedures. Further, DESI's spectrographs are housed in a temperature-controlled room, greatly improving the stability of DESI's optics and reducing the need for constant re-calibration. Combined with improved software and data reduction capabilities, DESI will thus be able to substantially improve the per-exposure SNR for a given object, as well as observing far greater numbers of objects than its predecessors.

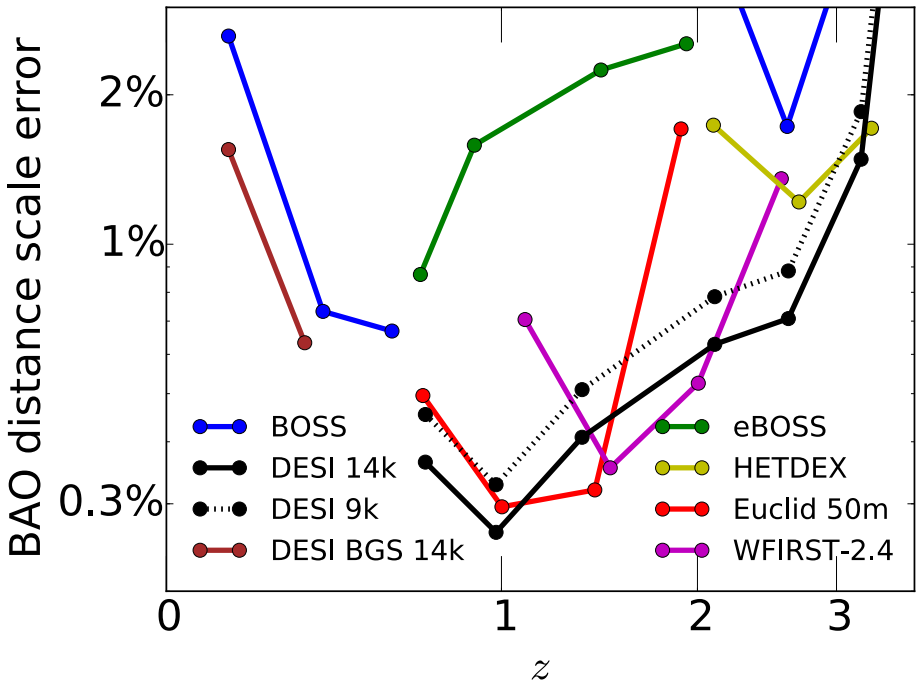

Figure 4.1. Forecast errors on the BAO distance scale from a range of current and future surveys, presented as a function of redshift. Results for DESI are shown in black, with the solid line representing the full survey of 14,000 deg^2, and the dotted line representing a reduced survey of 9,000 deg^2. *Figure credit: DESI Collaboration et al. (2016a).*

In conjunction, these developments will allow DESI to increase its survey volume by an order of magnitude over BOSS. This translates to significant improvements in DESI's ability to measure the BAO feature, as illustrated by Fisher forecast predictions in Figure 4.1. These suggest that DESI will be able to provide constrain the BAO scale to sub-percent levels of precision over the redshift range $z \sim$0.5–3.0. These contstraints are substantially tighter than those from BOSS and eBOSS, and cover a greater range of redshifts than the slightly tighter predictions for Euclid and WFIRST.

The precision of these constraints and the range of redshifts which they span will help to break further degeneracies when exploring models beyond ΛCDM, as was previously discussed in § 3.3. The ability of DESI to constrain these models can to some extent be quantified by a figure of merit (FoM). In particular, the Dark Energy Task Force FoM (Albrecht et al., 2006) is calculated by considering parameter errors when constraining a model extended to allow for a dark energy equation of state $w = w_p + (a_p - a)w_a$ for a pivot value of the scale factor a_p, and in which Ω_K is allowed to vary freely. To be precise, FoM $\propto (\sigma_{w_p}\sigma_{w_a})^{-1}$, where σ_X denotes the error on quantity X. In DESI Collaboration et al. (2016a), such a calculation is carried out combining forecast results from DESI with results from *Planck* and BOSS (in the range $0.45 < z < 0.6$, where

there was initially expected to be little overlap with DESI[3]). Using DESI's BAO results from galaxy clustering and the Lyα forest only, this yields a FoM of 169, approximately four times the BOSS-only value of 37. Extending to use the full shape of the galaxy power spectrum — thus capturing information from RSDs and the Alcock-Paczynski test (Alcock and Paczynski, 1979) — this increases further to 332 or 704, dependent on cuts in wavenumber. This marks a significant leap in cosmological constraining power, more than satisfying the threshold FoM of 110 set as the definition of a Stage-IV dark energy experiment set out in Albrecht et al. (2006).

Of course, the impact of DESI is not limited to measurements of the dark energy equation of state. It will also look to improve constraints of the spectral index of primordial perturbations n_s, as well as its "running" $\alpha_s = \mathrm{d}n_s/\mathrm{d}\ln k$. Combining results from DESI with existing constraints from *Planck* will enable us to differentiate between models of inflation, helping to constrain the number of inflaton fields and the nature of the inflationary potential. Simple inflationary models also predict that the initial fluctuations which seed the growth of large-scale structure are near-Gaussian in nature. This affects the scale-dependent bias of galaxies, and DESI will be able to build on measurements from SDSS (e.g. Castorina et al., 2019) to provide a measurement of primordial non-Gaussianities complementary to those from the CMB (Planck Collaboration et al., 2020c). Finally, DESI is projected to provide a measurement of the sum of neutrino masses with resolution of $\sigma_{\Sigma m_\nu} = 0.02$ eV when combined with *Planck* (DESI Collaboration et al., 2016a). This will provide the tightest such constraint available, and may enable a distinction to be made between the two possible hierarchies of neutrino masses (see Lesgourgues et al., 2013, for a review).

Evidently then, DESI represents a substantial step forward for measurements of large-scale structure, casting light into areas of the field that are poorly understood today and helping to push the boundaries of precision cosmology further still. In conjunction with other Stage-IV dark energy instruments, as well as next-generation dark matter and neutrino experiments, DESI is set to play an important role in an exciting age for cosmology.

[3]This is no longer the case, but for the purposes of this discussion it has no qualitative impact.

PART II

CLASSIFYING QUASAR TARGET SPECTRA

CHAPTER 5

OPTIMAL CLASSIFICATION STRATEGIES

In Chapter 4, we described the process of conducting a spectroscopic survey at a high level, dividing it into pre-, mid-, and post-observation stages. In this chapter, we focus on one part of the post-observation phase: the classification and redshift determination of quasar (QSO) targets from their spectra. We consider the tools currently available to carry out this dual task, and use them to test a selection of classification strategies when addressing two problems relevant to future surveys.

The work shown in this chapter is presented in Farr, Font-Ribera, and Pontzen (2020).

5.1 INTRODUCTION

The value of spectroscopic surveys hinges on their ability to deliver precise redshift determinations and confident classifications of the objects they observe. Precise redshifts allow clustering analyses to access additional information from modes in three dimensions, while confident classifications ensure that the biases of tracer samples can be accurately assessed. In particular, recent spectroscopic surveys have observed increasingly large populations of QSOs, using them to study the large-scale structure of the universe in 3D using the direct clustering of QSOs (e.g. Ata et al., 2018; Neveux et al., 2020; Hou et al., 2021) and the clustering of neutral hydrogen via the Lyman-α forest (e.g. de Sainte Agathe et al., 2019; Blomqvist et al., 2019; du Mas des Bourboux et al., 2020). Traditionally,

the joint task of redshift determination and classification of objects targeted as QSOs has relied on visual inspection (VI) by humans, requiring both substantial expertise and time in order to obtain large and reliable sets of classifications. In early surveys, this task was carried out exclusively by VI (e.g. Schmidt and Green, 1983). Subsequently, varying degrees of automation were introduced (Hewett, Foltz, and Chaffee, 1995; Croom et al., 2004), but VI was still integral to the success of these later surveys, with all spectra being inspected in order to eliminate the $\mathcal{O}(5\%)$ classification errors that early automatic pipelines introduced (Croom et al., 2001). This continued with the advent of the Sloan Digital Sky Survey (SDSS, Gunn et al., 2006; Schneider et al., 2007; Schneider et al., 2010), and in particular the Baryon Oscillation Spectroscopic Survey (BOSS, Dawson et al., 2013) of SDSS-III (Eisenstein et al., 2011). Over the course of its 5 years of operation, BOSS produced three QSO catalogues (Pâris et al., 2012; Pâris et al., 2014; Pâris et al., 2017), all of which relied on VI to classify spectra accurately. The final BOSS QSO catalogue came from the twelfth SDSS data release (referred to as DR12Q from here on), and consisted of 297,301 visually confirmed QSOs from a "Superset" of 546,856 QSO targets.

This enormous VI effort has provided a legacy product of immense value to the community due to the size of its sample and the reliability of its classifications. However, due to the increase in the number of QSO targets observed during the extended BOSS (eBOSS) programme of SDSS-IV (Dawson et al., 2016; Blanton et al., 2017), it was deemed infeasible to repeat such an extensive VI procedure. eBOSS produced two QSO catalogues from the fourteenth and sixteenth SDSS data releases (Pâris et al., 2018; Lyke et al., 2020), transitioning towards the use of automatic classifiers (e.g. Bolton et al., 2012; Hutchinson et al., 2016) in carrying out catalogue construction. The DR14Q and DR16Q catalogues relied on VI for only $\sim$3.7% and $\sim$2.9% of new spectra respectively, focusing on spectra for which automatic classifiers returned ambiguous results. These automatic classifiers were based primarily on the fitting of spectral templates to each QSO target spectrum, determining a class and redshift from the best fit solutions. This approach treats the problem of classification in a qualitatively different manner to VI, looking to find a minimum χ^2 value over a space of possible templates rather than identifying particular spectral features, as a human would do. In order to provide a complementary solution to template-fitting options, machine learning methods can be employed with the aim of replicating the "feature detection" approach of VI. Indeed, the set of VI classifications from BOSS DR12 provides a rich dataset on which to train and test such models. Since this dataset was released, two such tools have been developed, `QuasarNET` (Busca and Balland, 2018) and `SQUEzE` (Pérez-Ràfols et al., 2020; Pérez-Ràfols and Pieri, 2020), both of which are able to at least match the performance of template-fitting methods.

Looking to spectroscopic surveys of the future, the number of QSO targets is set to continue its dramatic increase. Over the course of the 5-year main survey of the Dark Energy Spectroscopic Instrument (DESI, DESI Collaboration et al., 2016a; DESI Collaboration et al., 2016b), approximately 3.6 million QSO targets will be observed, within which approximately 2.4 million QSOs are expected to be found. This number of objects will require automatic classifiers to further lead the process of classifying QSO target spectra, maintaining extremely high levels of accuracy while minimising any reliance on visual inspection. Ahead of surveys such as DESI, then, it is vital to understand how to make best use of the available QSO classification tools. In particular, with a range of qualitatively different classifiers now available, combining their classifications in order to take advantage of their differing strengths and weaknesses will be of great importance when designing optimal classification strategies.

In this chapter, we first assess the current landscape of publicly available QSO target spectra, and the tools available to classify them. In § 5.2, we provide brief descriptions of the existing data and classifiers alongside a comparison of classifier performance levels. We then look ahead to DESI in § 5.3, setting out the key classification tasks that will need to be addressed during its main survey, and describing broad prioritisations that should be made in each case. In § 5.3.1 and § 5.3.2, we address the two main QSO classification tasks in turn, considering various strategies to address them and quantifying the effectiveness of each strategy by testing on BOSS DR12Q Superset data.

5.2 Data and tools

In recent years, the largest sets of QSO spectra have been obtained by the BOSS (Dawson et al., 2013) and eBOSS (Dawson et al., 2016) programmes, which produced five major QSO catalogues (Pâris et al., 2012; Pâris et al., 2014; Pâris et al., 2017; Pâris et al., 2018; Lyke et al., 2020). These catalogues were all constructed in different ways, using different techniques to classify observed spectra of QSO targets. The main BOSS data we use in this chapter is described in § 5.2.1, while a summary of existing QSO classifiers is given in § 5.2.2.

5.2.1 BOSS DR12Q Superset data

The final data release from SDSS-III/BOSS was DR12, which included 627,751 spectra from 546,856 objects that were targeted as QSO candidates during the main survey (Ross et al., 2012). These spectra are each constructed from (on average) four, consecutively-taken, 900-second exposures, which were calibrated and coadded as part of the data reduction pipeline (Bolton et al.,

2012). At a later stage of the pipeline, each spectrum was automatically classified and assigned a redshift. However, they were subsequently re-classified via visual inspection (VI) by an expert (Pâris et al., 2017) in order to improve the quality of the classification. The full results of VI can be found in the DR12Q "Superset" catalogue[1], with a detailed explanation of the classification procedure available in Pâris et al. (2017). Each object was classified by VI as a star, galaxy or QSO (where data was of sufficient quality), and a VI redshift assigned. These redshifts consisted mostly of values from the automatic pipeline, with corrections applied by the expert during VI as appropriate. A confidence in the results of the classification/redshift determination procedure was also given, ranging from 1 to 3 for low to high confidences. This extensive VI effort makes DR12Q Superset spectra ideal for testing the predictions of automatic classifiers, and indeed we make use of this opportunity throughout this chapter. We consider only spectra from objects that were classified with the highest confidence, and thus can practically consider the VI results on these spectra as "true" classifications.

5.2.2 Existing classifiers

A number of automatic QSO classifiers have been developed in recent years, using a variety of different methods. Here, we outline some of the most prominent examples, dividing them into two sub-groups: those based on the concept of "template-fitting", and those based on machine learning methods. Finally, we compare the performance of four classifiers on BOSS DR12Q Superset data.

Template-fitting classifiers

In recent spectroscopic surveys, the most common spectral classification method has been to find the best fit to a given spectrum from a set of templates, framing the problem as one of χ^2 minimisation (Croom et al., 2001; Bolton et al., 2012). This necessitates a sound understanding of the data and its associated errors, as well as a broad set of templates that capture the full variety of features present in the data.

During the BOSS programme, such a classifier was developed as part of the `idlspec2d` data reduction pipeline (Bolton et al., 2012). This used a set of templates for each class (star, galaxy or QSO) constructed from well-understood spectra measured earlier on in SDSS. To construct the QSO templates, 1,000 QSOs were selected at random from the SDSS DR5 QSO catalogue (Schneider et al., 2007), enforcing a uniform redshift distribution. Of these, 568 defect-free spectra had been observed by BOSS, and were selected as the QSO template training sample. Due to the BOSS

[1] Publicly available at https://data.sdss.org/sas/dr12/boss/qso/DR12Q/Superset_DR12Q.fits.

target selection and observation strategies, the redshift distribution of these spectra was strongly weighted towards $z \geq 2.2$ (see Figure 4 of Bolton et al., 2012). Template source spectra for all classes were put through a PCA decomposition procedure, with leading principal components retained and used as a linear basis to fit to each DR12 spectrum. For each spectral type, linear combinations of the basis components were fitted to each spectrum at each redshift within a suitable range. From these fits, a best class and a best redshift was determined as the template class-redshift combination that resulted in the lowest reduced-χ^2. As an output, the classifier then provided this best-fit redshift and class, along with a 1σ statistical error in the redshift and a `ZWARNING` flag. This flag was raised for problematic fits, with a variety of possible bits to signify different possible issues. Most notably, the dominant source of warnings was from spectra where the difference in reduced-χ^2 value between the best and the next-best (with velocity difference greater than 1,000 $\mathrm{km\,s^{-1}}$) classification-redshift pair was below a threshold of 0.01. These spectra thus had two possible classification-redshift pairs which achieved approximately the same quality of fit, and so the classification and redshift were deemed to be insecure. It achieved little contamination at $z \geq 2.2$, but at low z it introduced substantial contamination from stars, as well as missing $\gtrsim 5\%$ of QSOs at all redshifts (see Figure 5.1). As such, its accuracy was not considered sufficient for final classifications, motivating the extensive VI programme discussed in § 5.2.1.

Recently, DESI has developed a new template-based classifier named `redrock` to improve upon several aspects of `idlspec2d`. When comparing the model to the data, `redrock` accounts for the wavelength-dependent spectral resolution of individual fibres. It resamples the model to the wavelength binning of the data rather than using an average resolution and rebinning the data to the model grid, a process which introduces covariances that were not fully modelled in `idlspec2d`. It also supports the ability to simultaneously fit individual exposures instead of requiring a coadded spectrum, which can also introduce covariances that were not modelled in `idlspec2d`. `redrock` optionally has the ability to compare the best fits to a suite of galaxy and stellar archetypes, aiming to identify and remove non-physical fits to those classes, such as those constructed to have negative emission lines. However, it does not currently include QSO archetypes due to their greater spectral diversity. Finally, `redrock` includes a new suite of galaxy and stellar templates. However, as of version 0.7.2 of `redrock-templates`[2] — as used in this chapter — the `redrock` QSO templates are the same as those used by `idlspec2d`. DESI is developing new QSO templates prior to its main survey, looking in particular to improve the $z < 2.2$ QSO performance where the `idlspec2d` QSO templates' training set are underrepresented, but these new `redrock` templates are not yet

[2]Publicly available at https://github.com/desihub/redrock-templates.

available.

Machine learning classifiers

The extensive VI effort carried out on the BOSS DR12Q Superset data (Pâris et al., 2017) provided a large, human-verified set of spectral classifications, ideal for use as training data for machine learning algorithms. Two QSO classifiers have since been developed separately, using very different approaches. Unlike the template-fitting methods described in § 5.2.2, these methods classify each spectrum as a binary QSO/non-QSO, and so they cannot fulfil the same all-round capabilities as their template-fitting counterparts. However, for the more specific problem of QSO template classification, they are able to offer complementary classification channels to the template-fitting methods described previously, with different areas of strength and weakness.

The first of these classifiers is called `QuasarNET` (Busca and Balland, 2018), which attempts to mimic human identification of emission lines. `QuasarNET` is a deep convolutional neural network classifier, taking a smoothed spectrum as an input before carrying out four layers of convolutions. The output from these convolutions is then passed to a fifth, fully-connected layer, before feeding into a number of "line finder" units. Each of these units consists of a fully-connected layer, trained to identify a particular emission line. This is carried out by first dividing each spectrum into a number of wide "boxes", equally spaced in log-wavelength. When training, each line finder unit is given a set of binary identifiers for each spectrum which indicate which box its emission line is in (when the line is present). When making a prediction, a line finder unit attempts to replicate these binary identifiers with a set of numbers between 0 and 1. We interpret these estimates as `QuasarNET`'s confidence as to whether a certain line is in a given box, and we take the largest of these as the confidence that the line has been found in the spectrum as a whole. The line finder unit is also trained to predict the offset of the relevant emission line within each box, and we use the offset within the most confident box to obtain a more precise estimate of the line's location. This then allows us to infer a redshift from each emission line. Throughout this chapter, we train `QuasarNET` models to detect the Lyα, CIV 1548, CIII 1909, MgII 2796, Hβ and Hα lines. These are chosen to ensure that at least two emission lines will be present in the BOSS spectrograph for reasonable QSO redshifts. From `QuasarNET`'s confidences and redshifts for these lines, we carry out classification via a simple procedure: if at least n_{detect} lines are found with confidence exceeding a "confidence threshold" c_{th}, then the spectrum is classified as a QSO. The most likely redshift is then taken from the most confidently identified line. Thus, keeping n_{detect} fixed, choosing a higher (lower) confidence threshold results in fewer (more) objects being classified as QSOs.

Equally, keeping c_{th} fixed, increasing (decreasing) n_{detect} has the same qualitative effect. The results of varying these two parameters is presented in Figure 2 of Busca and Balland (2018). `QuasarNET` is also trained to identify broad absorption line (BAL) QSO spectra, using the same method as described above for emission lines. It is able to identify BAL spectra with high success rates (see Figure 4 of Busca and Balland (2018) and surrounding discussion), though does not currently provide BAL properties such as "balnicity index" as more specialised tools do (e.g. Guo and Martini, 2019). We do not investigate `QuasarNET`'s BAL performance in this chapter, instead focusing on its ability to classify QSO spectra via emission line presence and location.

The classifier `SQUEzE` (Pérez-Ràfols et al., 2020; Pérez-Ràfols and Pieri, 2020) also attempts to mimic the human process of identifying QSO spectra, by looking for sequences of emission peaks. However, its methods are very different to those of `QuasarNET`. `SQUEzE` first smooths each spectrum to remove noise features, and searches for emission peaks above a certain significance threshold. Spectra for which no peaks are found are discarded, while those with significant peaks are retained. For each identified peak, `SQUEzE` then attempts to assign a "trial identity" in order to determine a redshift, analogous to the way that a VI expert would attempt to associate an observed peak with a particular emission line. For each possible identity a number of high-level metrics are computed, including the relative strength of the emission line above the continuum, and the slope of the continuum around the line (see eq. (1) – eq. (3) of Pérez-Ràfols et al. (2020)). These are then passed as features to a random forest classifier. By restricting the features seen by the random forest to high-level metrics, `SQUEzE` attempts to remove any tendency to learn from spurious features of the training set such as instrumental defects or pipeline reduction errors. The random forest then assesses the validity of each trial peak-identity pair, assigning a "confidence" of that pair being a correct identification. Finally, the spectrum is classified as a QSO if the largest of these probabilities meets a certain threshold value, which can be chosen according to the specifics of the classification task at hand.

Performance of classifiers

The classifiers described above all use different techniques, and one would expect variation in performance on different subsets of the data as a result. In order to summarise the performance of the classifiers and provide a comparison, we apply each classifier to BOSS DR12Q Superset data and present the results simultaneously.

We take the results of the DR12 pipeline from the publicly available data in the `spAll` file[3],

[3] https://data.sdss.org/sas/dr12/boss/spectro/redux/spAll-DR12.fits.

restricting to QSO targets by including only data corresponding to object identifiers (column name `THING_ID`) listed in the DR12Q Superset file. We obtain `redrock` results by running version 0.14.3 of the code[4], and using version 0.7.2 of the templates[5] on these same QSO target spectra. For both the DR12 pipeline and for `redrock`, we do not classify as QSOs spectra which raised a `ZWARN` flag[6], as classifications for these spectra are likely to be inaccurate (Pâris et al., 2017). `QuasarNET` results are obtained by applying models trained on 90% of DR12Q Superset data to all spectra from objects not included within this training sample. We train 10 such models, choosing training/testing splits such that all models have mutually exclusive test sets and thus allowing us to obtain `QuasarNET` classifications for 504,534 DR12Q Superset spectra. `SQUEzE` results were provided by its developers, and were obtained by applying a model trained on $\sim$3%[7] of DR12Q Superset data to all objects not included in the training set. Our initial test set consists of all DR12Q Superset spectra for which we have classifications from all four classifiers. To form our final test set, we restrict to those objects for which we have maximally confident VI results to ensure that we are comparing to "true" classifications. In total, then, the test set contains 481,201 spectra.

We quantify the classifiers' performance levels via purity and completeness, which we define in the same way as Busca and Balland (2018):

$$\text{purity} = \frac{\text{number of correctly predicted QSOs}}{\text{number of predicted QSOs}}, \tag{5.1}$$

$$\text{completeness} = \frac{\text{number of correctly predicted QSOs}}{\text{number of true QSOs}}. \tag{5.2}$$

In order for a classification to qualify as a "correctly predicted QSO", we require the classifier to have correctly identified a true QSO's spectrum as that of a QSO, and to have matched the VI redshift with a velocity error

$$\Delta v = \frac{c|z - z_{\rm VI}|}{1 + z_{\rm VI}} \leq 6000 \text{ kms}^{-1}. \tag{5.3}$$

This tolerance level rules out catastrophic failures — classifications which have mis-identified emission lines, for example — but does not require a highly accurate determination of the redshift as `QuasarNET` and `SQUEzE` were not designed to provide such a measurement. When calculating

[4] https://github.com/desihub/redrock/releases/tag/0.14.3.

[5] https://github.com/desihub/redrock-templates/releases/tag/0.7.2.

[6] Such flags are named `ZWARNING` in `idlspec2d` but `ZWARN` in `redrock`. We use the latter name from here on for simplicity.

[7] The convergence of the `SQUEzE` algorithm's performance with training set size was tested in Pérez-Ràfols et al. (2020). Using a training set size larger than this $\sim$ 3% was deemed not to yield any a better-performing model.

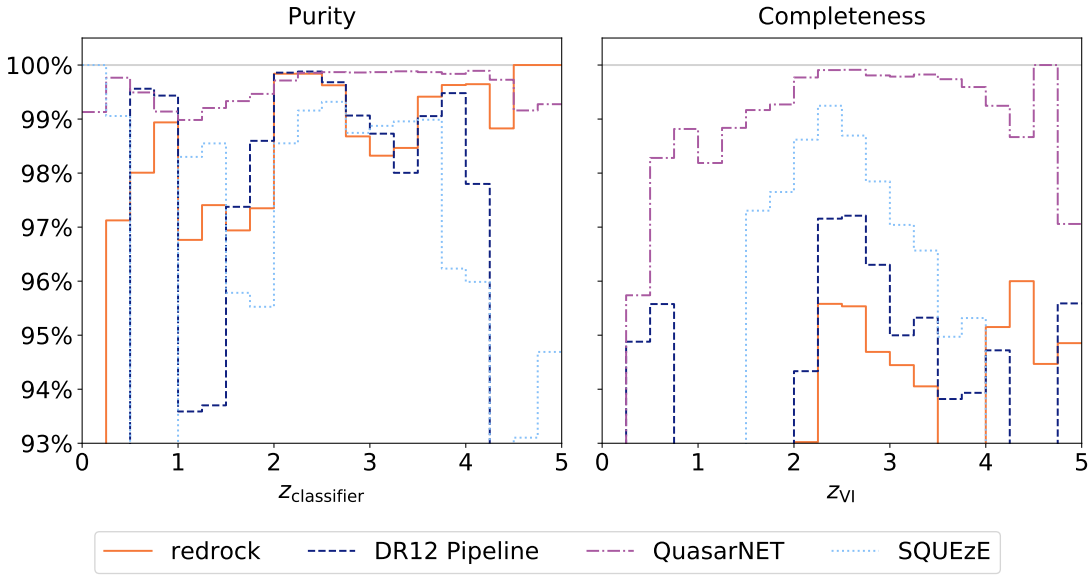

Figure 5.1. Purity as a function of classifier redshift, and completeness as a function of VI redshift for classifications from a selection of current QSO classifiers when applied to a subset of BOSS DR12Q Superset spectra. See equations eq. (5.1) and eq. (5.2) along with surrounding text for definitions of purity and completeness.

the purity, we also allow galaxy spectra identified by a classifier as QSOs to qualify as "correctly predicted QSOs", provided the classifier redshift matches that from VI to within 6,000 $\mathrm{km\,s^{-1}}$. This allows for the ambiguity as to whether a spectrum should be classified as a galaxy or QSO when a degree of broad-line emission is observed as a result of AGN activity.

The results of our classifications are presented in Figure 5.1, in which the left panel shows purity as a function of classifier redshift, and the right panel shows completeness as a function of VI redshift. It is clear that all classifiers are able to achieve very high levels of purity, with levels exceeding 97% for the majority of the redshift range. There is a slight dip in purity for all classifiers in the range $1 < z < 2$, perhaps because neither the Lyα nor the narrow OIII 5007 line is present to aid classification in this region (as noted in Bolton et al., 2012). This dip is less pronounced for `QuasarNET`, and indeed it appears that `QuasarNET` is able to achieve highest purity across the majority of the redshift range, showing very little variation and achieving $\gtrsim 99\%$ in all bins above $z = 0.5$. The other classifiers are able to achieve purities $> 99\%$ as well, but are only able to do so in more limited redshift ranges.

Completeness for all classifiers is also high, though is notably lower than purity and there is more variation in performance, both between classifiers and between redshift bins. `QuasarNET` achieves the highest completeness of the classifiers over the majority of the redshift range, exceeding

99.5% in the range $2 < z < 4$. This value drops slightly at lower redshifts but still exceeds the other classifiers. A rise in completeness at higher redshifts is also evident for the results from all other classifiers. `SQUEzE` shows a peak in completeness in the range $1.5 < z < 4$, perhaps related to the presence of the distinctive Lyα feature for these redshifts. A sharp increase in completeness is seen at $z \simeq 2.2$ for `redrock` and the DR12 Pipeline. This coincides with the redshift value above which the `idlspec2d` template training spectra increase significantly in number, enabling the template-fitting classifiers to identify a greater range of QSO spectra.

Whilst Figure 5.1 provides a suitable overview of the performance of the different classifiers, it is not helpful in assessing their complementarity. For example, while all classifiers miss $\gtrsim 1\%$ of QSOs below $z = 2$, Figure 5.1 does not indicate whether the different classifiers miss the same particular set of QSOs. If the intersection between two classifiers' sets of missed QSOs is small, then combining the classifiers' results in certain ways may enable us to find an even greater proportion of the QSOs in this redshift range. Combining classifications from different classifiers is discussed extensively in § 5.3.1 and § 5.3.2, taking into account different classification priorities in each case.

5.3 QSO TARGET CLASSIFICATION IN THE DARK ENERGY SPECTROSCOPIC INSTRUMENT

The Dark Energy Spectroscopic Instrument (DESI) will begin its main survey during 2021, and will run for 5 years. DESI will offer significant instrumental upgrades over (e)BOSS; its spectrograph has 5,000 fibres per exposure rather than 1,000, and it will be mounted to a 4m rather than a 2.5m telescope (Gunn et al., 2006; DESI Collaboration et al., 2016b). This will allow for a substantial increase in the number of objects observed: the forecasted number of QSOs in DESI is 2.4×10^6 (DESI Collaboration et al., 2016a), compared to 7.5×10^5 in the final combined (e)BOSS sample (Lyke et al., 2020). There will also be significant differences in the observational methods of DESI compared to (e)BOSS, with objects targeted using different photometric data and with different strategies employed to select from these targets. Both the increased dataset size and the different observing strategies will affect the QSO classification challenges facing DESI, and so new classification methods will be required.

Ahead of DESI's main survey, it will go through a period of "Survey Validation" (SV). During this time, the scientific capabilities of the instrument will be assessed, and decisions will be made about observational strategy during the main survey. It will involve the measurement of approximately 50,000–100,000 QSO target spectra, which will be classified via an extensive VI

effort. During this phase `redrock` will be run on all spectra, and comparison with VI results will be helpful for diagnosing and addressing any common failure modes. However, at this stage we will not yet be able to train new `QuasarNET` or `SQUEzE` models on DESI data, and so their use will be limited. It is possible that models trained on BOSS data will perform well on DESI data, and so our set of automatic classifiers could be used to "sense check" VI: if automatic classifiers confidently disagree with a first VI result, then that spectrum could be flagged to go through another round of VI. This would provide a useful aide to the VI effort, and could also prove instructive in understanding the classifiers themselves.

Once SV has been completed, we will be able to train new `QuasarNET` and `SQUEzE` models on DESI data, using the VI results from SV as a truth table. We check that these 50,000–100,000 spectra will be sufficient to train high-performing `QuasarNET` models in § A.1.1. Upon entering the main survey, QSO classifiers will be required for two main purposes: selecting high-z QSOs for reobservation, and constructing QSO catalogues. These are discussed in § 5.3.1 and § 5.3.2 respectively. We assess our classifiers' performance at these tasks using spectra from BOSS DR12Q Superset. The exact properties of these spectra will differ slightly to those from DESI for a variety of reasons. For example, DESI will observe QSO targets to approximately 1 magnitude fainter than BOSS, and will use different target selection strategies. However, the per-pixel signal-to-noise ratio (SNR) for the faintest objects in BOSS and DESI will be similar, and thus BOSS DR12Q Superset data can be considered a reasonable approximation to DESI data when testing QSO selection strategies.

5.3.1 Selecting high-z QSOs for reobservation

Whereas all QSO targets in (e)BOSS were allocated four, consecutively-taken, 900 second exposures on the same night, DESI will only carry out four observations of spectra considered likely to be high-z ($z \geq 2.1$) QSOs, and these will be distributed through the time period of the survey. Such objects benefit significantly from additional exposures as they are used for Lyα forest analyses, which use the values of individual pixels in each QSO's spectrum and so are directly sensitive to its signal-to-noise ratio (SNR). Other objects — contaminants and low-z QSOs — do not benefit from additional exposures provided classifications and redshifts can be accurately determined, and so will not generally be reobserved. This change allows for greater efficiency in DESI's data collecting, prioritising observations which will provide greatest scientific yield. However, it also introduces new challenges; we will need to be able to select high-z QSOs for reobservation from their first exposures, which will have lower SNR than the final, coadded spectra. This reduction in SNR

will make spectra harder to classify, necessitating careful thought about optimal selection methods. Further, the process of selecting for reobservation will need to be conducted entirely automatically, and should be built into the DESI pipeline in order to maximise efficiency and reduce the possibility of human error in generating selections. As such, no VI will be involved and we will need to rely entirely on automatic classifiers.

Ahead of DESI, we would like to assess which selection strategies perform best in this context. In order to approximately replicate the reduction in SNR, we construct a single-exposure dataset from BOSS DR12Q Superset data, choosing one exposure at random from each set of exposures that were coadded to make BOSS' final spectra. This ensures that our single-exposure dataset contains the same number of spectra and balance of contaminants as the coadded dataset. Using this single-exposure dataset, we obtain classifications from `redrock` and `QuasarNET` from which we can build selection strategies. We run `redrock` using the `andmask` option, which sets pixels' inverse variance to zero if they were masked during the BOSS data reduction procedure (due to the presence of strong sky lines, for example). We then train a `QuasarNET` model on 10% of spectra in our single-exposure dataset, a realistic training set size that we can expect from DESI SV (see § A.1.1 and § A.1.2 for analyses of the effects of size and SNR of `QuasarNET` training sets). We apply this model to all remaining single-exposure spectra from objects that were not included in the training set. We restrict our `redrock` classifications to this same set of spectra to ensure consistency.

In order to assess different classification strategies, we must consider what metrics are most relevant for the problem at hand. As a primary concern, any selection procedure must ensure that as great a proportion of true high-z QSOs as possible are chosen to be reobserved. As such, we certainly would like to measure the fraction of high-z QSOs that are selected by a strategy, a quantity related closely to completeness (see eq. (5.2) and surrounding discussion). However, unlike when computing completeness, we are now not interested in whether a strategy is able to determine an accurate redshift. Rather, we only care that a strategy can correctly identify whether an object is a QSO and whether it is at high z; we can consider a classification as correct provided these two criteria alone are met. Of course, any strategy must also avoid selecting too many contaminants for reobservation in order to ensure efficient use of available fibres. This is best assessed by calculating the total number density of objects selected rather than the purity, as it allows us to determine directly whether a strategy is recommending a number density of objects for reobservation that is feasible for DESI. Within current DESI plans, the number density of fibres that will be assigned to reobserving high-z QSOs will be approximately 50 sq deg^{-1}. There will a limited degree of

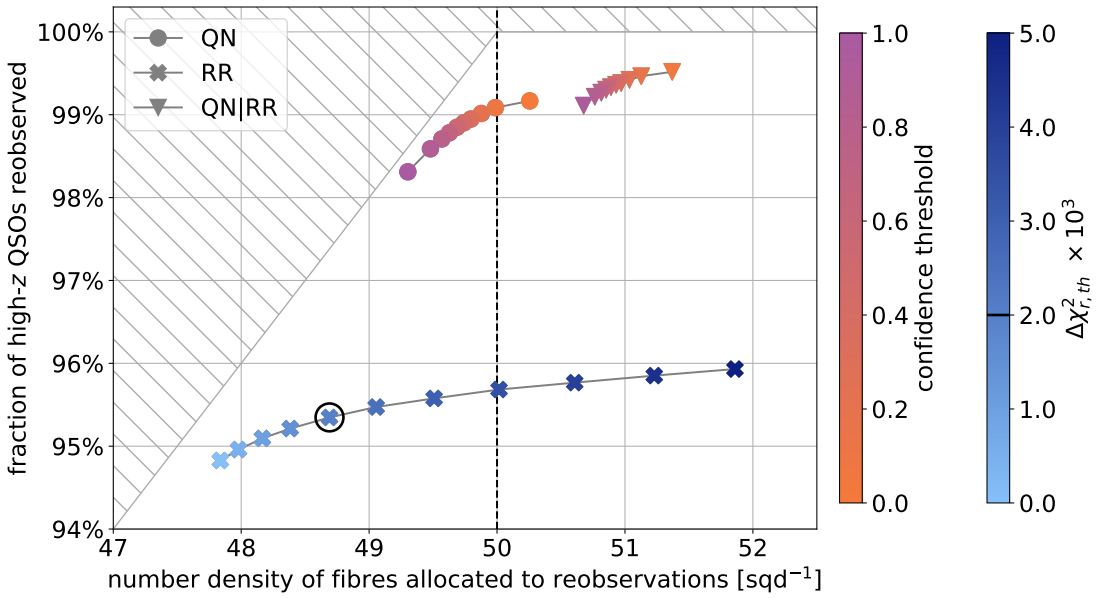

Figure 5.2. Performance of different strategies when selecting high-z QSOs for reobservation, using randomly-chosen single exposures from BOSS DR12Q Superset data. The x-axis has been normalised here to assume a fiducial true number density of high-z QSOs of 50 sq deg^{-1}, as is realistic for DESI (indicated by the vertical, dashed, black line). Here, the "QN" strategy selects objects for which `QuasarNET` identifies at least one emission line with confidence exceeding a confidence threshold, while the "RR" strategy selects objects for which `redrock` finds a high-z QSO template with $\Delta\chi^2_r < \Delta\chi^2_{r,th}$, as defined in eq. (5.4). In the "QN" and "RR" strategies, the variation in the colour of the points indicates the choice of confidence threshold and $\Delta\chi^2_{r,th}$ respectively. In the "QN|RR" strategy, we select objects identified as high-z QSOs either by `QuasarNET` *or* `redrock`. Here, we fix $\Delta\chi^2_{r,th} = 2 \times 10^{-3}$ — the corresponding point in RR is circled — and variation in the colour of the points once again indicates the choice of `QuasarNET` confidence threshold.

flexibility around this number, however a strategy that selects a substantially higher number density of objects than this will not be feasible, and so can be discarded.

In Figure 5.2, we place a number of strategies in the plane defined by these two quantities: the fraction of high-z QSOs selected on the y-axis, and the number density of fibres assigned to reobservations on the x-axis. Strategies can be defined by the results of a single classifier, or by combination of results from more than one classifier. The number density of fibres allocated to reobservations of potential high-z QSOs is normalised by assuming a true number density of high-z QSOs of 50 sq deg^{-1}, as indicated by the vertical, dashed line. As such, an ideal strategy would sit at the point $(50, 100\%)$: it would select 100% of high-z QSOs, while only recommending a number density of reobservations of 50 sq deg^{-1} (i.e. including no contaminants). Including some contaminants in the set of objects recommended for reobservation would then move this

strategy to the right, while missing high-z QSOs without including contaminants would move it diagonally downwards and to the left, along the grey line. The hashed region in the upper left corner of the plot denotes a section of the plane in which a strategy cannot sit: at a fixed number density of reobservations $N_{\text{reobs}} \leq 50$, the maximum fraction of high-z QSOs reobserved is $N_{\text{reobs}}/50$ (achieved when no contaminants are included). A realistic strategy includes both some degree of contamination and fails to select some high-z QSOs, and thus sits in the non-hashed region of the plane.

Performance of individual classifiers

In Figure 5.2, the results from a strategy using `redrock` classifications alone are represented by the set of blue crosses (labelled "RR"). The most straightforward strategy that could be defined from our `redrock` results would select all spectra for which `redrock` returned a high-z QSO as its best-fit template (i.e. that with the lowest χ^2). However, we find that we are able to improve performance by making two simple changes to this strategy. First, we select based on the reduced chi-squared value $\chi^2_r = \chi^2/\nu$ where $\nu = n_d - n_p$ is the number of degrees of freedom, defined as the difference between the number of data points n_d and the number of fit parameters n_p. This provides a more fair comparison between different spectral types which have different values of n_p. Second, we do not select based solely on the template with the lowest χ^2_r. Instead, we consider all templates whose χ^2_r values fall close to the minimum, selecting an object for reobservation if at least one such template corresponds to a high-z QSO. Formally, we consider templates with

$$\chi^2_r - \chi^2_{r,min} = \Delta\chi^2_r < \Delta\chi^2_{r,th}, \tag{5.4}$$

where $\chi^2_{r,min}$ is the minimum χ^2_r over all template-redshift combinations, and $\Delta\chi^2_{r,th}$ is a threshold value that can be chosen freely. We present results for a range of threshold values in Figure 5.2, from $\Delta\chi^2_{r,th} = 0$ to 5×10^{-3}. The value of $\Delta\chi^2_{r,th}$ at each point is indicated by the corresponding colour bar to the right of the main panel. The circled point corresponds to a choice of $\Delta\chi^2_{r,th} = 2 \times 10^{-3}$, and is the choice of threshold used when constructing the "QN|RR" strategy presented in § 5.3.1. One may also consider using `redrock`'s `archetypes` option. This compares stellar and galaxy template fits to a suite of archetypes, penalising fits which show non-physical features such as negative emission lines. This reduces the number of low-χ^2_r stellar and galaxy fits, resulting in more objects being classified as high-z QSOs. Subsequently, using the `archetypes` option selects a greater fraction of high-z QSOs as well as a greater number of contaminants, requiring 49.8

fibres sq deg^{-1} to select 96.6% of high-z QSOs. If using `redrock` alone, this may be deemed a better performing strategy than the RR points shown in Figure 5.2. However, when combined with results from `QuasarNET` (see § 5.3.1), using the `archetypes` option results in a greater number density of fibres being assigned to reobservations with negligible gain in the fraction of high-z QSOs selected. As such, we do not present these results here for clarity.

The results from a strategy using `QuasarNET` classifications alone are represented by the sequence of circular points and the line that joins them (labelled "QN"). Here, we select a spectrum if `QuasarNET` detects at least one emission line with confidence $c > c_{\rm th}$, where $c_{\rm th}$ is the confidence threshold (discussed in § 5.2.2). The colours of the QN points correspond to the choice of this confidence threshold (as indicated by the colour bar to the right of the main panel), with values varying from 0.05 to 0.95 in Figure 5.2. The value of $c_{\rm th}$ is free to be chosen: reducing $c_{\rm th}$ produces a less stringent selection strategy, resulting in a greater fraction of high-z QSOs being selected, as well as a greater number density of fibres being allocated to reobservations.

The QN strategy performs better than RR when selecting QSOs for reobservation. The proportion of high-z QSOs missed decreases from 4.1–5.2% for RR, to 0.9–1.7% for QN. For the threshold values shown, the number density of fibres allocated to reobservations depends varies within 47.8–51.8 sq deg^{-1} for RR and 49.3–50.3 sq deg^{-1} for QN, all of which are feasible values for DESI. As described, the values of $\Delta\chi^2_{r,th}$ and $c_{\rm th}$ can be varied to alter the properties of these two strategies. Indeed, values could be chosen automatically to suit the number of fibres available at any given sky location and on any given night.

Performance of combined strategies

We also define two simple strategies to combine classifications from `QuasarNET` and `redrock`. First, we define a strategy "QN&RR", which selects an object to be reobserved only if both the QN and RR strategies do so. This strategy represents a more stringent selection criteria than either the QN or RR strategies, and so selects fewer high-z QSOS as well as reducing the number density of fibres allocated to reobservations. In the context of selecting high-z QSOs for reobservation, we would like to prioritise increasing the fraction of high-z QSOs reobserved over reducing the number of contaminants reobserved (within reason). As such, the QN&RR strategy is not preferable to QN or RR as it results in at least 5.0% of high-z QSOs being missed and so we do not present this option in Figure 5.2.

Next, we define a strategy "QN|RR", which selects an object to be reobserved if either the QN or the RR strategy does so. The results for this strategy are represented in Figure 5.2 by the

coloured, downwards-pointing triangles and their adjoining line. These points sit further to the right and higher in the plane than either the QN or the RR strategies, reflecting the less stringent selection criteria of the QN|RR strategy and the greater number of objects selected. Combining `QuasarNET` and `redrock`'s classifications in this way results in a reduced proportion of high-z QSOs that are missed, with only 0.5–0.9% not selected. At the same time, it does not require a greater number of fibres be allocated to reobservations: between 50.6 and 51.3 sq deg^{-1} for different choices of `QuasarNET`'s confidence threshold.

As such, provided that 50.6–51.3 fibres sq deg^{-1} can be assigned to potential high-z QSOs, adopting a QN|RR strategy would appear to be the ideal strategy of those presented in Figure 5.2. However, the gain in the fraction of high-z QSOs selected over the QN strategy is not dramatic. Indeed, if fewer fibres are available for reobservations, then the QN strategy provides a suitable, high-performing alternative. These results demonstrate `QuasarNET` alone is able to provide a high-performing selection strategy, while combining it with results from `redrock` in a QN|RR strategy can boost performance further still.

5.3.2 Constructing a QSO catalogue

The second QSO classifier purpose during DESI's main survey is the construction of QSO catalogues. These catalogues will then be used to measure large-scale structure in two ways: via the clustering of QSOs at redshifts $0.9 < z < 2.1$, and via clustering of the Lyα forest at $z > 2.1$. In both cases, it is of great importance that the QSO catalogues contain minimal levels of contamination to ensure that the data used in these analyses truly traces the matter density at the relevant redshifts. At the same time, we need to ensure that we make maximal use of DESI's observations, without discarding significant numbers of spectra unnecessarily. As such, an optimal catalogue construction method should minimise contamination as a priority, while only maximising completeness as a secondary concern. In order to improve catalogue purity and completeness, DESI might be able to include a moderate level of visual inspection, as in the DR14Q (Pâris et al., 2018) and DR16Q (Lyke et al., 2020) catalogues from eBOSS. This VI can be targeted at specific spectra which automatic classifiers failed to classify confidently, and can be built into any catalogue construction strategy according to DESI's VI capabilities.

In order to test various strategies, we apply classifiers to coadded spectra from BOSS DR12Q Superset data. We apply `redrock` to all spectra in this dataset, using `redrock`'s `andmask` option as in § 5.3.1. In this case, as we are working with coadded spectra, this option sets a pixel's inverse variance to zero only when it was masked for all exposures used in constructing the coadded

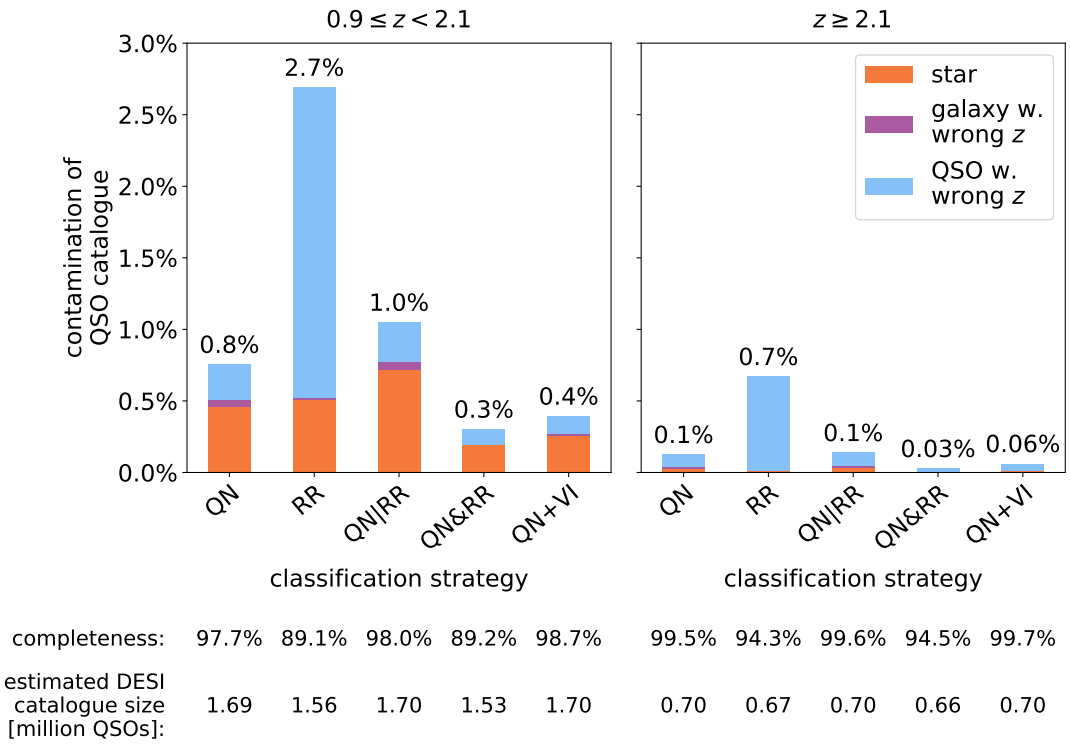

completeness:	97.7%	89.1%	98.0%	89.2%	98.7%	99.5%	94.3%	99.6%	94.5%	99.7%
estimated DESI catalogue size [million QSOs]:	1.69	1.56	1.70	1.53	1.70	0.70	0.67	0.70	0.66	0.70

Figure 5.3. Performance of different classification strategies when constructing a QSO catalogue. Each column shows the percentage level of contamination in the QSO catalogue, divided into different contaminating classes. In the left panel, we show performance on objects with predicted redshift $0.9 \leq z < 2.1$, and on the right on objects with predicted redshift ≥ 2.1. Below the panels, we also display the completeness of QSOs and an estimated DESI catalogue size for each strategy, in each z-bin. The estimated catalogue size is based on both the level of contamination and the level of completeness, and assumes that there will be 1.7×10^6 true QSOs at $z < 2.1$, and 0.7×10^6 true QSOs at $z \geq 2.1$ in DESI (DESI Collaboration et al., 2016a).

spectrum. When this is the case, a pixel's value cannot be considered reliable, and so should not be used in any fits. We then train a `QuasarNET` model on 10% of coadded spectra from DR12Q Superset, a realistic training set size to expect from DESI SV (see § A.1.1). We apply this model to all remaining coadded spectra from objects that were not included in the training set. Once again, we restrict our `redrock` classifications to this same set of spectra to ensure consistency.

In Figure 5.3 we show the level of contamination of QSO catalogues constructed using various strategies. Results are split into low ($0.9 < z < 2.1$) and high ($z \geq 2.1$) redshift bins in order to reflect the two separate uses of QSO catalogues in DESI, and to highlight differences in performance for each case. For each strategy, the level of contamination is broken down into three subsections:

contamination by stars, by galaxies with incorrect redshifts[8], and by QSOs with incorrect redshifts. Below each column, we also present the level of completeness achieved by each strategy, as well as combining contamination and completeness to estimate the size of a final DESI catalogue that each strategy would yield. This catalogue size is constructed by assuming that there will be 1.7×10^6 true QSOs at $0.9 < z < 2.1$, and 0.7×10^6 true QSOs at $z \geq 2.1$ amongst DESI's QSO targets (DESI Collaboration et al., 2016a).

Performance of individual classifiers

We first construct strategies using the classifiers `redrock` and `QuasarNET` individually. We define a strategy "RR" using only results from `redrock`. Unlike in § 5.3.1, we are now seeking to minimise contamination as a priority rather than maximising completeness, and so we do not consider template fits other than that with the lowest χ^2 (`redrock`'s best fit). However, the "RR" strategy presented here does incorporate `redrock`'s `ZWARN` flags — if a spectrum raises such a flag, it is deemed not to yield a reliable classification and so is discarded. This strategy results in 2.7% contamination and 89.1% completeness at low z, and 0.7% contamination and 94.3% completeness at high z. At both low and high z the majority of contaminants are QSOs with incorrect redshifts, with a moderate number of stars at low z as well. As in § 5.3.1, we could use `redrock`'s `archetypes` option when constructing the RR strategy in Figure 5.3. At low z, this results in a substantial increase in contamination to 4.8% (from 2.7%), while completeness increases to 94.3% (from 89.1%). At high z, there is no change in the contamination level (though the proportion of stellar contaminants increases slightly), and completeness increases to 95.8% (from 94.3%). As such, at low z using the `archetypes` option produces a substantially less effective strategy, whereas at high z it increases performance slightly. As we would like to prioritise low contamination when constructing QSO catalogues, we present results without using the `archetypes` option in Figure 5.3.

We then define a strategy using results from `QuasarNET` only, labelled "QN". Here, we include an object in our catalogue if `QuasarNET` is able to identify at least one emission line in the object's spectrum with confidence $c > c_{\text{th}} = 0.5$. This value could of course be varied, but we choose to present a single value for clarity. Using this value of c_{th} results in 0.8% contamination at low z, while achieving 97.7% completeness. Here, over half of the contaminants correspond to stars, while the remainder are mostly QSOs with incorrect redshifts. At high z the level of contamination drops to 0.1% while completeness rises to 99.5%. Contaminants here are mostly made up of QSOs

[8] As mentioned in § 5.2.2, the classification of a galaxy spectrum as belonging to a QSO is considered correct provided the classifier is able to determine the correct redshift.

with incorrect redshifts, with a small number of stars as well. It is also possible to classify using $n_{\text{detect}} = 2$, i.e. requiring that `QuasarNET` must find at least *two* emission lines with confidence greater than c_{th}. For a fixed choice of c_{th}, this results in a more stringent classification criteria, yielding less contamination and reducing the level of completeness as well. Again using $c_{\text{th}} = 0.5$, at low z this "2-line" strategy achieves 0.5% contamination at 96.9% completeness, whilst at high z it achieves 0.09% contamination at 99.4% completeness. We can also require that the two emission lines identified have consistent redshifts, discarding spectra for which `QuasarNET` confidently identifies two lines which would imply drastically different QSO redshifts. This can further reduce contamination in both high- and low-z catalogues, though the effects are small. Indeed, any number of more advanced interpretations of `QuasarNET`'s outputs are possible, and could be motivated by the properties of the dataset being classified.

These results show a significantly higher level of contamination and lower levels of completeness in catalogues produced by the RR strategy compared to those from QN, both in the low- and the high-z bins. In particular, there is much less contamination from QSOs with incorrect redshifts in the QN strategy, though it does introduce a greater level of contamination from galaxies with incorrect redshifts, and from stars in the high-z bin. This difference in the breakdown of contaminants suggests that the two classifiers may be able to act in a complementary manner when combined, as discussed in § 5.3.2. We can also see a substantial difference between the results in the low- and high-z bins: both classifiers perform significantly better at high redshift. This is certainly to be expected of the RR strategy as `redrock`'s templates are built from a high-z-dominated set of QSO spectra (see § 5.2.2), an issue which is currently being addressed by the developers. However, this z-dependence is also present in the QN strategy, likely due to the presence of the Lyα emission line in high-z QSO spectra, a strong and distinctive feature that makes classification more straightforward.

Performance of combined strategies

As in § 5.3.1, we also consider two strategies to combine outputs from `redrock` and `QuasarNET`. We define a "QN|RR" strategy which classifies a spectrum as a QSO if either `QuasarNET` has confidence $c > 0.5$ that it is a QSO, or `redrock` finds a QSO template as the best fit without any `ZWARN` flags. This strategy uses `QuasarNET`'s redshift if `QuasarNET` classifies as a QSO, and `redrock`'s otherwise. This redshift choice was found to be more effective than the alternative of using `redrock`'s redshift as default, and only using `QuasarNET`'s redshift where `QuasarNET` classified as a QSO but `redrock` did not. The QN|RR strategy introduces a greater degree of

contamination to the catalogue than the QN strategy at low z, but less than the RR strategy. This counterintuitive behaviour is due to the choice of redshift mentioned previously: some spectra which were correctly classified as QSOs but assigned an incorrect redshift by the RR strategy are now assigned a correct redshift from `QuasarNET` instead. At high z, the QN|RR strategy results in the same level of contamination as the QN strategy, and lower than the RR strategy. At the same time, it is able to achieve a higher level of completeness than either the RR or the QN strategies at both low and high z. The gain over the QN strategy in this respect is small, however. As we wish to prioritise low contamination in our QSO catalogues, it seems that the QN|RR strategy is less effective than then QN strategy at low z, and thus would not be an advisable strategy. At high z, it performs negligibly better than the QN strategy.

Similarly, we define a "QN&RR" strategy which classifies a spectrum as a QSO only if both `QuasarNET` has confidence $c > 0.5$ that it is a QSO, and `redrock` finds a QSO template as the best fit without any `ZWARN` flags. We also require that `QuasarNET` and `redrock` agree on the object's redshift to within 6,000 $\mathrm{km\,s^{-1}}$. Again, this strategy uses `QuasarNET`'s redshift for all objects classified as QSOs. As one would expect, this strategy achieves lower contamination levels than either then QN or RR strategies, reducing it to 0.3% at low z and 0.03% at high z. The completeness of the sample is also reduced relative to the individual classifier strategies, and as a result the predicted catalogue sizes are significantly lower than the true values. The very low levels of contamination achieved by the QN&RR strategy make it well suited for constructing QSO catalogues. Of course, it results in significant drop in completeness relative to the QN strategy, but this may be deemed a necessary sacrifice in order to construct a maximally pure catalogue.

It is also worth noting that both combined strategies perform significantly better at high z than low, suggesting that there may be benefit to using different classification strategies when constructing catalogues for different purposes. For example, when constructing a low-z QSO catalogue for QSO clustering analysis, one may judge that only the QN&RR strategy is able to construct a sufficiently pure catalogue, despite its modest level of completeness. At high z, the levels of contamination in either the QN or QN&RR strategies may be considered acceptable, and so the QN strategy may be favoured in order to make use of the 4.6 percentage point increase in completeness that this strategy offers. When measuring the clustering of QSOs at low z, however, understanding the completeness of the QSO sample more deeply is of great importance. The inclusion of any given QSO in a QSO catalogue depends on a number of other quantities both physical and observational, such as the photometric properties of the QSO and the position in the focal plane of the fibre used to observe it (e.g. Reid et al., 2016; Laurent et al., 2017; Ata et al.,

2018). As such, we emphasise that improvements to completeness must be complemented by future, detailed modelling of the entire pipeline when performing any clustering analysis.

Role of visual inspection in constructing QSO catalogues

As mentioned previously, DESI will have some capacity for VI when constructing QSO catalogues, which can be targeted towards spectra for which automatic classifiers were not able to yield a confident classification. In constructing the eBOSS DR14Q catalogue, approximately 3.7% of new spectra were visually inspected in this way, chosen via a simple decision tree (see § 3.2 of Pâris et al., 2018), while in constructing the DR16Q catalogue this dropped to approximately 2.9% (Lyke et al., 2020). Given the increase in data quantity that DESI will provide, it is likely that VI levels will need to be reduced further, perhaps to $\lesssim 1\%$. The results shown in this section thus far do not allow for any VI capacity, and could undoubtedly be improved by introducing a system to allow for spectra to be flagged for VI. There are several ways that this could be incorporated into any given strategy, and here we discuss two simple options.

Combining `QuasarNET` with VI, one could define two confidence thresholds $c_{\text{th}}^{\text{lo}}$ and $c_{\text{th}}^{\text{hi}}$, and classify each spectrum with confidence c as follows:

- $c \geq c_{\text{th}}^{\text{hi}}$: `QuasarNET` is sure that this spectrum is a QSO, thus classify as a QSO.
- $c_{\text{th}}^{\text{lo}} \leq c < c_{\text{th}}^{\text{hi}}$: `QuasarNET` unsure whether this spectrum is a QSO, thus send it to VI.
- $c < c_{\text{th}}^{\text{lo}}$: `QuasarNET` is sure that this spectrum is not a QSO, thus classify as a non-QSO.

This sends to VI those spectra which `QuasarNET` is least able to classify definitively. Adopting this classification strategy with values $c_{\text{th}}^{\text{lo}} = 0.04$ and $c_{\text{th}}^{\text{hi}} = 0.96$ results in VI being requested for 1.0% of spectra, an appropriate proportion for DESI. We can estimate its performance by simply assigning DR12Q Superset VI results to spectra that fall into this category; doing so yields the results labelled "QN+VI" in Figure 5.3. When compared with the QN strategy from the same figure, these results show a reduction in catalogue contamination levels by almost 50% in both redshift bins, alongside an increase in completeness of one percentage point at low z. This combination of `QuasarNET` and VI can thus be deemed highly effective, providing contamination rates in between those of the QN and QN&RR strategies, while exceeding the completeness of either strategy.

Equally, one could combine VI with `redrock`'s classification results. For example, one could send all spectra with `ZWARN` flags corresponding to fitting issues[9] to VI. Compared to the RR

[9] Some `ZWARN` flags correspond to issues such as broken fibres, and so VI would not be able to help with these issues.

strategy, this results in substantial improvements to completeness at both low and high z (3.7 and 1.6 percentage points respectively), and only small (< 0.1 percentage point) improvements to contamination. However, including VI in this way requires 6.7% of spectra to be visually inspected, a proportion that lies beyond the bounds of feasibility for DESI. Restricting VI to those spectra for which a `ZWARN` flag was raised and the best fit spectral type was "QSO", we require a far more manageable 0.6% of spectra to be inspected. Here, there is negligible reduction in contamination, but completeness increases by at least 1 percentage point in both low- and high-z bins.

Evidently, VI can also be built into more complex classification strategies as well. While in BOSS, all QSO target spectra were visually inspected, in eBOSS, decision trees were built to combine automatic classifier results with VI. In eBOSS DR14Q (Pâris et al., 2018), a decision tree was built based on the five templates from the pipeline (`idlspec2d`) with lowest reduced χ^2 values, highlighting spectra as requiring expert VI if the top 5 best-fit solutions were inconsistent, or if flags denoting low spectral quality were raised (see § 3 of Pâris et al., 2018, for further details). In DR16Q (Lyke et al., 2020), a similar decision tree was used, with `QuasarNET` employed to reduce the VI proportion further still. As a result of these decision trees, only 3.7% and 2.9% of new spectra were visually inspected in the construction of DR14Q and DR16Q respectively. These decision trees were built with prior knowledge of `idlspec2d`'s failure modes and the distribution of contaminant spectra. For example, applying `idlspec2d` to QSO target spectra from BOSS introduced a significant degree of contamination from stars, particularly at low z. As such, the decision tree was designed to carefully remove stellar spectra for which `idlspec2d` incorrectly returned "QSO" as its best fit. When constructing catalogues in DESI, a similar decision tree will be useful, constructed to take into account both the properties of `redrock`'s classifications and the contaminants in the DESI QSO target set. Such a decision tree could be constructed manually (as in eBOSS), or a simple machine learning approach could be used to assess in greater detail the full set of template fits from `redrock` and line identifications from `QuasarNET`.

5.4 Summary & conclusions

In this chapter, we have assessed problems of QSO classification relevant for future spectroscopic surveys such as DESI, and have demonstrated that existing automatic classifiers can be used to construct highly effective classification strategies. In § 5.2, we summarised the automatic classifiers currently available. We provided a simple comparison of their performance levels, demonstrating that `QuasarNET` is able to out-perform other classifiers over a range of redshifts. In § 5.3, we

identified the QSO classification tasks that will be relevant to DESI's main survey: selecting high-z QSOs for reobservation, and constructing QSO catalogues. We quantified how well `QuasarNET` and `redrock` perform at addressing these tasks by applying them to BOSS DR12Q Superset data, using `QuasarNET` models trained on appropriately-sized training sets, and using single-exposure spectra where necessary.

We then addressed the two classification tasks in turn, first presenting the performance of various strategies when selecting high-z QSOs for reobservation in § 5.3.1. We found that reobserving all objects selected by `QuasarNET` alone provides an effective solution, resulting in approximately 1% of high-z QSOs being lost. This loss can be reduced further — to 0.5% — by reobserving all objects selected by either `QuasarNET` or `redrock`, provided an additional 1 fibre sq deg^{-1} can be made available for reobservations. In § 5.3.2, we then used coadded spectra to consider various classification strategies in the context of constructing QSO catalogues. We showed that using `QuasarNET` alone for this task is able to offer sub-percent levels of contamination at both low and high z, simultaneously yielding high levels of completeness. We showed further that including in a catalogue only QSOs identified by both `QuasarNET` and `redrock` reduces levels of contamination by at least a factor of 2, though a substantial reduction in completeness is an unfortunate consequence. Alternatively, we considered combining `QuasarNET`'s outputs with a small visual inspection fraction, defining a subset of uncertain spectra using two confidence thresholds. Permitting 1% of spectra to be manually classified in this way reduced contamination to almost half that when using `QuasarNET` alone, improving levels of completeness at the same time.

In all, we have demonstrated that `QuasarNET` alone is able to suitably address the QSO classification tasks of DESI, and that combining its classifications with those from `redrock` in simple ways is able to boost performance further still. The exact performance levels achieved by the strategies we define may, of course, vary when applied to DESI data. They will be affected by the distribution of contaminant types through the set of QSO target spectra, and by the instrumental properties of DESI itself. The precise impact of these differences is not yet clear, though the similar pixel-level noise in BOSS and DESI's faintest QSOs indicates that BOSS data represents a suitable and well-understood proxy to use at the current moment. Further, both `QuasarNET` and `redrock` will evolve ahead of DESI; `redrock`'s QSO templates are actively being developed to yield improved performance, while there is potential to extract information more efficiently and more precisely from `QuasarNET`'s raw outputs. However, the results presented in this chapter can offer encouragement that the current range of automatic QSO classifiers are well suited to addressing the needs of forthcoming spectroscopic surveys, and can safely reduce the historical

reliance on visual inspection without adversely affecting science outcomes.

PART III

PRODUCING MOCK DATASETS

CHAPTER 6

MAKING MOCKS WITH LYACOLORE

Once quasar (QSO) catalogues have been constructed, they may be used in a number of different ways. In particular, high-z QSO spectra may be used to study the Lyα forest, which we may in turn use to measure baryon acoustic oscillations (BAO, see § 3.2.1 for details). The analysis process to carry out such measurements is non-trivial, and requires detailed tests to verify its reliability. In this chapter, we discuss the construction of mock datasets for this purpose, presenting the `LyaCoLoRe` package that we have developed to address this need.

The work shown in this chapter is presented in Farr et al. (2020). The mocks described in this section are constructed using two key pieces of software, `CoLoRe` *and* `LyaCoLoRe`*. Beyond a small modification relating to interpolation methods,* `CoLoRe` *was developed by collaborators rather than myself. The development of the* `LyaCoLoRe` *package, however, was led by myself.*

6.1 INTRODUCTION

The late-time measurement of the BAO "standard ruler" (Peebles and Yu, 1970) has become a key component of modern observational cosmology. As a precise method of measuring large-scale structure, it provides complementary information to the cosmic microwave background and thus is able to help constrain a number of extensions to the baseline ΛCDM model (e.g. eBOSS Collaboration et al., 2020). As discussed in § 3.2.1, the BAO signal has been measured most frequently using discrete objects such as galaxies at $z \sim 0.1 - 1.0$ and QSOs at $z \sim 1.2 - 1.7$ to

trace the matter density field. Beyond these redshifts, however, the number density of such objects observed in current surveys is too low to measure the BAO signal at any level of significance, and so alternative tracers must be found to extend BAO measurements to $z \sim 2$ and beyond.

In recent years, the Baryon Oscillation Spectroscopic Survey (BOSS, Dawson et al., 2013) and its extension eBOSS (Dawson et al., 2016) have enabled the use of such a tracer in the form of the Lyman-α (Lyα) forest: a sequence of absorption features that appears in the spectra of high-z QSOs as a result of Lyα absorption of light in the neutral hydrogen gas between QSO and observer. These spectral features thus trace the density of neutral hydrogen gas in the inter-galactic medium (IGM) along the line of sight (Bi, Boerner, and Chu, 1992). Indeed, analytical models developed during the 1990s showed that the Lyα forest absorption closely traces the distribution of dark matter on scales larger than the Jeans length (e.g. Cen et al., 1994; Petitjean, Muecket, and Kates, 1995; Miralda-Escudé et al., 1996). The Lyα forest should, then, provide a suitable means to extend measurements of cosmic expansion via BAO to earlier in the Universe's history. Measuring such a signal was first discussed in McDonald and Eisenstein (2007), while the 3D correlation of flux transmission was first studied in Slosar et al. (2011). The BAO signal was first detected from measurements of the Lyα auto-correlation using data from data release 9 (DR9) of BOSS (Busca et al., 2013; Slosar et al., 2013; Kirkby et al., 2013), with subsequent improvements in DR11 (Delubac et al., 2015) and DR12 (Bautista et al., 2017), as well as DR14 and DR16 of eBOSS (de Sainte Agathe et al., 2019; du Mas des Bourboux et al., 2020, respectively). The cross-correlation between the Lyα forest and QSOs was first measured in BOSS DR9 (Font-Ribera et al., 2013), with the first detection of BAO coming in DR11 (Font-Ribera et al., 2014), and improvements made in DR12 (du Mas des Bourboux et al., 2017) and eBOSS DR14 (Blomqvist et al., 2019) and DR16 (du Mas des Bourboux et al., 2020).

The upcoming Dark Energy Spectroscopic Instrument (DESI, DESI Collaboration et al., 2016a; DESI Collaboration et al., 2016b) will be able to advance these measurements greatly. Over the five years of its operation, it will measure approximately 700,000 QSO spectra with $z > 2.1$ (DESI Collaboration et al., 2016a), over three times the 210,005 spectra in the final eBOSS dataset (du Mas des Bourboux et al., 2020). Ahead of such an increase in statistical power, it is vital to be able to sufficiently test analysis pipelines to ensure that they do not introduce any biases. Equally, it is important to be able to quantify exactly how secondary astrophysical effects will impact upon BAO measurements. The best way to carry out both of these tests is through the development of mock datasets (e.g. Le Goff et al., 2011; Font-Ribera, McDonald, and Miralda-Escudé, 2012; Bautista et al., 2015) — synthetic realisations of a survey for which cosmological and astrophysical

parameters can be easily controlled. Producing such datasets must be computationally inexpensive in order to allow for generation of a large number of realisations, but the data must also provide realistic representations of the survey itself.

In this chapter, we introduce a package designed to produce mock datasets for current and future Lyα forest BAO analyses, `LyaCoLoRe`. In § 6.2, we describe the methods used to generate such datasets, including the use of a Gaussian random field to generate the 3D correlations and the subsequent post-processing to yield realistic skewers of transmitted flux fraction. The methods to determine the optimal values of parameters used in these transformations are detailed in § 6.3. We then verify that the datasets are able to fulfil their purpose for BAO analyses in § 6.4, measuring correlation functions in the same way as recent analyses from BOSS and eBOSS. In § 6.5, we introduce and briefly test additional astrophysical effects that `LyaCoLoRe` is able to include, before summarising and concluding in § 6.6.

6.2 Making the mocks

The requirement of mocks to be computationally inexpensive but also large in volume prohibits the use of hydrodynamical or N-body simulations in their construction. Instead, Gaussian random field methods can be used to generate a linear density field in a large box. This method does not capture non-linear evolution, generating data based solely on an initial power spectrum, but is orders of magnitude faster than state of the art simulations. Further, the presence of non-linear structure is not of vital importance to BAO measurements, particularly at $z \gtrsim 2$ where the Lyα forest is observed (Kirkby et al., 2013). As such, Gaussian random field methods are particularly well suited to the production of Lyα BAO mock datasets. Having generated such a box, tracers such as QSOs can be placed at peaks in the density field via Poisson sampling according to an input bias and number density, and line-of-sight skewers can be drawn by interpolating within the box.

Converting these skewers to mimic the transmitted flux fraction of the Lyα forest then requires a significant degree of post-processing. Despite the speed of Gaussian random field methods, resolution higher than $\mathcal{O}(1)$ Mpc h^{-1} is not possible within the computational bounds of mock production due to memory limitations. As a result, the 1D power spectrum of the skewers $P_{\rm 1D}(k_{\parallel})$ — the power spectrum measured only from modes lying along the line of sight of each skewer — is greatly suppressed. This subsequently affects the errors on our BAO measurements, as the 3D flux power spectrum of the Lyα forest has a significant contribution to its error that is proportional to the 1D power spectrum, known as aliasing noise (McDonald and Eisenstein, 2007). As such, we

must boost the 1D power spectrum by the addition of small-scale fluctuations in order to ensure that our BAO errors behave correctly. Further, we must convert from density to optical depth at each point of each skewer. The details of this relationship are complex, but in the context of Gaussian random field mocks we are constrained to using a simple approximation such as the fluctuating Gunn-Peterson approximation (FGPA, Croft et al., 1998). Finally, we must add redshift-space distortions to our skewers. These distortions occur as a result of peculiar velocities in the IGM, and we observe them as an anisotropy in measurements of power spectra and correlation functions.

In this chapter, we use `CoLoRe` (Alonso et al., in prep.) to generate our initial Gaussian skewers, as described in § 6.2.1. We then present the package `LyaCoLoRe`, which is able to convert `CoLoRe`'s output into realistic skewers of transmitted flux fraction. The methods used in this transformation are described in § 6.2.2. Finally, in § 6.2.3, we discuss the computational requirements of running both of these packages. The output skewers from `LyaCoLoRe` then require the addition of instrumental noise and combination with a QSO continuum before they can be considered realistic spectra. This can be carried out in the context of DESI by a package called `desisim`[1], which is not discussed here.

6.2.1 CoLoRe: Cosmological Lognormal Realisations

The `LyaCoLoRe` mocks originate from an existing program called `CoLoRe`[2], a highly parallelised code initially designed to produce large catalogues of multiple tracers with the same underlying density field (Alonso et al., in prep.). In this chapter, we use `CoLoRe`'s lognormal density model for speed, though first and second order Lagrangian perturbation theory methods (LPT/2LPT) are also available. From this density field, `CoLoRe` can produce a number of observables such as cosmic shear, intensity maps, CMB lensing and integrated Sachs-Wolfe maps. Most importantly in the context of this chapter, it is also able to draw line-of-sight skewers from each object to a central observer, interpolating the Gaussian field at intermediate points. This final functionality makes `CoLoRe` well suited for Lyα forest mocks. The basic steps that `CoLoRe` takes in computing such skewers are outlined in the 5-stage process below:

1. Generate a Gaussian random field δ_G at $z = 0$ in a Cartesian box according to an input power spectrum.

2. Compute a corresponding radial velocity in each cell using the gradient of the Newtonian

[1] Publicly available at https://github.com/desihub/desisim.
[2] Publicly available at https://github.com/damonge/CoLoRe.

gravitational potential ϕ:

$$v_r(z=0) = -\frac{2f_0}{3H_0^2\Omega_m}(\mathbf{e}_r \cdot \nabla)\phi(z=0), \tag{6.1}$$

where f_0 is the logarithmic growth rate at $z=0$, H_0 is the Hubble constant, Ω_m is the matter density parameter, and $\mathbf{e}_r$ is the radial unit vector.

3. Calculate the redshift of each cell (taking the centre of the box as the observer) using a given input cosmology, and evolve the fields to that redshift using the corresponding linear growth factor.

4. Carry out a lognormal transformation of the Gaussian field, and Poisson sample it using an input number density $n(z)$ and bias $b(z)$ to obtain a set of sources (QSOs in our case).

5. Compute line-of-sight skewers from each source to the centre of the box by interpolating the initial Gaussian field and the radial velocity field.

The final output from `CoLoRe` is a set of QSOs and corresponding Gaussian field skewers, as well as values of cosmological variables along the skewers. The QSOs have the correct two-point clustering properties in 3D on large scales, as demonstrated briefly in Appendix B.1 and in more detail in Alonso et al. (in prep.). The skewers also have the correct two-point 3D correlations, as demonstrated in § 6.4. Correlations beyond two-point statistics will not be reproduced by `CoLoRe`'s Gaussian methods. These are relevant to the Lyα forest as our lines-of-sight are not randomly placed, instead being determined by the locations of QSOs. Such higher-order correlations are unlikely to inhibit the determination of the BAO scale, but quantitative studies of such effects should be carried out ahead of DESI to ensure that this is the case. In the future, using `CoLoRe`'s LPT/2LPT methods will start to introduce higher-order correlations to our mocks, as discussed briefly in § 8.2.2.

6.2.2 LyaCoLoRe

While `CoLoRe` is able to produce skewers with 3D, large-scale correlations matching a given input in a short timeframe, its "raw" output requires significant post-processing before it can be considered a realistic representation of the Lyα forest. To implement these stages of processing, we have developed a `Python` module under the name `LyaCoLoRe`[3]. This code transforms `CoLoRe`'s output

[3] Publicly available at https://github.com/igmhub/LyaCoLoRe.

into realistic skewers of transmitted flux fraction. The following sections describe the key methods that LyaCoLoRe uses to do so, with each step represented visually in Figure 6.1.

Adding small-scale power

In order that the memory requirements of running CoLoRe do not become overwhelmingly large, we are limited to using a grid of 4096^3 cells. Requiring that this encloses the volume of a full Lyα survey limits us to using a low-resolution grid, with cells in CoLoRe's raw output of $\mathcal{O}(1)$ Mpc h^{-1}. In the context of the Lyα forest, we observe clustering on scales down to the Jeans Length, approximately 100 kpc h^{-1} (Walther et al., 2018) and an order of magnitude lower than the resolution we can feasibly achieve. While BAO is a large-scale phenomenon, imposing that the synthetic data has approximately the right small-scale properties ensures that the covariance matrices in our final analyses are realistic. We address this by first interpolating CoLoRe's Gaussian skewers — labelled as δ_C — to a smaller cell size, using nearest grid point (NGP) interpolation in order to avoid introducing additional smoothing.

We then generate a set of new, independent Gaussian skewers δ_ε on the grid of smaller cells according to an input 1D power spectrum. We take the k-dependence of this 1D power spectrum to follow that used in McDonald et al. (2006):

$$P_{1D}(k) \propto [1 + (k/k_1)^n]^{-1}, \tag{6.2}$$

where the normalisation is chosen to ensure unit variance. The additional skewers are then scaled by a common factor in order to control the variance in the extra power added. This factor is allowed to vary along the length of the skewers, effectively adding a redshift-dependency to the extra power. Hence, we write this factor as $\sigma_\varepsilon(z)$. The parameters n and k_1, as well as the function $\sigma_\varepsilon(z)$ are free, and we choose them according to the process described in § 6.3, aiming to achieve the correct 1D power spectrum across a range of redshifts. The new skewers are then simply added to each of the existing ones to form our final Gaussian skewers δ_G:

$$\delta_G(z,\mathbf{x}) = \delta_C(\mathbf{x}) + \sigma_\varepsilon(z)\delta_\varepsilon(\mathbf{x}). \tag{6.3}$$

The top panel of Figure 6.1 shows a sample skewer before and after the extra small-scale power is added. As the additional skewers are independent from one another, there are no correlations between the structures added to each of the skewers. When we measure the 3D correlation

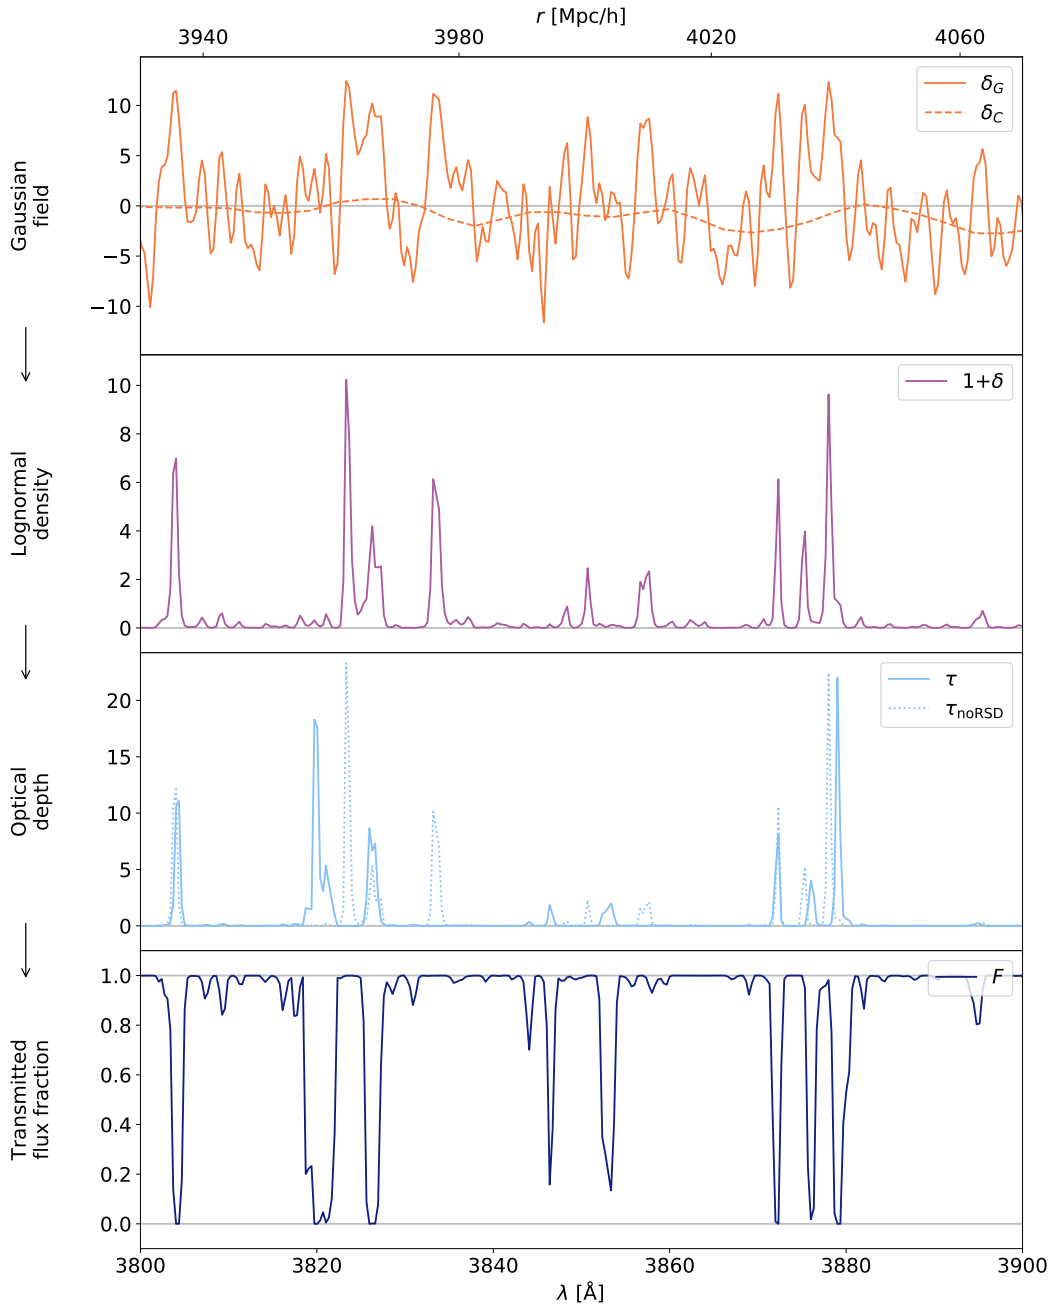

Figure 6.1. A sample skewer shown at the different stages of transformation from "raw" Gaussian `CoLoRe` output to a final `LyaCoLoRe` flux skewer. The top panel shows the addition of small-scale power to the skewer as described in § 6.2.2, converting δ_C (dashed) to δ_G (solid). The transition to the second panel shows the lognormal transformation from § 6.2.2, and moving to the dotted line of the third panel shows the fluctuating Gunn-Peterson approximation (FGPA) transformation from the same section. The application of redshift-space distortions (RSDs), as described in § 6.2.2, shifts the dotted line to the solid line in this third panel. The final transformation from optical depth to flux, as described in § 6.2.2, maps the third to the bottom panel. Here, the Hubble flow is used to map distances (top horizontal axis) to observed wavelengths (bottom horizontal axis).

function, we ignore contributions from pixel-pairs in the same skewer and so this process of adding small-scale power will not affect the 3D correlations of the Gaussian field beyond simply adding noise.

It is worth noting that we could have chosen to add extra small-scale fluctuations to the velocity field and achieved the same correct 1D power spectrum. However, allowing parameters describing extra small-scale velocities to vary freely would require the re-computation of the redshift-space distortions weights matrix (see p.113) at each step of the tuning process (as described in § 6.3). This is a considerably more time-consuming procedure than simply carrying out the inverse Fourier transform of eq. (6.2). As such, we choose to only add small-scale fluctuations to the Gaussian field and assign to each of the small cells the velocity of the nearest large CoLoRe cell.

Transformation to optical depth

In LyaCoLoRe, the transformation from skewers of the Gaussian field to ones of optical depth is governed by two equations. The first of these is known as a lognormal transformation. This approximates the density of the baryonic matter field closely by using a lognormally-distributed variable (Bi and Davidsen, 1997), introducing a degree of non-linearity. This is normalised so that we may define a zero-mean deviation δ from the mean density as:

$$1+\delta(z,\mathbf{x}) = \frac{\rho(z,\mathbf{x})}{\overline{\rho}(z)} = \exp\left[D(z)\delta_G(z,\mathbf{x}) - D^2(z)\frac{\sigma_G^2(z)}{2}\right], \tag{6.4}$$

where $D(z)$ is the linear growth factor at redshift z; $\delta_G(z,\mathbf{x})$ is the Gaussian field value from eq. (6.3); $\sigma_G(z)$ is the standard deviation of this Gaussian field, and $\rho(z,\mathbf{x})$ is the lognormal density at redshift z and position $\mathbf{x}$. This transformation is shown by the transition from the top to the second panel in Figure 6.1.

The second equation allows us to transform these deviations in density into an approximation of the optical depth at each point. Assuming adiabatic expansion implies a tight relationship between temperature and density of the form $\mathrm{d}\ln T/\mathrm{d}\ln\rho = \gamma - 1$ (Hui and Gnedin, 1997). If we further assume photoionization equilibrium, the temperature of the gas approximately determines the number of neutral hydrogen atoms $n_{\mathrm{HI}} \propto \rho^2 T^{-0.7}$ for a given baryonic matter density ρ (Hui, Gnedin, and Zhang, 1997). As the optical depth τ is proportional to n_{HI} (Gunn and Peterson, 1965, see § 3.2.1), these two assumptions allow us to provide an approximation for τ known as the

fluctuating Gunn-Peterson approximation (FGPA, Bi and Davidsen, 1997; Croft et al., 1998):

$$\tau(z,\mathbf{x}) = \tau_0(z)[1+\delta(z,\mathbf{x})]^{\alpha(z)}, \tag{6.5}$$

where $\tau_0(z)$ is a normalisation determined by the gas temperature and the photoionisation rate, and $\alpha(z) = 2 - 0.7[\gamma(z) - 1]$ is determined by the temperature-density relation. These parameter functions $\tau_0(z)$ and $\alpha(z)$ are free, and the method for choosing them is described in § 6.3. The transformation to optical depth is shown by the transition from the second panel to the dotted line of the third panel in Figure 6.1.

Adding redshift-space distortions

The Lyα forest exists as a sequence of absorption features due to the gradient in the recessional velocity of the IGM caused by the Universe's expansion. Features are redshifted according to their distance from the observer, appearing in a spectrum at an observed wavelength $\lambda_{\rm obs} = \lambda_\alpha(1+z)$ for λ_α the Lyα wavelength, and z the absorption redshift. However, peculiar velocities in a region of gas cause its redshift to differ from that due to expansion alone. These effects are known as redshift-space distortions (RSDs), and can be induced by a number of different effects. In particular, RSDs due to gravitationally-induced linear velocities in the IGM are calculated by `CoLoRe`: as mentioned in § 6.2.1, it produces velocity skewers quantifying this effect by calculating the gradient of the Newtonian gravitational potential.

The transition from real- to redshift-space in each skewer can be thought of as an integral over velocity space of the real-space optical depth field multiplied by a kernel K:

$$\tau(s) = \int \tau(x) K\Big(s - x - v_r\big[x|T(x)\big]\Big)\mathrm{d}x, \tag{6.6}$$

where x and s are velocity coordinates along the skewer in real- and redshift-space respectively, v_r is the radial peculiar velocity, and T is the temperature. The choice of K depends on the complexity of the physical effects that you wish to capture. Choosing a suitable Gaussian kernel allows the inclusion of thermal broadening effects: the apparent spreading of the gas's optical depth contribution in redshift-space due to random thermal velocities of the gas atoms. This is implemented as an option within `LyaCoLoRe`, the details of which are described in Appendix B.2. However, we find that the width σ_v of this Gaussian kernel is often smaller than the typical cell size used in `LyaCoLoRe` when adding small-scale fluctuations. Thus, the net effect of accounting

for this physical process is small, and so for the purposes of this chapter we choose the most straightforward option, setting $K(x) = \delta^D(x)$ for δ^D the Dirac delta function. This shifts the optical depth along each skewer according to the peculiar velocity, and does not attempt to include any further physical effects.

In order to implement eq. (6.6), we determine a matrix of weights W_{ij} for each skewer to map its real-space cells τ_j^x to redshift-space cells τ_i^s via the matrix equation $\tau_i^s = W_{ij}\tau_j^x$. The matrix W_{ij} depends on the velocities in the skewer as well as the choice of kernel K, and the details of its calculation can be found in Appendix B.2. Our implementation conserves the integrated optical depth along each line of sight (ignoring pixels which are shifted to un-observed wavelengths). The matrix W_{ij} is near-diagonal and filled mostly by zeros. It can thus be stored in the form of a sparse matrix, and applied to any additional absorption transitions (see § 6.5.2), reducing both the computation time and memory requirements of adding RSDs to the skewers.

The addition of RSDs (without thermal broadening) to a sample optical depth skewer is shown by the transition from the dotted to the solid line in the third panel of Figure 6.1.

Final transmission skewers

In one final stage, we convert from skewers of optical depth τ to transmitted flux fraction F via the equation:

$$F(s) = \exp\big[-\tau(s)\big], \tag{6.7}$$

and interpolate onto a wavelength grid of the user's choice to obtain $F(\lambda)$, where $\lambda = \lambda_\alpha(1+z)$. These skewers are then written to disc.

This final transformation can be seen in the transition between the solid lines in the third and fourth panels of Figure 6.1. It is worth noting that, while the signal in the lognormal density deviation $1+\delta$ and optical depth τ skewers is dominated by over-dense regions, the signal in flux F becomes saturated (equal to 0) at these points and does not carry a great deal of information. Rather, the intermediate density regions — where the density is high enough to cause some absorption but not so high that saturation occurs — are those from which the most information can be gleaned.

6.2.3 Computational requirements

In the realisations presented in this chapter, we specify that CoLoRe generates a 4096^3 cell box as a compromise between resolution and memory usage, given the large volume that we must cover in order to realistically represent a Lyα forest survey. Generating approximately 7.5M QSOs (across

the whole sky) and drawing subsequent skewers produces a dataset sufficient for a DESI-like survey, allowing for a significant degree of flexibility in the final survey strategy and number densities. The computational cost of producing one such dataset is relatively low, provided suitable multi-node, multi-core computational facilities are available. Running `CoLoRe` using the input data and options specified in § 6.4.1 in parallel across 32 Haswell compute nodes (each with 32 cores and 128GB of memory) on the National Energy Research Scientific Computing Centre's *Cori* machine requires approximately 18 minutes to run, equivalent to approximately 300 CPU hours. The large number of nodes is necessary to improve the speed of the code and to satisfy its memory requirements — a total of approximately 920 GB is needed for each run of this size. If such facilities are not available, then the box size must be reduced or the resolution lowered.

The precise requirements for running `LyaCoLoRe` depend strongly on the exact choices of input options. As an example, converting 800,000 skewers — similar to the number that will be observed by DESI — from `CoLoRe`'s Gaussian output to realistic transmission skewers including RSDs (though not thermal broadening effects) requires only 4 minutes when spread across the same 32 nodes mentioned previously. If such computational facilities are not available, then running `LyaCoLoRe` is still possible as its memory requirements are much lower than `CoLoRe`.

A very small test dataset of 1000 skewers is available within the `LyaCoLoRe` repository. It is straightforward to run `LyaCoLoRe` on this data on any standard laptop to generate sample skewers or to explore the functionality of the code.

6.3 Parameter tuning

A number of parameters are defined in the various transformations described in § 6.2.2, namely n, k_1, $\sigma_\varepsilon(z)$, $\tau_0(z)$ and $\alpha(z)$ (see eq. (6.2), eq. (6.3) and eq. (6.5) for definitions). These are all free parameters, and we would like to be able to choose their values so that our final skewers have particular properties. Specifically, we aim to match the 1D power spectrum $P_{\rm 1D}(k,z)$, mean transmitted flux fraction $\bar{F}(z)$ and large-scale Kaiser model bias $b_{\delta,F}(z)$ (Kaiser, 1987, see eq. (6.10) and surrounding discussion for further details) to literature values. Ignoring RSDs and the shape of the 1D power spectrum would allow the problem to be treated analytically, but unfortunately such simplifications are unrealistic. As such, it is not obvious how to choose our parameters correctly, and a more complex process is necessary.

We aim to solve this problem via a minimisation procedure. We first define a function that we will aim to minimise, and which takes the following steps:

1. Generate sample skewers in F corresponding to a given set of parameter values using the methods described in § 6.2.2.

2. Measure the 1D power spectrum, mean flux and large-scale bias of these skewers at a selection of redshift values.

3. Evaluate the deviation of each measurement at each redshift from literature results.

4. Quantify this deviation with a single number.

In step 2, we measure $P_{\rm 1D}$ and $\bar{F}$ straightforwardly, excluding cells that sit at a rest frame wavelength above 1200 Å. We measure $b_{\delta,F}$ by calculating the response of $\bar{F}$ to a small deviation in the average density field: $b_{\delta,F} = (1/\bar{F})\,\mathrm{d}\bar{F}/\mathrm{d}\delta$ (McDonald, 2003). The literature values referred to in step 3 are the fitting function from the BOSS DR9 $P_{\rm 1D}$ measurement from Palanque-Delabrouille et al. (2013), the fitting function of the mean flux measurement from Becker et al. (2013) and the bias value and redshift evolution determined by the BOSS DR12 combined Lyα auto- and cross-correlation analysis in du Mas des Bourboux et al. (2017). Using these literature results as targets, we compute a weighted error for each measurement at each redshift value. When computing the error on the $P_{\rm 1D}$, we prioritise the low-k modes by using a k-dependent error weighting. For $k < 0.02$ s km^{-1} — where the units indicate that we are working with wavenumbers in velocity space — this is proportional to $1/(1+(k/k_0)^2)$ where $k_0 = 0.01$ s km^{-1}. This ensures the modes most relevant for a BAO analysis — those with $k \lesssim 0.005$ s km^{-1} (McDonald and Eisenstein, 2007; McQuinn and White, 2011) — are prioritised over less important, high-k modes. Beyond $k = 0.02$ s km^{-1}, we ignore any errors as our finite cell size makes it unreasonable to expect realistic power at these scales, and these modes were not measured by BOSS. We sum the errors in quadrature over all k-modes using this weighting to produce an overall error on $P_{\rm 1D}$. In step 4, the errors on each measurement at each redshift value are summed in quadrature, and a single number produced. This number quantifies how well a given parameter set is able to produce realistic data, as measured by our specified properties. A standard minimisation routine can then be used to minimise it over the space of input parameters. We use `Minuit` (James and Roos, 1975), as implemented by the python module `iminuit`[4] to do so.

We introduce a number of simplifications to improve the speed of the minimisation. We assume that $\log\tau_0$ and $\log\sigma_\varepsilon$ follow the functional form:

$$\log(X) = \log A_0 + A_1 \log[(1+z)/(1+z_0)], \tag{6.8}$$

[4]Publicly available at https://github.com/iminuit/iminuit.

where $z_0 = 3.0$, the A_i are scalar parameters, and logarithms are calculated to base ten. In the case of $X = \tau_0$, we fix $A_1 = 4.5$ (Seljak, 2012). Further, we assume that $\alpha(z)$ takes a constant value of 1.65 across redshifts (McDonald et al., 2001). This is equivalent to a value of $\gamma = (2-\alpha)/0.7+1$ of 1.5, in reasonable agreement with literature results (e.g. Ricotti, Gnedin, and Shull, 2000; Hiss et al., 2018). With these simplifications, we end up with a five-parameter minimisation problem: one parameter describing the normalisation of $\tau_0(z)$; two describing the normalisation and z-dependence of $\sigma_\varepsilon(z)$; and two describing the shape of the 1D power of the small scale fluctuations (n and k_1). At each call of the routine, we produce sample skewers at a point in parameter space, and compute their 1D power spectra, mean flux and bias parameter values in seven redshift bins of width $\Delta z = 0.2$ centred at points evenly spaced between $z = 2.0$ and 3.2. We run this procedure using $\sim 55,000$ skewers to obtain an initial estimate, and increase this to $\sim 220,000$ skewers in order to fine tune the optimisation.

We also introduce a parameter a_v by which we multiply the velocities in our skewers in order to match the amount of anisotropy in the clustering of the Lyα forest to literature values. This is not because the velocities from `CoLoRe` are incorrect — when using `CoLoRe`'s unmodified velocities, we obtain the correct level of anisotropy in the QSO auto-correlation (see Appendix B.1) — but is a result of the approximations in our recipe to estimate F. We fix $a_v = 1.3$ when tuning; it is computationally costly to leave it free as a change in a_v requires re-computation of the RSD weights matrix W_{ij} (see § 6.2.2). The value is chosen on an ad hoc basis to match approximately the RSD parameter β (defined in § 6.4.3) measured from BOSS DR12 data (du Mas des Bourboux et al., 2017).

The final values of the transformation parameters are $\log[\tau_0(z)] = \log 1.48 + 4.5 \log x$, $\alpha(z) = 1.65$, $\log[\sigma_\varepsilon(z)] = \log 6.02 + 0.276 \log x$, $n = 0.732$, $k_1 = 0.0341$ and $a_v = 1.3$, where $x = [(1+z)/(1+z_0)]$ and numerical values are rounded to three significant figures where appropriate. These are the default values used by `LyaCoLoRe`. The tuning process is effective, matching literature values of P_{1D}, $\bar{F}$ and $b_{\delta,F}$ to within 10% at almost all relevant k-modes and z values. As an example, the P_{1D} measured across ~ 7.5M skewers is shown in Figure 6.2. We only plot 4 redshift bins and a limited number of k-modes here for clearer visualisation.

6.4 Verifying the mocks

The primary motivation for creating the `LyaCoLoRe` mocks is to provide realistic sets of test skewers for BAO analyses from Lyα forest surveys. Evidently then, it is important to verify that

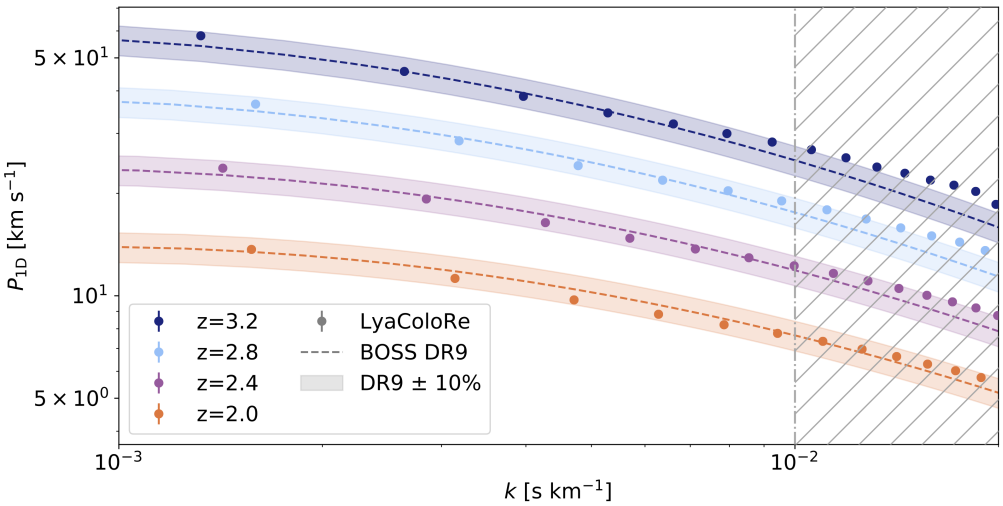

Figure 6.2. The 1D power spectrum as measured from one realisation of LyaCoLoRe mocks. The tuning process aims to match the measured P_{1D} to that from BOSS DR9 data (Palanque-Delabrouille et al., 2013) for k-modes that affect BAO analysis, as described in detail in § 6.3. Modes to the left of the dot-dash line at $k = 0.01$ s km^{-1} are the most important in this respect (McDonald and Eisenstein, 2007; McQuinn and White, 2011), and these all lie within 10% of our target P_{1D}, as indicated by the shaded areas. Modes to the right of the dot-dash line are not prioritised in our tuning procedureand so are not reproduced as accurately, though these are not important in the context of BAO.

the fundamental physical quantities studied by such analyses are correctly reproduced in the mock datasets. We thus seek to test that the BAO signal is present and unbiased in our mock datasets. § 6.4.1 describes the inputs we use in generating a collection of mock datasets; § 6.4.2 explains how we measure the correlation functions from each realisation, taking the skewers in F directly from LyaCoLoRe's output; and finally § 6.4.3 shows how we fit to a model. We do not visually compare the correlation functions measured from mocks to those from data, since our mock measurements are not affected by distortions from continuum fitting, and at this stage we wish to test the methods of LyaCoLoRe rather than techniques to quantify this distorting effect. Instead we compare fitted parameter values in order to assess the performance of our mock datasets.

6.4.1 Generating realisations

The input power spectrum that we use in step 1 of § 6.2.1 is generated by the Boltzmann solver CAMB (Lewis, Challinor, and Lasenby, 2000) using the *Planck* Collaboration's 2015 parameters for a flat, ΛCDM cosmology (see column 1 of Table 3 in Planck Collaboration et al., 2016). We generate the field in a box of 4096^3 cells, stipulating that this covers a redshift range $0.0 \leqslant z \leqslant 3.79$: a volume

large enough to contain a DESI-like survey. This results in a grid of total size $\sim (9.8 \text{ Gpc } h^{-1})^3$, with each cell $\sim (2.4 \text{ Mpc } h^{-1})^3$ in dimensions. The QSO number density function is based on estimates from SDSS-III data in Stripe 82 (Palanque-Delabrouille et al., 2016). This is considered to represent an optimistic estimate of the photometric capability of targeting for DESI, and results in ~ 3.7M QSOs[5] above $z = 1.8$ across the whole sky. We use as an input QSO bias the fitting function defined in equation 19 of Gontcho A Gontcho et al. (2018), which is based on clustering measurements from the BOSS DR12 QSO sample (Laurent et al., 2016). When running `LyaCoLoRe`, we use a cell size of 0.25 Mpc h^{-1}, and tune the parameters of our transformations according to the methods described in § 6.3.

For the purposes of this work we generate 10 such realisations, each with unique random seeds, and stack our results in order to test `LyaCoLoRe` as stringently as possible. This is approximately equivalent to 30 times the final number of Lyα QSOs with $z \geq 2.1$ that will be observed by DESI. It is worth noting that the signal to noise ratio will be significantly greater than 30 times that of DESI, as our skewers of $F(\lambda)$ do not include any instrumental noise, nor do they require any continuum fitting (as mentioned in § 6.2).

6.4.2 Measuring correlation functions

We test the BAO signal in our mock realisations in the standard way, by measuring correlation functions using the contrast in flux transmission:

$$\delta_F(\lambda) = \frac{F(\lambda)}{\bar{F}(\lambda)} - 1, \tag{6.9}$$

where $\bar{F}(\lambda)$ is the mean value of $F(\lambda)$ in each pixel over all skewers for which that cell corresponds to rest-frame wavelength $\lambda_r \in [1040, 1200]$ Å. The skewers of $F(\lambda)$ are taken straight from the processes described in § 6.2, with no further steps such as addition of continua or instrumental noise. This allows us to test the methods of § 6.2 to as high a degree of precision as possible, but consequently our covariance matrices may not necessarily be representative of true measurements. In order to assess the realism of the mocks' covariance matrices, these additional systematic effects must be added to our synthetic data, as was done in du Mas des Bourboux et al. (2020), and as is discussed briefly in § 7.2.

We would like to measure the 3D Lyα auto-correlation and the 3D Lyα-QSO cross-correlation,

[5] This is lower than the 7.5M quoted in § 6.2.3 as we no longer require the previously mentioned flexibility to adapt to different observing strategies in our realisations, and thus can reduce the QSO number density to more realistic values (approximately 59 QSOs per square degree).

the standard measurements made by recent Lyα BAO analyses from BOSS and eBOSS. Both are estimated using the `Package for IGM Cosmological-Correlations Analyses (picca)`[6]. We measure these correlations separately in 3,072 `HEALPix` (Górski et al., 2005) pixels on the sky for each of the 10 realisations, and treat the resultant measurements as a set of 30,720 independent subsamples. In order to compute the correlation functions more quickly, we rebin pixels in our final transmission skewers into larger pixels of width 3×10^{-4} log(Å) in log-wavelength. This enables us to use a larger number of skewers and thus reduce our errors, without compromising the large-scale properties of the correlations or incurring large computational costs.

Our computation of the 3D Lyα auto-correlation follows that of recent Lyα forest BAO analyses (Bautista et al., 2017; de Sainte Agathe et al., 2019; du Mas des Bourboux et al., 2020). We first define a grid of bins in parallel and perpendicular separation between pairs of pixels — $r_{\parallel}$ and $r_{\perp}$ respectively — where each bin is 4 Mpc h^{-1} × 4 Mpc h^{-1} in size, and the maximum separation is 200 Mpc h^{-1} in each direction. Pixel pairs are assigned to one of these bins by using a fiducial cosmology to convert from wavelength and angular separations to comoving distances parallel and perpendicular to the line-of-sight. The correlation is then computed as a weighted sum of products of pixel pairs of δ_F within each bin, where the weights depend on a number of factors, as described in eq. (4) and eq. (7) of du Mas des Bourboux et al. (2020). We restrict ourselves to include only contributions from the Lyα absorption in the Lyα region, ignoring delta pixels outside the rest-frame wavelength range [1040,1200] Å. The covariance matrix is estimated straightforwardly by calculating the scatter between our set of 30,720 subsamples.

The 3D Lyα-QSO cross-correlation is also computed in line with recent analyses of BOSS and eBOSS data (du Mas des Bourboux et al., 2017; Blomqvist et al., 2019; du Mas des Bourboux et al., 2020), as a weighted sum of pixels of δ_F within bins of parallel and perpendicular separation. We use the same bin size as in the auto-correlation, but are able to extend our minimum value of $r_{\parallel}$ to -200 Mpc h^{-1} as the pixel-pixel pair symmetry of the auto-correlation is not present in the pixel-QSO pairs of the cross-correlation. As for the Lyα auto-correlation, we restrict the rest-frame wavelength range of our δ_F pixels to [1040,1200] Å, and we estimate our covariance matrix from the scatter between our 30,720 subsamples.

6.4.3 Fitting the correlation functions

Having measured the 3D Lyα auto- and Lyα-QSO cross- correlations, we fit a model to our measurements to obtain the location of the BAO peak and check that no significant shift has been

[6]Publicly available at https://github.com/igmhub/picca.

introduced. We also seek to measure the bias parameters of our tracers: Lyα flux F and QSOs. These are defined by the relationship between the power spectra of the tracers, $P_F(\mathbf{k})$ and $P_{\rm QSO}(\mathbf{k})$, and the power spectrum of dark matter $P(\mathbf{k})$ (Kaiser, 1987):

$$P_F(\mathbf{k}) = [b_{\delta,F} + b_{\eta,F} f\mu^2]^2 P(\mathbf{k}) \tag{6.10}$$

$$P_{\rm QSO}(\mathbf{k}) = [b_{\delta,\rm QSO} + f\mu^2]^2 P(\mathbf{k}). \tag{6.11}$$

Here, the large-scale biases of flux and QSOs are $b_{\delta,F}$ and $b_{\delta,\rm QSO}$. The parameter $b_{\eta,F}$ is the velocity gradient bias of flux, which serves to quantify the effect of RSDs. This is often expressed alternatively using $\beta = f b_{\eta,F}/b_{\delta,F}$. The value of $b_{\eta,QSO}$ is 1 by default as QSOs are conserved under RSDs and so it is held fixed (Kaiser, 1987). This is not the case for Lyα flux F as the mapping to redshift space occurs in the Lyα optical depth τ, which is then related to our observable, F, by a non-linear relationship $F = e^{-\tau}$ (see McDonald et al., 2000; McDonald, 2003; Seljak, 2012, and references therein). The Lyα-QSO cross- power spectrum follows naturally from eq. (6.10) and eq. (6.11) as:

$$P_{F\times\rm QSO}(\mathbf{k}) = [b_{\delta,F} + b_{\eta,F} f\mu^2][b_{\delta,\rm QSO} + f\mu^2] P(\mathbf{k}). \tag{6.12}$$

We fit a model of the correlation functions to each of the measurements individually, and then to both correlations jointly. We use the same models as recent eBOSS analyses (de Sainte Agathe et al., 2019; Blomqvist et al., 2019) but ignore terms relating to systematics not present in our realisations, such as metal absorbers and high column density systems (HCDs). As we do not add continua to our skewers, we need not worry about the distortion of the correlations by the removal of long wavelength modes in the continuum fitting process, as occurs in real analyses. Thus, we do not need to consider distortion matrices, the standard method for taking these effects into account (introduced for the auto- and cross- correlations respectively in Bautista et al., 2017; du Mas des Bourboux et al., 2017). The relevant terms are described using Kaiser models (Kaiser, 1987), as described in § 4.1 of de Sainte Agathe et al. (2019) for the Lyα auto-correlation, and § 5.1 of Blomqvist et al. (2019) for the Lyα-QSO cross-correlation. We use the same cosmology as used to generate the input power spectrum of `CoLoRe` to produce the smooth and peak components of the fiducial model power spectrum.

The fit is carried out leaving free the parameters describing the position of the BAO peak in the

perpendicular and parallel directions:

$$\alpha_{\parallel} = \frac{D_H(z)/r_d}{[D_H(z)/r_d]_{\rm fid}}, \quad \alpha_{\perp} = \frac{D_A(z)/r_d}{[D_A(z)/r_d]_{\rm fid}}, \tag{6.13}$$

where $D_H(z) = c/H(z)$, as well as parameters describing the bias and RSDs of the Lyα-forest, $b_{\eta,F}$ and $\beta_F = f b_{\eta,F}/b_{\delta,F}$. We also leave free 2 parameters that describe the smoothing of the model power spectrum in the parallel and perpendicular directions, which help to account for the effects of the low-resolution of our `CoLoRe` grid. When fitting the Lyα-QSO cross-correlation individually, we fix the value of the QSO bias $b_{\delta,\rm QSO}$ to the input value in order to avoid degeneracies, though when we fit jointly with the Lyα auto-correlation we are able to leave it free.

Having defined our models, the fits are then carried out using `picca`. We fit only on separations $40 < r$ [Mpc h^{-1}] < 160 as the lognormal density approximation used in both `CoLoRe` and `LyaCoLoRe` begins to break down on scales smaller than this, and we are not able to fit the shape of the correlation function well at these separations. Further, the QSOs cannot be expected to be correctly clustered on the smallest scales due to the low-resolution of the `CoLoRe` box. To determine an effective redshift of our measurements, we consider pixel-pixel/pixel-QSO pairs which fall in bins A which satisfy $80 < r_A$ [Mpc h^{-1}] < 120, i.e. the bins that cover the BAO peak. The value of $z_{\rm eff}$ is then given by a weighted average of the redshifts of pairs in these bins.

The measured Lyα auto- and Lyα-QSO cross-correlations are shown in the left and right panels of Figure 6.3 respectively, along with the model from the combined fit[7]. We plot the correlations as $r^2\xi(r)$ in bins of $|\mu| = |r_{\parallel}|/r$, where $|\mu|$ close to 0 indicates correlations close to perpendicular to the line of sight, and $|\mu|$ close to 1 indicates correlations close to parallel to the line of sight. In each $|\mu|$ bin, we may clearly see the BAO peak at a separation of ~ 105 Mpc h^{-1}, on top of a broadband correlation. The anisotropy indicated by the difference between μ bins arises as a result of RSDs, which affect both the broadband and the peak of the correlation (e.g. Kaiser, 1987; Kirkby et al., 2013) due to the addition of apparent clustering in the parallel direction. The model appears to be a good fit to the measurement on the scales that we fit over, and the BAO peak is correctly placed. The two measurements deviate slightly from the model either side of the BAO bump in the highest $|\mu|$ bin, but this deviation is very small and is noticeable due to the extremely small error bars on our measurements.

The parameters from the individual and combined fits are shown in Table 6.1. The BAO peak location parameters $\alpha_{\perp}$ and $\alpha_{\parallel}$ for each fit are all consistent with 1 to within 1σ. Any deviation in

[7]Note that error bars are present for all points, but are often exceedingly small and thus obscured by the points themselves.

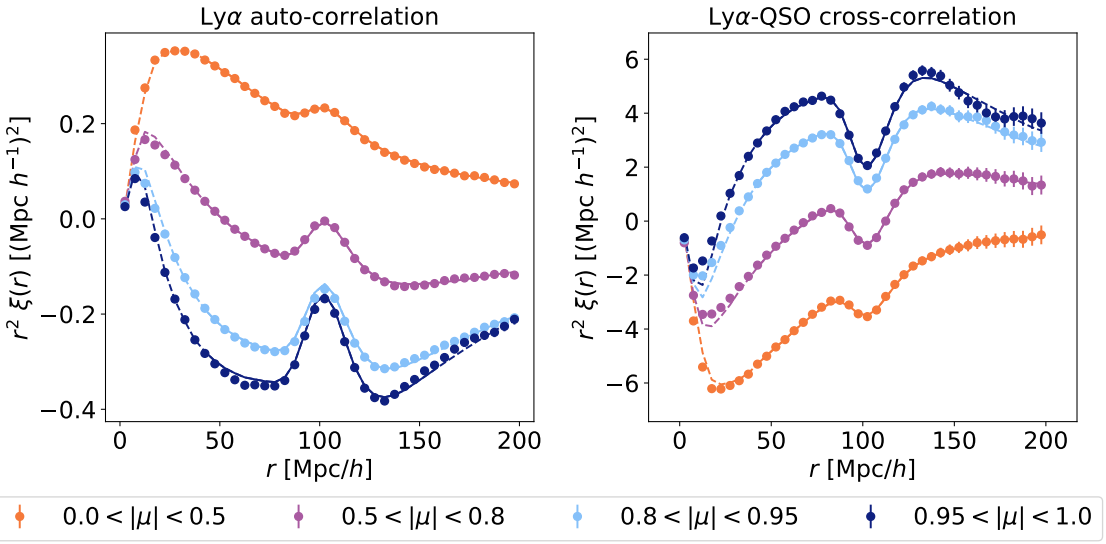

Figure 6.3. Correlation functions measured from 10 realisations of `LyaCoLoRe` datasets combined, and the best fit lines with parameters as described in the third column of Table 6.1. The left panel shows the Lyα auto-correlation, while the right panel shows the Lyα-QSO cross-correlation. Each plot panel shows the same 4 bins in $|\mu| = |r_\parallel|/r$. Note that the correlations presented here do not have any distortion from continuum-fitting and so should not be visually compared with the equivalent plots from recent BOSS/eBOSS data.

the α parameters from 1 is certainly less than 0.2%, and so can be considered insignificant in the context of a DESI-like survey. Thus, the mock production pipeline up to this stage can be said to introduce no clear systematic bias within the capabilities of a current or near-future instrument.

In order to compare the values of biases and βs to BOSS DR12 values (table 4 of du Mas des Bourboux et al., 2017), we first use the published functional forms of each parameter's redshift evolution to match the effective redshift of the BOSS DR12 measurements to that of our measurements. Having done so, we find that the two sets of values are very similar, with our measurements all lying within the 1σ errors on the BOSS DR12 values. In particular, the values of $b_{\delta,F}$ in each of our fits are almost identical to the BOSS DR12 value, demonstrating the effectiveness of our tuning of this parameter (see § 6.3). We do not compare the value of $\beta_{\rm QSO}$ to BOSS DR12 measurements as our input QSO bias takes a different value at this redshift. However, the value of $b_{\delta,{\rm QSO}}$ deduced from the joint fit is consistent with the input value (as shown at the bottom of the column showing the Lyα-QSO only fit). As such, we can consider the mocks to fulfil the basic criteria required of them, and thus they appear sufficient for a DESI-like survey. We do not assess the χ^2 of the fits as we do not expect our covariance matrices to be representative of those one would expect from a real survey given the lack of noise in our skewers.

	LyaCoLoRe			BOSS DR12
Parameter name	Lyα	Lyα-QSO	Lyα + Lyα-QSO	Lyα + Lyα-QSO
$\alpha_{\parallel}$	1.000 ± 0.002	1.001 ± 0.002	1.000 ± 0.001	
$\alpha_{\perp}$	0.998 ± 0.002	1.000 ± 0.002	0.999 ± 0.001	
$b_{\eta,F}$	-0.204 ± 0.0004	-0.201 ± 0.0009	-0.203 ± 0.0004	-0.206 ± 0.012
β_F	1.627 ± 0.008	1.624 ± 0.012	1.624 ± 0.007	1.650 ± 0.081
β_{QSO}		0.261	0.261 ± 0.0007	
$b_{\delta,F}$	-0.121 ± 0.0006	-0.120 ± 0.0009	-0.121 ± 0.0005	-0.121 ± 0.004
$b_{\delta,\mathrm{QSO}}$		*3.701*	3.700 ± 0.009	

Table 6.1. Parameters from model fits of the Lyα auto-correlation and Lyα-QSO cross-correlation functions measured from 10 realisations of LyaCoLoRe mocks combined. The relevant results from the BOSS DR12 combined fit (du Mas des Bourboux et al., 2017) — those to which our values of $b_{\delta,F}$ and $b_{\eta,F}$ are tuned — are presented in the rightmost column at the same effective redshift as our measurements. The parameters in the first segment of the table are those used in the minimisation process which determines the best fit to our correlations, while those in the second segment are calculated subsequently. The value of β_{QSO} is fixed when fitting the Lyα-QSO cross-correlation to avoid degeneracies, with its value chosen such that $b_{\delta,\mathrm{QSO}} = f/\beta_{\mathrm{QSO}}$ matches our input value.

6.5 Adding secondary astrophysical effects

A key purpose of creating mock datasets is to quantify the impact of secondary astrophysical effects on our measurements so that we may assess any biases that they could induce in our cosmological inference. When generating realisations of the synthetic data, we may choose to add or not to add different effects to each realisation, or to vary the strength of a given effect across a range of values. The resultant impact on BAO measurements can then be quantified. In Lyα forest analyses, two of the most pertinent effects are the presence of high column density systems (HCDs) and additional absorption transitions. LyaCoLoRe is able to compute both of these effects, and the methods it uses to do so are described in § 6.5.1 and § 6.5.2 respectively (alternative implementations of these effects are also possible). Once computed, LyaCoLoRe stores skewers of metal absorption and a table of HCDs in its output. These can then be added to the Lyα skewers during subsequent stages of the pipeline by packages such as desisim.

We do not present here a full study of the effects of HCDs and additional transitions on a BAO analysis. Rather, we simply illustrate in § 6.5.3 that their implementations within LyaCoLoRe are broadly correct and achieve the correct levels of clustering. We leave the study of these effects as systematics in a BAO analysis for future work.

In addition to HCDs and metal absorption, there are a number of further astrophysical effects

that we may wish to reproduce in our mock datasets. In particular, the presence of ionising sources such as galaxies may cause variations in the photoionisation rate in the IGM (e.g. McQuinn, Oh, and Faucher-Giguère, 2011; Pontzen, 2014; Gontcho A Gontcho, Miralda-Escudé, and Busca, 2014). Equally, on larger scales, the reionisation of helium at $z \sim 3$ may affect the temperature of the gas in large "bubbles" around ionising sources (McQuinn, Oh, and Faucher-Giguère, 2011; Meiksin and Tittley, 2012). Both of these processes induce spatial fluctuations in the Lyα optical depth that may yield additional contributions to measured correlations, though the impacts are small compared to HCDs and metal absorption, and we do not attempt to include them in our mock skewers at this stage.

6.5.1 Adding HCDs

HCDs occur in particularly dense regions of gas, and contain a number of subclasses determined by HI column density. Typically, we define regions with column density $N_{\rm HI} > 2 \times 10^{20}$ cm^{-2} as damped Lyα absorbers (DLAs), and regions with column density $10^{17.2} < N_{\rm HI} < 2 \times 10^{20}$ cm^{-2} as Lyman limit systems (LLSs, Wolfe et al., 1986). In detailed Lyα forest analyses, it is important to be able to identify HCDs as their high densities broaden their absorption profiles, impacting on inferred values of F over a significant wavelength range (Font-Ribera and Miralda-Escudé, 2012; Rogers et al., 2018). Further, HCDs are of scientific interest in and of themselves (e.g. Pettini et al., 1997; Prochaska et al., 2002; Padmanabhan, Choudhury, and Refregier, 2016; Pérez-Ràfols et al., 2018b; Pérez-Ràfols et al., 2018a). As such, being able to add HCDs to our mocks is important in maximising their realism.

We first determine potential HCD locations by computing a threshold value of the Gaussian field, set by an input bias $b_{\rm HCD}(z)$. In our realisations, we choose $b_{\rm HCD}(z) = 2.0$ to be constant with redshift, and in line with Pérez-Ràfols et al. (2018b). This picks out peaks in the field that are sufficiently dense to host an HCD. We then Poisson sample the potential locations according to an input number density $n_{\rm HCD}(z)$. This number density is imported from the default model of the IGM physics package `pyigm`[8] (Prochaska et al., 2017), which is fitted to a selection of literature results (summarised in Table 1 of Prochaska et al., 2014). The sampling is carried out before adding small-scale power (see § 6.2.2), as we would like the HCDs to correlate with the 3D fluctuations rather than the 1D extra power. A column density is then allocated to each HCD using a given redshift distribution — again from the default model of `pyigm` — and a radial velocity is determined using `CoLoRe`'s output. The resulting catalogue of HCDs can then be interpreted by

[8] Publicly available at https://github.com/pyigm/pyigm.

a package such as `desisim`, which is able to calculate the absorption profile of the HCD using a Voigt template, and insert it into the final spectrum.

6.5.2 Including additional absorption transitions

As with HCDs, absorption from additional transitions are an important level of detail to add to our mocks and are of significant scientific interest (e.g. Pieri et al., 2014; Blomqvist et al., 2018; Gontcho A Gontcho et al., 2018; du Mas des Bourboux et al., 2019). Additional transitions have a rest-frame absorption wavelength different to that of Lyα, and so absorption from gas at the same redshift appears at different observed wavelengths in spectra. Conversely, absorption from two different transitions can appear at the same observed wavelength even though the regions of gas hosting the absorbers are far apart physically. As a result, the presence of such absorption transitions acts to contaminate our measurements of Lyα flux, and thus our correlation functions and resultant BAO measurements. Such transitions include Lyman-β (Lyβ), as well as from silicon, oxygen and carbon gas, for example.

Similar to the method to add HCDs described in § 6.5.1, it would also be reasonable to place additional absorption transitions using a Poisson-sampled "density-peak" approach, as metals are typically produced in high-density regions of the Universe. However, we choose to follow the methods of previous works (Slosar et al., 2011; Bautista et al., 2015), assuming that the optical depth of each additional transition is proportional to that of the Lyα absorber. In the context of these mocks, the most important feature of these additional transitions that we seek to replicate is the strength of their 3D clustering, as it is this that will quantify any impact upon BAO measurements. In order to do so, we simply require an absorption strength (relative to Lyα) and a rest frame wavelength for each additional transition that we wish to include. Having calculated the skewers of optical depth in real space, we scale them differently for each absorption transition according to the transition's relative strength. For an additional transition X, we obtain $\tau_X = A_X \tau_\alpha$, where A_X is the relative strength and τ_α is the Lyα optical depth as defined in eq. (6.5). We then apply RSDs (using the same weights matrix W_{ij} as for Lyα), and convert to $F_X(\lambda)$ separately for each X according to its rest frame wavelength. For each line of sight, the separate $F_X(\lambda)$ skewers are then interpolated onto a common wavelength grid and are combined multiplicatively.

This method ensures that RSDs are correctly applied to each additional absorption transition, and we may tune the absorption strength in order to achieve the correct large-scale bias — and thus the correct 3D clustering — for each transition. A small selection of additional transitions and their relative strengths are shown in Table 6.2. These are the transitions most important to a Lyα BAO

Name	Rest frame wavelength [Å]	Relative absorption strength
Lyα	1215.67	1.0
Lyβ	1025.72	0.1901
SiII (1260)	1260.42	3.542×10^{-4}
SiIII (1207)	1206.50	1.8919×10^{-3}
SiII (1193)	1193.29	9.0776×10^{-4}
SiII (1190)	1190.42	1.28478×10^{-4}

Table 6.2. Details of a small selection of additional absorption transitions that can be used in `LyaCoLoRe`. The absorption strength for each absorber X has been tuned to match approximately the bias value $b_{\delta,X}$ found in literature (Bautista et al., 2017; du Mas des Bourboux et al., 2017; de Sainte Agathe et al., 2019). It is possible to add more absorbers straightforwardly, but the absorption strengths have not been calibrated beyond those listed above. These absorbers are those included in the skewers from which the correlation function in Figure 6.4 is measured.

analysis, though further transitions can be added straightforwardly if needed. These strengths have been tuned to approximately match bias values presented in the literature (Bautista et al., 2017; du Mas des Bourboux et al., 2017; de Sainte Agathe et al., 2019).

6.5.3 Testing astrophysical effects

We assess the methods of § 6.5.1 and § 6.5.2 by first computing the 3D Lyα-HCD cross-correlation. The methods used to do so are largely the same as used to compute the 3D Lyα-QSO cross-correlation, as described in § 6.4.2. One significant difference is that we restrict the HCDs in our calculation of the Lyα-HCD cross-correlation to lie in the rest frame wavelength range $[1040, 1100]$ Å, far from the background QSO. This restriction is necessary to prevent the correlation between Lyα flux and QSOs from significantly affecting our measurements close to the line of sight, as is discussed further in Appendix B.3. An effect can still be seen in the two μ-bins closest to the line of sight at large values of r, though this is mostly beyond the fitted range and so we are still able to measure the degree of clustering in the HCDs well. Future studies may prefer to model this effect in order to avoid reducing the HCD catalogue in this way, but such work is beyond the scope of this analysis. As in § 6.4, we measure correlations on each of our 10 realisations, and combine the measurements.

The measurement of the Lyα-HCD cross-correlation is shown in the right panel of Figure 6.4. Here, we fit for a model in the same way as for the Lyα-QSO cross-correlation. Carrying out a combined fit with the Lyα auto-correlation from § 6.4 allows us to measure the HCD bias

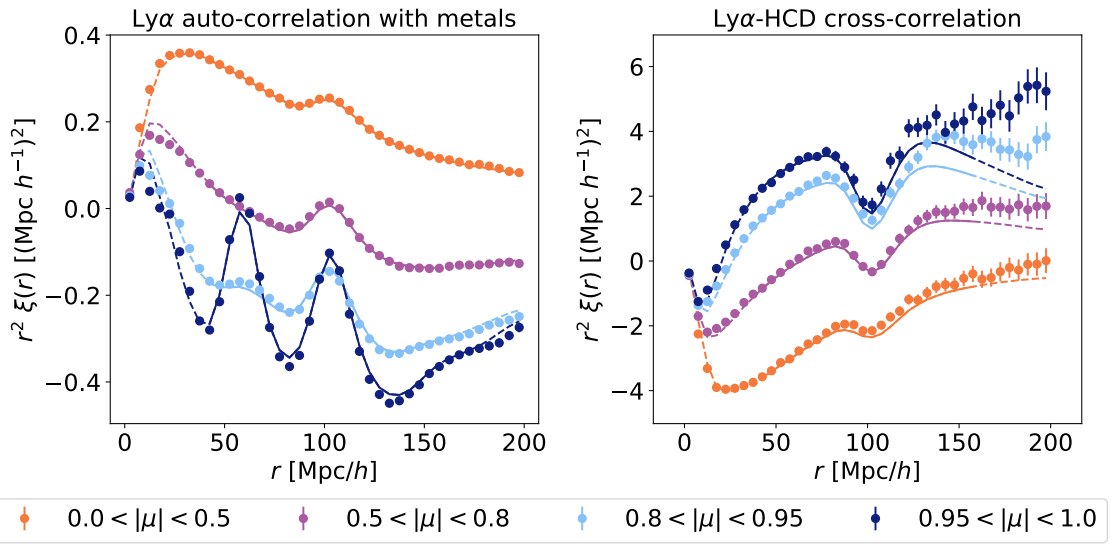

Figure 6.4. Correlation functions measured from 10 realisations of LyaCoLoRe combined, demonstrating the additional astrophysical effects that can be included in its skewers. The left panel shows the flux auto-correlation measured from skewers including metal absorption, from which we measure the metal absorber biases presented in Table 6.3. The right panel shows the Lyα-HCD cross-correlation. The subtleties of this measurement — including the discrepancy at large-r — are discussed in Appendix B.3. Note that the correlations presented here do not have any distortion from continuum-fitting and so should not be visually compared with the equivalent plots from recent BOSS/eBOSS data.

$b_{\delta,\mathrm{HCD}}(z_{\mathrm{eff}}) = 2.26 \pm 0.02$. Strictly, this is not consistent with the redshift-constant input value of $b_{\delta,\mathrm{HCD}} = 2.0$ (as motivated by Pérez-Ràfols et al., 2018b). There are a number of potential reasons for such a shift, but given the errors on current measurements of $b_{\delta,\mathrm{HCD}}$ from data (approximately 10% in Pérez-Ràfols et al., 2018b), we do not investigate the agreement further at this stage.

We then compute the 3D auto-correlation from skewers of F that include contributions from the additional absorption transitions in Table 6.2 (on top of Lyα absorption). The method used to do so is identical to that described for the 3D Lyα auto-correlation in § 6.4.2. We only include contributions from pixels that lie in the rest-frame wavelength range $[1040, 1200]$ Å, and so we do not include any absorption from the Lyβ absorber as its rest-frame wavelength is below this range. As such, from here on we refer to the additional absorption transitions as "metals". As in § 6.4, we measure correlations on each of our 10 realisations, and combine the measurements.

The measurement of the auto-correlation with metal absorbers is shown in the left panel of Figure 6.4. By comparison with Figure 6.3, the effect of including these metals in the skewers is clearly significant, particularly in the near-line of sight $0.95 < |\mu| < 1.0$ bin. Notably, we can

	bias $\times 10^3$		
Absorber	LyaCoLoRe	BOSS DR12	eBOSS DR14
SiII (1260 Å)	-1.70 ± 0.04	-1.5 ± 1.2	-2.5 ± 1.3
SiIII (1207 Å)	-3.3	-3.3 ± 1.3	-8.2 ± 1.0
SiII (1193 Å)	-3.28 ± 0.03	-3.5 ± 0.9	-4.6 ± 1.0
SiII (1190 Å)	-4.55 ± 0.03	-4.4 ± 0.9	-5.1 ± 1.0

Table 6.3. The biases of the metal absorbers measured from our realisations of LyaColoRe, along with the values from BOSS DR12 (Bautista et al., 2017) and eBOSS DR14 (de Sainte Agathe et al., 2019) for comparison. The bias of SiIII (1207) is held fixed to the DR12 value as the "bump" that it creates in the correlation function is at $r = 21$ Mpc h^{-1}, below the minimum separation of 40 Mpc h^{-1} used in our fits. The values from LyaColoRe are all within 1σ of those from BOSS DR12, indicating that the absorption strengths used in our realisations (see Table 6.2) result in the correct levels of large-scale clustering.

clearly see a peak at approximately 55-60 Mpc h^{-1} as a result of SiII (1190 Å) and SiII (1193 Å) absorption, as well as a peak at approximately 21 Mpc h^{-1} from SiIII (1207 Å). This final peak is not included in our fit as it is below the minimum separation. Less visually obvious, but more important to the BAO analysis, is the effect of absorption from SiII (1260 Å), which causes a bump at 105 Mpc h^{-1}, very close to the BAO peak.

In our fit of this correlation, we model the effect of metal absorbers in the same way as in de Sainte Agathe et al. (2019), summing contributions to the model power spectrum from each combination of pairs of absorbers. In Table 6.3, we show the biases for each of our metal absorbers, as well as the values from BOSS DR12 (Bautista et al., 2017) and eBOSS DR14 (de Sainte Agathe et al., 2019) for comparison. The bias of SiIII (1207 Å) is held fixed to the DR12 value as the peak that it creates in the correlation function is at $r = 21$ Mpc h^{-1}, below the minimum value used in our fits. The LyaColoRe values sit within 1σ of those from BOSS DR12, demonstrating that the levels of clustering given by the absorption strengths in Table 6.2 are similar to those found in data. Of course, each absorption strength can be tuned further so that the bias of the relevant absorber more closely matches any given value.

6.6 Summary & conclusions

In this chapter we have presented LyaCoLoRe, a tool for creating mock Lyα forest datasets when used in conjunction with a Gaussian random field code such as CoLoRe. We first use CoLoRe to generate skewers from a Gaussian random field, avoiding the use of N-body or hydrodynamical simulations due to the limited volume and high computational cost of such methods. LyaCoLoRe is

then able to transform the output into realistic skewers of transmitted flux fraction, with a number of properties defined by an automatic tuning process. The process is computationally efficient, making it suitable for generating large numbers of realisations of mocks with different input data and parameters.

We then demonstrate the effectiveness of LyaCoLoRe's output, generating a number of skewers equivalent to approximately 30 realisations of DESI and measuring the Lyα auto- and Lyα-QSO cross-correlations. Fitting these measurements with an appropriate model gives BAO peak positions that are consistent with the input cosmologies to within 0.2%, and certainly within the capabilities of an instrument such as DESI. In addition, the biases of the Lyα forest and of QSOs are shown to be very similar to those derived from BOSS DR12 data. As such, we conclude that the mock datasets generated by LyaCoLoRe are suitable for the BAO analyses of current and upcoming surveys such as eBOSS and DESI.

Finally, we demonstrate two additional capabilities of the LyaCoLoRe package in adding correlated high column density systems (HCDs) and additional absorption transitions to the skewers. We leave a full analysis on the impact of such features on a BAO analysis to a future work, but demonstrate that the HCDs are clustered approximately correctly on large scales, and that the additional transitions affect the Lyα auto-correlation in the expected manner.

Mock datasets such as those generated by LyaCoLoRe are of use to the BAO analyses of Lyα forest surveys in a number of ways. They are able to provide robust tests of analysis pipelines, while they can also help in assessing the impact of astrophysical effects — such as HCDs and additional absorption transitions — on BAO measurements. Finally, they can be used to provide evidence when making decisions regarding the planning of large surveys, such as in targeting and survey strategy. As such, we hope that LyaCoLoRe will be of use for Lyα BAO surveys both present and future.

Part IV

Applications, conclusions and further work

CHAPTER 7

APPLICATIONS

Thus far in this thesis, we have discussed strategies, techniques and datasets of direct relevance to spectroscopic QSO surveys present and future. With the eBOSS survey having recently concluded, and the first spectra from DESI having been taken in early 2020, we now present applications of the ideas introduced in earlier chapters to these datasets, both as part of ongoing work and from existing publications. In § 7.1, we discuss the initial results of applying the QSO selection strategies of Chapter 5 to DESI's first QSO target data. We provide a tentative insight into the classifiers' future performance levels, and suggest some areas for investigation once further data is available. Then, in § 7.2, we present the use of `LyaCoLoRe` mock datasets during the Lyα BAO analysis from the final eBOSS data release, highlighting the verification of the analysis pipeline that they enabled.

Finally, we go on to discuss the application of BAO results to constrain cosmological parameters, focusing on the Hubble parameter H_0. BAO is of particular use in this context as, in conjunction with data from big bang nucleosynthesis, it is able to constrain H_0 independent of other leading measurements, and may thus help to address the heavily debated "Hubble tension". In § 7.3, we explain how BAO is able to provide such a constraint, before demonstrating this approach by applying it to eBOSS data and DESI forecasts.

The discussion in subsequent sections is based partly on ongoing work, and partly on work presented in publications of which I am a co-author. At the start of each section, I describe in italics the status of the project in question, and describe the contributions I have made to it. In each case, I then summarise important methods and results of the project in my own words.

7.1 Applying QuasarNET to initial DESI data

Between March and September 2020, DESI was shut down due to the coronavirus pandemic, forcing a substantial delay to operations and preventing any observations from taking place. Fortunately, a small amount of data was taken before the shutdown, which has provided a tantalising initial test set for the strategies described in Chapter 5. The number of spectra available is small, but these early tests are still of great interest in understanding DESI as an instrument, as well as the data reduction pipeline and the classification tools available to us.

The work presented in this section is ongoing, and I highlight a number of questions which may be addressed once further DESI data is available. These provide numerous routes to explore ahead of a potential future publication.

7.1.1 Initial DESI data

During March 2020, DESI's "SV0" campaign was carried out, marking the beginning of the very first survey validation (SV) observations. Spectra were taken in ten areas of the sky, known as "tiles", in a range of different locations within DESI's footprint, as shown in Figure 7.1. Targets for observation were selected from data release 8[1] of the DESI Legacy Imaging Surveys (see Dey et al., 2019, for a summary), with three tile chosen to include quasar targets (alongside luminous red galaxy targets): 68000, 68001 and 68002, displayed in red in Figure 7.1. Within these tiles, QSO targets were selected in line with DESI's "baseline" QSO target selection strategy, which uses a random forest to select objects from photometric colour space. Further targets were then added to this to produce an "extended" strategy, in which secondary colour cuts were used to include additional objects, and restrictions were loosened on both the definition of point-source objects and the limiting magnitudes of objects (Yèche et al., 2020). This extended selection will be pursued during DESI's SV phase, and will inform decisions on target selection procedures for the main survey. It includes in its selection a number of fainter objects than in the baseline selection, resulting in lower-SNR spectra and potentially a greater fraction of contaminants. With this in mind, it is worth noting that tile 68002 lies close to the galactic plane, and overlaps the Saggitarius Stream. As such, it is expected that target selection for this tile may contain a greater level of stellar contamination than an average DESI tile, for the baseline and extended selections alike.

Observations of these targets were carried out successfully, with excellent observing conditions yielding above-average SNR in spectra. Subsequently, the spectra were processed by the DESI

[1] Detailed description available at https://www.legacysurvey.org/dr8/description/.

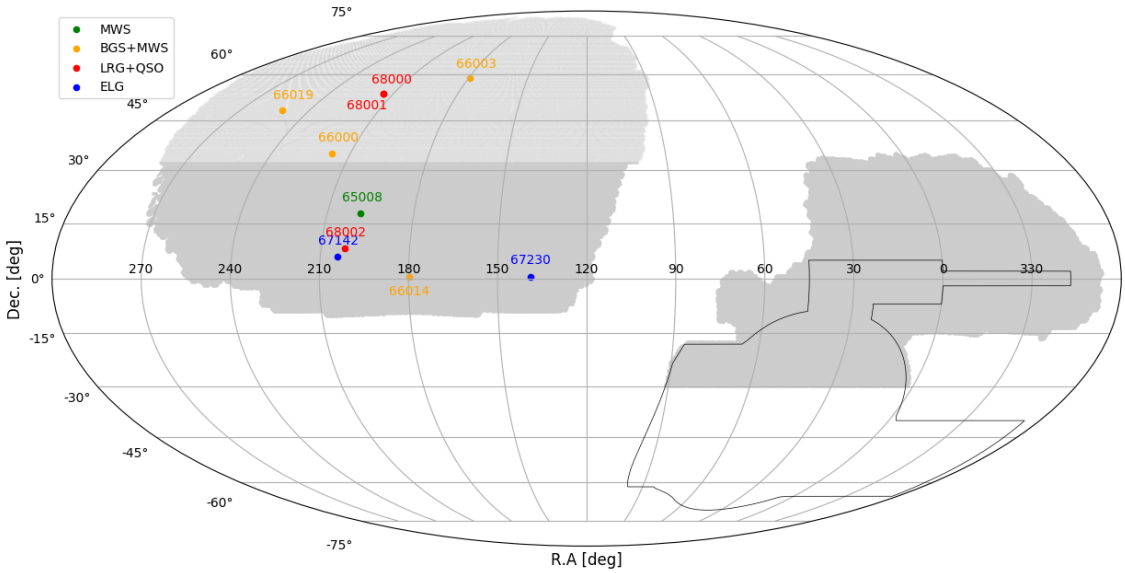

Figure 7.1. Sky map showing the locations of tiles observed by DESI during March 2020 as part of the "SV0" programme. Tiles in which QSOs were targeted are shown in red. Of particular note is tile 68002, spectra from which were put through an extensive visual inspection process, and which we use to test our QSO target classifiers. *Figure credit: DESI Internal*

pipeline, and released to the collaboration. Visual inspection (VI) was then carried out on tile 68002 by a team of DESI volunteers, with QSO experts used to check the classifications and redshifts provided by the team. Final VI results were shared with the collaboration, including a VI redshift confidence between 0 and 4, of which classifications with confidence of at least 2.5 were deemed "confident". From now on, we consider only spectra from tile 68002, and only those which met this VI confidence threshold. This left 696 confidently classified QSO targets, which we could use to test our classifiers and provide further insight into optimal classification strategies. For each QSO target, we could apply our classifiers either to coadded spectra, or to spectra from one of the three exposures taken of this tile.

7.1.2 Classification results

As discussed in Chapter 5, the classification of QSO targets from both single-exposure and coadded spectra are of direct relevance to DESI, for the purposes of selecting high-z QSOs for reobservation and constructing QSO catalogues respectively. When looking to classify the initial DESI spectra from tile 68002, we trained two `QuasarNET` models, one on coadded QSO target spectra from BOSS DR12 and one on randomly-chosen single exposures of the same objects. In each case, approximately 63,000 spectra were used for training, roughly equivalent to the number of QSO

target spectra which will be visually inspected in DESI's SV phase. As such, these models are directly equivalent to those used in § 5.3, the only difference being that the BOSS training spectra were truncated at 9800 Å in this instance, mimicking the upper wavelength limit of DESI. We then applied these two `QuasarNET` models and `redrock` to both the final coadded spectra from tile 68002, as well as the three sets of single exposures, and present our results below.

Coadded spectra

We first use coadded spectra from tile 68002 to provide a brief overview of the performance of `QuasarNET` and `redrock` on DESI data. Here, we use only the 343 spectra from objects selected by the "baseline" target selection procedure, and we use the `QuasarNET` model trained on coadded spectra from BOSS. In Figure 7.2, we show scatter plots in a plane defined by the classifier redshift on the y-axis, and the "true", VI-determined redshift on the x-axis. Results for `QuasarNET` are in the left panel while results for `redrock` are in the right panel. For `QuasarNET`, we classify objects as QSOs if at least one emission line is identified with confidence $c > 0.5$, while for `redrock` we do so if `redrock`'s best fit template is a QSO, and no `ZWARN` flags are raised. These selection strategies are the same as the "QN" and "RR" strategies used in § 5.3.2. For both classifiers, we display in this plane the 264 objects confidently identified by VI as QSOs, colouring the markers according to the success of the two classification strategies. The orange points represent true QSOs which were correctly identified as such, and for which the redshift was correctly estimated. Here, we allow for a deviation of up to 6,000 $\mathrm{km\,s^{-1}}$ from the VI redshift (as used throughout Chapter 5), which we indicate with a shaded grey region around a central dashed line in both panels. QSOs which were missed by each classifier are displayed as pale blue crosses, while QSOs which were correctly identified but received a catastrophic redshift failure are shown as dark blue crosses. In the bottom-right of each panel, we display the number of QSOs in each of these three classes. Further, each of these numbers is calculated for QSOs with low and high VI redshifts separately, with the dividing redshift between these categories set at $z = 2.1$ and indicated by the vertical, dotted grey line. These results are then presented in brackets below the full-sample values.

Clearly, the results presented here demonstrate reduced performance levels compared to those seen on BOSS data. We may calculate the completeness of the sample in each case via the ratio $n_{\mathrm{correct}}/n_{\mathrm{total}}$, yielding values of 94.3% for `QuasarNET` and 84.8% for `redrock`. This is noticeably lower than the equivalent results for BOSS data, where `QuasarNET` achieves 99.4% and `redrock` 91.3%. Of course, these new results based on a sample of only 264 QSOs, and so one would expect variation due to Poisson noise to be $1/\sqrt{N} \simeq 6\%$. As such, these reduced performance levels lie on

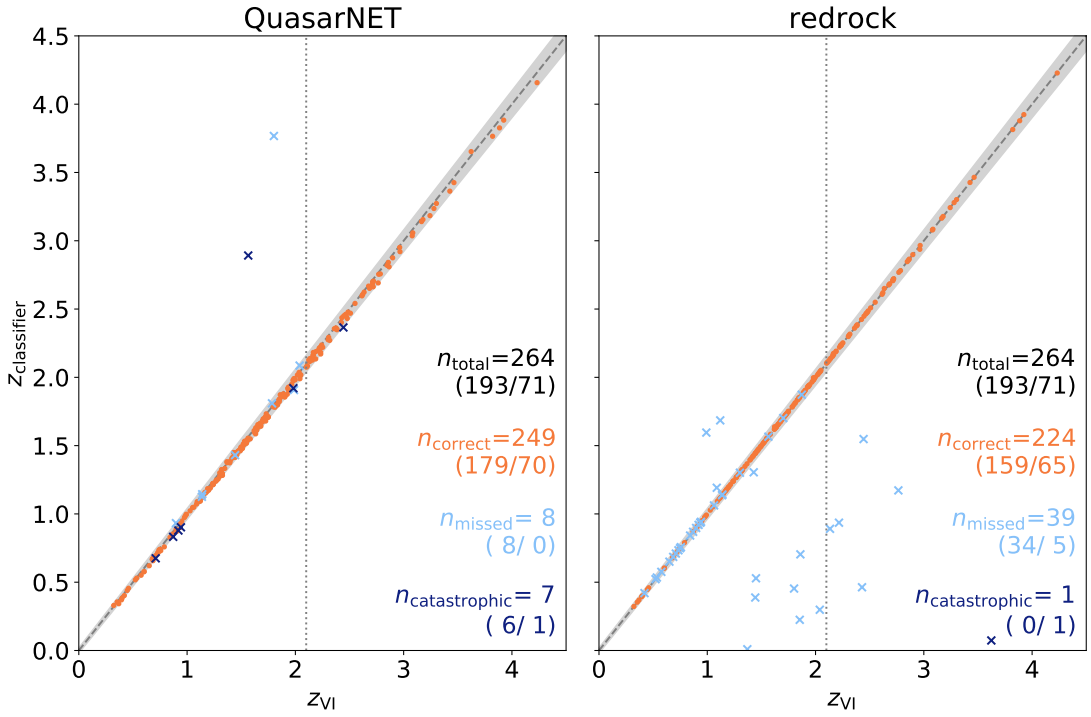

Figure 7.2. QSOs observed in tile 68002 of DESI's SV0 programme under the baseline target selection, shown as points in the $z_{\text{classifier}}$-z_{VI} plane for classifications from `QuasarNET` and `redrock` in the left and right panels respectively. In each panel, QSOs for which the classifier was able to provide a correct classification and redshift are shown as orange points, while those which the classifier missed are represented by pale blue crosses. Those which were correctly identified as QSOs by the classifier but for which the redshift estimate was a catastrophic failure ($\mathrm{d}v > 6,000\ \mathrm{km\,s^{-1}}$) are shown as dark blue crosses. The number of QSOs in each of these categories is shown in the bottom right of each panel, with low-/high-z splits in brackets beneath the overall numbers.

the boundary of significance, and should be further investigated upon the arrival of additional data. For a substantial proportion of these failures, `redrock` was able to provide the correct redshift, but classified the object as a galaxy rather than a QSO. This is likely due to the sparsity of `redrock`'s low-z QSO templates, providing further motivation for the development of improved templates in this redshift range. Equivalently, `QuasarNET` was able to correctly estimate the redshift of the majority of QSOs it missed, but did not identify any emission lines with confidence $c > 0.5$ and so did not classify these objects as QSOs. In addition to these objects, `redrock` also misses a number of QSOs at $1 \lesssim z \lesssim 3$ which for which it is unable to provide an accurate redshift, most of which have underestimated redshifts. This is not the case for `QuasarNET`, for which all bar one of its missed QSOs have accurate redshift estimates.

While `redrock` misses a far larger number of QSOs than `QuasarNET`, it also produces far

fewer catastrophic redshift failures. In particular, at low redshift, `QuasarNET` has six such failures whereas `redrock` has none. Looking at the `QuasarNET` results more closely, it appears that the majority of the catastrophic failures sit just below the lower edge of the shaded grey area, showing that `QuasarNET` underestimates these objects' redshifts with an error of magnitude slightly greater than 6000 $\mathrm{km\,s^{-1}}$. This is indicative of a broader trend within the `QuasarNET` results: we can see a general tendency for `QuasarNET` to underestimate QSOs' redshifts by looking at the orange points, a trend which is particularly visible at $z_{\mathrm{true}} \sim 3$. This bias in `QuasarNET`'s redshifts may be introduced by the difference in wavelength grids in BOSS and DESI data: while BOSS used a grid spaced evenly in $\log(\lambda)$, DESI uses a grid spaced evenly in λ. When training and testing our `QuasarNET` models, we first rebin all spectra onto a wavelength grid that is coarsely sampled in $\log(\lambda)$. For BOSS data, this means that each rebinned pixel receives contributions from the same number of instrument-pixels. For DESI data, however, the number of instrument-pixels contributing to each rebinned pixel increases as a function of wavelength. This may affect the wavelength-dependence of the SNR in the rebinned spectra, perhaps inducing a systematic biasing in `QuasarNET`'s redshift estimation. Certainly, training `QuasarNET` models on DESI SV data once available will help us to understand this issue further.

Single-exposure spectra

We now apply `redrock` and our single-exposure trained `QuasarNET` model to our single-exposure data from tile 68002. We do so separately for each of the exposures, further considering the two target selections described in § 7.1.1 separately as well. We present our results in the same way as in § 5.3.1, as if we are selecting high-z QSOs for reobservation. In each panel, we plot results for the same strategies considered in § 5.3.1. "QN" selects objects identified as $z \geq 2.1$ QSOs by `QuasarNET`; "RR" selects any objects for which `redrock` found a high-z QSO template with $\Delta\chi^2_r < 0.002$ relative to the best fit; and "QN|RR" selects objects identified as such by either of the QN or RR strategies. We apply each of these strategies to our three exposures separately, plotting the results in different shades. For simplicity, we choose to fix the confidence threshold used by `QuasarNET` in the QN and QN|RR strategies at $c_{th} = 0.35$ for the baseline target selection, and 0.4 for the extended target selection. We plot our results in Figure 7.3, quantifying strategies in terms of the fraction of high-z QSOs they select and the number density of fibres that they allocate to reobservations (again, as in § 5.3.1). This is carried out for the two target selections, with results for the baseline DESI selection in the left panel, and the extended selection in the right panel. These two targeting selections yield different number densities of true high-z QSOs, as indicated by the

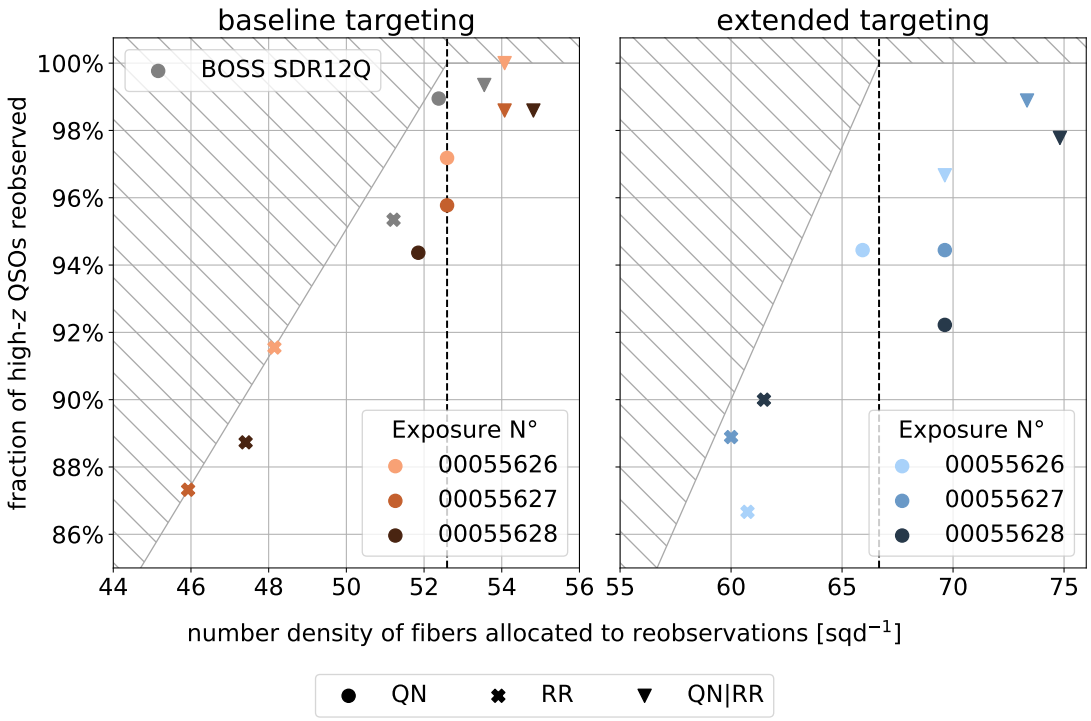

Figure 7.3. Performance of different strategies when selecting high-z QSOs for reobservation using single exposures from tile 68002, and with each exposure represented by a different tone. The strategies presented here are the same as those in Figure 5.2, only with QuasarNET's confidence threshold fixed at $c_{th} = 0.35$ in the left panel, and $c_{th} = 0.4$ in the right. We show the equivalent points from tests on BOSS SDR12Q data in grey in the left panel. The left panel shows results when only considering QSO targets included in DESI's baseline target selection, while those in the right panel also include additional QSO targets that are included in an experimental, "extended" target selection. This includes targets within extended colour cuts, as well as fainter objects that yield lower-SNR spectra. The dashed, vertical, black line here represents the true number density of high-z QSOs for each targeting procedure.

different locations of the vertical, dashed, black lines in the two panels.

Focusing on the left panel of Figure 7.3, we can see once again that performance levels on initial DESI data are somewhat lower than the equivalent results on BOSS data (see Figure 5.2), which are represented by the grey points. This is particularly noticeable for the RR strategy, somewhat less so for QN, and is not apparent for QN|RR. In this panel, our test samples contain 71 high-z QSO spectra for each of the three exposures used, and so we would expect the mean of these three points to vary by $\sim 7\%$ due to Poisson noise alone. As such, these differences are not particularly significant, but again mark points to be investigated upon the arrival of more data in future. Despite this possible drop in performance, the results of Figure 7.3 appear to align with those from Figure 5.2 in their overarching conclusions. In both cases, the QN strategy outperforms RR,

while combining the two strategies in QN|RR yields to the best performance, provided sufficient fibres are available. Indeed, this combined strategy is able to select at least 98.5% of high-z QSOs from DESI's baseline target selection, despite using a `QuasarNET` model which was not trained on DESI data and the limitations of `redrock`'s current targets. This high level of performance should offer further hope that the future versions of `QuasarNET` and `redrock` will be well suited to the task of selecting high-z QSOs for reobservation when combined in this way.

We may also compare the results of this left panel to those in the right panel, which correspond to applying our strategies to spectra in the extended target selection. Our strategies do not appear to perform as well on these spectra, as would be expected: this selection includes spectra from fainter objects, which have a lower average SNR than those included in the baseline strategy. It is thus, on average, more difficult to determine accurate classifications and redshifts for these objects. As mentioned previously, the extended targeting regime also uses different cuts in colour space to the baseline selection, which may impact upon the performance of our strategies beyond just the SNR of our spectra. Certainly, it would be useful to assess the efficacy of these classification strategies across a variety of sub-classes of QSO spectra, divided based on average SNR, colour properties or observing conditions, for example. Of course, the small number of spectra in our sample here are insufficient for these kinds of analyses, and we must wait for the full SV dataset before conducting such studies.

7.2 The use of LyaCoLoRe mocks in eBOSS

When conducting the Lyα BAO measurements from the final eBOSS data release — DR16 — the `LyaCoLoRe` mocks formed a key part of the analysis procedure. Numerous mock realisations were used to study the impact of systematic effects on the analysis pipeline, as well as to verify methods of calculating covariance methods. This use of mocks represented a substantial advancement over previous works, ultimately resulting in a more robustly-tested measurement of the BAO scale.

The full eBOSS DR16 Lyα BAO analysis is described in du Mas des Bourboux et al. (2020), of which I am a co-author. I provided 100 realisations of the `LyaCoLoRe` mocks for this analysis, which were then converted into mock QSO spectra and passed through the analysis pipeline by other members of the eBOSS collaboration. Here, I present in my own words the use of mocks in this analysis.

7.2.1 THE SDSS LYα SURVEY

The BOSS programme (Dawson et al., 2013) of SDSS-III (Eisenstein et al., 2011) and the eBOSS programme (Dawson et al., 2016) of SDSS-IV (Blanton et al., 2017) provided ten years of spectroscopic observations, covering a range of different tracers. In particular, both programmes observed QSOs at $z \geq 2.1$ from which the Lyα forest can be extracted, with eBOSS additionally observing tracer QSOs at $0.8 < z < 2.2$. The final eBOSS data release — DR16 — contained all QSO spectra observed in both BOSS and eBOSS, with the Lyα BAO analysis (du Mas des Bourboux et al., 2020) measuring correlations using Lyα absorption from 210,005 Lyα QSOs at $2.1 < z < 3.5$, and 341,468 QSOs at $z > 1.77$. This analysis measured Lyα absorption from the high-z QSOs in two spectral regions of each QSO: the "Lyα" region between the Lyα and Lyβ-O VI emission peaks, and the "Lyβ" region between the Lyβ-O VI emission peak and the rest-frame Lyman limit. The equivalent DR14 analyses (de Sainte Agathe et al., 2019; Blomqvist et al., 2019) also used the Lyβ region, though it was limited to the rest-frame wavelength range 974–1020 Å, considerably shorter than the range of 920–1020 Å used in DR16. In du Mas des Bourboux et al. (2020), the two sets of absorption regions were referred to as Lyα(Lyα) and Lyα(Lyβ) respectively, and are shown in the example (mock) spectrum of Figure 7.4 as blue and orange lines respectively. The two regions used to measure four different correlation functions:

- Lyα(Lyα)×Lyα(Lyα): the auto-correlation of Lyα absorption in the Lyα region
- Lyα(Lyα)×Lyα(Lyβ): the auto-correlation of Lyα absorption in the Lyα and Lyβ regions
- Lyα(Lyα)×QSO: the cross-correlation of Lyα absorption in the Lyα region with QSOs
- Lyα(Lyβ)×QSO: the cross-correlation of Lyα absorption in the Lyβ region with QSOs

These measurements consituted the first analysis of the full SDSS Lyα QSO sample, providing new insight into the impact of various calibration procedures, improving understanding of the effects of QSO redshift estimators, and expanding the Lyβ spectral region. In addition, it continued the development of Gaussian random field mocks, including those from `LyaCoLoRe`, and it is this use of mocks that we focus on now.

7.2.2 THE MOCKS USED

In conducting the analysis of the eBOSS DR16 Lyα forest dataset, 100 sets of `LyaCoLoRe` mocks[2] were used to validate the analysis pipeline, using v9.0 of `LyaCoLoRe`[3]. For each dataset, as in § 6.4.1, cosmological parameters for a flat ΛCDM Universe from Planck Collaboration et al. (2016) were used as inputs to `CoLoRe`, alongside a corresponding input power spectrum. Equally, the same QSO number density and bias as in § 6.4.1 were used for all runs of `CoLoRe`, and the same transformation parameters were used when running `LyaCoLoRe`, tuned to approximately match BOSS DR12 values for the bias and RSD parameter of the Lyα forest (as described in § 6.3). The only differences between each dataset were the random seeds used, with different seeds chosen for each dataset both when running `CoLoRe` and when adding small-scale fluctuations to the skewers. Thus the 100 mocks represent 100 independent realisations of the same Universe, allowing us to test our analysis methods to a high degree of precision.

While the output skewers of `LyaCoLoRe` are accurate representations of the Lyα forest, they are not sufficient in and of themselves to fully test the analysis pipeline. Indeed, there are several additional layers of realism which may be added in order to provide an approximation of real data:

- *High-column density systems (HCDs):* HCDs in the column density range $\log N_{\rm HI} = [17, 22.5]$ were generated when running `LyaCoLoRe`, correlated with the matter density field as described in § 6.5.1. These were added to the QSO spectra by simulating a Voigt profile for each HCD, and combining multiplicatively with the transmitted flux fraction.

- *Metal absorbers* While metal absorbers were included in the mocks, the method used to do so differed slightly from that described in in § 6.5.2 as the transmitted flux fraction for each absorber was calculated outside `LyaCoLoRe`. First, the Lyα optical depth τ_α was calculated from the transmitted flux fraction via $\tau_\alpha = -\log F_\alpha$, with the optical depth for each absorber then calculated by rescaling τ_α and shifting in wavelength. This is equivalent to rescaling and shifting the optical depth *after* applying redshift-space distortions (RSDs), rather than doing so *before* RSDs, as in § 6.5.2. This post-RSDs method is perhaps less realistic, as the velocity used to determine RSDs in the metal absorption should be calculated from the region where the metal lies. The difference between these methods is comparatively small, however, and at the time the rescaling values in Table 6.3 had not been finalised, so an alternative approach was required.

[2] In du Mas des Bourboux et al. (2020), these mocks are referred to as the "London mocks".
[3] https://github.com/igmhub/LyaCoLoRe/releases/tag/v9.0.

– *QSO continuum:* Each QSO was assigned a random magnitude according to the distributions measured in Ross et al. (2013), up to a limiting magnitude of $r = 21.3$. These were then used to generate unabsorbed QSO continua using the package `simqso`[4], which constructs continua by adding emission lines to a broken power-law (see McGreer et al., 2013, for details).

– *QSO redshift errors:* Random redshift errors were drawn from a Gaussian distribution with dispersion $\sigma_v = 400\,\mathrm{km\,s^{-1}}$, and then applied to the QSOs. This velocity error is an important feature of real data to replicate, as it causes additional smoothing in the parallel direction in the Lyα-QSO cross-correlation.

– *Instrumental noise:* While effects such as spectral resolution and pixelisation are important to small-scale measurements from the Lyα forest, in the context of a large-scale feature such as BAO, we need only worry about the level of noise in our spectra. As such, the readily available DESI simulator `specsim` (Kirkby et al., 2016) was used to provide an approximation of the instrumental noise of the eBOSS spectrograph, with an exposure time chosen to achieve a realistic SNR for eBOSS.

While, ultimately, we would like to test our analysis using maximally realistic mock spectra, it is helpful to isolate different elements of our methodology by running analyses on several different sets of mock spectra, each with different systematic effects. In du Mas des Bourboux et al. (2020), four such sets were used, with increasing levels of realism in each one:

– *raw (Lyα only):* "Spectra" are simply skewers of transmitted flux fraction, as was used in § 6.4.2. Analysis of these data provides a baseline, systematic-free measurement, and quantifies our methods in the case of "perfect" calculation of δ_F.

– *+continuum+noise:* QSO continua and instrumental noise are added to the transmitted flux fraction. We must now carry out continuum fitting — a process made more difficult by the presence of noise — and take into account in our modelling the distortions that this induces.

– *+metals:* Metal absorption is added, producing additional correlations near to the line-of-sight, and requiring further modelling terms to be included.

– *+HCDs+σ_v:* HCDs and random QSO redshift errors are introduced as well. HCDs with $\log N_{\mathrm{HI}} > 20.3$ are masked where pixel absorption is greater than 20%, approximately mimicking the process used in the analysis of real data.

[4]Publicly available at https://github.com/imcgreer/simqso.

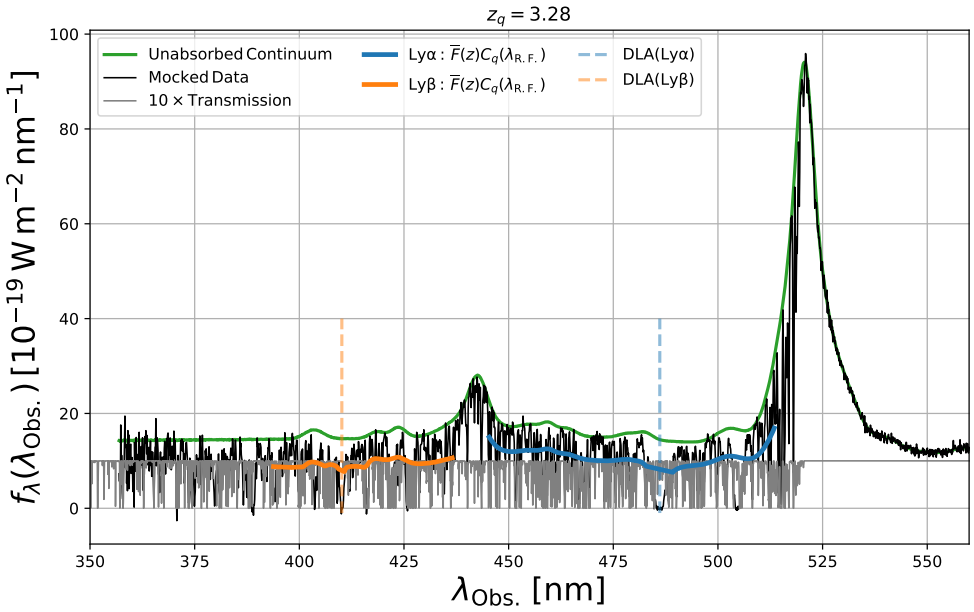

Figure 7.4. An example mock spectrum from a "+HCDs+σ_v" dataset. Here, the Lyα transmitted flux fraction from `LyaCoLoRe` is shown (multiplied by 10) in grey. Metal absorption is then added to this, and the combination is multiplied by the unabsorbed continuum in green. HCDs are then added to the spectrum, as indicated by the dashed vertical lines in the Lyα (blue) and Lyβ regions. Finally, instrumental noise is added to the spectrum, yielding the final spectrum in black. *Figure credit: du Mas des Bourboux et al. (2020)*

An example mock spectrum from this final set of spectra can be seen in Figure 7.4. Here, the transmitted flux fraction (multiplied by ten) from `LyaCoLoRe` is shown in grey, with the unabsorbed continuum shown in green and the final mock spectrum in black. Two HCD systems are also indicated by the vertical dashed lines.

7.2.3 Analyses on mocks

Full analyses were carried out on ten sets of mock spectra in each of these four varieties, with 90 additional analyses carried out on the second and fourth varieties described above to allow for greater precision. From each set of mock spectra, the four aforementioned correlation functions were calculated: the auto-correlations of the Lyα forest in the Lyα and Lyβ regions, and the cross correlations between QSOs and the Lyα forest in the Lyα and Lyβ regions. Each of these correlation functions was then fitted individually, as well as jointly across each set of mock spectra. In each case, the fitted model included factors relevant to the spectra in question, while discarding any unnecessary modelling. The measurement and fitting of all correlation functions was carried

out with picca, as in § 6.4.

In Figure 7.5, the Lyα auto-correlation measured from data is shown in blue, and the mean correlation from the 100 "+HCD+σ_v" mock datasets is shown in orange. Qualitatively, the mocks appear to be a fair representation of the data, with very similar overall amplitude, and similar shape as well. There does appear to be some difference between the mock- and data-based correlations in the bottom right panel of Figure 7.5, at separations $r \sim 25$ Mpc h^{-1}. This is mirrored by differences in some fitted parameters such as β_F, which is measured as 1.48 ± 0.051 from mocks, but 1.657 ± 0.088 from the data (du Mas des Bourboux et al., 2020). These deviations may be caused by differences in the effects of systematics such as metal absorption and HCDs, or by differences in the clustering at small separations. At this stage, we are not able to provide more definitive conclusions, and it is certainly an area to be investigated further ahead of DESI.

The shape of the mock auto-correlation function in Figure 7.5 is noticeably different to that shown in the left panel of Figure 6.3, largely due to the addition of systematic effects to the spectra used. In particular, the continuum fitting procedure suppresses power on scales larger than the typical length of Lyα regions used in our analyses, removing such fluctuations by setting the mean value of δ_F in each spectrum to zero artificially. The length of the Lyα region is typically around 520 Mpc h^{-1} (Blomqvist et al., 2015), and power is suppressed for $k \lesssim 0.01$ h Mpc^{-1}. This induces distortions to the broadband shape of two point correlation function (Blomqvist et al., 2015) that can be modelled with a "distortion matrix" (Bautista et al., 2017). Equally, astrophysical features of the Lyα forest such as HCDs and metal absorbers also induce distorting effects (Font-Ribera and Miralda-Escudé, 2012; Bautista et al., 2015), the latter of which are indicated in the left panel of Figure 6.4. Side-by-side illustrations of these effects can be seen in in Figures 11 and 12 of Bautista et al. (2017).

In Table 7.1 are the results of the analyses on mock datasets. It shows the mean value and mean standard deviation measured from the mock datasets for the two BAO parameters $\alpha_\parallel$ and $\alpha_\perp$, the Lyα flux velocity bias $b_{\eta,F}$, and the Lyα flux RSD parameter β_F[5]. This is presented for each of the different varieties of mocks, and for each of the different correlations measured. From these parameters, it is clear that, for all mock datasets and for all correlations, our analysis methods provide unbiased measurements of the BAO parameters to within the quoted errors, giving us confidence that our analysis pipeline is able to incorporate various systematic effects accurately. We can also see that the addition of systematic effects does not significantly alter our determination

[5]The parameters $b_{\eta,F}$ and β_F quantify the effect of RSDs on the Lyα flux power spectrum. They are defined and discussed in § 6.4.3.

Mock set	$\overline{\alpha_\parallel}$	$\overline{\sigma}$	$\overline{\alpha_\perp}$	$\overline{\sigma}$	$\overline{b_{\eta,F}}$	$\overline{\sigma}$	$\overline{\beta_F}$	$\overline{\sigma}$
Lyα(Lyα) × Lyα(Lyα)								
raw (Lyα only)	1.012	0.021	0.985	0.028	−0.200	0.001	1.568	0.021
+cont.+noise	1.003	0.027	0.995	0.040	−0.201	0.002	1.486	0.028
+metals	1.012	0.029	0.987	0.050	−0.202	0.002	1.485	0.030
+HCD+σ_v	1.004	0.029	1.001	0.041	−0.205	0.003	1.480	0.051
Lyα(Lyα) × QSO								
raw (Lyα only)	1.008	0.025	0.999	0.024	−0.189	0.003	1.568	0.041
+cont.+noise	1.008	0.029	0.992	0.033	−0.192	0.004	1.491	0.061
+metals	1.006	0.029	0.994	0.033	−0.193	0.004	1.510	0.063
+HCD+σ_v	1.003	0.033	0.998	0.033	−0.199	0.007	1.480	0.081
Lyα(Lyα) × Lyα(Lyβ)								
raw (Lyα only)	1.005	0.025	0.996	0.034	−0.200	0.002	1.588	0.026
+cont.+noise	1.014	0.049	0.983	0.069	−0.202	0.003	1.509	0.050
+metals	1.020	0.049	0.994	0.065	−0.203	0.004	1.528	0.054
+HCD+σ_v	1.009	0.054	1.019	0.087	−0.206	0.004	1.502	0.085
Lyα(Lyβ) × QSO								
raw (Lyα only)	1.028	0.042	1.009	0.044	−0.189	0.005	1.595	0.073
+cont.+noise	1.008	0.070	1.015	0.082	−0.193	0.010	1.527	0.146
+metals	0.994	0.071	1.002	0.093	−0.190	0.010	1.495	0.149
+HCD+σ_v	1.011	0.080	1.013	0.099	−0.192	0.015	1.447	0.186
all combined								
raw (Lyα only)	1.009	0.012	0.995	0.014	−0.203	0.001	1.628	0.015
+cont.+noise	1.005	0.017	0.992	0.022	−0.206	0.002	1.553	0.023
+metals	1.010	0.018	0.989	0.023	−0.206	0.002	1.558	0.025
+HCD+σ_v	1.005	0.019	0.998	0.023	−0.205	0.002	1.464	0.036

Table 7.1. Mean values and mean standard deviations of fits to correlations measured from mock datasets. We present these values for $\alpha_\parallel$, $\alpha_\perp$, $b_{\eta,F}$, and β_F, as well as the mean minimised value of χ^2 and the mean probability. Results are shown for "raw" mocks — where no QSO continuum or instrumental noise has been added to the transmitted flux fraction skewers — and three progressively more realistic mock types, which add continua and noise; metal absorption; and HCDs and QSO redshift errors to the spectra. The mean are calculated from 100 mocks for "+cont.+noise" and "HCD+σ_v" and 10 otherwise. *Table credit: eBOSS Collaboration et al. (2020).*

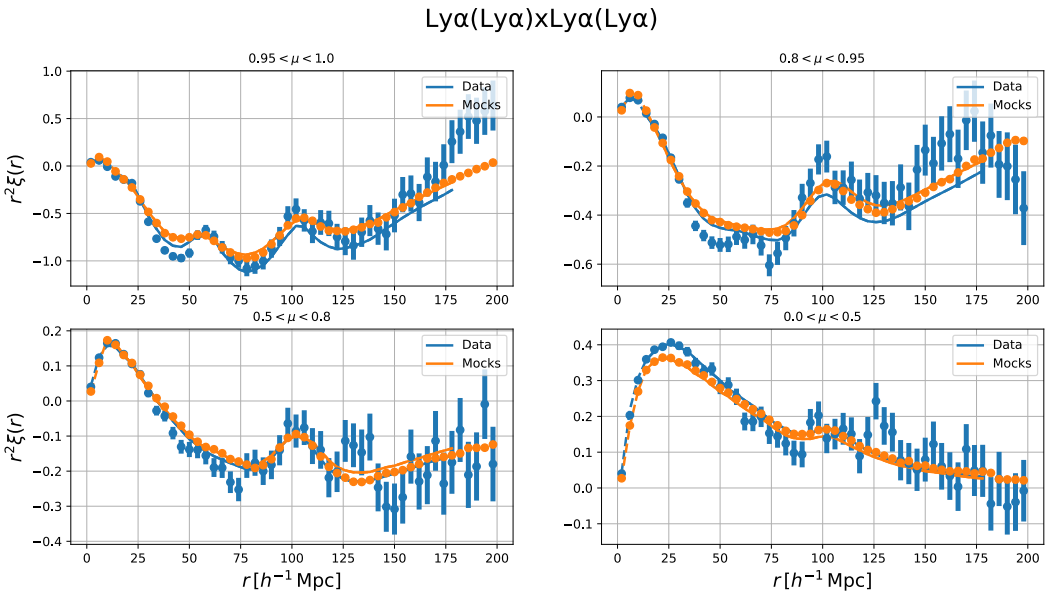

Figure 7.5. Comparison of the auto-correlation function of Lyα absorption in the Lyα region, measured from eBOSS DR16 data in blue, and from 100 mock datasets in orange. The correlations from these 100 mock realisations are averaged to produce the points and error bars shown. *Figure credit: du Mas des Bourboux et al. (2020).*

of $b_{\eta,F}$, though there does appear to be an impact upon β_F. Most notably, there appears to be a jump in the value of β_F when adding QSO continua and noise to our mocks, and a further jump when adding HCDs and QSO velocity errors. These suggest that our modelling of these systematics requires further development, though the minimal correlation between β_F and the BAO parameters ensures that our measurements of the BAO scale are not adversely affected by this.

Our suite of mock datasets also allows us to test our methods to measure covariances. As we have only one real Universe to measure, in du Mas des Bourboux et al. (2020) the covariance of the measurements from real data is calculated by subsampling across regions of the sky, as discussed in § 6.4.2. However, with 100 mock realisations available, we may calculate covariances from our mocks by looking at the "mock-to-mock" variation in correlation functions as well as via the subsampling approach. By then comparing these two different calculation methods, we may assess whether subsampling provides an accurate estimator of the true covariance. Figure 7.6 shows three different covariances for the Lyα auto- and the Lyα-QSO cross-correlations each: the measured subsampling covariance from DR16 data (dashed lines), the subsampling covariance from "+HCDs+σ_v" mock datasets (dotted lines), and the mock-to-mock covariance from the same mock datasets (solid lines). Comparing the solid and dotted lines, we can see that the subsampling approach yields an identical covariance to the mock-to-mock method at the 1% level, both for the

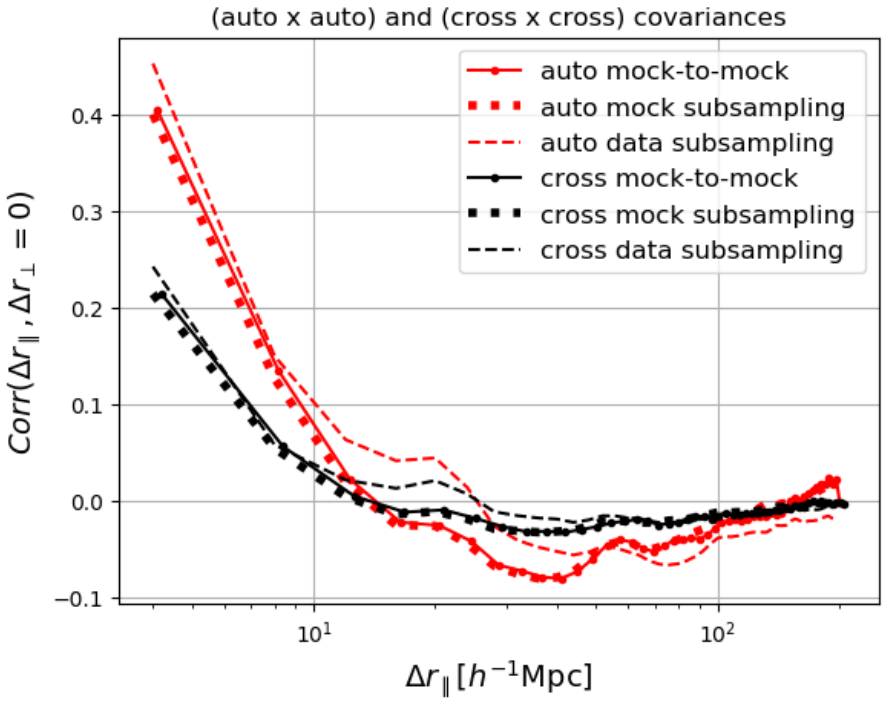

Figure 7.6. Covariances calculated for both Lyα auto- (black lines) and Lyα-QSO cross (red) correlations, and from both eBOSS DR16 data (dashed) and `LyaCoLoRe` mock data (dotted and solid). Here, we may compare the subsampling and mock-to-mock covariance methods on our mock datasets, while also comparing the covariances obtained via the subsampling method on both data and mocks. *Figure credit: du Mas des Bourboux et al. (2020)*.

auto- and cross-correlations. This validates the use of the subsampling method on real data.

We may also compare the mock covariances to those from data: while the two data sources yield covariances with similar overall shapes, there are noticeable differences between the two. This is at least in part due to the lesser contribution of metal absorption in the mocks, as is also visible in Figure 7.5. Further, this may be affected by the mock variances, which are measured to be lower than those from data, by a factor of 2 for the auto-correlation, and 1.5 for the cross-correlation, largely due to a lack of low-flux QSO spectra in the mocks. This difference in covariances highlights that there is still substantial potential for improving our mock datasets, both in `LyaCoLoRe` and in the subsequent stages to turn flux skewers into realistic spectra. Nonetheless, the use of these mocks to validate the subsampling covariance method provides a key example of their value as the arrival of DESI data draws ever closer.

7.3 H_0 AND BAO

The tension between measurements of the Hubble constant by *Planck* and by the SH0ES programme has proven a source of much debate in recent years. It has motivated the development of several new, independent approaches to measuring H_0 in an attempt to understand the source of the discrepancy. Here, we pursue one such method, combining BAO results from eBOSS' galaxy and Lyα tracers with results from big bang nucleosynthesis (BBN) to constrain H_0. We then provide forecasts for equivalent results from DESI.

This section is based on work published in Cuceu et al. (2019), of which I was the second author. I provided initial input on this project, developing a module to include recent Lyα constraints more accurately in the popular sampling tool `MontePython` *(Audren et al., 2013; Brinckmann and Lesgourgues, 2018), and contributing to the project's development via discussion in regular meetings. In this section, I describe the key details of this project as well as developments since the publication of Cuceu et al. (2019), all in my own words.*

7.3.1 COSMOLOGY FROM BAO

The BAO feature provides a powerful cosmological probe, helping to constrain cosmic expansion in the matter- and dark energy-dominated epochs. As mentioned in § 3.2.1, a number of different tracers may be used to measure the BAO scale, with galaxies, QSOs and the Lyα forest the most prominent. Here, we assess the relationship between BAO and the parameters of an underlying cosmological model, and address a recent tension between BAO measurements from different tracers.

Parameter dependencies

In a flat, ΛCDM cosmology — which we will assume for the remainder of § 7.3 — the dependency of the late-time BAO scale on cosmological parameters is relatively straightforward. As discussed in § 3.2.1, we may measure the BAO scale either perpendicular or parallel to the line of sight, constraining two different quantities:

$$\text{Perpendicular}: \quad \frac{d_M(z)}{r_d} = \frac{1}{r_d}\int_0^z \frac{c}{H(z')}\mathrm{d}z' \tag{7.1}$$

$$\text{Parallel}: \quad \frac{d_H(z)}{r_d} = \frac{c}{H(z)r_d}, \tag{7.2}$$

where

$$H(z)^2 = H_0^2 \left[\Omega_r(1+z)^4 + \Omega_m(1+z)^3 + \Omega_\Lambda\right], \tag{7.3}$$

and flatness imposes that $\Omega_\Lambda = 1 - \Omega_m - \Omega_r$. At late times, the contribution of the radiation sector is small, and we model it with a CMB temperature of $T_0 = 2.7255$ K from Fixsen (2009), and a fixed neutrino sector with $N_{\rm eff} = 3.046$ including two massless species. A third neutrino species is assumed to have a mass of 0.06 eV, thus is non-relativistic at the redshifts of interest to BAO, and contributes to Ω_m. With these choices, the two BAO scales depend on three parameters: Ω_m, H_0 and r_d.

At any given redshift, eq. (7.1) and eq. (7.2) clearly indicate that the BAO scale depends on Ω_m differently in the parallel and perpendicular directions. Equally, these dependencies vary with redshift for measurements in either direction. As such, measuring the anisotropic BAO signal at several redshifts allows us to constrain Ω_m tightly. However, these same equations also show that H_0 and r_d only appear as the product $H_0 r_d$ in our observable quantities, and thus these two parameters are completely degenerate. This makes intuitive sense, as H_0 and r_d both directly affect the normalisation of the late-time BAO scale, whereas Ω_m determines its evolution with variation in cosmic expansion rates through time.

Inter-tracer tension

The availability of numerous tracers from which to measure the BAO signal is key to BAO's cosmological power, enabling measurements to be made across a range of redshifts and tight constraints on Ω_m to be made. Measurements from these various tracers may be considered to be independent, and so we may straightforwardly compare their different constraints. However, when doing so, a mild tension emerges between measurements from galaxies and the Lyα forest: while results from galaxy BAO are consistent with *Planck*, the Lyα BAO constraints are in tension. We provide a visualisation of this apparent tension in Figure 7.7. In the left panel, we show constraints in the Ω_m-$H_0 r_d$ plane for BAO measurements from a number of different surveys and tracers: galaxies from 6dFGS (Beutler et al., 2011), SDSS MGS (Ross et al., 2015), BOSS (Alam et al., 2017) and eBOSS DR14 LRGs (Bautista et al., 2018); QSOs from eBOSS DR14 (Ata et al., 2018); and the Lyα auto-correlation (de Sainte Agathe et al., 2019) and Lyα-QSO cross-correlation (Blomqvist et al., 2019), also from eBOSS DR14. Then, in the right panel, we present a combined constraint from all galaxy BAO measurements, along with combined constraints from Lyα auto- and Lyα-QSO cross-correlation measurements from BOSS DR11 (Delubac et al., 2015;

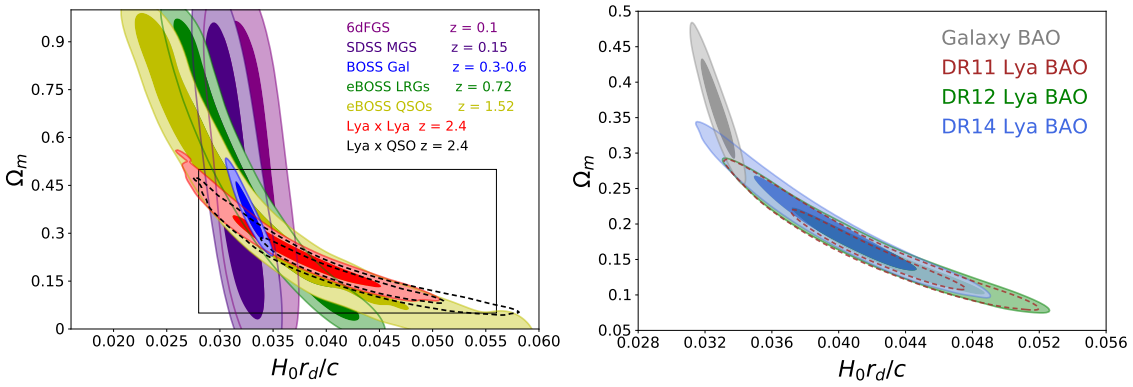

Figure 7.7. Contours in the Ω_m-$H_0 r_d$ plane from a selection of recent BAO measurements. In the left panel, results from numerous different sources are presented, across a range of different redshifts. Meanwhile, in the right panel we zoom in on the black rectangle of the left panel, highlighting the tension between combined galaxy BAO measurements and those from several recent Lyα measurements. *Figure credit: Cuceu et al. (2019)*

Font-Ribera et al., 2014) and DR12 (Bautista et al., 2017; du Mas des Bourboux et al., 2017), and eBOSS DR14. From this right panel, we can clearly see the aforementioned tension, which is more significant for the DR11 and DR12 Lyα results than those from DR14.

Quantifying the tension between pairs of complex datasets such as these is, however, a non-trivial task, and we may choose from a number of different methods to do so (e.g. Inman and Bradley Jr., 1989; Charnock, Battye, and Moss, 2017; Nicola, Amara, and Refregier, 2019; Adhikari and Huterer, 2019; Raveri and Hu, 2019). Most commonly used is the evidence ratio (Marshall, Rajguru, and Slosar, 2006; Trotta, 2008; Verde, Protopapas, and Jimenez, 2013), though this quantity is proportional to the prior volume shared by the two datasets and so may hide tension for certain prior choices (Handley and Lemos, 2019b). Alternatively, Handley and Lemos (2019b) defines the "suspiciousness" between two datasets, which attempts to remove the prior dependence of the evidence ratio while preserving its desirable characteristics. In conjunction with the Bayesian model dimensionality (Handley and Lemos, 2019a), this enables us to calculate the probability of two datasets being discordant by chance. We may thus quantify the tension between recent Lyα results from DR11 and DR12 of the BOSS programme and DR14 of the eBOSS programme. For the DR11 and DR12 results, the probability of the discordance being due to chance sits at $1.20 \pm 0.15\%$ and $1.31 \pm 0.16\%$ respectively. However, for the eBOSS DR14 results we calculate this probability to be $6.30 \pm 0.61\%$. As such, while the tension for the DR11 and DR12 results appears unlikely to be solely down to chance, that from DR14 is consistent with a purely statistical deviation.

Since the publication of Cuceu et al. (2019), the eBOSS DR16 Lyα BAO results have been

released (du Mas des Bourboux et al., 2020). The resultant constraints are in mild tension with *Planck* (and galaxy BAO), at a slightly reduced significance of $\sim 1.5\sigma$, compared to $\sim 2.5\sigma$ and $\sim 1.7\sigma$ for DR11/12 and DR14 respectively. In § 7 of du Mas des Bourboux et al. (2020), the reason for this evolution in the Lyα BAO constraints towards those from *Planck* and galaxy BAO is investigated. Re-analysing the BOSS DR12 results with the DR16 pipeline yields only small changes to the BAO constraints from the combined Lyα auto- and Lyα-QSO cross-correlations, of significance $< 0.25\sigma$. As such, the evolution from DR12 to DR16 can be taken to be due to differences in the underlying datasets rather than the analysis methodologies. du Mas des Bourboux et al. (2020) then consider whether this evolution is consistent with the increase of statistical power in the DR16 dataset: using the DR12 and DR16 covariance matrices, many mock DR12 and DR16 correlation functions are generated and the differences between the BAO constraints from pairs of correlations are assessed (see § 7 of du Mas des Bourboux et al., 2020, for details). These comparisons suggest that, due to statistical variation alone, one would expect to see differences in constraints larger than those observed between the real DR12 and DR16 results approximately 30% of the time. As such, it appears that both the tension observed between the DR12 Lyα BAO results and other cosmological probes, and the subsequent evolution of the Lyα BAO constraints are consistent with variation due to statistical fluctuations.

7.3.2 Measuring H_0 with BAO

As discussed in § 3.3, the most notable point of contention in modern cosmology lies around the current rate of expansion of the Universe, represented by the Hubble constant, H_0. Dispute over this value arises broadly between measurements from the local, late Universe at low z, and the distant, early Universe at high z. Most prominently, measurements using the local distance ladder from the SH0ES program have yielded $H_0 = 74.03 \pm 1.42$ km s^{-1} Mpc^{-1} (Riess et al., 2019), while measurements of CMB anisotropies from *Planck* have provided a value of $H_0 = 67.36 \pm 0.54$ km s^{-1} Mpc^{-1} within the ΛCDM model (Planck Collaboration et al., 2020b). Despite substantial efforts — investigating both possible systematic effects and new physics — no explanation for this tension has yet been agreed. Alternative methods to calculate the value of H_0 have also been developed in recent years, aiming to provide measurements independent of either SH0ES' or *Planck*'s data. A method using strong lensing time delays (Wong et al., 2020) agrees with SH0ES, as does an approach using the Tully-Fisher relation (Tully, Courtois, and Sorce, 2016) though at a much lower significance. Equally, calculations using the tip of the red giant branch (Freedman et al., 2020) have yielded a value between those from SH0ES and *Planck*, and a method based on gravitational wave sirens

(Abbott et al., 2017) does not yet have sufficient precision to provide any evidence either way.

At intermediate redshifts $z \sim 0.1$–4.0, a probe of H_0 is not immediately apparent. While BAO provides a tight constraint on Ω_m, it exhibits a fundamental degeneracy between H_0 and r_d. In order to measure H_0 with BAO, then, we must break this H_0–r_d degeneracy via the addition of non-BAO information. Typically, this is done via a measurement of r_d. In order to calculate r_d precisely, we must solve coupled Boltzmann equations with a programme such as `CAMB` (Lewis, Challinor, and Lasenby, 2000). For speed, however, we may instead use an approximation:

$$r_d \simeq \frac{55.154 \exp\left[-72.3(\omega_\nu + 0.0006)^2\right]}{\omega_m^{0.25351}\,\omega_b^{0.12807}}\ \text{Mpc}, \tag{7.4}$$

where $\omega_X = \Omega_X h^2$ for $X = \nu,\ m,\ b$ is proportional to the physical density field for neutrinos, matter and baryons respectively, and $h = H_0\ [\text{km s}^{-1}\ \text{Mpc}^{-1}]/100$. This approximation is accurate to within 0.021% (Aubourg et al., 2015). Thus to calculate r_d, we need values of Ω_m, Ω_ν and Ω_b. As mentioned previously, BAO measurements are able to constrain Ω_m tightly, and we assume a fixed neutrino sector with two massless and one massive species. In general, CMB anisotropy measurements from instruments such as *Planck* provide the best constraints on ω_b, and are able to constrain r_d directly. Here, though, as we seek a *Planck*-independent method, we instead follow the method first used in Addison, Hinshaw, and Halpern (2013) by using measurements of big bang nucleosynthesis (BBN) to constrain ω_b.

We may obtain such a constraint by considering BBN deuterium production, which is strongly related to ω_b (see Cyburt et al., 2016, for a review). As there are no known astrophysical sources of deuterium production (Epstein, Lattimer, and Schramm, 1976; Prodanović and Fields, 2003), measurements of deuterium abundance can provide an upper bound on primordial BBN production (Reeves et al., 1973; Gott et al., 1974). However, deuterium may be destroyed, and so in order to measure its abundance properly, we must find "pristine" regions of deuterium. These have not been "contaminated" by post-BBN astrophysical processes, and thus can be identified by their very low metallicities. Cooke, Pettini, and Steidel (2018) use seven damped Lyα absorbers (DLAs) for this purpose, measuring the deuterium abundance with 1% precision. A sequence of calculations is then required to convert this into a measurement of ω_b (Cooke et al., 2016), relying on measurements of various reaction cross sections in the BBN process. These are not straightforward to measure experimentally, with particular difficulty caused by the cross section of proton capture by deuterium to produce ^{3}He, denoted $d(p,\gamma)^3$He. As such, theoretical calculations are generally used, attaining 1% precision measurements (Cooke, Pettini, and Steidel, 2018) compared to empirical

measurements with approximately 7% precision (Adelberger et al., 2011). Notably, however, these two approaches yield conflicting values of ω_b:

$$100\omega_b = 2.166 \pm 0.015 \pm 0.011 \text{ (theoretical)}, \tag{7.5}$$

$$100\omega_b = 2.235 \pm 0.016 \pm 0.033 \text{ (empirical)}, \tag{7.6}$$

where the first error contribution comes from the measurement of deuterium abundance, and the second comes from the BBN calculations. Also of interest is the value derived from *Planck*'s 2018 data (Planck Collaboration et al., 2020b), which is consistent with empirical estimates yet in 2.9σ tension with the theoretical calculations:

$$100\omega_b = 2.237 \pm 0.015 \text{ (\textit{Planck})}. \tag{7.7}$$

We will consider both the theoretical and empirical methods to calculating ω_b from BBN, and will compare results from each value, referring back to this tension with *Planck*.

7.3.3 Cosmological results

eBOSS DR14 data

Using the two BBN-based measurements of ω_b in turn, we combine with BAO results from the combined galaxy sample presented in § 7.3.1, as well as those from eBOSS DR14 Lyα measurements, using `Polychord` (Handley, Hobson, and Lasenby, 2015a; Handley, Hobson, and Lasenby, 2015b) to sample parameter space. The constraints obtained from this are presented in Figure 7.8. In the left panel, results from galaxy and Lyα BAO are presented separately in the purple and red contours respectively, both of which have been combined with theoretical measurements of ω_b from BBN. Individually, either of these are consistent with the local distance ladder measurements from SH0ES (indicated by the vertical grey band Riess et al., 2019), though the Lyα results prefer a substantially lower value of Ω_m. Equally, both of these contours are also consistent with results from *Planck*. When the two sets of BAO results are combined, however, the resultant blue contours are no longer consistent with SH0ES at the 3.6σ level, but are highly consistent with *Planck*. If we instead combine our BAO results with the empirically determined value of ω_b from BBN, we obtain the green contours in the right panel. Here, again, the results are consistent with *Planck* but are in tension with SH0ES, now at the 3.3σ level. As such, no matter our choice of BBN estimate of ω_b, we derive values of H_0 consistent with *Planck*, but in substantial

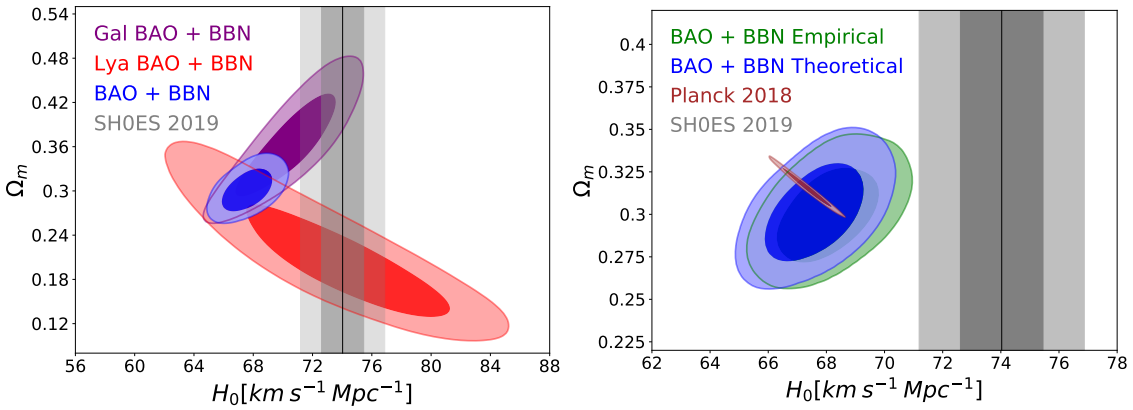

Figure 7.8. Contours in the Ω_m-H_0 plane from combining BAO results with constraints on ω_b from BBN. In the left panel, we show separately contours obtained by using BAO results from a combined galaxy sample in purple and from eBOSS DR14 Lyα measurments in red, combining each with a theoretical BBN measurement of ω_b separately. We then combine this BBN measurement with both the galaxy and Lyα BAO results to yield the blue contours. In the right panel, we show the same blue contours as well as equivalent contours in green when using the empirically-derived ω_m estimate from BBN, and results from *Planck* in dark red. In both panels, SH0ES constraints on H_0 are represented by a vertical grey bar, with the black line indicating the central value. A version of the left panel updated to use results from eBOSS DR16 is provided in Figure 4 of eBOSS Collaboration et al. (2020), but there is no qualitative difference with the results presented here. *Figure credit: Cuceu et al. (2019)*

tension with SH0ES.

This discrepancy may be recast into a tension in ω_b, and subsequently in the abundance of deuterium. Combining the BAO results with the value of H_0 from SH0ES and assuming ΛCDM and standard BBN physics, we may infer a primordial deuterium abundance of $10^5(D/H)_P = 1.38 \pm 0.25$. This is 4.5σ below the constraints from Cooke, Pettini, and Steidel (2018). Due to the aforementioned lack of significant astrophysical sources of deuterium, measurements of the primordial deuterium abundance have robust lower bounds, rendering the SH0ES+BAO value above highly unlikely.

Forecast DESI data

In the relatively near future, DESI will be able to provide substantially more precise BAO results than those available currently, using a range of different tracers. Evidently, this will have an impact upon BAO+BBN constraints on H_0, which we would like to assess. Using forecast uncertainties on DESI's BAO constraints for each individual tracer from DESI Collaboration et al. (2016a) and the "BAO + BBN Empirical" results from Figure 7.8 as a fiducial cosmology, we may calculate

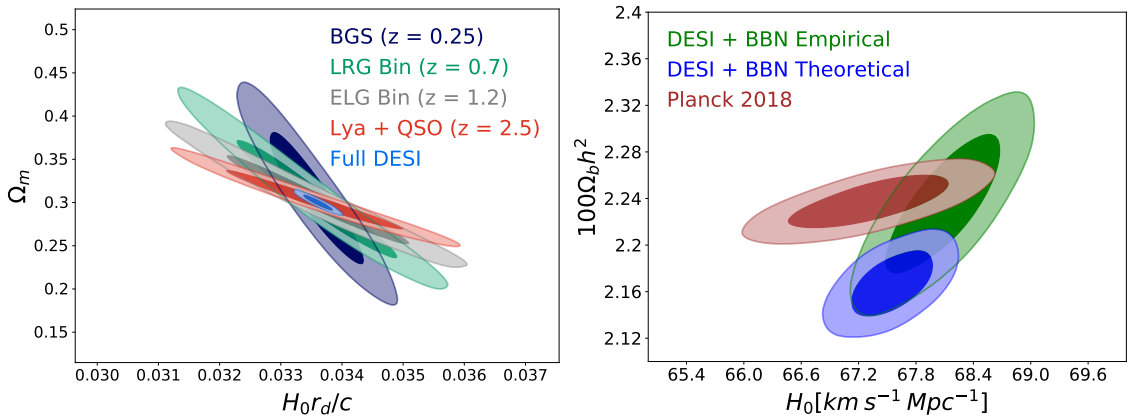

Figure 7.9. Forecast constraints from DESI's full survey. In the left panel are contours in the Ω_m-$H_0 r_d$ plane using BAO measurements from each of DESI's tracer classes individually, as well as a combined set. These contours are centred on values from Figure 7.8, with errors taken from the Fisher forecast estimates in DESI Collaboration et al. (2016a). In the right panel, contours in the $\Omega_b h^2$-H_0 plane are shown. The latest *Planck* constraints are shown in dark red, while the "Full DESI" contours of the left panel are combined with empirical and theoretical BBN constraints of ω_b to yield the green and blue contours respectively. *Figure credit: Cuceu et al. (2019)*

forecasts of DESI contours in the Ω_m-$H_0 r_d$ plane, as displayed in the left panel of Figure 7.9. These provide a clear demonstration on the power of multi-tracer BAO measurements, with the multiple redshift bins yielding contours at different angles in the plane, which may then be combined to great effect. Of course, these contours assume the same central point for each measurement, yielding very tight constraints when combining the tracers. In reality, the overlap between different tracers will likely not be so perfect, and there is a possibility that inter-tracer tensions will arise once again, similar to those shown in Figure 7.7.

In the right panel of Figure 7.9, we then present results from combining the forecast DESI results with the two BBN-based ω_b estimates used previously, showing contours in the ω_b-H_0 plane. While perfect alignment between BAO tracers is assumed here once again, the constraints in this panel suggest that BAO+BBN results from DESI will be able to provide constraints on H_0 comparable to those from *Planck* for either choice of BBN measurement. While the choice of BBN constraint does not significantly impact results when using current BAO data, this panel demonstrates that the increased statistical power of DESI will make the choice of BBN constraint important. Results from the two BBN methods presented here are in 1.2σ tension with each other, which, while by no means catastrophic, is certainly notable.

Since the publication of Cuceu et al. (2019), new laboratory measurements from the Laboratory for Underground Nuclear Astrophysics (Kochanek, 2016) have produced improved constraints on

the $d(p,\gamma)^3$He reaction rate (Mossa et al., 2020), which yield a value of ω_b given by

$$100\omega_b = 2.233 \pm 0.00036. \tag{7.8}$$

Notably, this remains consistent with the *Planck*-derived value of ω_b in eq. (7.7), and in tension with the theoretically derived value in eq. (7.5). This new result will help to improve the BAO+BBN constraints of H_0 in the near future, and also motivates further work into understanding the nature of the current discrepancy.

CHAPTER 8

CONCLUDING REMARKS

Since its first detection in 2013 (Busca et al., 2013; Slosar et al., 2013; Kirkby et al., 2013), the measurement of the baryon acoustic oscillation (BAO) feature using fluctuations traced by the Lyα forest has developed to become an important part of the cosmologist's inventory. It now provides a tight constraint on cosmic expansion history at a redshift not probed by other large-scale structure measurements, and thus plays an integral role in constraining cosmological models. With the arrival of the Dark Energy Spectroscopic Instrument (DESI, DESI Collaboration et al., 2016a; DESI Collaboration et al., 2016b), the power of Lyα BAO measurements will continue on its upwards trajectory, and there is ample motivation to ensure that this power is fully realised via a range of research opportunities. This will complement improved BAO measurements from DESI's other tracers as well as those from space-based experiments such as *Euclid* (Laureijs et al., 2011) and WFIRST (Spergel et al., 2015), ensuring that BAO will remain an integral part of cosmology for the foreseeable future.

In this thesis, we have focused on two key elements of conducting a Lyα forest BAO analysis from a spectroscopic QSO survey. First, we considered the classification of QSO target spectra, addressing problems relating to the observation of high-z QSO spectra and the construction of QSO catalogues. Second, we considered the construction of mock datasets to aid with the verification of Lyα BAO analysis pipelines, presenting and testing a new tool to produce such datasets. We then briefly described recent applications of the work in these two areas, along with a complementary study into the cosmological constraints imposed by BAO measurements. In this final chapter, we

summarise the work described above, providing a critical assessment and highlighting new areas into which it could be expanded.

8.1 Classifying QSO target spectra

To deliver scientific value, spectroscopic surveys must provide reliable classifications and accurate redshifts of the objects they observe. When targeting QSOs, sets of observed spectra may contain substantial contamination from stars and galaxies, which must be filtered out before conducting cosmological analyses. While traditionally this has been carried out by "visual inspection", the increasing size of QSO surveys has rendered this approach infeasible. Indeed, with the advent of DESI — which will observe ~ 2.4 million QSOs (DESI Collaboration et al., 2016a) — the need for automating this task will only grow, and making effective use of existing classification tools will be vital to successful measurements of large-scale structure.

8.1.1 Summary

In Chapter 5, we considered the problem of classifying QSO target spectra in the context of DESI, focusing on two key QSO target classification tasks: selecting high-z QSOs for reobservation, and constructing QSO catalogues. For each of these tasks in turn, we compared the performance of the official DESI template-fitter `redrock` and the convolutional neural network tool `QuasarNET`, using BOSS DR12 QSO target spectra (contained within the "DR12Q Superset", SDR12Q) as a test set in the absence of DESI data.

When selecting high-z QSOs for reobservation, we found that `QuasarNET` out-performs `redrock`, while a combined strategy that selects all objects chosen by either `QuasarNET` or `redrock` was able to provide further improved results. Similarly, when constructing QSO catalogues, `QuasarNET` provided catalogues with substantially reduced contamination compared to those from `redrock`, but a strategy which includes only QSOs identified by both `redrock` and `QuasarNET` was able to reduce contamination further still. In this second task, we also considered the inclusion of a visual inspection (VI) programme, demonstrating that catalogue contamination and completeness can be greatly improved by a moderate VI effort.

These analyses provide evidence that the outputs of available QSO classifiers may be used to construct QSO classification strategies that address DESI's QSO classification needs. These strategies yield high levels of performance, selecting $> 99\%$ of high-z QSOs for reobservation and producing QSO catalogues with $< 0.5\%$ contamination when tested on BOSS SDR12Q data. In

§ 7.1, we went on to test such strategies on the first DESI QSO target spectra, applying `QuasarNET` and `redrock` to 696 such observations with reliable visual inspections. Our conclusions from this analysis qualitatively agreed with those from Chapter 5, and we suggested some further tests to conduct once more data is available.

8.1.2 Outlook

The analyses presented in Chapter 5 are able to demonstrate the effectiveness of existing classifiers at addressing DESI's QSO classification needs, but they do not definitively determine the ideal strategies for use in DESI. While the SDR12Q data used in this chapter provides a reasonable approximation of DESI data, it is not a perfect substitution. For example, QSO targets will be chosen differently in DESI, altering the balance of contaminants in the QSO target set, and perhaps thus affecting the performance levels of classifiers on the dataset as a whole. The work presented in § 7.1 suggests that differences such as this will not cause major issues, but the number of DESI spectra currently available is too small to draw firm conclusions.

With the advent of DESI's SV phase, however, the number of DESI QSO target spectra will increase substantially, reaching a total of approximately 50,000–100,000. This will allow for more extensive tests of classification strategies to be carried out and, crucially, will enable `QuasarNET` models to be trained on DESI data. When applying these models to DESI spectra, the greater similarity between training and testing data ought to provide enhanced performance compared to models trained on BOSS data. By this point, we can also expect `redrock`'s set of templates to have been improved, and it will be instructive to re-test the strategies discussed in Chapter 5 using these new classification tools and the new DESI SV data.

The availability of these improved classifiers and a representative sample of DESI's QSO target spectra will then allow the development of more advanced classification strategies. These could follow a path analogous to those used in the construction of QSO catalogues from the 14$^{\text{th}}$ and 16$^{\text{th}}$ data releases from eBOSS (Pâris et al., 2018; Lyke et al., 2020), which used short decision trees to largely remove the contamination by stars introduced by the eBOSS pipeline classifications. Alternatively a more advanced, random forest approach could be taken to combine the outputs of our classifiers, seeking to learn more complex ways of translating sets of classifier results into a single, confident classification and redshift. This would then take advantage of the complementary nature of the existing classifiers in a maximally efficient way, diminishing the presence of any distinct failure modes.

Beyond the arrival of new data, there is substantial potential to develop `QuasarNET` further

as a tool. Despite the success of current QuasarNET models (see Figure 5.1), the existing system of "boxes" used to produce QuasarNET's features (see § 5.2.2) introduces artificial structure to QuasarNET's method. Suitable choices of activation functions appear to largely negate any negative consequences of this structure, but exploring more "natural" features may enable QuasarNET to make even more efficient use of its underlying architecture. Perhaps more fundamentally, in its current form QuasarNET behaves to some extent as a "black box" classifier. It identifies emission lines each with a certain confidence, but it is not clear how each line has been identified, nor exactly what this confidence value means. Providing a convincing explanation for the mapping between the input and output of machine learning models is a challenging problem that requires complex solutions (see Doshi-Velez and Kim, 2017, for a review), particularly in the context of deep neural networks such as QuasarNET (e.g. Ancona, Öztireli, and Gross, 2019; Angelov and Soares, 2019; Fan et al., 2020; Samek et al., 2020). While pixel-level explanations of QuasarNET's classifications may not be necessary in the context of large-scale structure surveys such as DESI, estimating quantities such as redshift errors would certainly prove helpful in understanding the reliability of QuasarNET's outputs.

8.2 Lyα BAO mock datasets

As measurements of BAO from the Lyα forest grow more precise, the complexity of modelling required to understand them increases too. In particular, potential sources of systematic errors must be studied with greater precision, while methods of calculating covariance matrices must be assessed in greater detail. In any survey, the development of a reliable analysis pipeline is of vital importance to maximising the scientific value of the data it will provide. This may be aided by the use of mock datasets: simple simulations constructed to mimic observational data with a minimal computational footprint. These provide data with a known cosmology and set of systematic effects, with which we may test certain parts of our analysis pipeline by checking that we are able to accurately extract cosmological information and associated errors. The small computational expense of each mock allows for a large number of datasets to be generated, and so pipelines can be tested to high levels of precision across a range of cosmologies and with varying systematic effects.

8.2.1 Summary

In Chapter 6, we presented a method to create such mock datasets. We first use an existing programme CoLoRe to generate a set of QSO positions and corresponding Gaussian density skewers,

with appropriate correlations imposed according to an input power spectrum. The Gaussian skewers are post-processed by a new tool `LyaCoLoRe`, which we developed for this purpose. This then adds small-scale fluctuations to the skewers before converting them into skewers of physical density and then optical depth, subsequently adding redshift-space distortions and finally converting to skewers of transmitted flux fraction. Within these various transformations are a number of tunable parameters, which we choose so that measurements of the Lyα forest bias, 1D power spectrum and mean flux approximately match literature measurements.

We then tested our mock-making process to ensure that the BAO signal remains unbiased throughout, generating ten sets of full-sky mock data from which we measure the Lyα auto-correlation and the Lyα-QSO cross correlation. We combined our measurements from the ten mocks and fitted the combined correlations with a simple model, finding that the position of the BAO peak remains unbiased to within 0.2%. This measurement is able to achieve far greater precision than DESI due to the larger footprint of each realisation, the use of numerous realisations, the high number density of QSOs, the lack of instrumental noise, and the lack of systematic effects needing to be modelled. As such, it provides a stringent test of our mocks' basic functionality, and demonstrates that they introduce no bias to the input BAO signal.

We further addressed two additional astrophysical effects that are implemented in `LyaCoLoRe`, and which are important to the Lyα forest: absorption from metals, and high-column density systems (HCDs). We first described the method used to introduce these effects, and then provided initial tests of their accuracy by computing the Lyα auto-correlation in the presence of metal absorption, and the Lyα-HCD cross correlation. These demonstrated that both `LyaCoLoRe`'s metals and HCDs have approximately the right large-scale bias, and thus will provide suitable tests of systematic modelling methods in Lyα analysis pipelines.

Having introduced and verified the methods of our mock datasets in Chapter 6, we then described their usage in the recent eBOSS DR16 Lyα forest analysis in § 7.2. In this context, they were used to test the analysis pipeline by computing correlations from mocks with successively greater numbers of systematic effects, checking that unbiased BAO results were recovered at each stage. Additionally, they helped to demonstrate the accuracy of the subsampling approach to covariance estimation, thus providing a key element of the analysis procedure and verifying the pipeline's integrity.

8.2.2 Outlook

In their current form, the `LyaCoLoRe` mocks are able to provide a useful tool for testing the Lyα analysis pipelines of BAO surveys. The analyses presented in this thesis, however, highlight some areas upon which we would like to improve. While the mocks are able to reproduce the correct BAO peak position with high precision, the full shape of the correlation function does not quite match that measured from real data, as shown in Figure 7.5. These deviations could be reduced by improving our implementation of metal absorbers and HCDs, as well as improving the small-scale realism of our skewers. The full shape of measured correlation functions contains more cosmological information than the BAO peak position alone, and surveys such as DESI would like to extract this information as effectively as possible. To achieve this goal, Lyα mocks with the correct shape of the full correlation function would be of great value.

Relatively straightforward improvements to the mocks could be brought about by taking advantage of an existing functionality within `CoLoRe`. Rather than providing initial skewers from a 3D Gaussian field, `CoLoRe` is able to apply second order Lagrangian perturbation theory (2LPT) to this field, displacing the density field along a trajectory determined by the gravitational potential and cosmological expansion (for a review of perturbation theories, including 2LPT, see Bernardeau et al., 2002). This then provides skewers of the physical density field that more accurately represent reality than using a lognormal transformation, and would certainly help to improve the realism of the mocks on smaller scales. In particular, the QSO clustering from the lognormal mocks shows deviations on separations $r \sim 20$ Mpc h^{-1} (see Figure B.1), and would be improved by the use of 2LPT. Using this approach would require changes to the existing `LyaCoLoRe` method: in using `CoLoRe`'s 2LPT functionality, we are no longer able to straightforwardly add Gaussian small scale fluctuations to our skewers as `CoLoRe`'s 2LPT output will approximate fluctuations in the physical density field (equivalent to $1+\delta$, as defined in eq. (6.4) for the Gaussian method). This is by no means an insurmountable challenge, but would certainly require substantial testing, as presented in the Gaussian case in Chapter 6.

Introducing the more advanced 2LPT approach would help to address some of the issues seen in the correlation functions of Chapter 6 at small separations. It would also introduce correlations beyond those produced by Gaussian fields alone, and would thus provide a suitable testing ground for large-scale measurements beyond the canonical two-point statistics. Historically, there have been several studies of the "1D bispectrum" (e.g. Mandelbaum et al., 2003; Viel et al., 2004), which measures correlations between triplets of flux pixels along the same line of sight. However, large-

scale, 3-dimensional statistics such as the three-point correlation function (3PCF, Tie et al., 2019) are yet to be measured from the Lyα forest. While it is not clear exactly what level of accuracy 2LPT mocks would provide for such measurements, they would offer a qualitative improvement over their Gaussian counterparts, and would enable more reliable tests of 3PCF methods to be carried out. Such tests could assess the impacts of systematic effects of 3PCF measurements using a realistic survey design, thus providing useful feedback ahead of attempting to measure such statistics from real data.

In either the Gaussian or 2LPT scenario, the `LyaCoLoRe` mocks may also prove useful in conducting a number of other tests ahead of DESI's Lyα BAO measurements. As yet, mocks have only been generated using a single cosmology, and thus producing mocks with varied cosmologies would constitute a natural extension of our work. This would then provide a more extensive testing regime for our analysis pipeline, enabling blinded analyses to be carried out and thus helping to ensure that no bias is introduced across a range of cosmological scenarios. Furthermore, one could extend the realism of mock Lyα BAO analyses by producing end-to-end survey simulations. These go beyond the types of spectra used in § 7.2 to simulate additional observational effects such as fibre assignment and variable observing conditions. This then can provide "snapshots" of DESI's data, either partway through the survey or upon its completion, and can enable a more accurate assessment of DESI's cosmological power at any stage of its survey.

8.3 Cosmological conclusions

In § 7.3, we described an approach to measure the Hubble constant, H_0, using BAO and constraints from big bang nucleosynthesis (BBN). This method is independent of both *Planck* and the local distance ladder, and so provides an interesting point of comparison when considering the tension between H_0 measurements from these two sources. With the advent of DESI, the BAO+BBN constraint on H_0 is set to tighten substantially, potentially yielding levels of precision similar to those from *Planck*.

This increase in precision will, no doubt, be mirrored across many facets of cosmology in the coming years. New instruments will provide unprecendented quantities of high-quality data, while advances in computational and theoretical areas will allow for ever more complex analytical efforts to be carried out. In conjunction, these will enable extensions to the current concordance ΛCDM model to be tested more stringently, including more complex extensions which cannot currently be constrained. It may further enable new detections to be made, providing insight into areas

such as neutrino mass, inflationary mechanisms and the nature of dark energy. Of course, whether these future measurements are in agreement the ΛCDM model, or whether they lead cosmology in a different direction remains to be seen. Certainly, though, the BAO feature will remain vital to cosmology in the immediate future, with measurements from the Lyα forest enabling cosmic history to be constrained across a broad range of redshifts, and providing an exciting approach to understanding our Universe.

PART V

APPENDICES

APPENDIX A

SUPPLEMENTARY MATERIAL FOR CHAPTER 5

A.1 TECHNICAL TESTS OF QUASARNET

A.1.1 DEPENDENCE OF QUASARNET'S PERFORMANCE ON TRAINING SET SIZE

In Busca and Balland (2018), QuasarNET models were trained on 80% of DR12Q Superset data (approximately 500,000 spectra), while in Figure 5.1, 90% was used (approximately 560,000 spectra). These training sets are significantly larger than the ∼50,000–100,000 spectra that will be visually inspected during DESI's survey validation (SV) period. In order to check that this reduced number of spectra will be sufficient to train high-performing QuasarNET models, we assess models trained on varying fractions of the DR12Q Superset data. We consider models trained on 10% and 5% of this data, approximately equivalent to 63,000 and 31,000 spectra respectively, and thus corresponding to "realistic" and "worst case scenario" training set sizes that we can expect from DESI SV. We compare the performances of models trained on datasets of these sizes to a "fiducial" training set size of 90% of DR12 data, as was used in Figure 5.1. For each training set size, we train 10 models on random subsets of DR12Q Superset data, and for each model, we test performance on spectra from all objects not in its training set. In Figure A.1, we then plot the mean performance across these 10 models for the 90%- and 10%-trained models, and the mean performance across 9 models for the 5%-trained models. We exclude one outlying 5%-trained model which showed

degraded performance relative to the remaining 9 models. It is not clear why this outlier exists, and it provides motivation to aim for an SV truth table of $\gtrsim 50,000$ QSO targets. If this is not feasible for any reason, then efforts should be made to investigate whether alternative `QuasarNET` architectures — fewer convolutional layers, for example — are better suited to smaller training set sizes.

The upper panel of Figure A.1 shows the levels of purity and completeness achieved by each model as a function of "confidence threshold" (as explained in 5.2.2). Here, we set $n_{\text{detect}} = 1$ for simplicity. Comparing the solid and dashed lines demonstrates that reducing the training set size from 90% (solid lines) to 10% (dashed lines) of DR12Q Superset has little effect on the model's purity, and results in only a small (0.25–0.5 percentage point) drop in the model's completeness at any given confidence threshold. The 10%-trained model achieves "optimal" performance — where purity and completeness are equal — at a confidence threshold of 0.08, where purity and completeness are both 99.4%. When compared to the 90%-trained model, which achieves purity and completeness of 99.5% at a confidence threshold of 0.15, it is clear that $\sim 62,000$ training spectra is sufficient for a `QuasarNET` model to achieve high purity and completeness in its test sample. Equally, considering the dash-dotted lines, reducing the training set size further to only 5% of DR12Q Superset results in an additional drop in "optimal" performance of only 0.2 percentage points, to 99.2% (at a confidence threshold of 0.07). Thus we can conclude that even $\sim 31,000$ spectra can be sufficient to achieve high levels of purity and completeness.

The lower panel of Figure A.1 shows a (normalised) histogram of the velocity errors (relative to VI redshifts) for spectra that `QuasarNET` correctly classified as belonging to QSOs, with velocity error less than 6000 $\mathrm{km\,s^{-1}}$. Once again comparing the solid and dashed lines, we can see relatively small differences between the results obtained by the 90%- and 10%-trained models. While predictions from the 90%-trained model have a median velocity error of -8 $\mathrm{km\,s^{-1}}$ and a standard deviation of 618 $\mathrm{km\,s^{-1}}$, results from the 10%-trained model have a median of 18 $\mathrm{km\,s^{-1}}$ and a standard deviation of 793 $\mathrm{km\,s^{-1}}$. As such, both models show no significant bias in the estimated redshifts, and the spread of velocity errors is 30% larger for the 10%-trained model. This marks a moderate increase, though it is important to note that `QuasarNET` is not designed to be a precision redshift fitter and so such values are not excessive. Finally, using a model trained on 5% of DR12 data results in further spreading of the velocity error distribution, achieving a median velocity error of -6 $\mathrm{km\,s^{-1}}$, with a standard deviation of 902 $\mathrm{km\,s^{-1}}$. Once again, there is no significant bias in the estimated redshifts for models using this training set size, and the velocity error spread does not increase to unacceptable levels.

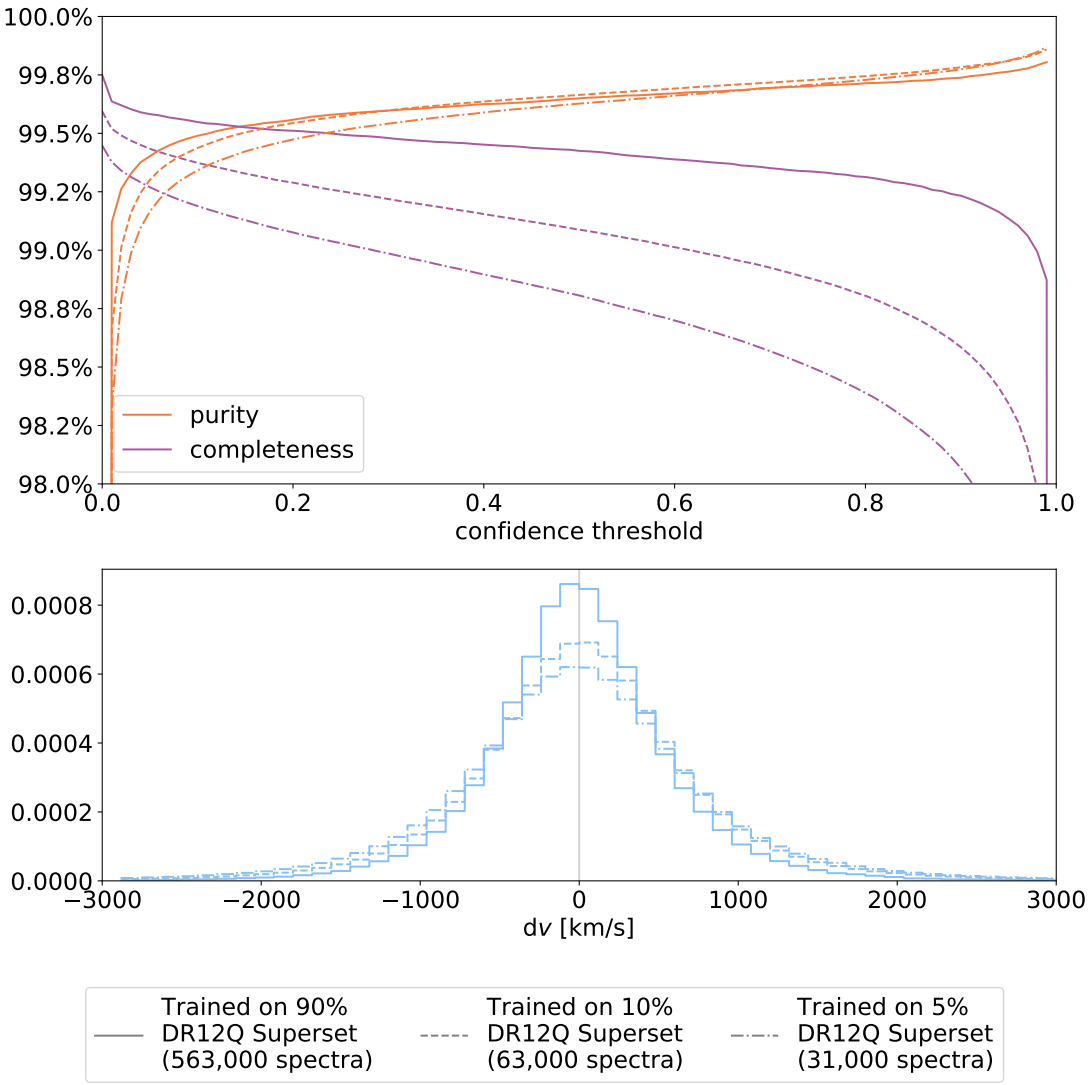

Figure A.1. Performance of QuasarNET models with varying training set sizes. The top panel shows purity and completeness as a function of the classification confidence threshold, while the bottom panel shows a histogram of velocity errors for successful QSO classifications. Here, the solid and dashed lines — corresponding to training sets made from 90% and 10% of BOSS DR12Q Superset data respectively — show the mean performance of 10 QuasarNET models. The dot-dashed line shows mean performance over 9 models with training sets made from 5% of DR12Q Superset data, excluding one outlying model which produced degraded performance levels.

Clearly, when training on DESI SV data, precise performance levels of QuasarNET models may vary slightly compared to those shown in Figure A.1 due to, for example, differences in targeting procedures between BOSS and DESI. However, it is reasonable to expect that a QuasarNET model trained on visually inspected data from DESI SV will be able to achieve similarly high levels of performance to the 10%-trained model presented here, and as such we can be reassured that the smaller training set provided by DESI SV will not inhibit our ability to train high-performing QuasarNET models. Equally, the small drop in performance when using the 5%-trained model suggests that even in a "worst case scenario" in which DESI SV is impaired, effective QuasarNET models would still be able to be trained.

A.1.2 Dependence of QuasarNET's performance on signal-to-noise ratio

As discussed in § 5.3.1, a key task for QSO classifiers in DESI will be to select high-z QSOs for reobservation. This introduces a number of new challenges, most notably the need to carry out classification on spectra obtained from single exposures. These spectra will have a lower signal-to-noise ratio (SNR) than coadded spectra, making classification more difficult. Ahead of DESI, we need to test how well QuasarNET models are able to classify single-exposure spectra, and also determine whether a model trained on single-exposures or coadded spectra is preferable when doing so. In order to answer these questions, we construct a single-exposure version of the DR12Q Superset data, taking spectra from one exposure chosen at random from each set of coadded exposures (excluding any low quality exposures). As such, our single-exposure dataset is of the same size as our coadded dataset, with each spectrum having a direct coadded counterpart, thus ensuring consistent balance of contaminants. Within each set of exposures used in BOSS' coadds, the BOSS pipeline identifies a "best" exposure, that with the highest average SNR. We could have used these best exposures rather than randomly chosen ones, but this simply improves performance of all models slightly, without affecting our qualitative conclusions.

In Figure A.2, we show the performance of QuasarNET models trained and tested on coadded and single-exposure data. In the left column are results for a model trained on coadded spectra — each made from on average 4 exposures — from BOSS DR12. This is the same model that was used in § 5.3.2 and § A.1.1, trained on 10% of DR12Q Superset as is appropriate given the estimated size of the DESI SV dataset. In the right column are results for a model trained on single exposures from DR12, using 10% of the single-exposure dataset for training, as was used in § 5.3.1. In the top row of Figure A.2 are results when models are applied to a coadded test dataset, while

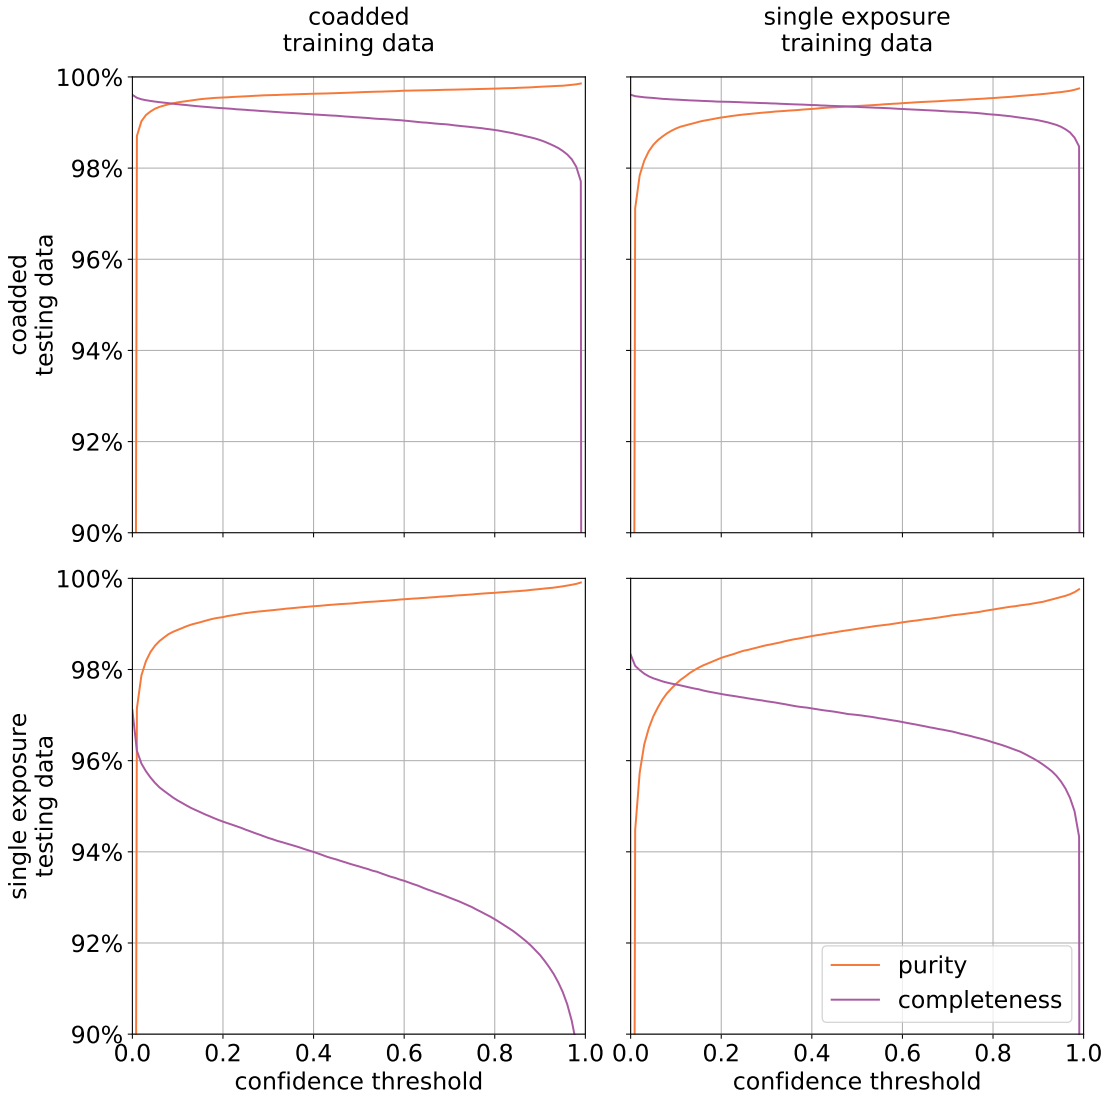

Figure A.2. Performance of QuasarNET models with different numbers of exposures in the training and testing data. The left column corresponds to a QuasarNET model trained on data consisting of coadded spectra, whereas the right column corresponds to a model trained on single-exposure spectra. Both models were trained on $\sim$10% of DR12Q Superset spectra ($\sim 62,000$ spectra), a similar quantity of visually inspected data to that which will be available from DESI SV. The top row corresponds to results obtained from applying each model to a coadded dataset, while the bottom row corresponds to a single-exposure test set. Each panel shows the purity and completeness of the classifications as a function of QuasarNET's confidence threshold.

in the bottom row are results when applied to a single-exposure test dataset. For each model, we use a test set of DR12Q Superset spectra from all objects that were not included in that model's training set. We also exclude from our test sets spectra which were classified with anything other than maximal confidence by the DR12Q Superset VI procedure. As such, the results shown in the top left panel are equivalent to the results shown by the dashed lines in the top panel of Figure A.1.

Moving from the top left to the top right panel, we are reducing the SNR in our training data while maintaining a high SNR in our test data. This change results in a slight increase in the completeness at a given confidence threshold, but a moderate drop in the purity of $\sim$0.5 percentage points. These differences are due to the single-exposure-trained model classifying more objects as QSOs for a given confidence threshold, of which most are incorrect. Conversely, moving from the top left to the bottom left panel, we are reducing the SNR in our test data while maintaining a high SNR in our training data. This results in a substantial drop in completeness ($\sim$5–8 percentage points) due to the model correctly identifying fewer single-exposure QSO spectra for a given confidence threshold. Equally, it results in a moderate drop in purity ($\sim$1 percentage point), mostly due to incorrect redshift determination in the single-exposure QSO spectra that are classified as QSOs. Finally, moving from the bottom left to the bottom right panel, we are reducing the SNR in our training data while maintaining a low SNR in our test data. This results in a $\sim$3–4 percentage point increase in completeness as the single-exposure-trained model is better able to identify low-SNR QSO spectra, but a $\sim$0.5–1 percentage point drop in purity due to contamination both by stellar spectra and incorrectly determined redshifts.

From these four sets of results, we can draw a number of conclusions. When classifying coadded spectra, a single-exposure-trained model is able to achieve a higher level of completeness than a coadd-trained model at a given confidence threshold. However, this comes at the expense of purity, and it is preferable to use a coadd-trained model with a lower confidence threshold if completeness is a priority. As such, a coadd-trained model is almost certainly preferable. A parallel conclusion can be drawn when classifying single-exposure spectra. A coadd-trained model is able to produce a more pure set of QSOs, but achieves substantially lower levels of completeness than a single-exposure-trained model when using the same confidence threshold. Again, it is almost always possible to match the coadd-trained model's purity with the single-exposure-trained model by simply increasing the confidence threshold, with a higher completeness being achieved at the same time. As such, it is preferable to use a single-exposure-trained model in this case.

Most importantly, we can be reassured that `QuasarNET` is able to produce models that perform well on single-exposure data. From the bottom right panel of Figure A.2, we can see that a single-

exposure-trained model achieves purity and completeness of 97.7% on single-exposure test data at a confidence threshold of 0.1. While there is a drop in performance compared to coadded data, this is to be expected: classifying low-SNR spectra is an inherently more difficult task. In particular, it is possible that single-exposure spectra classified incorrectly by QuasarNET would not have been confidently classified by a VI expert either, but we are able to assign that single-exposure spectrum a classification as our VI results were obtained using coadded data.

APPENDIX B

SUPPLEMENTARY MATERIAL FOR CHAPTER 6

B.1 THE QUASAR AUTO-CORRELATION

We measure the quasar (QSO) auto-correlation on 10 QSO catalogues from 10 realisations of `CoLoRe` and combine our results. Correlations are computed as the weighted sum of pairs of QSOs in a grid of parallel and perpendicular separation bins. We divide the sky into `HEALPix` pixels, computing "data-data", "data-random" and "random-random" correlations in each one using a random catalogue of QSOs. This random catalogue has the same number density distribution of QSOs as that in the mock data, and is generated by `LyaColoRe`. The different correlation types are then combined using the Landy-Szalay estimator (Landy and Szalay, 1993), and the covariance is estimated via sub-sampling across `HEALPix` pixelisations of all 10 realisations (as described in § 6.4.2). As in § 6.4.2, all correlations are computed using `picca`[1]. A Kaiser model (Kaiser, 1987) is then fitted to the measurement, leaving free parameters describing the location of the BAO peak and the QSO bias $b_{\delta,\mathrm{QSO}}$. As in § 6.4.3, we also leave free parameters describing the smoothing of the input power spectrum in the parallel and perpendicular directions. As in § 6.4.3, we fit only in the range $40 < r\ [\mathrm{Mpc}/h] < 160$ as the lognormal approximation begins to break down below this range. The resultant fit is very good in the fitted region, as shown in Figure B.1. We measure a

[1] Publicly available at https://github.com/igmhub/picca.

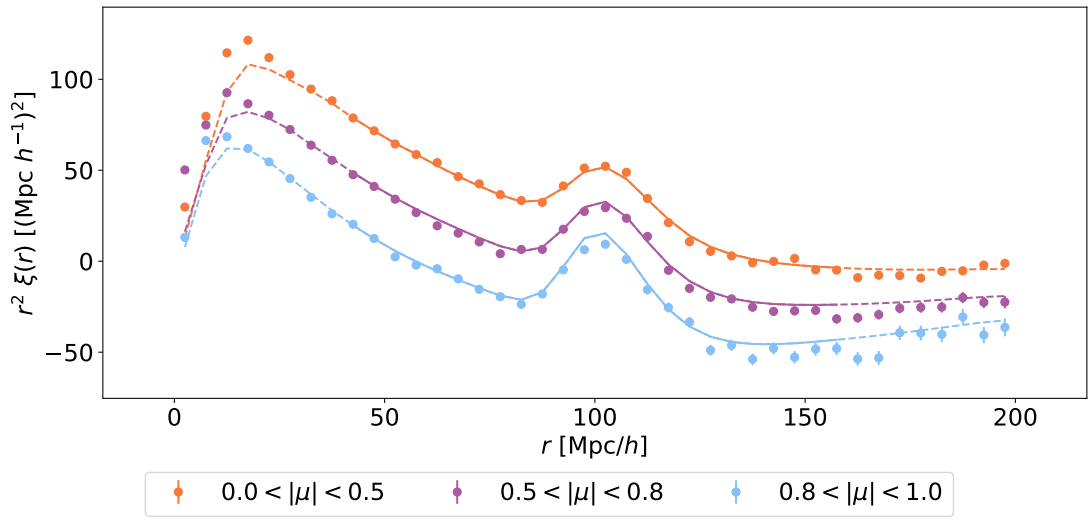

Figure B.1. The auto-correlation of QSOs, as measured from ten realisations of `CoLoRe`. The fit is generally good in the fitted region, though the correlation on smaller scales than this is evidently too high.

QSO bias of 3.57 ± 0.01 at an effective redshift of $z = 2.20$, consistent with the input value of 3.56 to within 1σ.

B.2 Redshift-space distortions: implementation details

As described in § 6.2.2, adding RSDs to our skewers requires the calculation of a matrix of weights W_{ij} to map each skewer's real-space cells τ^x_j to redshift-space cells τ^s_i via the matrix equation $\tau^s_i = W_{ij}\tau^x_j$. W_{ij} is determined by representing each cell as a top-hat function in real space, mapping this profile into redshift space according to the choice of kernel K, and calculating the overlap with each redshift-space cell:

$$W_{ij} = \int_{s^l_j}^{s^u_j} P(s - x_i - v_{r,i}|T_i, d_i)\mathrm{d}s, \tag{B.1}$$

where s^l_j and s^u_j are the lower and upper boundaries of cell j in redshift-space, and $P(x|T,d)$ describes the profile of the real-space cell when mapped into redshift space. $P(x|T,d)$ is dependent on the distance from the centre of the cell x, the temperature of the gas T and the half-width of the cell d. The form of P is determined by the choice of kernel, K, as defined in eq. (6.6):

$$P(x|T,d) = \frac{1}{2d}\int_{-d}^{d} K(x - y|T)\mathrm{d}y. \tag{B.2}$$

As such, in the case of K chosen to be a Dirac delta function, the redshift-space cell is represented by a top-hat function (as it was in real space).

In order to account for thermal broadening when adding RSDs to our skewers, we must instead choose our kernel K to be defined by

$$K(x|T) = \frac{1}{\sqrt{2\pi}\sigma_v(T)} \exp\left(-\frac{x^2}{2\sigma_v^2(T)}\right), \tag{B.3}$$

where $\sigma_v(T)$ is the thermal velocity dispersion, which we approximate as in McDonald et al. (2001) by

$$\sigma_v(T) = 9.1\left(\frac{T}{10,000\text{K}}\right)^{1/2} \text{km s}^{-1}, \tag{B.4}$$

for temperature $T(z,\mathbf{x}) = T_0(z)\rho(z,\mathbf{x})^{\gamma(z)-1}$. As described in § 6.3, for the purposes of this work we fix $\gamma = 1.5$. We also fix $T_0 = 10,000$ K in line with Slosar et al. (2011) and consistent with literature values (e.g. Ricotti, Gnedin, and Shull, 2000; McDonald et al., 2001; Hiss et al., 2018). Of course, these values can easily be updated to follow a more complex redshift dependence for any uses of `LyaCoLoRe` where thermal broadening effects become significant. Evaluating eq. (B.2) for this choice of K yields a cell profile in redshift space defined by

$$P(x|T,d) = \frac{1}{4d}\left[\text{erf}\left(\frac{x+d}{\sqrt{2}\sigma_v(T)}\right) - \text{erf}\left(\frac{x-d}{\sqrt{2}\sigma_v(T)}\right)\right], \tag{B.5}$$

and the matrix of weights can then be computed as per eq. (B.1).

B.3 The Lyα-HCD cross-correlation

In Figure 6.4, we showed the cross-correlation between Lyα absorption and high column density systems (HCDs) from 10 realisations of `LyaCoLoRe`, comparing it to a linear theory model similar to that used to describe the cross-correlation with QSOs. This model assumes that HCDs have the same clustering as dark matter halos, with a large-scale bias of approximately 2.0. However, in a QSO survey, HCDs are only detected when they are absorbing light from a background QSO, and this observational bias is not taken into account in our modelling. Here, we propose that this bias results in an asymmetry in the measured correlation function. We present a qualitative description of this effect and explain our choice to use only HCDs detected far away from the QSO in Figure 6.4 in this context.

In the left panel of Figure B.2 we show the same measurement of the Lyα-HCD cross-correlation as in the right panel of Figure 6.4, this time plotting the correlation against $r_\parallel$ in 3 narrow bins

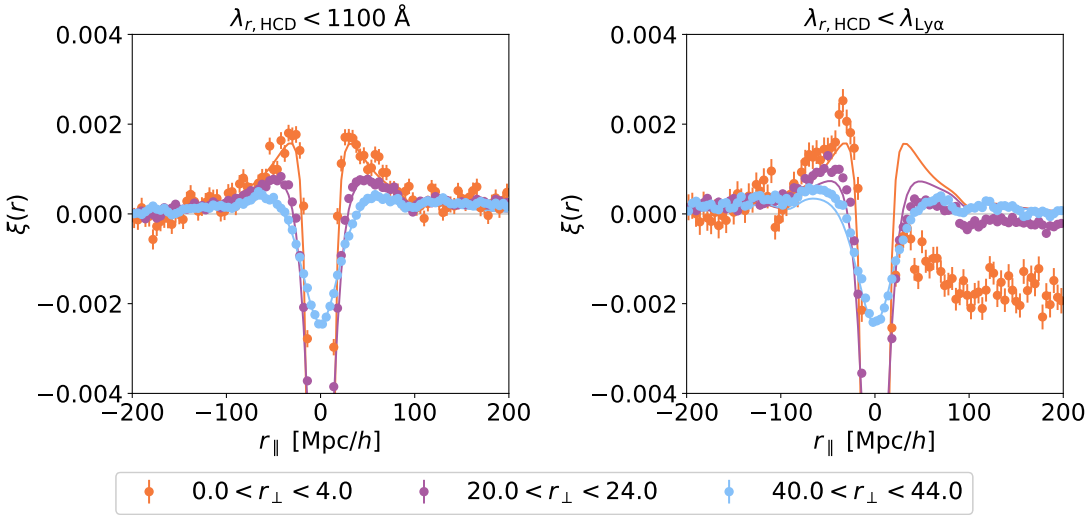

Figure B.2. The Lyα-HCD cross-correlation, plotted against $r_\parallel$ for different bins of $r_\perp$. The left panel shows the combined measurement from ten realisations using an HCD catalogue that only includes HCDs with rest-frame wavelength less than 1100 Å (as in the right panel of Figure 6.4). The right panel shows the correlation measured from one realisation when using an HCD catalogue that includes HCDs in the full rest-frame wavelength range, up to $\lambda_{\mathrm{Ly}\alpha} = 1215.67$ Å. The solid lines in both panels show the same fitted correlation as in Figure 6.4: the joint fit of the Lyα auto-correlation and Lyα-HCD cross-correlation from ten realisations of `LyaCoLoRe`.

of $r_\perp$. The solid lines show the model obtained by fitting this measurement jointly with the Lyα auto-correlation. The model is generally able to fit the measurement well, though some small residuals remain at large $r_\parallel$. These are visible at large separations in the right panel of Figure 6.4, accentuated by plotting $r^2\xi(r)$ in that figure.

In the right panel of Figure B.2 we plot the Lyα-HCD cross-correlation measured on one realisation of `LyaCoLoRe`, this time using a full HCD catalogue (with no maximum rest-frame wavelength). The solid lines are the exact same lines as in the left panel. It is clear from this plot that there is a strong asymmetry in the data, and the model used to fit the data in the left panel does not fit this measurement well.

We propose that this asymmetry is a consequence of the observational bias that is inherently present in our HCD sample, and the dependence on the Lyα-QSO cross-correlation that this induces. According to the density-QSO cross-correlation, a QSO q will tend to have dense regions of gas around it. In relation to an HCD X in q's spectrum, these dense regions will be located at small $r_\perp$ and $r_\parallel \simeq r_{Xq}$, the distance between X and q (as X is constrained to lie directly along the line of sight between q and the observer). This preferential location of dense regions of gas will imprint a

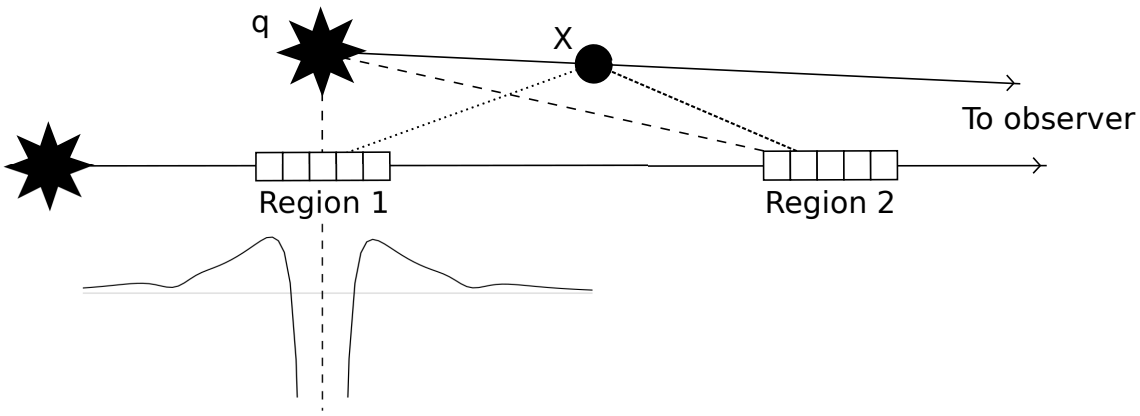

Figure B.3. Diagram showing the geometry of the setup involved when measuring the Lyα-HCD cross-correlation between two near-parallel skewers. Given the proximity of the QSO q to "Region 1" of the lower skewer, we expect to measure biased values of δ_F for cells corresponding to that region. The cells' values will tend to be reduced or boosted according to the Lyα-QSO cross correlation, as indicated beneath "Region 1". This biasing is then imprinted on the correlation between an HCD X and the skewer.

feature in the correlation between X and neighbouring skewers of δ_F at these specific separations. Referring to the diagram in Figure B.3, we can see that the cells of δ_F in Region 1 will tend to be significantly biased according to the Lyα-QSO cross-correlation. Thus, we will see a feature in the correlation between HCD X and its neighbouring skewer corresponding to this region. The shape of this feature is determined by the shape of the Lyα-QSO cross-correlation at small r_t, as shown beneath Region 1. The cells in Region 2 will not be significantly affected by the presence of QSO q, and so we would not expect to see a feature here.

Summing over HCD-pixel pairs in order to compute the full Lyα-HCD cross-correlation will average out most of the signal, but a small, asymmetric residual will remain, as seen in the right panel of Figure B.2. The contribution from each HCD will carry a similar signature but the signature will be centred at different values of $r_\parallel$ due to the different values of r_{Xq} for each X-q (HCD-QSO) pair. Certainly though, the sign of r_{Xq} will always be the same as an HCD is always less distant than its host QSO. Using `picca`'s definition of the sign of $r_\parallel$, this means that $r_{Xq} > 0$ for all X and q. As a result, we will see a reduction of the Lyα-HCD cross-correlation for all $r_\parallel > 0$, due to the strong reduction in δ_F at the centre of regions such as Region 1 in Figure B.3. This is only apparent for $r_\perp$ small as the reduced area shown in Figure B.3 is narrow. We also see a secondary effect: a boost in the Lyα-HCD cross-correlation for small, negative $r_\parallel$. This is a result of the small boost in δ_F on the right-hand side of Region 1 in Figure B.3, which appears at $r_\parallel < 0$ for HCDs that are very close to their host QSOs. This effect extends to larger values of $r_\perp$ due to the greater width of

the boosted area (relative to that which is reduced).

Whilst interesting, these effects are very small. In order to assess their visibility in current/future studies, we would need to carry out tests using a more realistic mock dataset. This would involve using the entire data reduction pipeline — including continuum fitting and the use of a distortion matrix — and is beyond the scope of this work. As an approximate comparison, we observe that the size of the deviation of points in the $0.0 < r_\perp < 4.0$ bin in the right panel of Figure B.2 is approximately an order of magnitude smaller than the size of the error bars in the uppermost two panels of Figure 2 of Pérez-Ràfols et al. (2018b)[2].

In order to mitigate this effect in the right panel of Figure 6.4, we measure the Lyα-HCD cross-correlation using only HCDs in the rest-frame wavelength range $[1040, 1100]$ Å. This imposes a minimum value of r_{Xq} in our measurements, and thus pushes the effect described above to large separations, beyond those over which we fit. Of course, making such an extreme cut in rest-frame wavelength greatly reduces the number of HCDs in our catalogue. In Chapter 6 we use approximately 30 times the number of skewers as DESI will have, and so this reduction does not cause us any concern. For studies from real surveys, however, maximising the scientific value of their data will be of much greater importance. As such, we would recommend the development of a new model to account for the effects described above using the measured Lyα-QSO cross-correlation. Alternatively, a catalogue of random HCDs, uncorrelated with the Lyα forest, could be generated and used to quantify these effects before accounting for them appropriately. Either way, further tests are needed in order to understand more fully the effect described in this Appendix, particularly if new modelling is required for future Lyα-HCD cross-correlation measurements.

[2]It should be noted that Pérez-Ràfols et al. (2018b) includes only HCDs at least 5,000 km s^{-1} away from their host quasar, equivalent to a rest-frame wavelength cut of approximately 1195 Å. We choose to use $\lambda_{r,\mathrm{HCD}} < \lambda_{\mathrm{Ly}\alpha}$ in the right panel of Figure B.2 in order to explain the relationship between the geometry of the problem and the observed effect more clearly.

BIBLIOGRAPHY

Abbott, B. P. et al. (Nov. 2017). "A gravitational-wave standard siren measurement of the Hubble constant". In: *Nature* 551.7678, pp. 85–88. DOI: 10.1038/nature24471. arXiv: 1710.05835 [astro-ph.CO] (cit. on p. 153).

Addison, G. E., G. Hinshaw, and M. Halpern (Dec. 2013). "Cosmological constraints from baryon acoustic oscillations and clustering of large-scale structure". In: *MNRAS* 436.2, pp. 1674–1683. DOI: 10.1093/mnras/stt1687. arXiv: 1304.6984 [astro-ph.CO] (cit. on pp. 57, 153).

Addison, G. E. et al. (Feb. 2018). "Elucidating ΛCDM: Impact of Baryon Acoustic Oscillation Measurements on the Hubble Constant Discrepancy". In: *ApJ* 853.2, 119, p. 119. DOI: 10.3847/1538-4357/aaa1ed. arXiv: 1707.06547 [astro-ph.CO] (cit. on p. 57).

Adelberger, E. G. et al. (Jan. 2011). "Solar fusion cross sections. II. The pp chain and CNO cycles". In: *Reviews of Modern Physics* 83.1, pp. 195–246. DOI: 10.1103/RevModPhys.83.195. arXiv: 1004.2318 [nucl-ex] (cit. on p. 154).

Adhikari, Saroj and Dragan Huterer (Jan. 2019). "A new measure of tension between experiments". In: *JCAP* 2019.1, 036, p. 036. DOI: 10.1088/1475-7516/2019/01/036. arXiv: 1806.04292 [astro-ph.CO] (cit. on p. 151).

Aihara, Hiroaki et al. (Jan. 2018). "The Hyper Suprime-Cam SSP Survey: Overview and survey design". In: *PASJ* 70, S4, S4. DOI: 10.1093/pasj/psx066. arXiv: 1704.05858 [astro-ph.IM] (cit. on p. 65).

Alam, Shadab et al. (Sept. 2017). "The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample". In: *MNRAS* 470.3, pp. 2617–2652. DOI: 10.1093/mnras/stx721. arXiv: 1607.03155 [astro-ph.CO] (cit. on pp. 58, 150).

Albrecht, Andreas et al. (Sept. 2006). "Report of the Dark Energy Task Force". In: *arXiv e-prints*, astro-ph/0609591, astro–ph/0609591. arXiv: astro-ph/0609591 [astro-ph] (cit. on pp. 69, 75, 76).

Alcock, C. and B. Paczynski (Oct. 1979). "An evolution free test for non-zero cosmological constant". In: *Nature* 281, p. 358. DOI: 10.1038/281358a0 (cit. on p. 76).

Alonso, David et al. (in prep.). *CoLoRe: Cosmological Lognormal Realisations* (cit. on pp. 108, 109).

Ancona, Marco, Cengiz Öztireli, and Markus Gross (Mar. 2019). "Explaining Deep Neural Networks with a Polynomial Time Algorithm for Shapley Values Approximation". In: *arXiv e-prints*. arXiv: 1903.10992 [cs.LG] (cit. on p. 162).

Anderson, Lauren et al. (Dec. 2012). "The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Release 9 spectroscopic galaxy sample". In: *MNRAS* 427.4, pp. 3435–3467. DOI: 10.1111/j.1365-2966.2012.22066.x. arXiv: 1203.6594 [astro-ph.CO] (cit. on p. 59).

Angelov, Plamen and Eduardo Soares (Dec. 2019). "Towards Explainable Deep Neural Networks (xDNN)". In: *arXiv e-prints*. arXiv: 1912.02523 [cs.LG] (cit. on p. 162).

Ata, Metin et al. (Feb. 2018). "The clustering of the SDSS-IV extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample: first measurement of baryon acoustic oscillations between redshift 0.8 and 2.2". In: *MNRAS* 473.4, pp. 4773–4794. DOI: 10.1093/mnras/stx2630. arXiv: 1705.06373 [astro-ph.CO] (cit. on pp. 58, 79, 98, 150).

Aubourg, Éric et al. (Dec. 2015). "Cosmological implications of baryon acoustic oscillation measurements". In: *Phys.Rev.D* 92.12, 123516, p. 123516. DOI: 10.1103/PhysRevD.92.123516. arXiv: 1411.1074 [astro-ph.CO] (cit. on p. 153).

Audren, Benjamin et al. (Feb. 2013). "Conservative constraints on early cosmology with MONTE PYTHON". In: *JCAP* 2013.2, 001, p. 001. DOI: 10.1088/1475-7516/2013/02/001. arXiv: 1210.7183 [astro-ph.CO] (cit. on p. 149).

Bahcall, John N. and E. E. Salpeter (Nov. 1965). "On the Interaction of Radiation from Distant Sources with the Intervening Medium." In: *ApJ* 142, pp. 1677–1680. DOI: 10.1086/148460 (cit. on p. 60).

Barkana, R. and A. Loeb (July 2001). "In the beginning: the first sources of light and the reionization of the universe". In: *physrep* 349.2, pp. 125–238. DOI: 10.1016/S0370-1573(01)00019-9. arXiv: astro-ph/0010468 [astro-ph] (cit. on p. 39).

Bartelmann, M. and P. Schneider (Jan. 2001). "Weak gravitational lensing". In: *physrep* 340.4-5, pp. 291–472. DOI: 10.1016/S0370-1573(00)00082-X. arXiv: astro-ph/9912508 [astro-ph] (cit. on p. 65).

Bautista, Julian E. et al. (May 2015). "Mock Quasar-Lyman-α forest data-sets for the SDSS-III Baryon Oscillation Spectroscopic Survey". In: *JCAP* 2015.5, 060, p. 060. DOI: 10.1088/1475-7516/2015/05/060. arXiv: 1412.0658 [astro-ph.CO] (cit. on pp. 106, 126, 145).

Bautista, Julian E. et al. (June 2017). "Measurement of baryon acoustic oscillation correlations at z = 2.3 with SDSS DR12 Lyα-Forests". In: *A&A* 603, A12, A12. DOI: 10.1051/0004-6361/201730533. arXiv: 1702.00176 [astro-ph.CO] (cit. on pp. 63, 106, 120, 121, 127, 129, 145, 151).

Bautista, Julian E. et al. (Aug. 2018). "The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations at Redshift of 0.72 with the DR14 Luminous Red Galaxy Sample". In: *ApJ* 863.1, 110, p. 110. DOI: 10.3847/1538-4357/aacea5. arXiv: 1712.08064 [astro-ph.CO] (cit. on p. 150).

Bautista, Julian E. et al. (Jan. 2021). "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". In: *MNRAS* 500.1, pp. 736–762. DOI: 10.1093/mnras/staa2800. arXiv: 2007.08993 [astro-ph.CO] (cit. on pp. 58, 64).

Becker, George D. et al. (Apr. 2013). "A refined measurement of the mean transmitted flux in the Lyα forest over $2 < z < 5$ using composite quasar spectra". In: *MNRAS* 430.3, pp. 2067–2081. DOI: 10.1093/mnras/stt031. arXiv: 1208.2584 [astro-ph.CO] (cit. on p. 116).

Bennett, C. L. et al. (June 1996). "Four-Year COBE DMR Cosmic Microwave Background Observations: Maps and Basic Results". In: *ApJ* 464, p. L1. DOI: 10.1086/310075. arXiv: astro-ph/9601067 [astro-ph] (cit. on p. 54).

Bennett, C. L. et al. (Jan. 2003). "The Microwave Anisotropy Probe Mission". In: *ApJ* 583.1, pp. 1–23. DOI: 10.1086/345346. arXiv: astro-ph/0301158 [astro-ph] (cit. on p. 54).

Bernal, José Luis, Licia Verde, and Adam G. Riess (Oct. 2016). "The trouble with H_0". In: *JCAP* 2016.10, 019, p. 019. DOI: 10.1088/1475-7516/2016/10/019. arXiv: 1607.05617 [astro-ph.CO] (cit. on p. 68).

Bernardeau, F. et al. (Sept. 2002). "Large-scale structure of the Universe and cosmological perturbation theory". In: *physrep* 367.1-3, pp. 1–248. DOI: 10.1016/S0370-1573(02)00135-7. arXiv: astro-ph/0112551 [astro-ph] (cit. on pp. 51, 164).

Beutler, Florian et al. (Oct. 2011). "The 6dF Galaxy Survey: baryon acoustic oscillations and the local Hubble constant". In: *MNRAS* 416.4, pp. 3017–3032. DOI: 10.1111/j.1365-2966.2011.19250.x. arXiv: 1106.3366 [astro-ph.CO] (cit. on pp. 58, 150).

Bi, H. G., G. Boerner, and Y. Chu (Dec. 1992). "An alternative model for the Ly-alpha absorption forest." In: *A&A* 266, pp. 1–5 (cit. on p. 106).

Bi, Hongguang and Arthur F. Davidsen (Apr. 1997). "Evolution of Structure in the Intergalactic Medium and the Nature of the Lyα Forest". In: *ApJ* 479.2, pp. 523–542. DOI: 10.1086/303908. arXiv: astro-ph/9611062 [astro-ph] (cit. on pp. 112, 113).

Bianchi, Davide et al. (Dec. 2018). "Unbiased clustering estimates with the DESI fibre assignment". In: *MNRAS* 481.2, pp. 2338–2348. DOI: 10.1093/mnras/sty2377. arXiv: 1805.00951 [astro-ph.CO] (cit. on p. 73).

BICEP2 Collaboration et al. (Nov. 2018). "Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season". In: *Phys.Rev.Lett* 121.22, 221301, p. 221301. DOI: 10.1103/PhysRevLett.121.221301. arXiv: 1810.05216 [astro-ph.CO] (cit. on p. 65).

Blake, Chris et al. (Dec. 2011). "The WiggleZ Dark Energy Survey: mapping the distance-redshift relation with baryon acoustic oscillations". In: *MNRAS* 418.3, pp. 1707–1724. DOI: 10.1111/j.1365-2966.2011.19592.x. arXiv: 1108.2635 [astro-ph.CO] (cit. on p. 58).

Blandford, R. D. et al. (Aug. 1991). "The distortion of distant galaxy images by large-scale structure." In: *MNRAS* 251, p. 600. DOI: 10.1093/mnras/251.4.600 (cit. on p. 64).

Blanton, Michael R. et al. (July 2017). "Sloan Digital Sky Survey IV: Mapping the Milky Way, Nearby Galaxies, and the Distant Universe". In: *AJ* 154.1, 28, p. 28. DOI: 10.3847/1538-3881/aa7567. arXiv: 1703.00052 [astro-ph.GA] (cit. on pp. 58, 73, 80, 141).

Blomqvist, Michael et al. (Nov. 2015). "Broadband distortion modeling in Lyman-α forest BAO fitting". In: *JCAP* 2015.11, 034, p. 034. DOI: 10.1088/1475-7516/2015/11/034. arXiv: 1504.06656 [astro-ph.CO] (cit. on p. 145).

Blomqvist, Michael et al. (May 2018). "The triply-ionized carbon forest from eBOSS: cosmological correlations with quasars in SDSS-IV DR14". In: *JCAP* 2018.5, 029, p. 029. DOI: 10.1088/1475-7516/2018/05/029. arXiv: 1801.01852 [astro-ph.CO] (cit. on p. 126).

Blomqvist, Michael et al. (Sept. 2019). "Baryon acoustic oscillations from the cross-correlation of Lyα absorption and quasars in eBOSS DR14". In: *A&A* 629, A86, A86. DOI: 10.1051/0004-6361/201935641. arXiv: 1904.03430 [astro-ph.CO] (cit. on pp. 63, 79, 106, 120, 121, 141, 150).

Bolton, Adam S. and David J. Schlegel (Feb. 2010). "Spectro-Perfectionism: An Algorithmic Framework for Photon Noise-Limited Extraction of Optical Fiber Spectroscopy". In: *PASP* 122.888, p. 248. DOI: 10.1086/651008. arXiv: 0911.2689 [astro-ph.IM] (cit. on pp. 73, 74).

Bolton, Adam S. et al. (Nov. 2012). "Spectral Classification and Redshift Measurement for the SDSS-III Baryon Oscillation Spectroscopic Survey". In: *AJ* 144.5, 144, p. 144. DOI: 10.1088/0004-6256/144/5/144. arXiv: 1207.7326 [astro-ph.CO] (cit. on pp. 80–83, 87).

Bovy, Jo et al. (Mar. 2011). "Think Outside the Color Box: Probabilistic Target Selection and the SDSS-XDQSO Quasar Targeting Catalog". In: *ApJ* 729.2, 141, p. 141. DOI: 10.1088/0004-637X/729/2/141. arXiv: 1011.6392 [astro-ph.CO] (cit. on p. 72).

Bovy, Jo et al. (Apr. 2012). "Photometric Redshifts and Quasar Probabilities from a Single, Data-driven Generative Model". In: *ApJ* 749.1, 41, p. 41. DOI: 10.1088/0004-637X/749/1/41. arXiv: 1105.3975 [astro-ph.CO] (cit. on p. 72).

Brinckmann, Thejs and Julien Lesgourgues (Apr. 2018). "MontePython 3: boosted MCMC sampler and other features". In: *arXiv e-prints*. arXiv: 1804.07261 [astro-ph.CO] (cit. on p. 149).

Bromm, Volker and Naoki Yoshida (Sept. 2011). "The First Galaxies". In: *ARA&A* 49.1, pp. 373–407. DOI: 10.1146/annurev-astro-081710-102608. arXiv: 1102.4638 [astro-ph.CO] (cit. on p. 39).

Burden, Angela et al. (Mar. 2017). "Mitigating the impact of the DESI fiber assignment on galaxy clustering". In: *JCAP* 2017.3, 001, p. 001. DOI: 10.1088/1475-7516/2017/03/001. arXiv: 1611.04635 [astro-ph.CO] (cit. on p. 73).

Busca, N. G. et al. (Apr. 2013). "Baryon acoustic oscillations in the Lyα forest of BOSS quasars". In: *A&A* 552, A96, A96. DOI: 10.1051/0004-6361/201220724. arXiv: 1211.2616 [astro-ph.CO] (cit. on pp. 63, 106, 159).

Busca, Nicolas and Christophe Balland (Aug. 2018). "QuasarNET: Human-level spectral classification and redshifting with Deep Neural Networks". In: *arXiv e-prints*. arXiv: 1808.09955 [astro-ph.IM] (cit. on pp. 80, 84–86, 169).

Castorina, Emanuele et al. (Sept. 2019). "Redshift-weighted constraints on primordial non-Gaussianity from the clustering of the eBOSS DR14 quasars in Fourier space". In: *JCAP* 2019.9, 010, p. 010. DOI: 10.1088/1475-7516/2019/09/010. arXiv: 1904.08859 [astro-ph.CO] (cit. on p. 76).

Cen, Renyue et al. (Dec. 1994). "Gravitational Collapse of Small-Scale Structure as the Origin of the Lyman-Alpha Forest". In: *ApJ* 437, p. L9. DOI: 10.1086/187670. arXiv: astro-ph/9409017 [astro-ph] (cit. on p. 106).

Charnock, Tom, Richard A. Battye, and Adam Moss (June 2017). "Planck data versus large scale structure: Methods to quantify discordance". In: *Phys.Rev.D* 95.12, 123535, p. 123535. DOI: 10.1103/PhysRevD.95.123535. arXiv: 1703.05959 [astro-ph.CO] (cit. on p. 151).

Colless, Matthew et al. (Dec. 2001). "The 2dF Galaxy Redshift Survey: spectra and redshifts". In: *MNRAS* 328.4, pp. 1039–1063. DOI: 10.1046/j.1365-8711.2001.04902.x. arXiv: astro-ph/0106498 [astro-ph] (cit. on pp. 58, 70).

Cooke, Ryan J., Max Pettini, and Charles C. Steidel (Mar. 2018). "One Percent Determination of the Primordial Deuterium Abundance". In: *ApJ* 855.2, 102, p. 102. DOI: 10.3847/1538-4357/aaab53. arXiv: 1710.11129 [astro-ph.CO] (cit. on pp. 153, 155).

Cooke, Ryan J. et al. (Oct. 2016). "The Primordial Deuterium Abundance of the Most Metal-poor Damped Lyman-α System". In: *ApJ* 830.2, 148, p. 148. DOI: 10.3847/0004-637X/830/2/148. arXiv: 1607.03900 [astro-ph.CO] (cit. on p. 153).

Crocce, Martín and Román Scoccimarro (Jan. 2008). "Nonlinear evolution of baryon acoustic oscillations". In: *Phys.Rev.D* 77.2, 023533, p. 023533. DOI: 10.1103/PhysRevD.77.023533. arXiv: 0704.2783 [astro-ph] (cit. on p. 59).

Croft, Rupert A. C. et al. (Mar. 1998). "Recovery of the Power Spectrum of Mass Fluctuations from Observations of the Lyα Forest". In: *ApJ* 495.1, pp. 44–62. DOI: 10.1086/305289. arXiv: astro-ph/9708018 [astro-ph] (cit. on pp. 108, 113).

Croom, S. et al. (Oct. 2004). "The 2DF QSO Redshift Survey". In: *Multiwavelength AGN Surveys*. Ed. by Raúl Mújica and Roberto Maiolino, pp. 57–62. DOI: 10.1142/9789812702432_0015 (cit. on p. 80).

Croom, S. M. et al. (Apr. 2001). "The 2dF QSO Redshift Survey - V. The 10k catalogue". In: *MNRAS* 322.4, pp. L29–L36. DOI: 10.1046/j.1365-8711.2001.04474.x. arXiv: astro-ph/0104095 [astro-ph] (cit. on pp. 80, 82).

Cuceu, Andrei et al. (Oct. 2019). "Baryon Acoustic Oscillations and the Hubble constant: past, present and future". In: *JCAP* 2019.10, 044, p. 044. DOI: 10.1088/1475-7516/2019/10/044. arXiv: 1906.11628 [astro-ph.CO] (cit. on pp. 149, 151, 155, 156).

Cyburt, Richard H. et al. (Jan. 2016). "Big bang nucleosynthesis: Present status". In: *Reviews of Modern Physics* 88.1, 015004, p. 015004. DOI: 10.1103/RevModPhys.88.015004. arXiv: 1505.01076 [astro-ph.CO] (cit. on p. 153).

Dawson, Kyle S. et al. (Jan. 2013). "The Baryon Oscillation Spectroscopic Survey of SDSS-III". In: *AJ* 145.1, 10, p. 10. DOI: 10.1088/0004-6256/145/1/10. arXiv: 1208.0022 [astro-ph.CO] (cit. on pp. 20, 70, 80, 81, 106, 141).

Dawson, Kyle S. et al. (Feb. 2016). "The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data". In: *AJ* 151.2, 44, p. 44. DOI: 10.3847/0004-6256/151/2/44. arXiv: 1508.04473 [astro-ph.CO] (cit. on pp. 58, 70, 80, 81, 106, 141).

de Jong, J. T. A. et al. (Dec. 2013). "The Kilo-Degree Survey". In: *The Messenger* 154, pp. 44–46 (cit. on p. 65).

de Mattia, Arnaud et al. (Dec. 2020). "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". In: *MNRAS*. DOI: 10.1093/mnras/staa3891. arXiv: 2007.09008 [astro-ph.CO] (cit. on pp. 58, 64).

de Sainte Agathe, Victoria et al. (Sept. 2019). "Baryon acoustic oscillations at z = 2.34 from the correlations of Lyα absorption in eBOSS DR14". In: *A&A* 629, A85, A85. DOI: 10.1051/0004-6361/201935638. arXiv: 1904.03400 [astro-ph.CO] (cit. on pp. 63, 79, 106, 120, 121, 127, 129, 141, 150).

Dehnen, W. and J. I. Read (May 2011). "N-body simulations of gravitational dynamics". In: *European Physical Journal Plus* 126, 55, p. 55. DOI: 10.1140/epjp/i2011-11055-3. arXiv: 1105.1082 [astro-ph.IM] (cit. on p. 51).

Dekker, Hans et al. (Aug. 2000). "Design, construction, and performance of UVES, the echelle spectrograph for the UT2 Kueyen Telescope at the ESO Paranal Observatory". In: *Optical and IR Telescope Instrumentation and Detectors*. Ed. by Masanori Iye and Alan F. Moorwood. Vol. 4008. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 534–545. DOI: 10.1117/12.395512 (cit. on p. 71).

Delubac, Timothée et al. (Feb. 2015). "Baryon acoustic oscillations in the Lyα forest of BOSS DR11 quasars". In: *A&A* 574, A59, A59. DOI: 10.1051/0004-6361/201423969. arXiv: 1404.1801 [astro-ph.CO] (cit. on pp. 63, 106, 150).

DESI Collaboration et al. (Oct. 2016a). "The DESI Experiment Part I: Science,Targeting, and Survey Design". In: *arXiv e-prints*. arXiv: 1611.00036 [astro-ph.IM] (cit. on pp. 21, 59, 71, 72, 75, 76, 81, 88, 95, 96, 106, 155, 156, 159, 160).

— (Oct. 2016b). "The DESI Experiment Part II: Instrument Design". In: *arXiv e-prints*. arXiv: 1611.00037 [astro-ph.IM] (cit. on pp. 59, 71–73, 81, 88, 106, 159).

Desjacques, Vincent, Donghui Jeong, and Fabian Schmidt (Feb. 2018). "Large-scale galaxy bias". In: *physrep* 733, pp. 1–193. DOI: 10.1016/j.physrep.2017.12.002. arXiv: 1611.09787 [astro-ph.CO] (cit. on p. 51).

Dey, Arjun et al. (Aug. 2016). "Mosaic3: a red-sensitive upgrade for the prime focus camera at the Mayall 4m telescope". In: *Ground-based and Airborne Instrumentation for Astronomy VI*. Ed. by Christopher J. Evans, Luc Simard, and Hideki Takami. Vol. 9908. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 99082C. DOI: 10.1117/12.2231488 (cit. on p. 72).

Dey, Arjun et al. (May 2019). "Overview of the DESI Legacy Imaging Surveys". In: *AJ* 157.5, 168, p. 168. DOI: 10.3847/1538-3881/ab089d. arXiv: 1804.08657 [astro-ph.IM] (cit. on pp. 72, 134).

Dicke, R. H. et al. (July 1965). "Cosmic Black-Body Radiation." In: *ApJ* 142, pp. 414–419. DOI: 10.1086/148306 (cit. on p. 54).

Dodelson, Scott (2003). *Modern cosmology* (cit. on pp. 25, 43, 45).

Doshi-Velez, Finale and Been Kim (Feb. 2017). "Towards A Rigorous Science of Interpretable Machine Learning". In: *arXiv e-prints*. arXiv: 1702.08608 [stat.ML] (cit. on p. 162).

du Mas des Bourboux, Hélion et al. (Dec. 2017). "Baryon acoustic oscillations from the complete SDSS-III Lyα-quasar cross-correlation function at z = 2.4". In: *A&A* 608, A130, A130. DOI: 10.1051/0004-6361/201731731. arXiv: 1708.02225 [astro-ph.CO] (cit. on pp. 63, 106, 116, 117, 120, 121, 123, 124, 127, 151).

du Mas des Bourboux, Hélion et al. (June 2019). "The Extended Baryon Oscillation Spectroscopic Survey: Measuring the Cross-correlation between the Mg II Flux Transmission Field and Quasars and Galaxies at z = 0.59". In: *ApJ* 878.1, 47, p. 47. DOI: 10.3847/1538-4357/ab1d49. arXiv: 1901.01950 [astro-ph.CO] (cit. on p. 126).

du Mas des Bourboux, Hélion et al. (Oct. 2020). "The Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations with Lyα Forests". In: *ApJ* 901.2, 153, p. 153. DOI: 10.3847/1538-4357/abb085. arXiv: 2007.08995 [astro-ph.CO] (cit. on pp. 63, 79, 106, 119, 120, 140–145, 147, 148, 152).

eBOSS Collaboration et al. (July 2020). "The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological Implications from two Decades of Spectroscopic Surveys at the Apache Point observatory". In: *arXiv e-prints*. arXiv: 2007.08991 [astro-ph.CO] (cit. on pp. 33–36, 64, 66, 67, 105, 146, 155).

Efstathiou, G., W. J. Sutherland, and S. J. Maddox (Dec. 1990). "The cosmological constant and cold dark matter". In: *Nature* 348.6303, pp. 705–707. DOI: 10.1038/348705a0 (cit. on p. 33).

Efstathiou, George and Steven Gratton (May 2020). "The evidence for a spatially flat Universe". In: *MNRAS* 496.1, pp. L91–L95. DOI: 10.1093/mnrasl/slaa093. arXiv: 2002.06892 [astro-ph.CO] (cit. on p. 66).

Einstein, Albert (Jan. 1915). "Die Feldgleichungen der Gravitation". In: *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin*, pp. 844–847 (cit. on pp. 19, 25).

Eisenstein, Daniel J., Hee-Jong Seo, and Martin White (Aug. 2007). "On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter". In: *ApJ* 664.2, pp. 660–674. DOI: 10.1086/518755. arXiv: astro-ph/0604361 [astro-ph] (cit. on pp. 48, 50).

Eisenstein, Daniel J. et al. (Aug. 2007). "Improving Cosmological Distance Measurements by Reconstruction of the Baryon Acoustic Peak". In: *ApJ* 664.2, pp. 675–679. DOI: 10.1086/518712. arXiv: astro-ph/0604362 [astro-ph] (cit. on p. 59).

Eisenstein, Daniel J. et al. (Sept. 2011). "SDSS-III: Massive Spectroscopic Surveys of the Distant Universe, the Milky Way, and Extra-Solar Planetary Systems". In: *AJ* 142.3, 72, p. 72. DOI: 10.1088/0004-6256/142/3/72. arXiv: 1101.1529 [astro-ph.IM] (cit. on pp. 80, 141).

Epstein, R. I., J. M. Lattimer, and D. N. Schramm (Sept. 1976). "The origin of deuterium". In: *Nature* 263, pp. 198–202. DOI: 10.1038/263198a0 (cit. on p. 153).

Fan, Fenglei et al. (Jan. 2020). "On Interpretability of Artificial Neural Networks: A Survey". In: *arXiv e-prints*. arXiv: 2001.02522 [cs.LG] (cit. on p. 162).

Farr, James, Andreu Font-Ribera, and Andrew Pontzen (Nov. 2020). "Optimal strategies for identifying quasars in DESI". In: *JCAP* 2020.11, 015, p. 015. DOI: 10.1088/1475-7516/2020/11/015. arXiv: 2007.10348 [astro-ph.CO] (cit. on p. 79).

Farr, James et al. (Mar. 2020). "LyaCoLoRe: synthetic datasets for current and future Lyman-α forest BAO surveys". In: *JCAP* 2020.3, 068, p. 068. DOI: 10.1088/1475-7516/2020/03/068. arXiv: 1912.02763 [astro-ph.CO] (cit. on p. 105).

Feeney, Stephen M., Daniel J. Mortlock, and Niccolò Dalmasso (May 2018). "Clarifying the Hubble constant tension with a Bayesian hierarchical model of the local distance ladder". In: *MNRAS* 476.3, pp. 3861–3882. DOI: 10.1093/mnras/sty418. arXiv: 1707.00007 [astro-ph.CO] (cit. on p. 68).

Feng, Jonathan L. (Sept. 2010). "Dark Matter Candidates from Particle Physics and Methods of Detection". In: *ARA&A* 48, pp. 495–545. DOI: 10.1146/annurev-astro-082708-101659. arXiv: 1003.0904 [astro-ph.CO] (cit. on p. 32).

Fixsen, D. J. (Dec. 2009). "The Temperature of the Cosmic Microwave Background". In: *ApJ* 707.2, pp. 916–920. DOI: 10.1088/0004-637X/707/2/916. arXiv: 0911.1955 [astro-ph.CO] (cit. on pp. 38, 54, 150).

Flaugher, B. et al. (Nov. 2015). "The Dark Energy Camera". In: *AJ* 150.5, 150, p. 150. DOI: 10.1088/0004-6256/150/5/150. arXiv: 1504.02900 [astro-ph.IM] (cit. on p. 72).

Flaugher, Brenna (Jan. 2005). "The Dark Energy Survey". In: *International Journal of Modern Physics A* 20.14, pp. 3121–3123. DOI: 10.1142/S0217751X05025917 (cit. on p. 65).

Follin, B. and L. Knox (July 2018). "Insensitivity of the distance ladder Hubble constant determination to Cepheid calibration modelling choices". In: *MNRAS* 477.4, pp. 4534–4542. DOI: 10.1093/mnras/sty720. arXiv: 1707.01175 [astro-ph.CO] (cit. on p. 68).

Font-Ribera, Andreu, Patrick McDonald, and Jordi Miralda-Escudé (Jan. 2012). "Generating mock data sets for large-scale Lyman-α forest correlation measurements". In: *JCAP* 2012.1, 001, p. 001. DOI: 10.1088/1475-7516/2012/01/001. arXiv: 1108.5606 [astro-ph.CO] (cit. on p. 106).

Font-Ribera, Andreu and Jordi Miralda-Escudé (July 2012). "The effect of high column density systems on the measurement of the Lyman-α forest correlation function". In: *JCAP* 2012.7, 028, p. 028. DOI: 10.1088/1475-7516/2012/07/028. arXiv: 1205.2018 [astro-ph.CO] (cit. on pp. 125, 145).

Font-Ribera, Andreu et al. (May 2013). "The large-scale quasar-Lyman α forest cross-correlation from BOSS". In: *JCAP* 2013.5, 018, p. 018. DOI: 10.1088/1475-7516/2013/05/018. arXiv: 1303.1937 [astro-ph.CO] (cit. on pp. 63, 106).

Font-Ribera, Andreu et al. (May 2014). "Quasar-Lyman α forest cross-correlation from BOSS DR11: Baryon Acoustic Oscillations". In: *JCAP* 2014.5, 027, p. 027. DOI: 10.1088/1475-7516/2014/05/027. arXiv: 1311.1767 [astro-ph.CO] (cit. on pp. 63, 106, 151).

Freedman, Wendy L. et al. (Mar. 2020). "Calibration of the Tip of the Red Giant Branch". In: *ApJ* 891.1, 57, p. 57. DOI: 10.3847/1538-4357/ab7339. arXiv: 2002.01550 [astro-ph.GA] (cit. on p. 152).

Gil-Marín, Héctor et al. (Aug. 2020). "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". In: *MNRAS* 498.2, pp. 2492–2531. DOI: 10.1093/mnras/staa2455. arXiv: 2007.08994 [astro-ph.CO] (cit. on pp. 58, 64).

Glazebrook, Karl et al. (Nov. 2005). "Monster redshift surveys through dispersive slitless imaging: The Baryon Oscillation Probe [review article]". In: *NewAR* 49, pp. 374–378. DOI: 10.1016/j.newar.2005.08.007. arXiv: astro-ph/0410037 [astro-ph] (cit. on p. 59).

Gontcho A Gontcho, Satya, Jordi Miralda-Escudé, and Nicolás G. Busca (July 2014). "On the effect of the ionizing background on the Lyα forest autocorrelation function". In: *MNRAS* 442.1, pp. 187–195. DOI: 10.1093/mnras/stu860. arXiv: 1404.7425 [astro-ph.CO] (cit. on p. 125).

Gontcho A Gontcho, Satya et al. (Oct. 2018). "Quasar - CIV forest cross-correlation with SDSS DR12". In: *MNRAS* 480.1, pp. 610–622. DOI: 10.1093/mnras/sty1817. arXiv: 1712.09886 [astro-ph.CO] (cit. on pp. 119, 126).

Górski, K. M. et al. (Apr. 2005). "HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere". In: *ApJ* 622.2, pp. 759–771. DOI: 10.1086/427976. arXiv: astro-ph/0409513 [astro-ph] (cit. on p. 120).

Gott J. R., III et al. (Dec. 1974). "An unbound universe?" In: *ApJ* 194, pp. 543–553. DOI: 10.1086/153273 (cit. on p. 153).

Gunn, James E. and Bruce A. Peterson (Nov. 1965). "On the Density of Neutral Hydrogen in Intergalactic Space." In: *ApJ* 142, pp. 1633–1636. DOI: 10.1086/148444 (cit. on pp. 60, 112).

Gunn, James E. et al. (Apr. 2006). "The 2.5 m Telescope of the Sloan Digital Sky Survey". In: *AJ* 131.4, pp. 2332–2359. DOI: 10.1086/500975. arXiv: astro-ph/0602326 [astro-ph] (cit. on pp. 80, 88).

Guo, Zhiyuan and Paul Martini (July 2019). "Classification of Broad Absorption Line Quasars with a Convolutional Neural Network". In: *ApJ* 879.2, 72, p. 72. DOI: 10.3847/1538-4357/ab2590. arXiv: 1901.04506 [astro-ph.GA] (cit. on p. 85).

Handley, W. J., M. P. Hobson, and A. N. Lasenby (June 2015a). "polychord: nested sampling for cosmology." In: *MNRAS* 450, pp. L61–L65. DOI: 10.1093/mnrasl/slv047. arXiv: 1502.01856 [astro-ph.CO] (cit. on p. 154).

— (Nov. 2015b). "POLYCHORD: next-generation nested sampling". In: *MNRAS* 453.4, pp. 4384–4398. DOI: 10.1093/mnras/stv1911. arXiv: 1506.00171 [astro-ph.IM] (cit. on p. 154).

Handley, Will and Pablo Lemos (July 2019a). "Quantifying dimensionality: Bayesian cosmological model complexities". In: *Phys.Rev.D* 100.2, 023512, p. 023512. DOI: 10.1103/PhysRevD.100.023512. arXiv: 1903.06682 [astro-ph.CO] (cit. on p. 151).

Handley, Will and Pablo Lemos (Aug. 2019b). "Quantifying tensions in cosmological parameters: Interpreting the DES evidence ratio". In: *Phys.Rev.D* 100.4, 043504, p. 043504. DOI: 10.1103/PhysRevD.100.043504. arXiv: 1902.04029 [astro-ph.CO] (cit. on p. 151).

Hewett, Paul C., Craig B. Foltz, and Frederic H. Chaffee (Apr. 1995). "The Large Bright Quasar Survey.VI.Quasar Catalog and Survey Parameters". In: *AJ* 109, p. 1498. DOI: 10.1086/117380 (cit. on p. 80).

Hiss, Hector et al. (Sept. 2018). "A New Measurement of the Temperature-density Relation of the IGM from Voigt Profile Fitting". In: *ApJ* 865.1, 42, p. 42. DOI: 10.3847/1538-4357/aada86. arXiv: 1710.00700 [astro-ph.CO] (cit. on pp. 117, 179).

Hou, Jiamin et al. (Jan. 2021). "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". In: *MNRAS* 500.1, pp. 1201–1221. DOI: 10.1093/mnras/staa3234. arXiv: 2007.08998 [astro-ph.CO] (cit. on pp. 58, 64, 79).

Hu, Wayne and Scott Dodelson (Jan. 2002). "Cosmic Microwave Background Anisotropies". In: *ARA&A* 40, pp. 171–216. DOI: 10.1146/annurev.astro.40.060401.093926. arXiv: astro-ph/0110414 [astro-ph] (cit. on p. 55).

Hu, Wayne, Naoshi Sugiyama, and Joseph Silk (Mar. 1997). "The physics of microwave background anisotropies". In: *Nature* 386.6620, pp. 37–43. DOI: 10.1038/386037a0. arXiv: astro-ph/9504057 [astro-ph] (cit. on p. 55).

Hu, Wayne and Martin White (Nov. 1996). "Acoustic Signatures in the Cosmic Microwave Background". In: *ApJ* 471, p. 30. DOI: 10.1086/177951. arXiv: astro-ph/9602019 [astro-ph] (cit. on p. 48).

— (Apr. 1997). "The Damping Tail of Cosmic Microwave Background Anisotropies". In: *ApJ* 479.2, pp. 568–579. DOI: 10.1086/303928. arXiv: astro-ph/9609079 [astro-ph] (cit. on p. 54).

Hubble, E. P. (Dec. 1926). "Extragalactic nebulae." In: *ApJ* 64, pp. 321–369. DOI: 10.1086/143018 (cit. on p. 19).

Hubble, Edwin (Jan. 1929). "A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae". In: *Contributions from the Mount Wilson Observatory* 3, pp. 23–28 (cit. on pp. 19, 28, 29).

Hui, Lam and Nickolay Y. Gnedin (Nov. 1997). "Equation of state of the photoionized intergalactic medium". In: *MNRAS* 292.1, pp. 27–42. DOI: 10.1093/mnras/292.1.27. arXiv: astro-ph/9612232 [astro-ph] (cit. on p. 112).

Hui, Lam, Nickolay Y. Gnedin, and Yu Zhang (Sept. 1997). "The Statistics of Density Peaks and the Column Density Distribution of the Lyα Forest". In: *ApJ* 486.2, pp. 599–622. DOI: 10.1086/304539. arXiv: astro-ph/9608157 [astro-ph] (cit. on p. 112).

Hutchinson, Timothy A. et al. (Dec. 2016). "Redshift Measurement and Spectral Classification for eBOSS Galaxies with the redmonster Software". In: *AJ* 152.6, 205, p. 205. DOI: 10.3847/0004-6256/152/6/205. arXiv: 1607.02432 [astro-ph.IM] (cit. on p. 80).

Inman, Henry F. and Edwin L. Bradley Jr. (1989). "The overlapping coefficient as a measure of agreement between probability distributions and point estimation of the overlap of two normal densities". In: *Communications in Statistics - Theory and Methods* 18.10, pp. 3851–3874. DOI: 10.1080/03610928908830127. eprint: https://doi.org/10.1080/03610928908830127. URL: https://doi.org/10.1080/03610928908830127 (cit. on p. 151).

Iršič, Vid et al. (July 2017). "New constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman-α forest data". In: *Phys.Rev.D* 96.2, 023522, p. 023522. DOI: 10.1103/PhysRevD.96.023522. arXiv: 1702.01764 [astro-ph.CO] (cit. on p. 32).

Jain, Bhuvnesh and Uroš Seljak (July 1997). "Cosmological Model Predictions for Weak Lensing: Linear and Nonlinear Regimes". In: *ApJ* 484.2, pp. 560–573. DOI: 10.1086/304372. arXiv: astro-ph/9611077 [astro-ph] (cit. on p. 64).

James, F. and M. Roos (Dec. 1975). "Minuit - a system for function minimization and analysis of the parameter errors and correlations". In: *Computer Physics Communications* 10.6, pp. 343–367. DOI: 10.1016/0010-4655(75)90039-9 (cit. on p. 116).

Jeans, J. H. (Jan. 1902). "The Stability of a Spherical Nebula". In: *Philosophical Transactions of the Royal Society of London Series A* 199, pp. 1–53. DOI: 10.1098/rsta.1902.0012 (cit. on p. 46).

Jones, D. Heath et al. (Dec. 2004). "The 6dF Galaxy Survey: samples, observational techniques and the first data release". In: *MNRAS* 355.3, pp. 747–763. DOI: 10.1111/j.1365-2966.2004.08353.x. arXiv: astro-ph/0403501 [astro-ph] (cit. on p. 70).

Kaiser, Nick (July 1987). "Clustering in real space and in redshift space". In: *MNRAS* 227, pp. 1–21. DOI: 10.1093/mnras/227.1.1 (cit. on pp. 63, 115, 121, 122, 177).

— (Apr. 1992). "Weak Gravitational Lensing of Distant Galaxies". In: *ApJ* 388, p. 272. DOI: 10.1086/171151 (cit. on p. 64).

Kirkby, David et al. (Mar. 2013). "Fitting methods for baryon acoustic oscillations in the Lyman-α forest fluctuations in BOSS data release 9". In: *JCAP* 2013.3, 024, p. 024. DOI: 10.1088/1475-7516/2013/03/024. arXiv: 1301.3456 [astro-ph.CO] (cit. on pp. 63, 106, 107, 122, 159).

Kirkby, David et al. (Sept. 2016). *Quick Simulations Of Fiber Spectrograph Response V0.5*. Version v0.5. DOI: 10.5281/zenodo.154130 (cit. on p. 143).

Kochanek, Izabela (Apr. 2016). "Towards the study of ^{2}H(p, γ)^{3}He reaction in the Big Bang Nucleosynthesis energy range in LUNA". In: *Journal of Physics Conference Series*. Vol. 703. Journal of Physics Conference Series, p. 012023. DOI: 10.1088/1742-6596/703/1/012023 (cit. on p. 156).

Kosowsky, Arthur and Michael S. Turner (Aug. 1995). "CBR anisotropy and the running of the scalar spectral index". In: *Phys.Rev.D* 52.4, R1739–R1743. DOI: 10.1103/PhysRevD.52.R1739. arXiv: astro-ph/9504071 [astro-ph] (cit. on p. 37).

Landy, Stephen D. and Alexander S. Szalay (July 1993). "Bias and Variance of Angular Correlation Functions". In: *ApJ* 412, p. 64. DOI: 10.1086/172900 (cit. on p. 177).

Laureijs, R. et al. (Oct. 2011). "Euclid Definition Study Report". In: *arXiv e-prints*. arXiv: 1110.3193 [astro-ph.CO] (cit. on pp. 59, 159).

Laurent, Pierre et al. (Nov. 2016). "A 14 h^{-3} Gpc^{3} study of cosmic homogeneity using BOSS DR12 quasar sample". In: *JCAP* 2016.11, 060, p. 060. DOI: 10.1088/1475-7516/2016/11/060. arXiv: 1602.09010 [astro-ph.CO] (cit. on p. 119).

Laurent, Pierre et al. (July 2017). "Clustering of quasars in SDSS-IV eBOSS: study of potential systematics and bias determination". In: *JCAP* 2017.7, 017, p. 017. DOI: 10.1088/1475-7516/2017/07/017. arXiv: 1705.04718 [astro-ph.CO] (cit. on p. 98).

Le Goff, J. M. et al. (Oct. 2011). "Simulations of BAO reconstruction with a quasar Ly-α survey". In: *A&A* 534, A135, A135. DOI: 10.1051/0004-6361/201117736. arXiv: 1107.4233 [astro-ph.CO] (cit. on p. 106).

Lesgourgues, Julien et al. (2013). *Neutrino Cosmology* (cit. on pp. 32, 76).

Lewis, Antony and Anthony Challinor (June 2006). "Weak gravitational lensing of the CMB". In: *physrep* 429.1, pp. 1–65. DOI: 10.1016/j.physrep.2006.03.002. arXiv: astro-ph/0601594 [astro-ph] (cit. on p. 55).

Lewis, Antony, Anthony Challinor, and Anthony Lasenby (Aug. 2000). "Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models". In: *ApJ* 538.2, pp. 473–476. DOI: 10.1086/309179. arXiv: astro-ph/9911177 [astro-ph] (cit. on pp. 118, 153).

Lyke, Brad W. et al. (Sept. 2020). "The Sloan Digital Sky Survey Quasar Catalog: Sixteenth Data Release". In: *ApJS* 250.1, 8, p. 8. DOI: 10.3847/1538-4365/aba623. arXiv: 2007.09001 [astro-ph.GA] (cit. on pp. 80, 81, 88, 94, 99, 100, 161).

Mandelbaum, R. et al. (Sept. 2003). "Precision cosmology from the Lyman α forest: power spectrum and bispectrum". In: *MNRAS* 344.3, pp. 776–788. DOI: 10.1046/j.1365-8711.2003.06859.x. arXiv: astro-ph/0302112 [astro-ph] (cit. on p. 164).

Marshall, Phil, Nutan Rajguru, and Anže Slosar (Mar. 2006). "Bayesian evidence as a tool for comparing datasets". In: *Phys.Rev.D* 73.6, 067302, p. 067302. DOI: 10.1103/PhysRevD.73.067302. arXiv: astro-ph/0412535 [astro-ph] (cit. on p. 151).

McDonald, Patrick (Mar. 2003). "Toward a Measurement of the Cosmological Geometry at z ~2: Predicting Lyα Forest Correlation in Three Dimensions and the Potential of Future Data Sets". In: *ApJ* 585.1, pp. 34–51. DOI: 10.1086/345945. arXiv: astro-ph/0108064 [astro-ph] (cit. on pp. 62, 116, 121).

McDonald, Patrick and Daniel J. Eisenstein (Sept. 2007). "Dark energy and curvature from a future baryonic acoustic oscillation survey using the Lyman-α forest". In: *Phys.Rev.D* 76.6, 063009, p. 063009. DOI: 10.1103/PhysRevD.76.063009. arXiv: astro-ph/0607122 [astro-ph] (cit. on pp. 63, 106, 107, 116, 118).

McDonald, Patrick et al. (Nov. 2000). "The Observed Probability Distribution Function, Power Spectrum, and Correlation Function of the Transmitted Flux in the Lyα Forest". In: *ApJ* 543.1, pp. 1–23. DOI: 10.1086/317079. arXiv: astro-ph/9911196 [astro-ph] (cit. on pp. 62, 121).

McDonald, Patrick et al. (Nov. 2001). "A Measurement of the Temperature-Density Relation in the Intergalactic Medium Using a New Lyα Absorption-Line Fitting Method". In: *ApJ* 562.1, pp. 52–75. DOI: 10.1086/323426. arXiv: astro-ph/0005553 [astro-ph] (cit. on pp. 117, 179).

McDonald, Patrick et al. (Mar. 2006). "The Lyα Forest Power Spectrum from the Sloan Digital Sky Survey". In: *ApJS* 163.1, pp. 80–109. DOI: 10.1086/444361. arXiv: astro-ph/0405013 [astro-ph] (cit. on p. 110).

McGreer, Ian D. et al. (May 2013). "The z = 5 Quasar Luminosity Function from SDSS Stripe 82". In: *ApJ* 768.2, 105, p. 105. DOI: 10.1088/0004-637X/768/2/105. arXiv: 1212.4493 [astro-ph.CO] (cit. on p. 143).

McQuinn, Matthew (Sept. 2016). "The Evolution of the Intergalactic Medium". In: *ARA&A* 54, pp. 313–362. DOI: 10.1146/annurev-astro-082214-122355. arXiv: 1512.00086 [astro-ph.CO] (cit. on pp. 32, 61).

McQuinn, Matthew, S. Peng Oh, and Claude-André Faucher-Giguère (Dec. 2011). "On Lyman-limit Systems and the Evolution of the Intergalactic Ionizing Background". In: *ApJ* 743.1, 82, p. 82. DOI: 10.1088/0004-637X/743/1/82. arXiv: 1101.1964 [astro-ph.CO] (cit. on p. 125).

McQuinn, Matthew and Martin White (Aug. 2011). "On estimating Lyα forest correlations between multiple sightlines". In: *MNRAS* 415.3, pp. 2257–2269. DOI: 10.1111/j.1365-2966.2011.18855.x. arXiv: 1102.1752 [astro-ph.CO] (cit. on pp. 116, 118).

Mehta, Kushal T. et al. (June 2011). "Galaxy Bias and Its Effects on the Baryon Acoustic Oscillation Measurements". In: *ApJ* 734.2, 94, p. 94. DOI: 10.1088/0004-637X/734/2/94. arXiv: 1104.1178 [astro-ph.CO] (cit. on p. 59).

Meiksin, Avery and Eric R. Tittley (June 2012). "The impact of helium reionization on the structure of the intergalactic medium". In: *MNRAS* 423.1, pp. 7–25. DOI: 10.1111/j.1365-2966.2011.20380.x. arXiv: 1109.5037 [astro-ph.CO] (cit. on p. 125).

Meiksin, Avery A. (Oct. 2009). "The physics of the intergalactic medium". In: *Reviews of Modern Physics* 81.4, pp. 1405–1469. DOI: 10.1103/RevModPhys.81.1405. arXiv: 0711.3358 [astro-ph] (cit. on p. 32).

Menci, N., F. Fiore, and A. Lamastra (Apr. 2012). "Galaxy formation in warm dark matter cosmology". In: *MNRAS* 421.3, pp. 2384–2394. DOI: 10.1111/j.1365-2966.2012.20470.x. arXiv: 1201.1617 [astro-ph.CO] (cit. on p. 32).

Miralda-Escude, Jordi (Oct. 1991). "The Correlation Function of Galaxy Ellipticities Produced by Gravitational Lensing". In: *ApJ* 380, p. 1. DOI: 10.1086/170555 (cit. on p. 64).

Miralda-Escudé, Jordi et al. (Nov. 1996). "The Ly alpha Forest from Gravitational Collapse in the Cold Dark Matter + Lambda Model". In: *ApJ* 471, p. 582. DOI: 10.1086/177992. arXiv: astro-ph/9511013 [astro-ph] (cit. on p. 106).

Mo, H. J. and S. D. M. White (Sept. 1996). "An analytic model for the spatial clustering of dark matter haloes". In: *MNRAS* 282.2, pp. 347–361. DOI: 10.1093/mnras/282.2.347. arXiv: astro-ph/9512127 [astro-ph] (cit. on p. 51).

Mossa, V. et al. (Nov. 2020). "The baryon density of the Universe from an improved rate of deuterium burning". In: *Nature* 587.7833, pp. 210–213. DOI: 10.1038/s41586-020-2878-4 (cit. on p. 157).

Myers, Adam D. et al. (Dec. 2015). "The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Quasar Target Selection". In: *ApJS* 221.2, 27, p. 27. DOI: 10.1088/0067-0049/221/2/27. arXiv: 1508.04472 [astro-ph.CO] (cit. on p. 72).

Naab, Thorsten and Jeremiah P. Ostriker (Aug. 2017). "Theoretical Challenges in Galaxy Formation". In: *ARA&A* 55.1, pp. 59–109. DOI: 10.1146/annurev-astro-081913-040019. arXiv: 1612.06891 [astro-ph.GA] (cit. on p. 52).

Neveux, Richard et al. (Sept. 2020). "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". In: *MNRAS* 499.1, pp. 210–229. DOI: 10.1093/mnras/staa2780. arXiv: 2007.08999 [astro-ph.CO] (cit. on pp. 58, 64, 79).

Nicola, Andrina, Adam Amara, and Alexandre Refregier (Jan. 2019). "Consistency tests in cosmology using relative entropy". In: *JCAP* 2019.1, 011, p. 011. DOI: 10.1088/1475-7516/2019/01/011. arXiv: 1809.07333 [astro-ph.CO] (cit. on p. 151).

O'Meara, J. M. et al. (Oct. 2015). "The First Data Release of the KODIAQ Survey". In: *AJ* 150.4, 111, p. 111. DOI: 10.1088/0004-6256/150/4/111. arXiv: 1505.03529 [astro-ph.CO] (cit. on p. 71).

O'Meara, J. M. et al. (Sept. 2017). "The Second Data Release of the KODIAQ Survey". In: *AJ* 154.3, 114, p. 114. DOI: 10.3847/1538-3881/aa82b8. arXiv: 1707.07905 [astro-ph.GA] (cit. on p. 71).

Ostriker, J. P. and E. T. Vishniac (July 1986). "Generation of Microwave Background Fluctuations from Nonlinear Perturbations at the ERA of Galaxy Formation". In: *ApJ* 306, p. L51. DOI: 10.1086/184704 (cit. on p. 54).

Padmanabhan, Hamsa, T. Roy Choudhury, and Alexandre Refregier (May 2016). "Modelling the cosmic neutral hydrogen from DLAs and 21-cm observations". In: *MNRAS* 458.1, pp. 781–788. DOI: 10.1093/mnras/stw353. arXiv: 1505.00008 [astro-ph.CO] (cit. on p. 125).

Padmanabhan, Nikhil and Martin White (Sept. 2009). "Calibrating the baryon oscillation ruler for matter and halos". In: *Phys.Rev.D* 80.6, 063508, p. 063508. DOI: 10.1103/PhysRevD.80.063508. arXiv: 0906.1198 [astro-ph.CO] (cit. on p. 59).

Padmanabhan, Nikhil et al. (Dec. 2012). "A 2 per cent distance to z = 0.35 by reconstructing baryon acoustic oscillations - I. Methods and application to the Sloan Digital Sky Survey". In: *MNRAS* 427.3, pp. 2132–2145. DOI: 10.1111/j.1365-2966.2012.21888.x. arXiv: 1202.0090 [astro-ph.CO] (cit. on p. 59).

Palanque-Delabrouille, N. et al. (Mar. 2016). "The extended Baryon Oscillation Spectroscopic Survey: Variability selection and quasar luminosity function". In: *A&A* 587, A41, A41. DOI: 10.1051/0004-6361/201527392. arXiv: 1509.05607 [astro-ph.CO] (cit. on p. 119).

Palanque-Delabrouille, Nathalie et al. (Nov. 2013). "The one-dimensional Lyα forest power spectrum from BOSS". In: *A&A* 559, A85, A85. DOI: 10.1051/0004-6361/201322130. arXiv: 1306.5896 [astro-ph.CO] (cit. on pp. 116, 118).

Palanque-Delabrouille, Nathalie et al. (Apr. 2020). "Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman-α and Planck full-survey data". In: *JCAP* 2020.4, 038, p. 038. DOI: 10.1088/1475-7516/2020/04/038. arXiv: 1911.09073 [astro-ph.CO] (cit. on p. 32).

Pâris, I. et al. (Dec. 2012). "The Sloan Digital Sky Survey quasar catalog: ninth data release". In: *A&A* 548, A66, A66. DOI: 10.1051/0004-6361/201220142. arXiv: 1210.5166 [astro-ph.CO] (cit. on pp. 80, 81).

Pâris, Isabelle et al. (Mar. 2014). "The Sloan Digital Sky Survey quasar catalog: tenth data release". In: *A&A* 563, A54, A54. DOI: 10.1051/0004-6361/201322691. arXiv: 1311.4870 [astro-ph.CO] (cit. on pp. 80, 81).

Pâris, Isabelle et al. (Jan. 2017). "The Sloan Digital Sky Survey Quasar Catalog: Twelfth data release". In: *A&A* 597, A79, A79. DOI: 10.1051/0004-6361/201527999. arXiv: 1608.06483 [astro-ph.GA] (cit. on pp. 80–82, 84, 86).

Pâris, Isabelle et al. (May 2018). "The Sloan Digital Sky Survey Quasar Catalog: Fourteenth data release". In: *A&A* 613, A51, A51. DOI: 10.1051/0004-6361/201732445. arXiv: 1712.05029 [astro-ph.GA] (cit. on pp. 80, 81, 94, 99, 100, 161).

Peebles, P. J. E. and J. T. Yu (Dec. 1970). "Primeval Adiabatic Perturbation in an Expanding Universe". In: *ApJ* 162, p. 815. DOI: 10.1086/150713 (cit. on pp. 48, 105).

Penzias, A. A. and R. W. Wilson (July 1965). "A Measurement of Excess Antenna Temperature at 4080 Mc/s." In: *ApJ* 142, pp. 419–421. DOI: 10.1086/148307 (cit. on p. 54).

Percival, Will J. et al. (Feb. 2010). "Baryon acoustic oscillations in the Sloan Digital Sky Survey Data Release 7 galaxy sample". In: *MNRAS* 401.4, pp. 2148–2168. DOI: 10.1111/j.1365-2966.2009.15812.x. arXiv: 0907.1660 [astro-ph.CO] (cit. on p. 58).

Pérez-Ràfols, Ignasi and Matthew M. Pieri (June 2020). "Spectroscopic QUasar extractor and redshift (z) estimator SQUEZE - II. Universality of the results". In: *MNRAS* 496.4, pp. 4941–4950. DOI: 10.1093/mnras/staa1786. arXiv: 1911.04891 [astro-ph.GA] (cit. on pp. 80, 85).

Pérez-Ràfols, Ignasi et al. (Nov. 2018a). "The cosmological bias factor of damped Lyman alpha systems: dependence on metal line strength". In: *MNRAS* 480.4, pp. 4702–4709. DOI: 10.1093/mnras/sty2158. arXiv: 1805.00943 [astro-ph.GA] (cit. on pp. 63, 125).

Pérez-Ràfols, Ignasi et al. (Jan. 2018b). "The SDSS-DR12 large-scale cross-correlation of damped Lyman alpha systems with the Lyman alpha forest". In: *MNRAS* 473.3, pp. 3019–3038. DOI: 10.1093/mnras/stx2525. arXiv: 1709.00889 [astro-ph.CO] (cit. on pp. 63, 125, 128, 182).

Pérez-Ràfols, Ignasi et al. (June 2020). "Spectroscopic QUasar Extractor and redshift (z) Estimator SQUEZE - I. Methodology". In: *MNRAS* 496.4, pp. 4931–4940. DOI: 10.1093/mnras/stz3467. arXiv: 1903.00023 [astro-ph.GA] (cit. on pp. 80, 85, 86).

Perlmutter, S. et al. (June 1999). "Measurements of Ω and Λ from 42 High-Redshift Supernovae". In: *ApJ* 517.2, pp. 565–586. DOI: 10.1086/307221. arXiv: astro-ph/9812133 [astro-ph] (cit. on pp. 19, 33, 40).

Petitjean, Patrick, J. P. Muecket, and Ronald E. Kates (Mar. 1995). "The Lyα forest at low redshift: tracing the dark matter filaments." In: *A&A* 295, pp. L9–L12. arXiv: astro-ph/9502100 [astro-ph] (cit. on p. 106).

Pettini, Max et al. (Sept. 1997). "The Metallicity of High-Redshift Galaxies: The Abundance of Zinc in 34 Damped Lyα Systems from z = 0.7 to 3.4". In: *ApJ* 486.2, pp. 665–680. DOI: 10.1086/304564. arXiv: astro-ph/9704102 [astro-ph] (cit. on p. 125).

Pieri, Matthew M. et al. (June 2014). "Probing the circumgalactic medium at high-redshift using composite BOSS spectra of strong Lyman α forest absorbers". In: *MNRAS* 441.2, pp. 1718–1740. DOI: 10.1093/mnras/stu577. arXiv: 1309.6768 [astro-ph.CO] (cit. on pp. 63, 126).

Pinol, Lucas et al. (Apr. 2017). "Imprint of DESI fiber assignment on the anisotropic power spectrum of emission line galaxies". In: *JCAP* 2017.4, 008, p. 008. DOI: 10.1088/1475-7516/2017/04/008. arXiv: 1611.05007 [astro-ph.CO] (cit. on p. 73).

Planck Collaboration et al. (Nov. 2014). "Planck 2013 results. I. Overview of products and scientific results". In: *A&A* 571, A1, A1. DOI: 10.1051/0004-6361/201321529. arXiv: 1303.5062 [astro-ph.CO] (cit. on p. 54).

Planck Collaboration et al. (Sept. 2016). "Planck 2015 results. XIII. Cosmological parameters". In: *A&A* 594, A13, A13. DOI: 10.1051/0004-6361/201525830. arXiv: 1502.01589 [astro-ph.CO] (cit. on pp. 118, 142).

Planck Collaboration et al. (Sept. 2020a). "Planck 2018 results. I. Overview and the cosmological legacy of Planck". In: *A&A* 641, A1, A1. DOI: 10.1051/0004-6361/201833880. arXiv: 1807.06205 [astro-ph.CO] (cit. on pp. 55, 56).

Planck Collaboration et al. (Sept. 2020b). "Planck 2018 results. VI. Cosmological parameters". In: *A&A* 641, A6, A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO] (cit. on pp. 33, 49, 55, 65, 66, 152, 154).

Planck Collaboration et al. (Sept. 2020c). "Planck 2018 results. X. Constraints on inflation". In: *A&A* 641, A10, A10. DOI: 10.1051/0004-6361/201833887. arXiv: 1807.06211 [astro-ph.CO] (cit. on pp. 65, 76).

Pontzen, Andrew (Apr. 2014). "Scale-dependent bias in the baryonic-acoustic-oscillation-scale intergalactic neutral hydrogen". In: *Phys.Rev.D* 89.8, 083010, p. 083010. DOI: 10.1103/PhysRevD.89.083010. arXiv: 1402.0506 [astro-ph.CO] (cit. on p. 125).

Press, William H. and Paul Schechter (Feb. 1974). "Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation". In: *ApJ* 187, pp. 425–438. DOI: 10.1086/152650 (cit. on p. 51).

Prochaska, J. Xavier et al. (Feb. 2014). "Towards a unified description of the intergalactic medium at redshift z $\approx$ 2.5". In: *MNRAS* 438.1, pp. 476–486. DOI: 10.1093/mnras/stt2218. arXiv: 1310.0052 [astro-ph.CO] (cit. on p. 125).

Prochaska, J. Xavier et al. (Nov. 2017). *pyigm/pyigm: Initial release for publications*. DOI: 10.5281/zenodo.1045480. URL: https://doi.org/10.5281/zenodo.1045480 (cit. on p. 125).

Prochaska, Jason X. et al. (Sept. 2002). "The UCSD HIRES/Keck I Damped Lyα Abundance Database. IV. Probing Galactic Enrichment Histories with Nitrogen". In: *PASP* 114.799, pp. 933–954. DOI: 10.1086/342354. arXiv: astro-ph/0206296 [astro-ph] (cit. on p. 125).

Prodanović, Tijana and Brian D. Fields (Nov. 2003). "On Nonprimordial Deuterium Production by Accelerated Particles". In: *ApJ* 597.1, pp. 48–56. DOI: 10.1086/378272. arXiv: astro-ph/0307183 [astro-ph] (cit. on p. 153).

Raichoor, Anand et al. (Jan. 2021). "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". In: *MNRAS* 500.3, pp. 3254–3274. DOI: 10.1093/mnras/staa3336. arXiv: 2007.09007 [astro-ph.CO] (cit. on p. 58).

Raveri, Marco and Wayne Hu (Feb. 2019). "Concordance and discordance in cosmology". In: *Phys.Rev.D* 99.4, 043506, p. 043506. DOI: 10.1103/PhysRevD.99.043506. arXiv: 1806.04649 [astro-ph.CO] (cit. on p. 151).

Rees, M. J. and D. W. Sciama (Feb. 1968). "Large-scale Density Inhomogeneities in the Universe". In: *Nature* 217.5128, pp. 511–516. DOI: 10.1038/217511a0 (cit. on p. 54).

Reeves, Hubert et al. (Feb. 1973). "On the Origin of Light Elements". In: *ApJ* 179, pp. 909–930. DOI: 10.1086/151928 (cit. on p. 153).

Reid, Beth et al. (Jan. 2016). "SDSS-III Baryon Oscillation Spectroscopic Survey Data Release 12: galaxy target selection and large-scale structure catalogues". In: *MNRAS* 455.2, pp. 1553–1573. DOI: 10.1093/mnras/stv2382. arXiv: 1509.06529 [astro-ph.CO] (cit. on p. 98).

Richards, Gordon T. et al. (Jan. 2009a). "Efficient Photometric Selection of Quasars from the Sloan Digital Sky Survey. II. ~1,000,000 Quasars from Data Release 6". In: *ApJS* 180.1, pp. 67–83. DOI: 10.1088/0067-0049/180/1/67. arXiv: 0809.3952 [astro-ph] (cit. on p. 72).

Richards, Gordon T. et al. (Apr. 2009b). "Eight-Dimensional Mid-Infrared/Optical Bayesian Quasar Selection". In: *AJ* 137.4, pp. 3884–3899. DOI: 10.1088/0004-6256/137/4/3884. arXiv: 0810.3567 [astro-ph] (cit. on p. 72).

Ricotti, Massimo, Nickolay Y. Gnedin, and J. Michael Shull (May 2000). "The Evolution of the Effective Equation of State of the Intergalactic Medium". In: *ApJ* 534.1, pp. 41–56. DOI: 10.1086/308733. arXiv: astro-ph/9906413 [astro-ph] (cit. on pp. 117, 179).

Riess, Adam G. et al. (Sept. 1998). "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant". In: *AJ* 116.3, pp. 1009–1038. DOI: 10.1086/300499. arXiv: astro-ph/9805201 [astro-ph] (cit. on pp. 19, 33, 40).

Riess, Adam G. et al. (May 2019). "Large Magellanic Cloud Cepheid Standards Provide a 1% Foundation for the Determination of the Hubble Constant and Stronger Evidence for Physics beyond ΛCDM". In: *ApJ* 876.1, 85, p. 85. DOI: 10.3847/1538-4357/ab1422. arXiv: 1903.07603 [astro-ph.CO] (cit. on pp. 66, 152, 154).

Rogers, Keir K. et al. (May 2018). "Correlations in the three-dimensional Lyman-alpha forest contaminated by high column density absorbers". In: *MNRAS* 476.3, pp. 3716–3728. DOI: 10.1093/mnras/sty603. arXiv: 1711.06275 [astro-ph.CO] (cit. on p. 125).

Roos, Matts (Jan. 2010). "Dark Matter: The evidence from astronomy, astrophysics and cosmology". In: *arXiv e-prints*. arXiv: 1001.0316 [astro-ph.CO] (cit. on p. 32).

Ross, Ashley J. et al. (May 2015). "The clustering of the SDSS DR7 main Galaxy sample - I. A 4 per cent distance measure at z = 0.15". In: *MNRAS* 449.1, pp. 835–847. DOI: 10.1093/mnras/stv154. arXiv: 1409.3242 [astro-ph.CO] (cit. on p. 150).

Ross, Nicholas P. et al. (Mar. 2012). "The SDSS-III Baryon Oscillation Spectroscopic Survey: Quasar Target Selection for Data Release Nine". In: *ApJS* 199.1, 3, p. 3. DOI: 10.1088/0067-0049/199/1/3. arXiv: 1105.0606 [astro-ph.CO] (cit. on pp. 72, 81).

Ross, Nicholas P. et al. (Aug. 2013). "The SDSS-III Baryon Oscillation Spectroscopic Survey: The Quasar Luminosity Function from Data Release Nine". In: *ApJ* 773.1, 14, p. 14. DOI: 10.1088/0004-637X/773/1/14. arXiv: 1210.6389 [astro-ph.CO] (cit. on p. 143).

Sachs, R. K. and A. M. Wolfe (Jan. 1967). "Perturbations of a Cosmological Model and Angular Variations of the Microwave Background". In: *ApJ* 147, p. 73. DOI: 10.1086/148982 (cit. on p. 54).

Salvato, Mara, Olivier Ilbert, and Ben Hoyle (June 2019). "The many flavours of photometric redshifts". In: *Nature Astronomy* 3, pp. 212–222. DOI: 10.1038/s41550-018-0478-0. arXiv: 1805.12574 [astro-ph.GA] (cit. on p. 57).

Samek, Wojciech et al. (Mar. 2020). "Toward Interpretable Machine Learning: Transparent Deep Neural Networks and Beyond". In: *arXiv e-prints*. arXiv: 2003.07631 [cs.LG] (cit. on p. 162).

Scheuer, P. A. G. (Aug. 1965). "A Sensitive Test for the Presence of Atomic Hydrogen in Intergalactic Space". In: *Nature* 207.5000, p. 963. DOI: 10.1038/207963a0 (cit. on p. 60).

Schmidt, M. and R. F. Green (June 1983). "Quasar evolution derived from the Palomar bright quasar survey and other complete quasar surveys." In: *ApJ* 269, pp. 352–374. DOI: 10.1086/161048 (cit. on p. 80).

Schmidt, Maarten (Apr. 1965). "Large Redshifts of Five Quasi-Stellar Sources." In: *ApJ* 141, p. 1295. DOI: 10.1086/148217 (cit. on p. 60).

Schneider, Donald P. et al. (July 2007). "The Sloan Digital Sky Survey Quasar Catalog. IV. Fifth Data Release". In: *AJ* 134.1, pp. 102–117. DOI: 10.1086/518474. arXiv: 0704.0806 [astro-ph] (cit. on pp. 80, 82).

Schneider, Donald P. et al. (June 2010). "The Sloan Digital Sky Survey Quasar Catalog. V. Seventh Data Release". In: *AJ* 139.6, 2360, p. 2360. DOI: 10.1088/0004-6256/139/6/2360. arXiv: 1004.1167 [astro-ph.CO] (cit. on p. 80).

Seljak, Uroš (Mar. 2012). "Bias, redshift space distortions and primordial nongaussianity of nonlinear transformations: application to Ly-α forest". In: *JCAP* 2012.3, 004, p. 004. DOI: 10.1088/1475-7516/2012/03/004. arXiv: 1201.0594 [astro-ph.CO] (cit. on pp. 62, 117, 121).

Seo, Hee-Jong and Daniel J. Eisenstein (Nov. 2005). "Baryonic Acoustic Oscillations in Simulated Galaxy Redshift Surveys". In: *ApJ* 633.2, pp. 575–588. DOI: 10.1086/491599. arXiv: astro-ph/0507338 [astro-ph] (cit. on p. 58).

— (Aug. 2007). "Improved Forecasts for the Baryon Acoustic Oscillations and Cosmological Distance Scale". In: *ApJ* 665.1, pp. 14–24. DOI: 10.1086/519549. arXiv: astro-ph/0701079 [astro-ph] (cit. on p. 59).

Seo, Hee-Jong et al. (Oct. 2008). "Nonlinear Structure Formation and the Acoustic Scale". In: *ApJ* 686.1, pp. 13–24. DOI: 10.1086/589921. arXiv: 0805.0117 [astro-ph] (cit. on p. 58).

Seo, Hee-Jong et al. (Sept. 2010). "High-precision Predictions for the Acoustic Scale in the Nonlinear Regime". In: *ApJ* 720.2, pp. 1650–1667. DOI: 10.1088/0004-637X/720/2/1650. arXiv: 0910.5005 [astro-ph.CO] (cit. on p. 59).

Shapley, Harlow and Heber D. Curtis (May 1921). "The Scale of the Universe". In: *Bulletin of the National Research Council* 2.11, pp. 171–217 (cit. on p. 19).

Sheth, Ravi K., H. J. Mo, and Giuseppe Tormen (May 2001). "Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes". In: *MNRAS* 323.1, pp. 1–12. DOI: 10.1046/j.1365-8711.2001.04006.x. arXiv: astro-ph/9907024 [astro-ph] (cit. on p. 51).

Sheth, Ravi K. and Giuseppe Tormen (Jan. 2002). "An excursion set model of hierarchical clustering: ellipsoidal collapse and the moving barrier". In: *MNRAS* 329.1, pp. 61–75. DOI: 10.1046/j.1365-8711.2002.04950.x. arXiv: astro-ph/0105113 [astro-ph] (cit. on p. 51).

Silk, Joseph (Feb. 1968). "Cosmic Black-Body Radiation and Galaxy Formation". In: *ApJ* 151, p. 459. DOI: 10.1086/149449 (cit. on p. 54).

Slosar, Anže et al. (Sept. 2011). "The Lyman-α forest in three dimensions: measurements of large scale flux correlations from BOSS 1st-year data". In: *JCAP* 2011.9, 001, p. 001. DOI: 10.1088/1475-7516/2011/09/001. arXiv: 1104.5244 [astro-ph.CO] (cit. on pp. 63, 106, 126, 179).

Slosar, Anže et al. (Apr. 2013). "Measurement of baryon acoustic oscillations in the Lyman-α forest fluctuations in BOSS data release 9". In: *JCAP* 2013.4, 026, p. 026. DOI: 10.1088/1475-7516/2013/04/026. arXiv: 1301.3459 [astro-ph.CO] (cit. on pp. 63, 106, 159).

Smith, Alex et al. (Mar. 2019). "Correcting for fibre assignment incompleteness in the DESI Bright Galaxy Survey". In: *MNRAS* 484.1, pp. 1285–1300. DOI: 10.1093/mnras/stz059. arXiv: 1809.07355 [astro-ph.CO] (cit. on p. 73).

Smoot, G. F. et al. (Sept. 1992). "Structure in the COBE Differential Microwave Radiometer First-Year Maps". In: *ApJ* 396, p. L1. DOI: 10.1086/186504 (cit. on p. 54).

Spergel, D. et al. (Mar. 2015). "Wide-Field InfrarRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report". In: *arXiv e-prints*. arXiv: 1503.03757 [astro-ph.IM] (cit. on pp. 59, 159).

Springel, Volker (Sept. 2010). "Smoothed Particle Hydrodynamics in Astrophysics". In: *ARA&A* 48, pp. 391–430. DOI: 10.1146/annurev-astro-081309-130914. arXiv: 1109.2219 [astro-ph.CO] (cit. on p. 51).

Sunayama, Tomomi et al. (June 2020). "Mitigating the impact of fiber assignment on clustering measurements from deep galaxy redshift surveys". In: *JCAP* 2020.6, 057, p. 057. DOI: 10.1088/1475-7516/2020/06/057. arXiv: 1912.06583 [astro-ph.CO] (cit. on p. 73).

Sunyaev, R. A. and Ia. B. Zeldovich (Jan. 1980). "Microwave background radiation as a probe of the contemporary structure and history of the universe". In: *ARA&A* 18, pp. 537–560. DOI: 10.1146/annurev.aa.18.090180.002541 (cit. on p. 54).

Sunyaev, R. A. and Ya. B. Zeldovich (Apr. 1970). "Small-Scale Fluctuations of Relic Radiation". In: *Ap&SS* 7.1, pp. 3–19. DOI: 10.1007/BF00653471 (cit. on p. 54).

Tao, Dominique et al. (July 2018). "Priority coordination of fiber positioners in multi-objects spectrographs". In: *Ground-based and Airborne Instrumentation for Astronomy VII*. Ed. by Christopher J. Evans, Luc Simard, and Hideki Takami. Vol. 10702. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 107028K. DOI: 10.1117/12.2313962 (cit. on p. 73).

Teyssier, Romain (Aug. 2015). "Grid-Based Hydrodynamics in Astrophysical Fluid Flows". In: *ARA&A* 53, pp. 325–364. DOI: 10.1146/annurev-astro-082214-122309 (cit. on p. 51).

Tie, Suk Sien et al. (Aug. 2019). "UV background fluctuations and three-point correlations in the large-scale clustering of the Lyman α forest". In: *MNRAS* 487.4, pp. 5346–5362. DOI: 10.1093/mnras/stz1632. arXiv: 1905.02208 [astro-ph.CO] (cit. on p. 165).

Trotta, Roberto (Mar. 2008). "Bayes in the sky: Bayesian inference and model selection in cosmology". In: *Contemporary Physics* 49.2, pp. 71–104. DOI: 10.1080/00107510802066753. arXiv: 0803.4089 [astro-ph] (cit. on p. 151).

Tully, R. Brent, Hélène M. Courtois, and Jenny G. Sorce (Aug. 2016). "Cosmicflows-3". In: *AJ* 152.2, 50, p. 50. DOI: 10.3847/0004-6256/152/2/50. arXiv: 1605.01765 [astro-ph.CO] (cit. on p. 152).

Verde, Licia, Pavlos Protopapas, and Raul Jimenez (Sept. 2013). "Planck and the local Universe: Quantifying the tension". In: *Physics of the Dark Universe* 2.3, pp. 166–175. DOI: 10.1016/j.dark.2013.09.002. arXiv: 1306.6766 [astro-ph.CO] (cit. on p. 151).

Viel, M. et al. (Jan. 2004). "The bispectrum of the Lyman α forest at z~2-2.4 from a large sample of UVES QSO absorption spectra (LUQAS)". In: *MNRAS* 347.2, pp. L26–L30. DOI: 10.1111/j.1365-2966.2004.07404.x. arXiv: astro-ph/0308151 [astro-ph] (cit. on p. 164).

Viel, Matteo et al. (Mar. 2005). "Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman-α forest". In: *Phys.Rev.D* 71.6, 063534, p. 063534. DOI: 10.1103/PhysRevD.71.063534. arXiv: astro-ph/0501562 [astro-ph] (cit. on p. 32).

Vogt, S. S. et al. (June 1994). "HIRES: the high-resolution echelle spectrometer on the Keck 10-m Telescope". In: *Instrumentation in Astronomy VIII*. Ed. by David L. Crawford and Eric R. Craine. Vol. 2198. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 362. DOI: 10.1117/12.176725 (cit. on p. 71).

Walther, Michael et al. (Jan. 2018). "A New Precision Measurement of the Small-scale Line-of-sight Power Spectrum of the Lyα Forest". In: *ApJ* 852.1, 22, p. 22. DOI: 10.3847/1538-4357/aa9c81. arXiv: 1709.07354 [astro-ph.CO] (cit. on p. 110).

Weinberg, David H. et al. (Sept. 2013). "Observational probes of cosmic acceleration". In: *physrep* 530.2, pp. 87–255. DOI: 10.1016/j.physrep.2013.05.001. arXiv: 1201.2434 [astro-ph.CO] (cit. on p. 65).

Weinberg, Steven (Jan. 1989). "The cosmological constant problem". In: *Reviews of Modern Physics* 61.1, pp. 1–23. DOI: 10.1103/RevModPhys.61.1 (cit. on p. 33).

— (Dec. 2002). "Cosmological Fluctuations of Small Wavelength". In: *ApJ* 581.2, pp. 810–816. DOI: 10.1086/344441. arXiv: astro-ph/0207375 [astro-ph] (cit. on p. 47).

— (2008). *Cosmology* (cit. on p. 25).

Williams, G. Grant et al. (Sept. 2004). "90prime: a prime focus imager for the Steward Observatory 90-in. telescope". In: *Ground-based Instrumentation for Astronomy*. Ed. by Alan F. M. Moorwood and Masanori Iye. Vol. 5492. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 787–798. DOI: 10.1117/12.552189 (cit. on p. 72).

Wolfe, A. M. et al. (June 1986). "Damped Lyman-Alpha Absorption by Disk Galaxies with Large Redshifts. I. The Lick Survey". In: *ApJS* 61, p. 249. DOI: 10.1086/191114 (cit. on p. 125).

Wong, Kenneth C. et al. (Oct. 2020). "H0LiCOW – XIII. A 2.4 per cent measurement of H_0 from lensed quasars: 5.3σ tension between early- and late-Universe probes". In: *MNRAS* 498.1, pp. 1420–1439. DOI: 10.1093/mnras/stz3094. arXiv: 1907.04869 [astro-ph.CO] (cit. on p. 152).

Wright, Edward L. et al. (Dec. 2010). "The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance". In: *AJ* 140.6, pp. 1868–1881. DOI: 10.1088/0004-6256/140/6/1868. arXiv: 1008.0031 [astro-ph.IM] (cit. on p. 72).

Wright, Thomas (1750). *An original theory or new hypothesis of the universe founded upon the laws of nature, and solving by mathematical principles the general phaenomena of the visible creation and particularly the Via Lactea*. DOI: 10.3931/e-rara-28672 (cit. on p. 20).

Yeche, Ch. et al. (Oct. 2009). "QSO Selection and Photometric Redshifts with Neural Networks". In: *arXiv e-prints*. arXiv: 0910.3770 [astro-ph.CO] (cit. on p. 72).

Yèche, Christophe et al. (Oct. 2020). "Preliminary Target Selection for the DESI Quasar (QSO) Sample". In: *Research Notes of the American Astronomical Society* 4.10, 179, p. 179. DOI: 10.3847/2515-5172/abc01a. arXiv: 2010.11280 [astro-ph.CO] (cit. on pp. 72, 134).

York, Donald G. et al. (Sept. 2000). "The Sloan Digital Sky Survey: Technical Summary". In: *AJ* 120.3, pp. 1579–1587. DOI: 10.1086/301513. arXiv: astro-ph/0006396 [astro-ph] (cit. on pp. 58, 70).

Zaldarriaga, Matias and Uroš Seljak (July 1998). "Gravitational lensing effect on cosmic microwave background polarization". In: *Phys.Rev.D* 58.2, 023003, p. 023003. DOI: 10.1103/PhysRevD.58.023003. arXiv: astro-ph/9803150 [astro-ph] (cit. on p. 55).

Zel'Dovich, Y. B. (Mar. 1970). "Reprint of 1970A&A.....5...84Z. Gravitational instability: an approximate theory for large density perturbations." In: *A&A* 500, pp. 13–18 (cit. on p. 59).

Zhang, Bonnie R. et al. (Oct. 2017). "A blinded determination of H_0 from low-redshift Type Ia supernovae, calibrated by Cepheid variables". In: *MNRAS* 471.2, pp. 2254–2285. DOI: 10.1093/mnras/stx1600. arXiv: 1706.07573 [astro-ph.CO] (cit. on p. 68).