UNIVERSITY COLLEGE LONDON

Faculty of Mathematics and Physical Sciences

Department of Physics & Astronomy

PRECISION COSMOLOGY FROM THE CLUSTERING OF LARGE-SCALE STRUCTURES

ANDREI ALEXANDRU CUCEU

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

SUPERVISORS:

DR. ANDREU FONT-RIBERA

PROF. BENJAMIN JOACHIMI

EXAMINERS:

PROF. JOHN PEACOCK

PROF. AMELIE SAINTONGE

Saturday 9th April, 2022

I, Andrei Alexandru Cuceu, 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.

ANDREI CUCEU

ABSTRACT

Cosmology is going through an exciting period characterized by increasingly large and diverse surveys of the Universe. The resulting datasets contain both cosmological and astrophysical information, and provide many different perspectives on the components of the Universe and its evolution. Despite this wealth of data, there are still many open questions and unknowns. These vary from broad questions about the nature of dark energy and dark matter, to more concrete ones, such as how fast exactly is the Universe expanding. In this thesis, we will focus on spectroscopic surveys of the large-scale structure (LSS) in the Universe, and how they are used to study these problems.

We begin in Chapters 1-4 with an introduction of modern cosmology, focusing on the topics relevant to LSS surveys and the analysis tools used to extract cosmological information from these datasets. After that, in Chapter 5 we use the latest baryon acoustic oscillation (BAO) measurements and baryon density constraints based on big bang nucleosynthesis (BBN) to measure the expansion rate of the Universe, through the Hubble constant. This is an independent measurement that contributes to the ongoing Hubble tension debate. In Chapter 6, we perform for the first time a Bayesian analysis of the Lyman-α (Lyα) forest correlation functions in order to measure BAO. Finally, in Chapter 7 we study the possibility of an analysis of the full shape of the Lyα forest correlation functions. This would expand on BAO analyses by providing cosmological information from a broader range of scales. The next generation of cosmological surveys is just starting with the Dark Energy Spectroscopic Instrument (DESI). Throughout this thesis, we performed multiple studies that are relevant for future cosmological analyses with DESI. These analyses will advance our physical understanding of the Universe by precisely mapping its evolution to higher redshifts than ever before.

Impact Statement

This thesis presents novel work in the field of observational cosmology. It is designed to pave the way for future spectroscopic surveys, such as the Dark Energy Spectroscopic Instrument (DESI). The methods presented here will be helpful in robustly extracting cosmological information from future analyses of the Lyman-α (Lyα) forest.

The work in Chapter 5 contributes directly to the ongoing Hubble tension, by providing an estimate of the Hubble constant that is independent of cosmic microwave background (CMB) anisotropy measurements and cosmic distance ladder analyses. Our measurement is consistent with CMB measurements from the Planck satellite, reinforcing the robustness of that result.

The Bayesian analysis we performed in Chapter 6 represents a novel approach to extracting baryon acoustic oscillations (BAO) information from Lyα forest correlations. We have also implemented an interface to the `PolyChord` sampler in the community code `picca`. This allows Bayesian sampling to be used by future Lyα forest analyses. The work in Chapters 5 and 6 has already been directly relevant to the final extended Baryon Oscillation Spectroscopic Survey (eBOSS) analysis as outlined in Chapter 8.

The work in Chapter 7 is directly relevant to future analyses of the 3D distribution of the Lyα forest as it provides techniques for extracting more cosmological information from existing datasets. As we have shown, this could lead to the first measurement of the growth rate of cosmic structure at redshifts $z > 2$.

The work in this thesis has led to three first-author publications, and has been presented at multiple collaboration meetings relating to eBOSS and DESI. As part of the work presented here, I have developed the `Vega` library for modelling and fitting Lyα forest correlation functions. This package is public, and is meant to replace the old fitting tool used in eBOSS.

Beyond academia, the discussion on Bayesian versus frequentist statistics from Chapter 6 is relevant not only for science at large, but for other fields as well. The Bayesian framework and

many associated tools have become mainstream in the field of machine learning, with wide-ranging applications in all parts of society. Furthermore, the work in this thesis is meant to advance the tools used to study and understand the Universe we live in and its physical laws. This has the potential to inspire the next generation of physicists, and lead to increased interest in science.

ACKNOWLEDGEMENTS

To begin, I would like to thank Andreu Font-Ribera for his invaluable guidance along this journey, and for providing me with encouragement along the way. Your mentorship has made these last few years both enjoyable and productive. Secondly, I would like to thank Benjamin Joachimi for his advice and insightful discussion, and for often helping me see the bigger picture.

I would like to thank Seshadri Nadathur, Pablo Lemos, Jim Rich, James Farr and Paul Martini for their contributions in supporting my research, and for their valuable feedback and suggestions. I would also like to thank Bruno Moraes, Arthur Loureiro, Lorne Whiteway, Filipe Abdalla and Michael McLeod for their guidance in the early stages of my PhD. Thank you to everyone in the astrophysics group at UCL for providing such a friendly and interesting atmosphere. Thank you also to Rammstein for providing the soundtrack for my PhD.

I would like to acknowledge the support and encouragement of my parents, and to thank them along with my brother for the many insightful conversations that led to my interest in science. I also have to thank Lucian Stoian for his guidance and for nurturing my curiosity for physics and astronomy. And finally, thank you to Blanka for being patient with me along this journey. It has been undoubtedly easier because of you, and it has been a true pleasure to share it with you.

Contents

List of Figures

List of Tables

PART I

INTRODUCTION AND BACKGROUND

CHAPTER 1

GRAVITY AND THE HOMOGENEOUS UNIVERSE

We begin the introduction with an overview of the basis of the standard cosmological model, known as Λ Cold Dark Matter (ΛCDM). In Section 1.1, I introduce general relativity and the core assumptions used in modern cosmology. After that, I introduce the components of the Universe and the equations that govern them in Section 1.2. Finally, in Section 1.3 I give an overview of how distances behave and how to measure them in an evolving Universe.

1.1 GENERAL RELATIVITY

The theory of General Relativity (GR) was constructed by Albert Einstein (Einstein, 1915), starting from the equivalence principle. The weak version states that: *the motions of freely-falling particles are the same in a gravitational field and a uniformly accelerated frame, in small enough regions* (Carroll, 2004). This means that we cannot locally distinguish between the effect of gravity and that of uniform acceleration. This led Einstein to formulate a more general version of the equivalence principle, stating that: *In small enough regions of spacetime, the laws of physics reduce to those of special relativity* (Carroll, 2004).

The equivalence principle leads us to look for a geometry that has curvature but appears locally flat. A manifold is an object from the mathematical field of topology that describes a space that may be curved, but locally resembles Euclidean space. We also need the structure of the manifold

to be differentiable in order to be able to use calculus. Therefore, spacetime is defined as a four-dimensional (4D) differentiable manifold. Further, we define the metric $g_{\mu\nu}$ which allows us to measure distances by converting a set of coordinates X^{μ}[1] into the line element:

$$ds^2 = g_{\mu\nu} dX^{\mu} dX^{\nu}, \tag{1.1}$$

where we use the Einstein summation convention.

A general metric for the Universe can be created by using the cosmological principle. This is based on the a priori expectation that we are not in a special location in the Universe. Firstly, our inferences about the large-scale structure of the Universe should not change by moving to a different position which is arbitrarily far away. This means that the Universe should be *homogeneous* on large scales. Secondly, our inferences should not depend on which area of the sky we survey, which means the Universe should be *isotropic*.

Robertson and Walker showed that a metric first used by Friedman (Friedman, 1922) and Lemaître (Lemaître, 1931) is the most general metric with a time-dependent spatial component under the constraints of homogeneity and isotropy (Robertson, 1935; Robertson, 1936a; Robertson, 1936b; Walker, 1937). The Friedman-Lemaître-Robertson-Walker (FLRW) metric using a metric signature $(+,-,-,-)$ is given by:

$$ds^2 = dt^2 - a^2(t) d\sigma^2, \tag{1.2}$$

where $a(t)$ is the scale factor and $d\sigma^2$ is a metric of a 3D spatial slice with uniform curvature. We also use the convention where we set the speed of light $c = 1$ throughout this introduction. In polar coordinates, the most general $d\sigma^2$ (assuming homogeneity and isotropy) is given by:

$$d\sigma^2 = \frac{dr^2}{1-kr^2} + r^2 d\theta^2 + r^2 \sin^2\theta d\phi^2, \tag{1.3}$$

where k is a constant that sets the curvature of space. Negative values of k correspond to an open Universe, positive values correspond to a closed Universe, and $k = 0$ corresponds to a flat (Euclidean) Universe.

Having introduced the metric, which describes the geometry of spacetime, we turn our attention to the contents of the Universe. On large scales we use the concept of a perfect fluid, that is, a fluid completely characterized by its rest frame energy density ρ and isotropic pressure P. Imposing

[1] For example in Cartesian space: X^{μ}=(t,x,y,z)

homogeneity and isotropy allows us to define the energy-momentum tensor in the frame of a comoving observer:

$$T^{\mu}_{\nu} = \begin{pmatrix} \rho & 0 & 0 & 0 \\ 0 & -P & 0 & 0 \\ 0 & 0 & -P & 0 \\ 0 & 0 & 0 & -P \end{pmatrix}. \tag{1.4}$$

In General Relativity the conservation of energy-momentum is expressed as:

$$\Delta_{\mu} T^{\mu}_{\nu} = 0, \tag{1.5}$$

with the covariant derivative defined as:

$$\Delta_{\mu} F^{\nu}_{\gamma} = \partial_{\mu} F^{\nu}_{\gamma} + \Gamma^{\nu}_{\sigma\mu} F^{\sigma}_{\gamma} - \Gamma^{\sigma}_{\gamma\mu} F^{\nu}_{\sigma}, \tag{1.6}$$

where $\partial_{\alpha} \equiv \frac{\partial}{\partial X^{\alpha}}$ and the Christoffel symbols are given by:

$$\Gamma^{\mu}_{\alpha\beta} = \frac{g^{\mu\nu}}{2}(\partial_{\beta} g_{\alpha\nu} + \partial_{\alpha} g_{\beta\nu} - \partial_{\nu} g_{\alpha\beta}). \tag{1.7}$$

Using equations 1.4-1.7 and the FLRW metric, we arrive at the conservation equation:

$$\dot{\rho} + 3\frac{\dot{a}}{a}(\rho + P) = 0, \tag{1.8}$$

where the dot notation represents derivatives with respect to the time coordinate ($\dot{a} = \partial a/\partial t$). Note that Equation 1.8 can also be derived in a more direct way from the mass continuity equation.

Finally, in order to study the evolution of the Universe, we need the Einstein field equations which tie together the geometry of spacetime with its contents (in our case a perfect fluid). These are given by:

$$\begin{aligned} G_{\mu\nu} + \Lambda g_{\mu\nu} &= 8\pi G T_{\mu\nu}, \\ \text{with } G_{\mu\nu} &= R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu}, \end{aligned} \tag{1.9}$$

where $G_{\mu\nu}$ is the Einstein tensor, $R_{\mu\nu}$ and R are the Ricci curvature tensor and scalar respectively, G is the Newtonian gravitational constant and Λ is the cosmological constant.

It is common to absorb the Λ term into the energy-momentum tensor, by interpreting it as an energy density. This is preferred when interpreting Λ as the quantum energy of vacuum, or as an extension to known physics meant to explain the observed accelerated expansion of the Universe.

However, it is important to mention that the cosmological constant appears naturally when using the principle of least action to derive the Einstein field equations (Lovelock, 1971; Lovelock, 1972).

The Einstein field equations using the FLRW metric and the energy-momentum tensor described above give us only two equations (with the others vanishing due to homogeneity and isotropy), known as the *Friedman equations*:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{k}{a^2}, \tag{1.10}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3P). \tag{1.11}$$

The first Friedman equation relates density and curvature to the speed of expansion. The second equation gives us the acceleration. At first glance, it may appear that $\ddot{a}$ is always negative as the density and pressure should be positive, leading to a decelerated expansion. However, as we shall see, the cosmological constant Λ behaves as a perfect fluid with negative pressure, leading to accelerated expansion when it becomes dominant. Finally, the Hubble parameter is defined as:

$$H \equiv \frac{\dot{a}}{a}. \tag{1.12}$$

The Friedman equations (Equations 1.10 and 1.11) combined with the conservation equation (Equation 1.8) describe the evolution of the Universe under the assumptions of homogeneity and isotropy. The perfect fluid we considered has multiple ingredients that behave in qualitatively different ways, so we next turn our attention to the components of the Universe.

1.2 Cosmic ingredients

We have introduced the contents of the Universe as a perfect fluid that is homogeneous and isotropic. However, this fluid has different components that come to dominate at different times and have different impacts on the evolution of the Universe. Before introducing these cosmic ingredients, we introduce a common framework to describe them, based on the equation of state:

$$P = w\rho, \tag{1.13}$$

where the state parameter w is assumed to be uniform. Plugging the equation of state into Equation 1.8 and integrating leads to:

$$\rho \propto a^{-3(1+w)}. \tag{1.14}$$

This simple formalism combined with the Friedman equations allows us to study the behaviour and evolution of different components based on their state parameter w.

When describing the components of the Universe, we generally use their density parameter Ω_i, instead of their physical density ρ_i. The density parameter is given by:

$$\Omega_i = \frac{8\pi G}{3H^2}\rho_i = \frac{\rho_i}{\rho_c}, \tag{1.15}$$

$$\text{with } \rho_c = \frac{3H^2}{8\pi G}, \tag{1.16}$$

where ρ_c is the critical density. This allows us to express the first Friedman equation in a very compact form:

$$\sum_i \Omega_i(a) + \Omega_k(a) = 1, \tag{1.17}$$

where the sum is performed over the different components i, and the density parameter of curvature is given by:

$$\Omega_k(a) \equiv -\frac{k}{a^2H^2(a)}. \tag{1.18}$$

We will also use the notation $\Omega_{i,0} = \Omega_i(t_0)$ to refer to the present value of the density parameters.

Finally, we have the tools necessary to describe the components of the Universe, and their impact on its evolution:

- **Radiation:** A majority of the radiation contribution today is made up of free-streaming photons that are part of the cosmic microwave background (CMB). They have an equation of state $w_r = 1/3$, meaning they exert a positive pressure, $P_r = \rho_r/3$. When compared to the total energy budget today, photons make up a very small fraction $\Omega_{r,0} \approx 10^{-5}$. However, from Equation 1.14 we can see that their density evolves as $\rho_r \propto a^{-4}$, meaning that they used to have a much bigger contribution in the past. In the early Universe, radiation was a dominant component and was coupled to baryons due to scattering by charged particles. This led to an oscillatory behaviour between the opposing forces of radiation pressure and gravity (see Chapter 3), that were left imprinted in the CMB.

- **Baryons:** It is standard in cosmology to refer to all ordinary matter (i.e. protons, neutrons electrons) as baryons. Even though electrons are actually leptons, they are included for simplicity, as their mass is tiny compared to nucleons. Today, baryons can be found either in collapsed objects such as stars and planets, or in diffuse gas in (interstellar medium), around (circum-galactic medium) or between (inter-galactic medium) galaxies. They make

up a relatively small percentage of the cosmic energy budget, with $\Omega_{b,0} \approx 0.05$. On the other hand, in the early Universe baryons were part of the primordial plasma and coupled to radiation. This coupling led to an oscillatory behaviour that gave rise to baryon acoustic oscillations (BAO; see Chapter 3). At recombination the first atoms formed, which led to a drastic decrease in the number of free charged particles and thus the decoupling from radiation. Since then, the only large-scale interaction relevant for baryons is due to gravity, leading to an equation of state $w_b = 0$, and a density evolution $\rho_b \propto a^{-3}$.

- **Neutrinos:** Neutrinos were initially thought to be massless. However, neutrino oscillation experiments (Fukuda et al., 1998) showed that they do have a total mass of at least 0.06 eV[2]. In the early Universe, neutrinos were relativistic and can be treated similarly to radiation. They turn non-relativistic around redshift ~ 100, which is after the emission of the CMB, but well before the low redshift ($z < 4$) measurements discussed in this thesis. Once that happens, neutrinos become pressureless, with $w_\nu = 0$. The interactions of neutrinos with other particles are very weak, and therefore the strongest constraints on their masses come from their gravitational interaction with the rest of the Universe. Neutrinos are an active area of research, with only upper bounds on their masses currently available from cosmology (e.g. Planck Collaboration et al., 2018; eBOSS Collaboration et al., 2021).

- **Dark Matter:** We have known for a long time that there is a type of matter in the Universe that we do not see directly, yet it dominates the gravitational potential on scales from the size of a galaxy to that of the entire Universe. Strong observational evidence for dark matter started to appear in the 1970s from galaxy rotation curves (Rubin and Ford, 1970), however, there were earlier hints of is existence (e.g. Kapteyn, 1922; Zwicky, 1933; Zwicky, 1937). Today, the best evidence for the existence of matter beyond the familiar baryonic type comes from the CMB and large-scale clustering (e.g. Planck Collaboration et al., 2018; eBOSS Collaboration et al., 2021; DES Collaboration et al., 2021). Dark matter either does not interact with other standard model particles (except through gravity), or interacts so weakly we have not yet detected any such events. In the early Universe, dark matter was free to collapse gravitationally in a period when baryons were still coupled to radiation, meaning they dominated early gravitational potential wells and left a strong imprint in CMB anisotropies. Therefore, the best current estimate of the total matter density comes from the CMB: $\Omega_{m,0} = 0.315 \pm 0.007$ (Planck Collaboration et al., 2018). When compared to

[2]Oscillation experiments only measure mass differences, not the absolute masses, so we can only get a lower bound.

the best measurements of the total baryon density given above, dark matter accounts for approximately 80% of the total matter in the Universe. After recombination, baryons were free to collapse into the dark matter potential wells, and today the two components behave similarly on large scales with $w_m = 0$ and $\rho_m \propto a^{-3}$. Because the CMB points to dark matter being able to collapse gravitationally in the early Universe, it must have been *cold*. This means that dark matter particles had small velocities that allowed them to efficiently clump together. Therefore, we generally refer to this component as *cold dark matter* (CDM).

- **Dark Energy:** The final component we consider is also the most abundant one in the Universe today, with $\Omega_{\Lambda,0} \approx 0.7$. Dark energy is the most recent addition to the standard cosmological model, only becoming an established component about twenty years ago when two groups independently showed that the expansion of the Universe is accelerating (Riess et al., 1998; Perlmutter et al., 1999). The simplest explanation for this comes from the cosmological constant we have introduced in Section 1.1. The cosmological constant can be described by an equation of state $w_\Lambda = -1$, and its energy density is constant: $\rho_\Lambda \propto$ const. However, we still use the term dark energy to refer to this mysterious component because Λ is just the simplest possibility, but by no means the only one. A common generalization considers dark energy as a scalar field with a constant equation of state $w \neq -1$, or a time evolving equation of state. Furthermore, there have been attempts to identify dark energy as the quantum energy of vacuum, however, this led to the famous 120 orders of magnitude difference between theory and cosmological measurements (Carroll, 2001). The best observational constraints currently available (see Figure 1.1) are consistent with a value of $w_\Lambda = -1$ (i.e. a cosmological constant; e.g. Scolnic et al., 2018; Planck Collaboration et al., 2018; eBOSS Collaboration et al., 2021), however, this is still an active area of research.

I have not included curvature in the list above because it is a geometric property of spacetime rather than a physical component of the Universe (it is on the left-hand side of the Einstein field equations). This may also be true for dark energy if it is described by the cosmological constant, however, that has not been clearly established yet. Current observational evidence (Figure 1.1) is consistent with the Universe being flat (i.e. $k = 0$, e.g. Planck Collaboration et al., 2018; eBOSS Collaboration et al., 2021).

The model presented in Section 1.1 together with the components described here form the basis for the currently accepted standard model of Cosmology, referred to as Λ cold dark matter

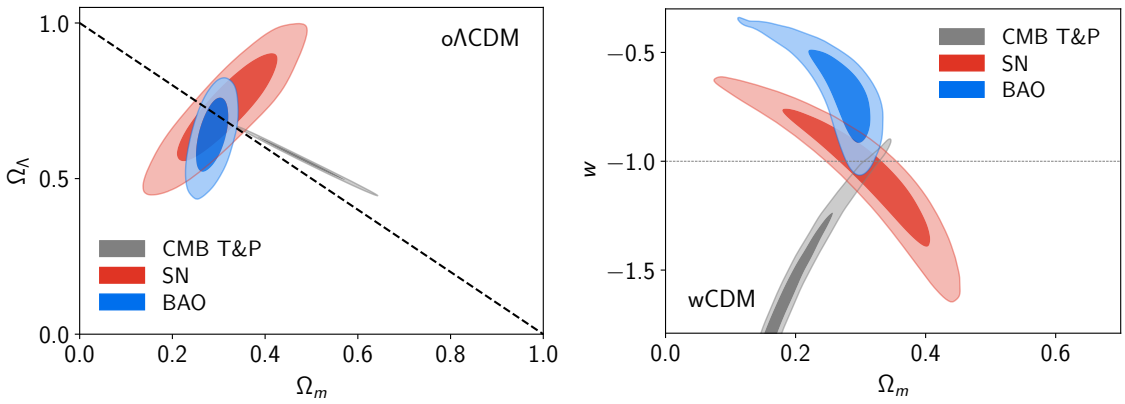

Figure 1.1. Constraints of the dark energy density parameter (left) and the dark energy equation of state (right) versus the matter density parameter from various cosmological probes. These measurements overlap at $\Omega_{\Lambda,0} \approx 0.7, \Omega_{m,0} \approx 0.3, w = -1$, consistent with a flat Universe (dashed line on the left) and a cosmological constant (dashed line on the right). From eBOSS Collaboration et al. (2021).

(ΛCDM)[3]. Finally, we reformulate the Friedman equation to explicitly include the components described above in terms of their density parameters today:

$$H^2 = H_0^2 \left(\Omega_{r,0} a^{-4} + \Omega_{m,0} a^{-3} + \Omega_{k,0} a^{-2} + \Omega_{\Lambda,0} \right), \tag{1.19}$$

where $H_0 = H(t_0)$ is the Hubble constant. The matter density parameter $\Omega_{m,0}$ is the sum of the density parameters of all non-relativistic species today, which includes baryons, dark matter and neutrinos. Note however, that when neutrinos were relativistic, their density parameter is included in Ω_r because they behave like radiation. While so far we have used $\Omega_{i,0}$ to denote the density parameter of species i at the present time, from now we will simplify the notation. Therefore, whenever we refer to Ω_i, we mean the value of the density parameter at the present time.

1.3 Distances

In order to tie the theoretical understanding of the structure of spacetime to observations, we need the ability to measure distances across the Universe. However, this is non-trivial due to the expansion of the spatial component of the metric and also due to curvature. In general relativity free-falling particles move along geodesics, with paths $X^\mu(\tau)$ obeying the geodesic equation:

$$\frac{d^2 X^\mu}{d\tau^2} + \Gamma^\mu_{\alpha\beta} \frac{dX^\alpha}{d\tau} \frac{dX^\beta}{d\tau} = 0, \tag{1.20}$$

[3]The final key ingredient of ΛCDM is inflation, which we discuss in Section 2.1.

where τ is an affine parameter (monotonically increases along the path).

Photons, our main way of observing the Universe, are massless, and thus their four momentum $p_\mu = (E, p^i)$ has zero magnitude: $p_\mu p^\mu = 0$. Combined with the Friedman equations, this means that the energy of photons evolves as $E \propto a^{-1}$, or equivalently, the wavelength $\lambda \propto a(t)$. This matches with the evolution of their energy density, $\rho \propto a^{-4}$, that we found in Section 1.2. Besides the dilution due to the expansion of space which gives an a^{-3} term, photons also have redshift effect a^{-1}. Therefore, we introduce the cosmological redshift, z, to track the redshift due to the expansion of the Universe. It is given by:

$$z \equiv \frac{\lambda_0 - \lambda_e}{\lambda_e} = \frac{a(t_0)}{a(t_e)} - 1, \tag{1.21}$$

where the photon was emitted at time t_e with wavelength λ_e and observed at time t_0 with wavelength λ_0. Very often, the wavelength of the emitted photons is known due to their relation to known features in an object's spectrum. This makes the redshift one of the most important quantities in observational cosmology. It is convention in cosmology to set the scale factor at the present time to equal 1: $a(t_0) = 1$. This means we can directly relate the redshift to the scale factor of the Universe at the time when the photon was emitted:

$$1 + z = \frac{1}{a(t_e)}. \tag{1.22}$$

As the scale factor is the only unknown function in the metric, in the absence of peculiar velocities the redshift acts as both a measure of time and distance.

A very useful concept for measuring distances in an evolving Universe is that of a comoving distance. This is a rescaling of physical distance by the scale factor, meaning two objects initially at rest with respect to each other remain equidistant in comoving coordinates. An object that emitted light at time t_e, is today at a comoving distance given by[4]:

$$D_C(t_e) = \int_{t_e}^{t_0} \frac{dt'}{a(t')} = \int_0^{z_e} \frac{dz}{H(z)}, \tag{1.23}$$

where $z_e = z(t_e)$.

Besides the comoving distance, cosmologists often use the angular diameter distance and the luminosity distance.

The angular diameter distance is used when we want to relate the physical size, s, of an object

[4]Where we keep to the convention of $c = 1$.

to its observed angular size θ. The angular size is commonly assumed to be small, such that the angular diameter distance can be expressed as $D_A = s/\theta$. In terms of the comoving distance, this is given by:

$$D_A(z_e) = \frac{S_k(D_C)}{1+z_e}, \tag{1.24}$$

where the function S_k depends on the curvature parameter Ω_k, and is given by:

$$S_k(D_C) = \begin{cases} \sinh(D_C H_0\sqrt{\Omega_k})/(H_0\sqrt{\Omega_k}), & \Omega_k > 0 \\ D_C, & \Omega_k = 0 \\ \sin(D_C H_0\sqrt{|\Omega_k|})/(H_0\sqrt{|\Omega_k|}), & \Omega_k < 0 \end{cases}. \tag{1.25}$$

An often used closely related quantity is the comoving angular diameter distance, given by:

$$D_M(z_e) = (1+z_e)D_A(z_e) = S_k(D_C). \tag{1.26}$$

Finally, we consider observations of a distant object with absolute luminosity, L. If the Universe was flat and static, the observed flux at a distance d from the source would be given by $F = L/(4\pi d^2)$. When translating this into the FLRW spacetime, we first have to change the area of the spherical shell to comoving coordinates such that it is given by $4\pi D_M^2$. Furthermore, photons are redshifted on their way to us, and their rate of arrival changes due to the expansion of the Universe, each giving an extra term of a^{-1}. Therefore, the observed flux is given by:

$$F = \frac{L}{4\pi D_M^2 (1+z)^2} \equiv \frac{L}{4\pi D_L^2}, \tag{1.27}$$

where we have defined the luminosity distance:

$$D_L(z_e) = (1+z_e)D_M(z_e) = (1+z_e)S_k(D_C). \tag{1.28}$$

The distances introduced in this section allow us to work with familiar physical quantities such as flux, angles and the physical size of objects. Therefore, they allow us to tie cosmological observables to theory and study the geometry of a constantly evolving Universe.

CHAPTER 2

THE INHOMOGENEOUS UNIVERSE

So far we have relied on the cosmological principle, by assuming homogeneity and isotropy, to build the framework of ΛCDM. However, the Universe around us is clearly inhomogeneous and has plenty of structure. Therefore, in this chapter we turn our attention to the formation of this structure. We start in Section 2.1 by introducing a mechanism to seed small primordial inhomogeneities while still maintaining the homogeneous background. After that, in Sections 2.2 and 2.3 we will study the evolution of these perturbations through the different epochs and show how they lead to the Universe we observe today.

2.1 INFLATION

The FLRW framework we introduced in Chapter 1 has a singularity as $a \to 0$. This Big Bang singularity points to either the model being incorrect in the very early Universe, or at best incomplete. Cosmic inflation was introduced to solve some of the perceived problems with the standard Big Bang cosmological model. Here we will briefly mention these problems, however our main focus in this chapter is on the inhomogeneities in the Universe. As we shall see, one of the main features of inflation is that it provides a framework for the Universe to be very homogeneous on large scales, and yet have small deviations from homogeneity that will act as the seeds for the formation of cosmic structure.

The problems that motivated the development of cosmic inflation are:

- **Flatness problem:** We have previously defined the curvature density parameter $\Omega_k = -k/(a^2H^2)$. As far as we know, the early Universe was first dominated by radiation and after that by matter. Both of these have a Hubble radius, $r_H = 1/(aH)$, that grows with time, leading to any small deviations from $\Omega_k = 0$ being amplified. In other words, $\Omega_k = 0$ is an unstable equilibrium. Therefore, if the Universe we observe today is very close to being flat, it must have been even closer in the distant past. This fine-tuning issue can only be addressed if the early Universe underwent an epoch of accelerated expansion such that r_H shrinks over time (similar to the effect of dark energy today).

- **Horizon problem:** The Universe appears to be very homogeneous on large scales. The CMB is isotropic to about one part in 10^5. For two different patches of the CMB to have the same temperature, they must have been in causal contact in the distant past. In the standard big bang cosmology, the comoving horizon, that is the largest distance between two points such that they had time to communicate since the big bang, is growing monotonically. Therefore, regions that are currently entering the horizon appear to have never been in causal contact before, and yet the Universe still appears homogeneous on the largest scales. This leads us to assume that these patches must have been in causal contact in the early Universe, but there was a mechanism that pushed them out of causal contact. As with the last point, the only way to explain this would be for the early Universe to undergo a period of accelerated expansion, such that the comoving horizon shrinks, pushing regions that were previously in causal contact outside the horizon.

- **Magnetic-monopole problem:** The final problem has to do with the absence of magnetic monopoles in the Universe today. A majority of Grand Unified Theories (GUT) predict that a large number of stable magnetic monopoles should have been produced in the early Universe. As we have not detected any such particles, a mechanism to explain their absence is required.

We note that for of all of these problems, one could make the case that the Universe is just the way it is. Especially with the last problem, introducing a theory to explain the absence of hypothetical particles appears dubious [1]. However, the attraction of inflation is that it simultaneously provides an explanation to all of these problems, and furthermore has some predictive power, as we shall see.

In order to address the problems described above, we need the early Universe to undergo an

[1] As was eloquently put by Martin Rees: *Skeptics about exotic physics might not be hugely impressed by a theoretical argument to explain the absence of particles that are themselves only hypothetical.* Rees (1997)

epoch of exponential expansion, such that the Hubble radius and the comoving horizon shrink. Such an expansion would also dilute magnetic monopoles in the Universe, potentially making them extremely rare. This is achieved if the Universe is effectively dominated by a fluid with equation of state $w < -1/3$. To achieve this, we start by assuming that the accelerated expansion is driven by a scalar field φ called the inflaton. The energy momentum tensor for a scalar field is given by:

$$T_\nu^\mu = g^{\mu\beta}\partial_\beta\varphi\partial_\nu\varphi - \delta_\nu^\mu\left(\frac{1}{2}g^{\lambda\alpha}\partial_\alpha\varphi\partial_\lambda\varphi - V(\varphi)\right), \tag{2.1}$$

where $V(\phi)$ is the potential. Assuming homogeneity, we can derive the energy density and pressure for this field:

$$\rho_\varphi = \frac{1}{2}\dot{\varphi}^2 + V(\varphi), \tag{2.2}$$

$$P_\varphi = \frac{1}{2}\dot{\varphi}^2 - V(\varphi), \tag{2.3}$$

where the first term represents the kinetic energy, and $\dot{\varphi} = \partial_t\varphi$. Therefore, the equation of state is given by:

$$w = \frac{\frac{1}{2}\dot{\varphi}^2 - V(\varphi)}{\frac{1}{2}\dot{\varphi}^2 + V(\varphi)}. \tag{2.4}$$

This means that we can obtain accelerated expansion ($w < -1/3$) if the potential energy dominates over the kinetic energy: $V(\varphi) \gg \dot{\varphi}^2$. This is known as the slow-roll limit. We also want this condition to be maintained long enough for inflation to solve the problems we discussed (Liddle and Lyth, 2000). Therefore, we also require the second derivative of ϕ to be small: $|\ddot{\varphi}| \ll |3H\dot{\varphi}|$ and $|\ddot{\varphi}| \ll |\partial_\phi V|$. These conditions are usually expressed by requiring two slow roll parameters be small. The parameters are given by (Peacock, 1999):

$$\varepsilon_V = \frac{m_P^2}{16\pi}\left(\frac{\partial_\varphi V}{V}\right)^2, \tag{2.5}$$

$$\eta_V = \frac{m_P^2}{8\pi}\frac{\partial_\varphi\partial_\varphi V}{V}. \tag{2.6}$$

As long as these conditions are satisfied, the field is slowly rolling down the potential V, and the Universe undergoes accelerated expansion. When the field reaches the minimum potential, it enters a period of reheating when it oscillates away its energy into Standard Model fields (Dodelson, 2003). The equation of motion for φ in the slow-roll limit is given by (Peacock, 1999):

$$3H\dot{\varphi} = -\partial_\varphi V. \tag{2.7}$$

When building an inflation model, we generally require two important conditions. The first is that inflation stops at some point, while the second is that it must go on for a certain amount of e-folds[2]. In particular, scales that enter the horizon today, ~ 60 e-folds after the end of inflation, must have been in causal contact before inflation started. This puts a lower limit of ~ 60 e-folds on the duration of inflation. There is no required upper limit for the duration of inflation, and in the limit of infinite duration there are models of eternal inflation that provide a possible answer to the problem of the initial singularity.

Finally, inflation provides a mechanism for the creation of initial perturbations from homogeneity that go on to act as the seeds of structure formation in the Universe. This is achieved by magnifying quantum fluctuations to macroscopic scales, by pushing them outside the horizon during inflation (effectively freezing them). After inflation, these perturbations re-enter the horizon and begin evolving.

We can represent the statistics of these perturbations through the variance of their harmonic modes, given by the power spectrum. For a homogeneous and isotropic overdensity field $\delta(\vec{x}) \equiv \rho(\vec{x})/\bar{\rho} - 1$ (where $\bar{\rho}$ is the mean density) the power spectrum is given by:

$$\langle \hat{\delta}(\vec{k}) \hat{\delta}(\vec{k}')^* \rangle = (2\pi)^3 \delta_D(\vec{k} - \vec{k}') P(k), \tag{2.8}$$

where $\hat{\delta}(\vec{k})$ is the Fourier transform of $\delta(\vec{x})$.

These perturbations are generally assumed to be very close to Gaussian, as predicted by most inflationary models. However, there are also models which predict deviations from a Gaussian random field (e.g. Allen, Grinstein, and Wise, 1987; Kofman and Pogosyan, 1988). The search for non-gaussianities is an active area of research due to its potential to distinguish between different inflation models.

Inflation predicts a primordial power spectrum that is very close to scale invariant. This can be expressed as (Dodelson, 2003):

$$P_{\mathcal{R}}(k) = 2\pi^2 k^{-3} A_s \left(\frac{k}{k_p}\right)^{n_s - 1}, \tag{2.9}$$

where $\mathcal{R}$ represents curvature perturbations, n_s is the scalar spectral index, k_p is the pivot scale and A_s is the amplitude of scalar perturbations. The prediction of a nearly scale-invariant spectrum corresponds to a value of n_s that is close but not equal to 1. Currently, the best measurements of this parameter come from the CMB: $n_s = 0.9649 \pm 0.0042$ (Planck Collaboration et al., 2018), and thus

[2] One e-fold is expansion by a factor of e.

verify the prediction. The parameters A_s and n_s are a core part of the standard ΛCDM cosmological model.

2.2 Scalar perturbations

In general relativity we have three types of perturbations: scalar, vectorial and tensorial. In linear theory these evolve independently (Dodelson, 2003). This is an important result because it means that we can study them separately, without having to consider e.g. tensor perturbations initiating scalar perturbation during their evolution. Scalar perturbations are the most studied as they are the main driver of structure formation in the Universe. Vector perturbations only have decaying modes and are not sourced in appreciable amounts in most cosmological scenarios (Dodelson, 2003). On the other hand, tensor perturbations form gravitational waves, which are an important topic in cosmology. However, they are not directly relevant to this thesis. Therefore, we will only focus on scalar perturbations as these are the primary cause of structure formation.

In order to study the evolution of perturbations, we have to introduce deviations from homogeneity in our formalism. We do this by introducing small perturbations (inhomogeneities) to the FLRW metric. The perturbed metric is given by:

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + h_{\mu\nu}, \tag{2.10}$$

where $\bar{g}_{\mu\nu}$ is the FLRW metric and $h_{\mu\nu}$ represents perturbations, assumed to be small. In particular, we assume that the perturbative part of the metric does not change the global evolution of the background described by $\bar{g}_{\mu\nu}$. The choice of coordinates, known as gauge, is very important when it comes to perturbation theory as certain gauges can significantly simplify the solutions. Here, we will work in the Newtonian gauge, and write the perturbed line element:

$$ds^2 = a^2(\eta)[(1+2\Psi)d\eta^2 - (1-2\Phi)\delta_{ij}dX^i dX^j], \tag{2.11}$$

where Φ and Ψ define our perturbations, and $\eta(t) = \int_0^t dt'/a(t')$ is the conformal time. We can identify Ψ with the Newtonian gravitational potential while Φ represents the perturbation to the spatial curvature. In the absence of anisotropic stress, we have $\Phi = \Psi$.

Perturbation modes behave differently depending on the cosmological era. In the early Universe, we have two relevant epochs, that of radiation domination and that of matter domination, while in the late Universe we have dark energy domination. Furthermore, modes also behave differently

depending on their scale relative to the horizon at that time. In the early Universe, perturbations are super-horizon (pushed there by inflation), meaning their length is much larger than the Hubble radius $2\pi/k \gg r_H$. As the horizon expands, perturbations transition to the sub-horizon regime ($2\pi/k \ll r_H$).

At early times, large-scale and small-scale modes behave differently depending on when they enter the horizon, while at late times, all modes evolve the same. As we observe these perturbations in the late Universe, it is useful to break down the potential into three different components, as shown by Dodelson (2003):

$$\Phi(\mathbf{k},a) \sim \mathcal{R}(\mathbf{k}) \times \{\text{Transfer Function}(k)\} \times \{\text{Growth Factor}(a)\}. \tag{2.12}$$

The first component represents the primordial curvature perturbations $\mathcal{R}$ that we introduced in Section 2.1. The second component is a scale dependent transfer function that encapsulates the evolution of different modes in the early Universe, defined as:

$$T(k) \equiv \frac{\Phi(\mathbf{k}, a_{\text{late}})}{\Phi_{\text{large-scale}}(\mathbf{k}, a_{\text{late}})}, \tag{2.13}$$

where a_{late} is an epoch well after the matter-radiation equality, and the large-scale solution denotes modes that enter the horizon deep into matter domination. The third component is a time dependant growth factor that describes the evolution of all modes in the late Universe, defined by:

$$\frac{\Phi(\mathbf{k},a)}{\Phi(\mathbf{k},a_{\text{late}})} \equiv \frac{D_+(a)}{a}, \text{ with } a > a_{\text{late}}. \tag{2.14}$$

Precise descriptions on how to build the transfer function and growth factor can be found in Peacock (1999) and Dodelson (2003). Here we will only give a qualitative description of the evolution of the potential and different density modes which are used to build the two objects.

2.2.1 The evolution of the potential

Assuming a perfect fluid with adiabatic perturbations, the evolution of the gravitational potential is described by:

$$\Phi'' + 3(1+w)\mathcal{H}\Phi' + wk^2\Phi = 0, \tag{2.15}$$

where $\mathcal{H} = aH$ and $'$ denotes a derivative with respect to η. For super-horizon modes, the potential is constant during both radiation and matter domination, $\Phi = \Phi(0)$. However, during the transition

through the epoch of matter-radiation equality, the potential for these modes drops by a factor of $9/10$ such that $\Phi \to (9/10)\Phi(0)$ (Dodelson, 2003).

For sub-horizon modes we start with the era of radiation domination where $w = 1/3$ and the potential is described by:

$$\Phi'' + \frac{4}{\eta}\Phi' + \frac{k^2}{3}\Phi = 0. \tag{2.16}$$

This has solutions of the form $\Phi \propto \cos(k\eta/\sqrt{3})/(k\eta)^2$, corresponding to decaying oscillations. Therefore, modes that enter the horizon during radiation domination oscillate and decay away. On the other hand, during matter domination $w = 0$ and we have $\Phi'' + 3\mathcal{H}\Phi' = 0$. This leads to solutions $\Phi = \text{const}$, meaning that during matter domination the potential is constant on all scales. Note that this assumes negligible curvature. If the Universe has negative curvature, or dark energy starts to become important, structure growth will be suppressed and the potential starts to decay (Peacock, 1999).

2.2.2 The evolution of single fluids

Having considered the evolution of the potential, we turn our attention to density perturbations. We formally define the overdensity:

$$\delta_i = \frac{\delta\rho_i}{\bar{\rho}_i}, \tag{2.17}$$

where $\delta\rho_i$ is a (small) perturbation to the mean energy density $\bar{\rho}_i$ of fluid i. Combining the continuity and Euler equations in Fourier space with the Poisson equation, we can describe the evolution of density perturbations by:

$$\ddot{\delta}_i + 2H\dot{\delta}_i + \left(\frac{c_s^2 k^2}{a^2} - 4\pi G\bar{\rho}\right)\delta_i = 0, \tag{2.18}$$

where the speed of sound is given by $c_s^2 = d\bar{P}_i/d\bar{\rho}_i$. This equation includes the Hubble drag, fluid pressure and the effect of gravity. It has a critical scale called the Jeans scale, given by $k_J = \sqrt{4\pi G\bar{\rho}/c_s^2}$. On large scales, $k \ll k_J$ and modes have growing solutions. On small scales, $k \gg k_J$ and the solutions describe a damped oscillator. Therefore, on small scales, perturbations cannot grow due to the effect of pressure. The transition between growth and oscillations is determined by the properties of the fluid through w and c_s. In the case of radiation, only the largest scales grow because c_s is very large. On the other hand, dark matter has a small c_s, and most perturbations grow.

For pressureless dark matter, $w = 0$ and we can ignore c_s. We can use the substitution $y = a/a_{\text{eq}}$,

where $a_{\rm eq}$ is the scale factor at the matter-radiation equality, to obtain the Mészáros equation (Meszaros, 1974):

$$\frac{d^2\delta_c}{dy^2} + \frac{2+3y}{2y(y+1)}\frac{d\delta_c}{dy} - \frac{3}{2y(y+1)}\delta_c = 0. \tag{2.19}$$

The dominant solution during radiation domination is logarithmic growth, $\delta_c \propto \ln a$. This means that dark matter can collapse, but this is significantly slowed by the dominant radiation component, which cannot collapse on most scales. During matter domination, $y \gg 1$, dark matter can collapse freely and the dominant solution is $\delta_c \propto a \propto t^{2/3}$. On the other hand, when dark energy becomes dominant in the late Universe, growth is suppressed again on all scales and $\delta_c \propto$ const.

2.2.3 Photon-baryon perturbations

Before the epoch of recombination photons and baryons were tightly coupled due to photon scattering on charged particles. We therefore have to treat these two components with different properties as a single fluid in order to understand their evolution. Again using the Euler equation, we can describe the evolution of a Fourier mode of the baryon density perturbation as (Peebles and Yu, 1970; Doroshkevich, Zel'dovich, and Syunyaev, 1978; Hu and White, 1996):

$$\frac{d}{d\eta}[(1+R)\delta_b'] + \frac{k^2}{3}\delta_b = -k^2(1+R)\Psi - \frac{d}{d\eta}[3(1+R)\Phi'], \tag{2.20}$$

where $R = 3\rho_b/4\rho_\gamma$ with the photon density ρ_γ. This is the equation of a driven oscillator with frequency $c_s k$, where the speed of sound in the photon-baryon fluid is $c_s = 1/\sqrt{3(1+R)}$. The acoustic waves produced by these oscillations leave an imprint in the baryon and photon distributions, known as *baryon acoustic oscillations* (BAO).

Qualitatively, we can understand this behaviour by considering an initial overdensity which corresponds to a region that is overpressured compared to its surroundings (Eisenstein, Seo, and White, 2007). A spherical sound wave is driven out in order for the fluid to equilibrate. Dark matter is left behind at the centre of the perturbation because it is pressureless. The wave travels outwards at the speed of sound c_s, which is determined by the opposing forces of gravity and pressure. After recombination, photons are allowed to free-stream, and the baryons lose the pressure driving force. The wave becomes frozen in the baryon distribution at a radius of:

$$r_d = \int_{z_d}^{\infty} \frac{c_s(z)}{H(z)} dz, \tag{2.21}$$

where $z_d \simeq 1060$ is the redshift of the *drag epoch*, when the baryons decouple from radiation. This

is slightly later than the epoch of recombination which is at a redshift of $z_r \simeq 1090$. In Chapter 5 we will go in more detail into how r_d depends on cosmological parameters, and its usefulness for observational cosmology.

Once the baryons are free from the effects of pressure, perturbations can grow again as $\delta_b \propto a$, similarly to dark matter perturbations. However, we are left with a dark matter overdensity at the centre and an echo of baryons at radius r_d. As these two evolve together due to gravity, baryons fall into the central dark matter overdensity and dark matter falls into the baryon acoustic overdensity. Over the whole Universe, we have many such perturbations overlapping, which leads to increased clustering at a typical separation of $r_d \simeq 150$ Mpc. This appears as a peak in the correlation function and a harmonic sequence in the power spectrum. See Figure 2.1 for a visualization of the evolution of BAO.

On small scales the acoustic oscillations are damped due to photon-baryon diffusion which leads to a broadening of the acoustic peak. This effect is known as Silk damping (Silk, 1968). Furthermore, the peak is distorted by small-scale non-linear evolution in the late Universe. We will discuss the effect of this in Section 2.3.

2.2.4 The growth factor

In order to build the growth factor in the late Universe we have to take into account the effect of dark energy on the evolution of matter perturbations. As baryons and dark matter behave similarly after recombination, we treat them as one component with the total matter perturbation $\bar{\rho}_m \delta_m = \bar{\rho}_c \delta_c + \bar{\rho}_b \delta_b$. Neglecting the effect of massive neutrinos, their evolution can be described by:

$$\frac{d^2\delta_m}{da^2} + \frac{d\ln(a^3H)}{da}\frac{d\delta_m}{da} - \frac{3\Omega_m H_0^2}{2a^5H^2}\delta_m = 0. \tag{2.22}$$

In general this has to be solved numerically, however, we can obtain an analytic solution by assuming that dark energy is the cosmological constant, Λ, and ignoring the small radiation component. The growth factor is then given by (Dodelson, 2003):

$$D_+(a) = \frac{5\Omega_m}{2}\frac{H(a)}{H_0}\int_0^a \frac{da'}{[a'H(a')/H_0]^3}. \tag{2.23}$$

This describes the growth of all matter perturbations at late times. When matter dominates, the overdensity grows as $\delta_c \propto D_+$. Once Λ starts to become relevant, the growth slows down due to accelerated expansion.

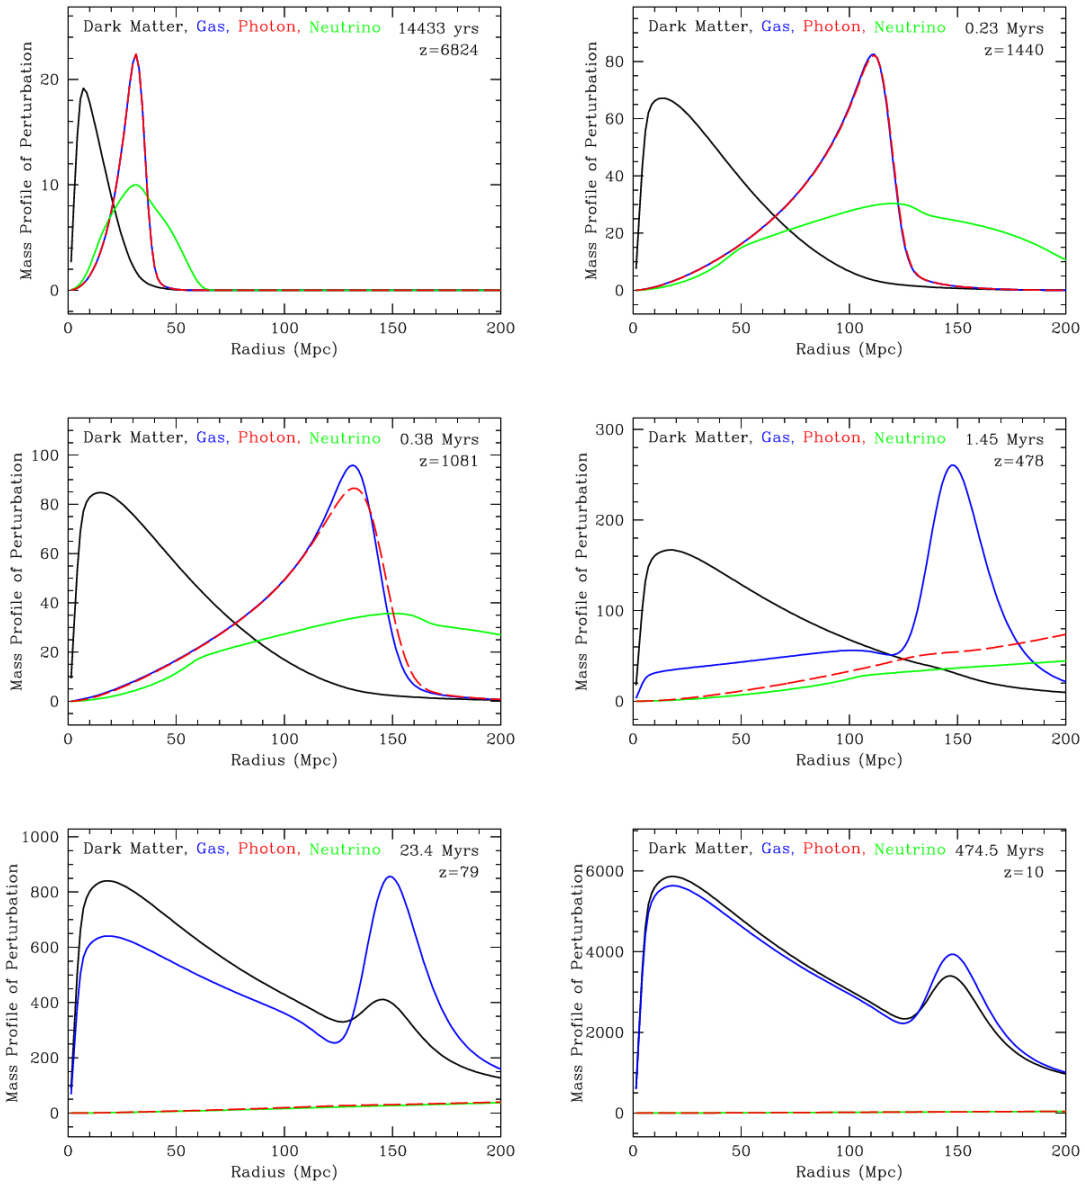

Figure 2.1. The evolution of the mass profile of a perturbation in different components. In the top plots, the baryon-photon acoustic wave is propagating outwards, while the dark matter perturbation grows towards the centre, with a wake following the acoustic wave. In the middle plots, the photons decouple from the gas after recombination and are free-streaming away. The acoustic peak becomes frozen in the gas distribution at around 150 Mpc. In the bottom two plots, gravitational collapse takes over, and all matter follows the gravitational potentials set by the dark matter perturbation in the centre and the baryonic perturbation in a shell. Figures from Eisenstein, Seo, and White (2007).

Finally, we define the *growth rate*, f, as the logarithmic derivative of the growth factor:

$$f(a) \equiv \frac{d\ln D_+(a)}{d\ln a} \tag{2.24}$$

This has an empirical fitting under GR, $f(a) \simeq \Omega_m(a)^{0.55}$, which is remarkably precise even for dynamical models of dark energy (Dodelson, 2003). In linear theory, the growth rate relates the velocity field, $\theta(k)$, to the density field: $\theta(k) = -f\delta(k)$. This makes it useful when working with galaxy surveys in redshift space as we shall see in Chapter 3.

2.3 Non-linear growth

After recombination all matter is free to collapse under gravity. The initial small perturbations we observe in the CMB grow into the stars, galaxies and clusters of galaxies we observe today. During this evolution, the linear theory approximation we have used so far starts to break down, especially on small scales. In the late Universe the effects of small-scale non-linear evolution have an impact even on large-scale features. Here, we briefly describe some of the methods used to study and model non-linear growth, and we focus on its impact on the BAO feature.

A natural extension to the linear theory framework is to go beyond first order in perturbation theory. These techniques are especially useful on large scales in the late Universe, the so-called quasi-linear regime. On these scales, we can use modifications to linear theory that account for the impact of non-linear growth. For a review of such techniques, see Bernardeau et al. (2002) or the more recent Desjacques, Jeong, and Schmidt (2018).

On small scales, perturbation theory approaches are generally not sufficient, and so we must rely on numerical simulations of structure formation. Due to recent advances in computational power, we can now simulate the behaviour of large-scale structure. These simulations are typically seeded with a large number of particles distributed as a Gaussian random field that matches the initial conditions of our Universe (as best as we currently understand them). In N-body simulations, the evolution of structure under the effect of gravity is simulated (we do not distinguish between baryons and dark matter). These simulations are very useful, especially on medium to large scales as dark matter makes up $\sim 85\%$ of matter in the Universe (for a review see Dehnen and Read, 2011). In hydrodynamic simulations, we model baryons more accurately by simulating their hydrodynamic behaviour, which allows us to study the evolution of gas and its properties such as pressure and temperature. Furthermore, we can simulate the complex interaction of galaxies with their surrounding environment, which is composed of a large range of astrophysical effects. This

allows simulations to be accurate down to small scales, and therefore incredibly useful for both cosmology and astrophysics.

When it comes to cosmological analyses, both higher order perturbation theory and simulations have been used to inform simple empirical models of non-linear effects, for example in the case of BAO (Eisenstein, Seo, and White, 2007). We already described the broadening of the acoustic peak due to pre-recombination effects, however, another major source of broadening is the non-linear evolution of structure. Even though this is a small scale effect, it has an impact on a large scale feature. A common approach in the literature is to perform a reconstruction of the density field by either directly smoothing the density, or by using the linear velocity field to approximately reverse the growth (Eisenstein et al., 2007). This results in sharper measurements of the peak, and thus tighter constrains on its scale. A slightly simpler approach is to just model the broadening of the peak. Eisenstein, Seo, and White (2007) used a simple theoretical argument and simulations to show that this broadening can be modelled as a simple Gaussian smoothing given by:

$$P(\mathbf{k}) = P_{\rm lin}(\mathbf{k}) \exp\Big(-\frac{k_{||}^2}{2\sigma_{||}^2} - \frac{k_{\perp}^2}{2\sigma_{\perp}^2}\Big), \tag{2.25}$$

where $P_{\rm lin}(\mathbf{k})$ is the linear theory power spectrum, $\mathbf{k}$ is the wavenumber with components $(k_{||}, k_{\perp})$ along and across the line of sight respectively and $(\sigma_{||}, \sigma_{\perp})$ are two parameters that correspond to the smoothing scales along and across the line of sight respectively. Perturbation theory can be used to infer the expected values of these parameters, and simulations to test how realistic the model is. We will use this simple model in Chapter 7.

Another non-linear effect that appears in large-scale structure surveys are *fingers-of-god* (FOG). This effect appears because we generally build our maps in redshift space (see Chapter 3). Non-linear structure formation leads to large peculiar velocities that induce red and blue-shifts comparable or even larger than the one from the Hubble flow. This only affects line-of-sight correlations, but leads to very prominent elongated features known as fingers-of-god. These can generally be modelled using simple analytical formulas (see e.g. Percival and White, 2009) informed by simulations.

CHAPTER 3

COSMOLOGICAL PROBES

Over the last few decades, cosmology has entered an era of unprecedented precision. We now have many different observables that probe different epochs and regimes through the evolution of the Universe. These measurements allow us to test the theoretical framework we have introduced in the last two chapters, and search for missing components. Much of the interest in observational cosmology is focused on studying the elusive dark matter and dark energy, as well as understanding the origins of the Universe and its evolution. Here we introduce some of the major probes currently used in cosmology, focusing on the ones directly relevant to this thesis.

3.1 THE COSMIC MICROWAVE BACKGROUND

Described as the *first light in the Universe*, the Cosmic Microwave Background (CMB) is composed of the photons that were finally allowed to free-stream once the first atoms formed at the epoch of recombination. Before that epoch, scattering of photons on charged particles effectively coupled the baryon fluid to the photon fluid. After the charged particles were locked into the first atoms, the mean free path of CMB photons grew larger than the size of the horizon. This happened at a redshift of $z_r \simeq 1090$.

The CMB is currently the most powerful cosmological probe to which we have access, in the sense that it provides the most precise measurements on most cosmological parameters in ΛCDM and extensions. This is due to the fact that the CMB probes an epoch before the formation of non-linear structure, when the Universe was remarkably uniform and its physics relatively easy

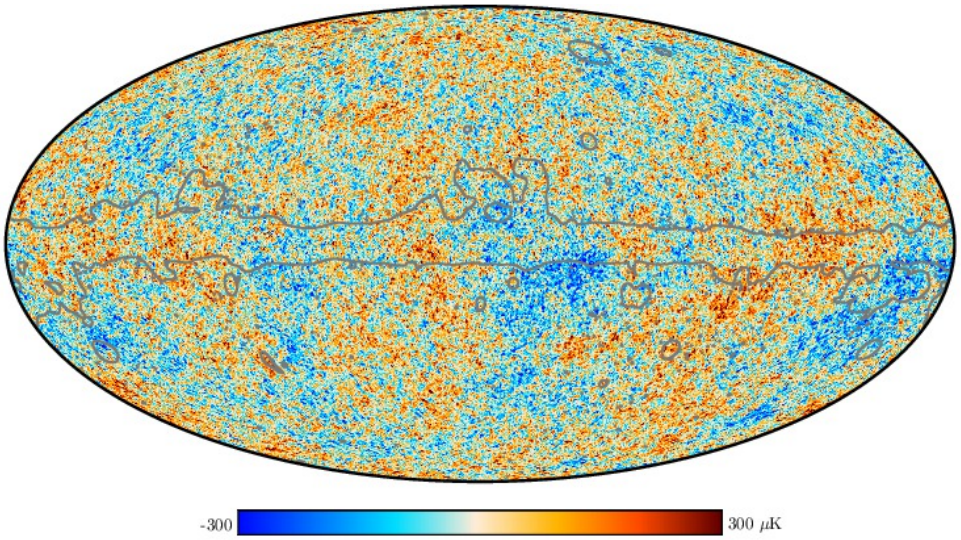

Figure 3.1. Map of CMB temperature anisotropies measured by the Planck space mission (Planck Collaboration et al., 2020a). The gray lines show the region around the Milky Way galaxy which has been in-painted.

to derive analytically. In fact, the framework we discussed in Section 2.2 accurately describes the core features observed in the CMB, with just a few modifications necessary to account for the interactions of CMB photons after recombination.

We currently observe CMB photons arriving from an apparent 2D surface with a radius determined by the distance travelled by photons after decoupling. This signal was famously discovered by accident in 1964 (Penzias and Wilson, 1965; Dicke et al., 1965). The Cosmic Background Explorer (COBE) satellite accurately measured the current temperature of the CMB at $T = 2.72548 \pm 0.00057$ K (Fixsen et al., 1996; Fixsen, 2009), as well as detecting for the first time anisotropies at the 10^{-5} level (Smoot et al., 1992; Bennett et al., 1996). These anisotropies were measured with increasing precision by subsequent space missions, first the Wilkinson Microwave Anisotropy Probe (WMAP; Bennett et al., 2003) and then the Planck space satellite (Planck Collaboration et al., 2014a).

In order to constrain cosmology, the CMB anisotropy maps (Figure 3.1) are generally used to measure the angular power spectrum. The primary features observed in this spectrum are the acoustic oscillations of the baryon-photon fluid we described in Section 2.2.3. These features allow us to precisely constrain the baryon and dark matter densities. This is because the relative densities of baryons and radiation affects the relative amplitudes of the peaks due to baryon drag, and dark matter acts as an extra source of gravitational potential. Furthermore, observing the large scales in the CMB from full-sky maps allows us to measure the scalar spectral index n_s, and the

pristine nature of the perturbations gives us the amplitude of primordial perturbations A_s. Finally, we can use BAO as a standard ruler to get precise measurements of Ω_m, Ω_b and H_0. For a recent summary of current constraints see Planck Collaboration et al. (2020b) and Planck Collaboration et al. (2020c).

After recombination, CMB photons still have a few interactions with the rest of the Universe. These interactions have to be taken into account when modelling the CMB power spectrum, but they also provide an extra source of information on structure in the Universe and its evolution. The most important of these interactions is due to gravitational potentials. Firstly, the photons are red or blue-shifted depending on their position in an overdensity or underdensity at decoupling, which is known as the Sachs-Wolfe effect (Sachs and Wolfe, 1967). Secondly, they are affected by time-varying gravitational potentials in the late Universe (especially due to dark energy), which is known as the integrated Sachs-Wolfe effect (Rees and Sciama, 1968). CMB photons can also interact with high-energy electrons via inverse Compton scattering, which is known as the Sunyaev-Zel'dovich effect (Sunyaev and Zeldovich, 1970; Sunyaev and Zel'Dovich, 1980). Finally, photons are gravitationally lensed by intervening matter, leading to a lensing signal in the CMB that allows us to study the evolution of structure (Lewis and Challinor, 2006).

3.2 Galaxy clustering

The structures we observe in the CMB anisotropies form the seeds for structure formation in the Universe. Therefore, mapping the large-scale structure of the Universe provides a rich source of information on its evolution and components. The simplest way to map this structure is to build maps of galaxies. Measuring the position of galaxies on the sky is a relatively easy endeavour, however, measuring how far away they are is much harder (especially for large data sets).

Spectroscopic surveys are the most common way to build three-dimensional (3D) maps of galaxies. They measure spectra for each object and use the features in the spectra, such as emission lines, to infer the redshift to the galaxy. These redshift maps can then be used to constrain cosmology, based on the fact that galaxies trace the dark underlying density field. The relationship between the density field and the galaxy field is a biased one (Peacock, 1999; Dodelson, 2003). This is due to the fact that only peaks in the density distribution go on to collapse into galaxies. At linear order, we can express this relationship as (Bardeen et al., 1986; Dekel and Lahav, 1999; Sheth and Tormen, 1999):

$$\delta_g(\mathbf{x}) = b\delta(\mathbf{x}), \tag{3.1}$$

where b is the linear bias parameter. In general this bias could be a function of both redshift and scale, however, we usually assume it is only a function of redshift.

Beyond the bias, we build the galaxy map in redshift space, not in real space. This means that peculiar velocities may change the position of galaxies, an effect known as redshift space distortions (RSD). The position in redshift space, $\mathbf{s}$, can be related to the position in real space, $\mathbf{x}$, through:

$$\mathbf{s} = \mathbf{x} - u_z(\mathbf{x})\hat{\mathbf{z}}, \tag{3.2}$$

where $u_z(\mathbf{x})$ is the galaxy velocity along the line-of-sight (which is taken to be the $\mathbf{z}$ direction). Assuming an irrotational velocity field, we write $u_z = \partial/\partial z \nabla^{-2}\theta$ where ∇^{-2} is the inverse Laplacian and $\theta \equiv \nabla \cdot \mathbf{u}$ is the velocity divergence. In the linear regime, we can express the galaxy density field in the Fourier space as (Percival and White, 2009):

$$\delta_g^s(k) = \delta_g(k) - \mu^2\theta(k), \tag{3.3}$$

where μ is the cosine of the line-of-sight angle. Again assuming linear theory for the velocity field, we have $\theta(k) = -f\delta(k)$, where f is the linear growth rate defined in Equation 2.24. Using this and the linear bias above, we can express the linear power spectrum of galaxies as (Kaiser, 1987):

$$P_g^s(k,\mu) = (b + f\mu^2)^2 P_m(k) = b^2(1+\beta\mu^2)^2 P_m(k), \tag{3.4}$$

where $P_m(k)$ is the isotropic linear power spectrum of matter, and we have defined $\beta \equiv f/b$. This relation can be used to confront theory with data, because we can measure P_g^s from galaxy redshift maps, and we can compute a model $P_m(k)$ based on the framework introduced in Chapter 2. Libraries such as CAMB (Lewis, Challinor, and Lasenby, 2000) and CLASS (Lesgourgues, 2011; Blas, Lesgourgues, and Tram, 2011) are used to solve the Boltzmann equations and compute accurate theoretical power spectra. The bias parameter is generally treated as a nuisance and marginalized out.

The problem with this simple linear approach is that galaxies form in the strongly non-linear regime, which means that the linear bias approach is not a very good description of the relationship between the galaxy and density fields (Gaztanaga, 1992). Instead, higher order perturbation theory approaches have been employed. For a recent review on this topic see Desjacques, Jeong, and Schmidt (2018). However, this simple description can be useful for tracers that are still in the quasi-linear regime, for example the Lyman-α forest which we discuss in the next section.

Even though we can use a model power spectrum to fit cosmological parameters given a model (e.g. ΛCDM) directly from the measured data, that is not the approach favoured in the literature so far. Instead, the *template fitting* approach is taken, where we use a template $P(k)$ to fit for a few well-defined physical quantities which tend to be useful in testing a large variety of cosmological models. The most commonly used such quantity is the size of the acoustic peak, which as we have seen in Section 2.2 can be used as a standard ruler. The BAO feature was first detected in the distribution of galaxies by Cole et al. (2005) and Eisenstein et al. (2005), although there were earlier hints (Percival et al., 2001). Since then, it has turned into one of the most robust cosmological measurements, with the latest results from the Baryon Oscillation Spectroscopic Survey (BOSS) and extended BOSS (eBOSS) at below 2% precision (Alam et al., 2017; eBOSS Collaboration et al., 2021).

Another commonly used feature are redshift space distortions. As we have seen above, RSD are caused by peculiar velocities because we construct our maps in redshift space. The quantity measured in RSD analyses is $f\sigma_8$, where the rms overdensity in a sphere of comoving radius R is given by (Dodelson, 2003):

$$\sigma_R^2 \equiv \langle \delta_{m,R}^2(\mathbf{x}) \rangle, \tag{3.5}$$

$$\text{with } \delta_{m,R}(\mathbf{x}) \equiv \int d^3x' \delta_m(\mathbf{x}') W_R(|\mathbf{x}-\mathbf{x}'|), \tag{3.6}$$

where $W_R(x)$ is a tophat window function with radius R equal to 8 h^{-1}Mpc by convention. Note that σ_8 is a linear-theory extrapolation. The $f\sigma_8$ term appears as the normalization of the velocity power spectrum, and we effectively write $P_\theta = f\sigma_8 P_\theta^t / (f\sigma_8)^t$, where P_θ^t and $(f\sigma_8)^t$ are the template velocity power spectrum and its normalization respectively. In linear theory $P_\theta = f^2 P_m$, therefore, when performing an RSD analysis we effectively measure $b\sigma_8$ and $f\sigma_8$, with only the second quantity being of cosmological interest. This type of analysis has also established itself as a standard cosmological tool (see e.g. Peacock et al., 2001; Blake et al., 2012; Alam et al., 2017; eBOSS Collaboration et al., 2021).

The next generation of spectroscopic surveys is spearheaded by the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al., 2016), which began survey operations in May 2021. DESI is aiming to measure the spectra of roughly 33 million objects over 14000 square degrees. The objects are split in a few different target classes. These are luminous red galaxies (LRGs) at redshifts $0.4 < z < 1.0$, emission line galaxies (ELGs) at $0.6 < z < 1.6$, quasars at $0.6 < z < 4.0$ and bright galaxies at $0 < z < 0.5$. The last class is meant to take advantage of the times when the

sky is not fully dark (e.g. due to moonlight) and it is hard to observe faint targets.

Besides spectroscopic surveys, galaxies can also be mapped by photometric surveys. These surveys take images of the sky in a few different bands, which are used to compute very rough estimates of the redshifts of the galaxies observed. Even though the redshift measurements are less accurate, these surveys are very powerful because they map a much larger number of galaxies. For example, while DESI will precisely measure the redshifts of about 33 million objects, the next generation of photometric surveys, spearheaded by the Euclid satellite (starting in 2022; Laureijs et al., 2011) and the Vera C. Rubin Legacy Survey of Space and Time (LSST, starting in 2023; LSST Science Collaboration et al., 2009) will map approximately one billion galaxies.

3.3 Lyman-α forest

Over the last decade, the Lyman-α (Lyα) forest became an important cosmological tracer at high redshift ($z > 2$). First observed in the 1960s (Schmidt, 1965; Scheuer, 1965; Bahcall and Salpeter, 1965; Gunn and Peterson, 1965), the Lyα forest consists of a series of absorption features in spectra of high redshift quasars (QSO).

These features can be qualitatively understood by considering photons with rest-frame wavelength $\lambda_{\rm rest} \leq \lambda_\alpha$, with the Ly$\alpha$ transition at $\lambda_\alpha = 1215.67\text{Å}$, that are emitted by a QSO at redshift z_q. As the Universe is expanding, these photons are redshifted while they travel towards us, meaning their wavelength increases. At some point their wavelength will be equal to λ_α, and they may be absorbed by a hydrogen atom (H I) in the inter-galactic medium (IGM). A fraction of these photons will be absorbed depending on the interaction probability. This probability depends on the density and kinematics of the neutral hydrogen gas in that region. Once the remaining photons reach us, we will observe an absorption line at $\lambda_{\rm obs} = \lambda_{\rm rest}(1+z_q)$. If we can measure the redshift of the quasar through other features in the spectrum (e.g. the Lyα emission line), we can compute the redshift of the gas that leads to the absorption line through its relation to the observed wavelength of the Lyα line.

Extrapolating to an entire spectrum of photons with different wavelengths, we will observe a series of such absorption lines blue-ward of the Lyα line (see Figure 3.2). These features trace the underlying density field between us and the quasar. We can express the Lyα optical depth through gas at a fixed density by (McQuinn, 2016):

$$\tau_\alpha(z) = 1.3\delta_b \left(\frac{x_{\rm H\,I}}{10^{-5}}\right)\left(\frac{1+z}{4}\right)^{3/2}\left(\frac{dv/dx}{H(z)/(1+z)}\right)^{-1}, \tag{3.7}$$

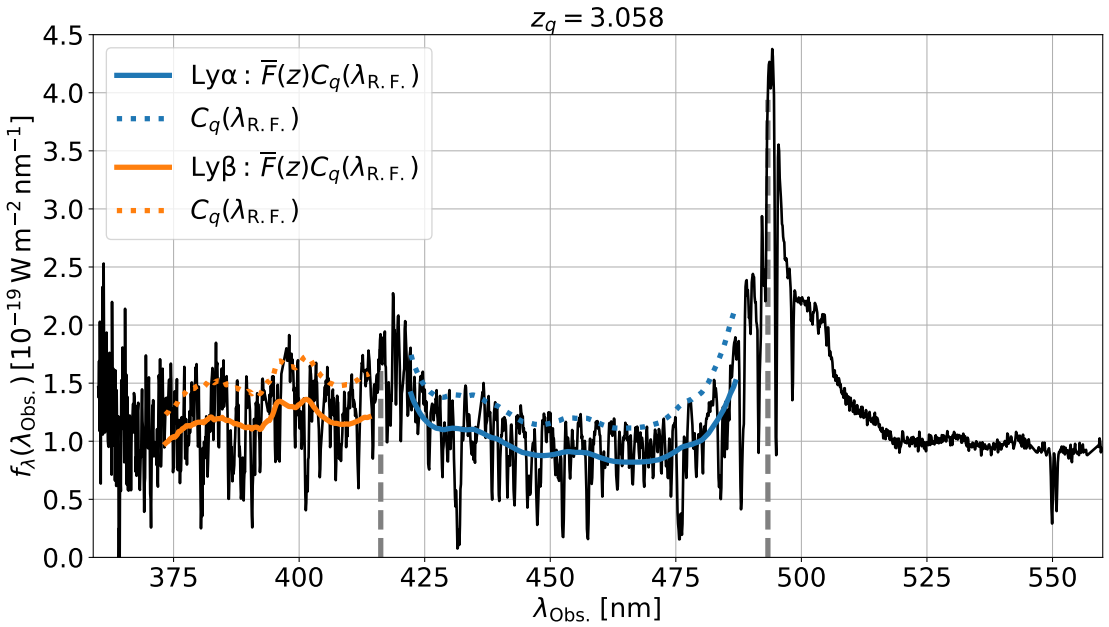

Figure 3.2. Quasar spectrum showing the absorption forest blue-ward of the Lyman-α emission line. The colours represent the different parts of the forest, with blue being the section between the Lyman-α and Lyman-β emission peaks, and orange being left of the Lyman-β peak. The solid lines are the fitted mean transmission times the quasar continuum. The dotted lines show the extrapolated quasar continuum, assuming a mean transmission function. Figure from du Mas des Bourboux et al. (2020).

where $x_{\text{H I}}$ is the fraction of hydrogen that is neutral, and we have assumed a smooth line-of-sight velocity gradient, dv/dx. This shows that the optical depth is directly proportional to the baryon density and the neutral hydrogen fraction. However, it also depends on the peculiar velocities of the absorbing gas. This will shift the inferred redshift of the absorption lines from their nominal values given by the Hubble flow, the same RSD effect we have discussed in Section 3.2.

The relevant tracer for us is the transmitted flux fraction, F, which is related to the optical depth via the non-linear mapping $F \propto \exp(-\tau_\alpha)$. We can compute fluctuations in the transmitted flux fraction from the directly observed absolute flux, $f(\lambda)$, through:

$$\delta_F(\lambda) = \frac{f(\lambda)}{\bar{F}(\lambda)C(\lambda)} - 1, \tag{3.8}$$

where $\bar{F}(\lambda)$ is the mean transmitted flux fraction, and $C(\lambda)$ is the quasar continuum. $\bar{F}(\lambda)$ and $C(\lambda)$ are generally unknown and have to be measured from the data, which can be a very difficult process, especially for low signal-to-noise spectra. Any error made in the continuum estimation will cause large-scale fluctuations (of the size of the forest) to be erased along the line-of-sight, and

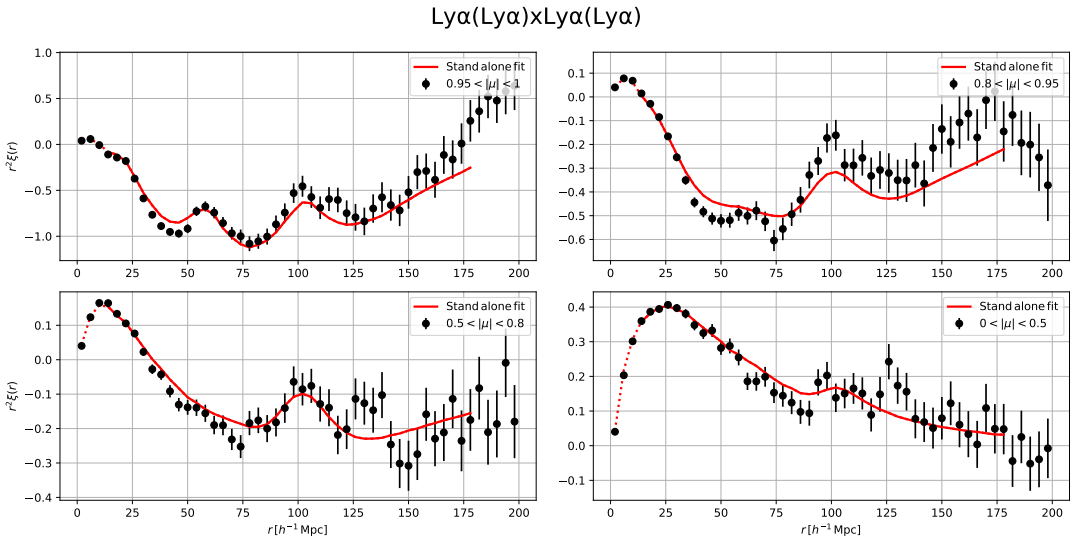

Figure 3.3. Plots of the 3D Lyα forest auto-correlation function in bins of the cosine of the line-of-sight angle, μ, along with the best-fit model. The BAO feature can be seen at separations of about 100 h^{-1}Mpc. The other visible peaks in the bins close to the line-of-sight are due to contamination by metal absorbers. Figure from du Mas des Bourboux et al. (2020).

will introduce correlated errors in the inferred δ_F from the same spectrum. Therefore, we generally ignore correlations within the same spectrum, and only use correlations between pixels in different spectra.

The 3D correlation of the Lyα forest was first measured from Sloan Digital Sky Survey (SDSS) data by Slosar et al. (2011). The most recent measurement from SDSS data release 16 (DR16) along with the best-fit model are shown in Figure 3.3. As the correlation function is the Fourier transform of the power spectrum, we can use a framework similar to the one we introduced for galaxies, in order to build a model of the correlation. However, in this case the tracer goes through a non-linear transformation, which introduces more complexity (McDonald et al., 2000; McDonald, 2003). In linear theory, the power spectrum of the Lyα transmitted flux fraction is given by:

$$P_F^s(k,\mu) = (b_F + b_{\eta,F} f \mu^2)^2 P(k), \tag{3.9}$$

where we now have an extra velocity divergence bias $b_{\eta,F}$ which multiplies the growth rate f. This extra bias appears because we work with the two-point statistics of transmitted flux, which has a non-linear mapping to the directly distorted field of optical depth (Seljak, 2012; Arinyo-i-Prats et al., 2015). As this new bias is also unknown, it means we cannot use the auto-correlation of the forest to directly measure f because the two are degenerate. We will study this issue in Chapter 7.

Beyond this simple prescription, there are a number of more complex processes that affect the forest, which must be taken into account. For example, not all the absorption we measure is caused by neutral hydrogen. Some of it is due to heavier elements, which we refer to as metals[1]. Another source of complexity are regions with high column density, which leave broad absorption features in the spectra. Even though these processes are of interest for astrophysics (e.g. Pieri et al., 2014; Pérez-Ràfols et al., 2018a; Pérez-Ràfols et al., 2018b), for the purposes of cosmological analyses they are treated as systematics.

For cosmology, the 3D distribution of the Lyα forest has so far only been used to measure BAO. The first analyses to detect the BAO feature in the Lyα auto-correlation used SDSS data release 9 (DR9; Slosar et al., 2013; Busca et al., 2013; Kirkby et al., 2013). Since then, it has been measured multiple times, with major improvements in the analysis, especially through improved modelling of contaminants and other systematic effects (Delubac et al., 2015; Bautista et al., 2017; de Sainte Agathe et al., 2019; du Mas des Bourboux et al., 2020). Furthermore, BAO has also been measured in the cross-correlation of the forest with quasar positions (Font-Ribera et al., 2014; du Mas des Bourboux et al., 2017; Blomqvist et al., 2019; du Mas des Bourboux et al., 2020).

The next generation of Lyα forest large-scale structure analyses will also come from DESI. The quasar sample measured by DESI will include roughly 0.7 million quasars at high redshift ($z > 2.1$), where we can measure the Lyα forest.

3.4 Other probes

There are many other cosmological probes, with even more constantly being developed. Here I will list some of them, and add a few details that are relevant for this thesis. We start with one of the first probes to be used in cosmology, supernovae (SNe) explosions of type Ia. These are termed *standardizable candles* because we can relate their peak luminosity to their luminosity distance (Whelan and Iben, 1973), with some corrections (see e.g. Betoule et al., 2014; Scolnic et al., 2018). They are thought to be caused by mass transfer onto a white dwarf star from a companion (Whelan and Iben, 1973). Once the white dwarf reaches a critical mass, it explodes as a SNe Ia. These events were famously used for the first detection of the accelerated expansion of the Universe (Riess et al., 1998; Perlmutter et al., 1999), and have remained one of the best probes of dark energy and the late Universe (e.g. Scolnic et al., 2018; Abbott et al., 2019). However, SNe Ia cannot be used to measure the Hubble constant directly because it is degenerate with their absolute magnitude.

[1] Following the convention in astronomy to refer to all elements besides hydrogen and helium as metals.

Therefore, they have to be combined with other probes in order to directly measure the speed of expansion.

A cosmic *distance ladder* is commonly used to calibrate SNe Ia in order to measure H_0. This generally uses another type of standardizable candles called cepheid variables, a type of pulsating stars (Leavitt and Pickering, 1912; Eddington, 1917). A small number of cepheids are close enough that we can measure their distance directly through parallax, or other geometric means. The period-luminosity relation of these cepheids is then calibrated and used to infer the distance to other nearby galaxies where we can observe cepheids and also happened to have had a SNe Ia explosion. This allows us to calibrate the absolute magnitude of SNe Ia and by using a large sample of these infer H_0 (see e.g. Freedman et al., 2001; Riess et al., 2016).

Over the last decade, an apparent *Hubble tension* has appeared primarily between H_0 measurements from the distance ladder (Riess et al., 2009; Riess et al., 2011; Riess et al., 2016; Riess et al., 2018; Riess et al., 2019), and those from the CMB (Planck Collaboration et al., 2014b; Planck Collaboration et al., 2016; Planck Collaboration et al., 2020c). The latest local measurement gives $H_0 = 73.2 \pm 1.3$ km s^{-1} Mpc^{-1}(Riess et al., 2021), while the best CMB constraint is $H_0 = 67.36 \pm 0.54$ km s^{-1} Mpc^{-1}(Planck Collaboration et al., 2020c), $\sim 4.2\sigma$ away. This tension is the topic of Chapter 5.

Jumping to the beginning of the Universe, another useful cosmological probe termed Big Bang Nucleosynthesis (BBN) uses known nuclear physics to model the relative abundances of different elements produced in the early Universe (see Cyburt et al., 2016 for a recent review). By measuring some of these abundances in the late Universe, we can infer the primordial baryon density, as well as the properties of neutrinos (see e.g. Cooke et al., 2016; Cooke, Pettini, and Steidel, 2018; Grohs et al., 2019).

The final probe we consider is *weak gravitational lensing* (WL). The gravitational lensing effect was predicted by general relativity and first discovered by Dyson, Eddington, and Davidson (1920). Weak lensing involves small changes in the shapes (shear) and sizes (magnification) of background galaxies due to intervening gravitational potentials. The shear power spectrum of background galaxies is the most commonly used way of detecting this effect. WL is a very powerful probe because it probes the gravitational potential directly, without having to deal with galaxy biasing (Weinberg et al., 2013). Over the last decade, a number of photometric surveys have been used to perform weak lensing analyses, with the most recent results coming from the Dark Energy Survey (DES; DES Collaboration et al., 2021), the Kilo-Degree Survey (KiDS; Heymans et al., 2021) and the Hyper Suprime-Cam Subaru Strategic Program (HSC; Hikage et al., 2019).

CHAPTER 4

INFERENCE

The precision era of cosmology was brought about by increasingly large surveys of the Universe, facilitated by an increase in computational power and technological advancement. However, with these large data sets came renewed interest into the statistical tools used to distil vast amounts of astronomical data into a few cosmological parameters. In Chapter 6 we will take an in-depth look at different statistical frameworks used in cosmology today. Here, I will give a brief overview of Bayesian inference, which provides a unified framework for quantifying belief and uncertainty. We will start in Section 4.1 with an introduction of this framework. After that, in Section 4.2 I will introduce the Nested Sampling algorithm, which is the most commonly used computational tool throughout this thesis.

4.1 THE BAYESIAN FRAMEWORK

The foundations of Bayesian statistics were set by Bayes (Bayes, 1763) and Laplace (Laplace, 1820). They built a framework for dealing with uncertainty by quantifying degrees of belief through the mathematical concept of probability. For a long time this framework was ignored in favour of the frequentist approach where probabilities are defined in terms of the frequency of an event in repeated experiments. This view is often perceived as being more objective compared to the Bayesian approach because of the common misconception that degrees of belief are inherently subjective (Sivia and Skilling, 2006). In fact, the power of Bayesian statistics comes from the fact that it provides a coherent framework for quantifying beliefs based on all available information.

Differences only appear when the background assumptions are not clearly stated. However, this is an issue for all types of inference, even if they are often not as obvious as in the Bayesian case.

Both frameworks are still in use today across different sciences, however, modern cosmology is largely performed using the Bayesian framework. This is firstly due to the inherently limited nature of experimentation in this field. We only have one Universe, and we do not have the luxury of running it again and again in order to measure frequencies of different events. Secondly, we have a limited perspective as observers, which often makes it very hard to disentangle different effects across the evolution of the Universe. This means we often have to build a coherent model for different complex phenomena and test it with limited data. An important point to note is that the use of the Bayesian framework in this field was only made possible by the recent advancements in computational power over the last few decades.

In order to introduce the Bayesian framework, we start with two simple rules, the *sum rule* and *product rule*. These were derived by Richard Cox (Cox, 1946) from the rules of logic and algebra, based on two simple assertions. The first is that belief in the probability of an event X implies belief in its opposite, that is, the probability of not observing X, denoted $\bar{X}$. The second is that knowing about X, and given that X is true knowing about a second event Y, implies knowledge of the joint probability of X and Y. The two rules can be expressed as:

$$P(X|\mathcal{M}) + P(\bar{X}|\mathcal{M}) = 1, \tag{4.1}$$

$$P(Y,X|\mathcal{M}) = P(Y|X,\mathcal{M})P(X|\mathcal{M}), \tag{4.2}$$

where we always condition probabilities on the assumed model, $\mathcal{M}$, which encompasses all relevant background information. Note that from the very start we refer to these probabilities as measures of belief or knowledge, rather than frequencies of events (see Chapter 6 for a more detailed comparison with the frequentist approach).

Using the sum and product rules, we can derive two important results in Bayesian statistics:

$$P(X|Y,\mathcal{M}) = \frac{P(Y|X,\mathcal{M})P(X|\mathcal{M})}{P(Y|\mathcal{M})}, \tag{4.3}$$

$$P(Y|\mathcal{M}) = \sum_{X_i} P(X_i,Y|\mathcal{M}), \tag{4.4}$$

where the sum is performed over the set of possible X_i outcomes. The first equation is known as Bayes' theorem, and the second is marginalization.

The usefulness of these equations becomes apparent if we identify Y with the observed experi-

mental data, D, and X with the parameters $\vec{\theta} = (\theta_1, ..., \theta_n)$ of the theoretical model we wish to test ($\mathcal{M}$). We can then express Bayes' theorem as:

$$P(\vec{\theta}|D,\mathcal{M}) = \frac{P(D|\vec{\theta},\mathcal{M})P(\vec{\theta}|\mathcal{M})}{P(D|\mathcal{M})}. \tag{4.5}$$

The probability distribution of the parameters given experimental data, $P(\vec{\theta}|D,\mathcal{M})$, is known as the *posterior* distribution (denoted $\mathcal{P}$), while the probability of observing a data set given some parameter values, $P(D|\vec{\theta},\mathcal{M})$, is known as the *likelihood* (denoted $\mathcal{L}$). The power of Bayes' theorem is that it allows us to compute the object of interest, the posterior distribution, from the object we have much better chances of being able to assign, the likelihood (Sivia and Skilling, 2006).

The *prior* probability of the parameters, $P(\vec{\theta}|\mathcal{M})$ (denoted π), is a common source of criticism of Bayesian statistics, because it encompasses the prior knowledge of the parameters values. Different choices of priors may lead to different posterior results, which makes it crucial to clearly state them in any inference. Jeffreys (1946) pioneered the use of so-called *uninformative priors*, which are designed to have minimal impact on the inference when we do not want to include any concrete prior knowledge.

The final term in Equation 4.5 is the marginal likelihood or the *evidence*, $P(D|\mathcal{M})$ (denoted $\mathcal{Z}$). This is obtained by marginalizing over the parameter values:

$$P(D|\mathcal{M}) = \int P(D,\vec{\theta}|\mathcal{M})d\vec{\theta} = \int P(D|\vec{\theta},\mathcal{M})P(\vec{\theta}|\mathcal{M})d\vec{\theta}, \tag{4.6}$$

which is the continuous version of Equation 4.4. This term is the normalization of the posterior and is often ignored when performing inferences due to it being very difficult to compute in general. However, if it can be computed, it is very useful in Bayesian model selection. We can derive this by plugging the evidence back into Bayes' theorem to compute the posterior probability of the model given the data, $P(\mathcal{M}|D)$. If we have two models, $\mathcal{M}_1$ and $\mathcal{M}_2$, we can compute how much more likely $\mathcal{M}_1$ is compared to $\mathcal{M}_2$ in light of experimental data, using the ratio of the two posterior probabilities:

$$\frac{P(\mathcal{M}_1|D)}{P(\mathcal{M}_2|D)} = \frac{P(D|\mathcal{M}_1)}{P(D|\mathcal{M}_2)}\frac{P(\mathcal{M}_1)}{P(\mathcal{M}_2)}. \tag{4.7}$$

If we then assign equal prior probabilities to the two models,[1] $P(\mathcal{M}_1) = P(\mathcal{M}_2)$, this is simply

[1] Note that we do not have to assign equal priors to the two models. This choice can be made on a case by case basis. However, the preference for equal priors comes from a desire to not interfere with the analysis through an arguably subjective choice.

given by the ratio of the two evidences. As we shall see in Chapter 6, the evidence also has other uses, however, we now turn our attention to the problem of computing the quantities presented here.

4.2 Computation

In cosmology, we generally have high dimensional parameter spaces, with ΛCDM having 6 free parameters, and many more being added in extensions or as nuisance parameters for different analyses. As most posterior distributions are analytically intractable due to the complexities of the model, we most often have to infer them by sampling over the multidimensional parameter-space. See Trotta (2008) and Hobson et al. (2010) for reviews of Bayesian methods used in cosmology.

Markov Chain Monte Carlo (MCMC) methods (e.g. Metropolis et al., 1953; Hastings, 1970) have proved very successful at sampling these spaces and inferring the shape of the posterior. Therefore, they are commonly used in cosmology (e.g. Lewis and Bridle, 2002; Audren et al., 2013). However, here we will discuss a second class of algorithms, called *Nested Samplers*, that can compute both the shape of the posterior and its normalization (the evidence).

Nested sampling was first introduced by Skilling (2004). Its main purpose is to compute the evidence, given by Equation 4.6, which for simplicity we re-label as:

$$\mathcal{Z} = \int \mathcal{L}(\vec{\theta})\pi(\vec{\theta})d\vec{\theta}. \tag{4.8}$$

This multidimensional integral quickly becomes infeasible to compute through brute force integration, and therefore a change of variables is used to cast it into a one-dimensional problem.

We define $X(\lambda)$ to be the fraction of the prior contained within the iso-likelihood contour $\mathcal{L}(\vec{\theta}) = \lambda$. This can be expressed as:

$$X(\lambda) = \int_{\mathcal{L}(\vec{\theta})>\lambda} \pi(\vec{\theta})d\vec{\theta}, \tag{4.9}$$

where the element of prior mass is $dX = \pi(\vec{\theta})d\vec{\theta}$. This is a strictly decreasing function of the likelihood, with extremes at $X = 1$ when the iso-likelihood contour encompasses the entire prior ($\lambda \geq 0$), and at $X = 0$ when $\lambda = \mathcal{L}_{\max}$. See Figure 4.1 for a visualization. Casting the likelihood as a function of X, we have $\mathcal{L}(X(\lambda)) \equiv \lambda$, which means we can write the evidence as:

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X)dX. \tag{4.10}$$

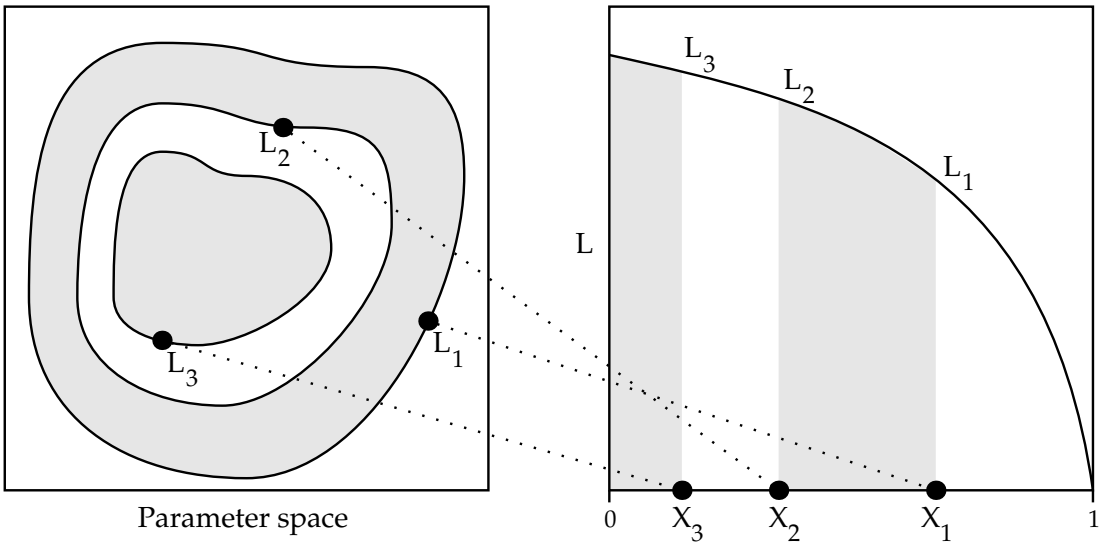

Figure 4.1. Visualization of the change of variables performed in Nested Sampling. The N-dimensional parameter space is transformed into a one dimensional space, where the points are ordered by the value of their likelihood. The new parameter X (x-axis in the right-hand plot) corresponds to the prior volume encased by the iso-likelihood contour of the corresponding $\mathcal{L}$ in the plot on the left. Figure from Sivia and Skilling (2006).

This is a one dimensional integral that can be computed numerically.

The sampling algorithm starts by generating N random samples, called *live points*, from the full prior distribution. New points are then randomly sampled from the prior under an evolving constraint $\mathcal{L}(\vec{\theta}) > \mathcal{L}^*$. Initially, $\mathcal{L}^* = 0$ corresponding to $X^* = 1$, and therefore the points are generated under the constraint $X < X^*$. At each step of the iteration, the point with the worst likelihood $\mathcal{L}$, corresponding to the largest volume X, is discarded and $(X, \mathcal{L})$ becomes the new constraint $(X^*, \mathcal{L}^*)$. A new point is then generated from the now smaller volume X^*, hence each new domain is nested within the old.

If the new random samples are uniformly generated from the prior volume X^*, at iterate k, the volume X_k will have mean and standard deviation:

$$\log X_k = (-k \pm \sqrt{k})/n. \tag{4.11}$$

Therefore, at each step, the volume shrinks approximately by a factor $\exp(1/N)$. The evidence integral is computed by summing the thin area strips, $A = w\mathcal{L}$, of each point. The simplest way to

define the width is $w_k = X_{k-1} - X_k$[2]. Therefore, the evidence is given by:

$$\mathcal{Z} \approx \sum_k A_k = \sum_k w_k \mathcal{L}_k, \tag{4.12}$$

and the iteration is usually stopped when most of $\mathcal{Z}$ has been found. This is based on the fact that after we find the region that contains most of the posterior volume, each new element of the $\mathcal{Z}$ sum will be an increasingly smaller fraction of the total (see Sivia and Skilling, 2006 for more details).

Finally, we note that the set of samples produced by this method also gives us the shape of the posterior distribution, by weighting each sample as:

$$p_k = \frac{w_k \mathcal{L}_k}{\mathcal{Z}}. \tag{4.13}$$

Therefore, nested sampling is an efficient algorithm both for computing the evidence and for performing parameter estimation.

The last detail needed is a method to uniformly generate new samples from the prior under the constraint $\mathcal{L}(\vec{\theta}) > \mathcal{L}^*$. This is a non-trivial problem, especially in high-dimensional spaces. The solution to this problem is what differentiates the many implementations of nested sampling algorithms available in the literature.

Originally, Skilling (2004) proposed a simple MCMC method using the Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970). Feroz and Hobson (2008) introduced the popular *Multinest* package (subsequently improved in Feroz, Hobson, and Bridges, 2009 and Feroz et al., 2019), which uses D-dimensional ellipsoids to define the iso-likelihood contours. The package used throughout this thesis, called *Polychord* (Handley, Hobson, and Lasenby, 2015a; Handley, Hobson, and Lasenby, 2015b), uses slice sampling to generate new points. This method takes each dimension separately and generates a new sample from a proposal distribution across a slice of the likelihood. More complex methods have also been recently proposed, for example using a varying number of live points (Higson et al., 2019) or deep learning (Moss, 2020).

[2] In practice, this is often computed using the trapezoid rule.

PART II

PRECISION COSMOLOGY FROM THE CLUSTERING OF LARGE-SCALE STRUCTURES

CHAPTER 5

BAO AND THE HUBBLE CONSTANT: PAST, PRESENT AND FUTURE

We investigate constraints on the Hubble constant (H_0) using Baryon Acoustic Oscillations (BAO) and baryon density measurements from Big Bang Nucleosynthesis (BBN). We start by investigating the tension between galaxy BAO measurements and those using the Lyman-α forest, within a Bayesian framework. Using the latest results from eBOSS DR14, we find that the probability of this tension being statistical is $\simeq 6.3\%$ assuming flat ΛCDM. We measure $H_0 = 67.6 \pm 1.1$ km s^{-1} Mpc^{-1}, with a weak dependence on the BBN prior used, in agreement with results from *Planck* Cosmic Microwave Background (CMB) results and in strong tension with distance ladder results. Finally, we forecast the future of BAO + BBN measurements of H_0, using the Dark Energy Spectroscopic Instrument (DESI). We find that the choice of BBN prior will have a significant impact when considering future BAO measurements from DESI.

The work shown in this chapter is presented in Cuceu et al. (2019). The BBN computation mentioned in Section 5.3.2, which was used to turn our measured baryon density into a constraint on the deuterium abundance, was performed by collaborators. The rest of the work presented here was led by myself.

5.1 Introduction

Over the last twenty years, a clear picture of the Universe has started to emerge, with Lambda Cold Dark Matter (ΛCDM) becoming the standard cosmological model. However, with the improved precision of the latest surveys, tensions between different measurements of some parameters have also started to appear. Perhaps none have been debated more than the discrepant values of the Hubble constant, H_0, that measures the expansion rate of the Universe. The cosmic distance ladder has long been used to directly measure H_0 (e.g. Riess et al., 2009; Riess et al., 2011; Riess et al., 2016; Riess et al., 2018; Riess et al., 2019), and the latest value from the Supernova, H_0, for the Equation of State of Dark energy (SH0ES) program is $H_0 = 74.03 \pm 1.42$ km s^{-1} Mpc^{-1}(Riess et al., 2019). On the other hand, indirect constraints using Cosmic Microwave Background (CMB) anisotropy measurements from the *Planck* satellite (Astropy Collaboration et al., 2013; Planck Collaboration et al., 2016; Planck Collaboration et al., 2018) give a significantly different value: $H_0 = 67.36 \pm 0.54$ km s^{-1} Mpc^{-1}(Planck Collaboration et al., 2018) (assuming ΛCDM).

Possible explanations for this tension are systematic errors in one or both datasets, or problems with the standard model and the need for new physics. Reanalyses of the distance ladder data (e.g. Cardona, Kunz, and Pettorino, 2017; Zhang et al., 2017; Feeney, Mortlock, and Dalmasso, 2018; Follin and Knox, 2018) still prefer high values of H_0, while using most subsets of the *Planck* data yields lower values (e.g. Planck Collaboration et al., 2018; Bernal, Verde, and Riess, 2016). The 4.4σ difference between the two H_0 measurements is also hard to reconcile with extensions to the standard ΛCDM model. A promising prospect is a higher value of the effective number of neutrinos, $N_{\rm eff}$. However, the tension is only slightly reduced ($\sim 3.9\sigma$), as CMB constraints rule out very high values for this parameter (Planck Collaboration et al., 2018).

Baryon Acoustic Oscillations (BAO) provide a standard ruler, which has been evolving with the Universe since recombination. As such, probing the BAO scale at different times is a powerful tool in constraining cosmology. The best measurements of the BAO scale come from CMB anisotropy measurements at redshift $z \approx 1100$ (e.g. Planck Collaboration et al., 2018). BAO are also present in the distribution of matter, and there are measurements at low redshifts using the clustering of galaxies (e.g. Beutler et al., 2011; Ross et al., 2015; Alam et al., 2017). It has also been detected in the correlation function of the Lyα forest at $z \sim 2.4$ and in its cross-correlation with quasar positions (e.g. Delubac et al., 2015; Font-Ribera et al., 2014; Bautista et al., 2017; du Mas des Bourboux et al., 2017; Blomqvist et al., 2019; de Sainte Agathe et al., 2019).

BAO data can only constrain a combination of the size of the sound horizon and the expansion

rate of the Universe (H_0). Therefore, a constraint on H_0 requires extra data to calibrate the size of the sound horizon; usually CMB anisotropy measurements are used. Recently, Addison et al. (2018) used an alternative method, introduced by Addison, Hinshaw, and Halpern (2013), that uses deuterium abundance measurements and the Big Bang Nucleosynthesis (BBN) theory. This BAO + BBN method assumes standard pre-recombination physics and gives a value of H_0 consistent with the *Planck* value using a flat ΛCDM model. Addison et al. (2018) emphasized the importance of this method in providing a constraint on H_0 independent of CMB anisotropy measurements and the distance ladder. The focus of this work is to discuss past results of this data combination, compute the latest constraints, and investigate future implications.

The BAO measurements used by Addison et al. (2018) come from galaxy clustering analyses (Beutler et al., 2011; Ross et al., 2015; Alam et al., 2017), and the Lyα forest (Delubac et al., 2015; Font-Ribera et al., 2014). Questions arise, however, when considering the $\sim 2.5\sigma$ tension between Galaxy BAO and Lyα BAO in the 11th and 12th data release of the Sloan Digital Sky Survey (SDSS DR11 and DR12, e.g. du Mas des Bourboux et al., 2017; Bautista et al., 2017; Aubourg et al., 2015). The question of consistency between datasets, especially when it comes to combining them, has long been debated (e.g. Inman and Jr, 1989; Charnock, Battye, and Moss, 2017; Nicola, Amara, and Refregier, 2019; Adhikari and Huterer, 2019; Raveri and Hu, 2019). Recently, a new method was proposed by Handley and Lemos (2019b) to quantify tension using a new statistics they call suspiciousness. As such, in Section 5.2 we use this method to investigate the tension between Galaxy BAO and Lyα BAO for the purpose of testing the reliability of their combined results.

In Section 5.3 we update the constraint from BAO + BBN using the latest BAO and BBN results. Compared to Addison et al. (2018), we add the latest BAO measurements from the Extended Baryon Oscillation Spectroscopic Survey (eBOSS) using QSO clustering and the Lyα forest (Ata et al., 2018; de Sainte Agathe et al., 2019; Blomqvist et al., 2019). We also use the latest primordial deuterium abundance results (Cooke, Pettini, and Steidel, 2018). In Section 5.4, we forecast future BAO + BBN measurements of H_0 using the Dark Energy Spectroscopic Instrument (DESI), and discuss the role of BBN priors on future results.

5.2 Galaxy BAO vs Lyα Forest BAO

When combining different BAO measurements, Addison et al. (2018) split the data in two types: Galaxy BAO and Lyα BAO, that includes both the Lyα auto-correlation and its cross-correlation with quasars. BOSS DR11 Lyα BAO measurements were in $\approx 2.3\sigma$ tension with CMB predictions

BAO Measurement	Dataset	Reference	Tracer	$z_{\rm eff}$
6dFGS	6dFGS	Beutler et al. (2011)	galaxies	0.106
SDSS MGS	SDSS DR7	Ross et al. (2015)	galaxies	0.15
BOSS Gal	SDSS DR12	Alam et al. (2017)	galaxies	0.38, 0.51, 0.61
eBOSS QSO	SDSS DR14	Ata et al. (2018)	QSO	1.52
eBOSS Ly$\alpha\times$Lyα	SDSS DR14	de Sainte Agathe et al. (2019)	Ly$\alpha\times$Lyα	2.34
eBOSS Ly$\alpha\times$QSO	SDSS DR14	Blomqvist et al. (2019)	Ly$\alpha\times$QSO	2.35

Table 5.1. Datasets measuring the BAO peak that are used in our Hubble constant analysis. We have also used other past results such as Lyα DR11 and DR12 for our tension analysis. We assume Gaussian likelihoods for the galaxy BAO measurements, but we use the full χ^2 tables provided by the Lyα forest analyses (see Section 5.6).

from the *Planck* Collaboration (Font-Ribera et al., 2014; du Mas des Bourboux et al., 2017), while the samples that go into Galaxy BAO were all consistent with CMB predictions. This translated into a tension between Lyα BAO and Galaxy BAO that can clearly be seen in the bottom panel of Figure 5.1 (red dashed contours).

Recently, the eBOSS collaboration published the latest Lyα BAO measurements using DR14 data (de Sainte Agathe et al., 2019; Blomqvist et al., 2019). They use $\sim 15\%$ more quasar spectra than the previous DR12 results, and, for the first time, Lyα absorbers in the Lyβ region are used. With these new measurements, the tension with CMB predictions has gone down to $\sim 1.7\sigma$. In this section, we discuss the internal tensions of the latest BAO results, listed in Table 5.1.

5.2.1 BAO COSMOLOGY

Studies of the BAO feature in the transverse direction provide a measurement of $D_M(z)/r_d$, while BAO studies along the line of sight measure the combination $D_H(z)/r_d = c/H(z)r_d$, where D_M is the comoving angular diameter distance, c is the speed of light in vacuum, z is the redshift and $r_d \equiv r_s(z_d)$ is the size of the sound horizon at the drag epoch (z_d).

In a flat ΛCDM cosmology, D_M is given by:

$$D_M(z) = c\int_0^z \frac{dz'}{H(z')}. \tag{5.1}$$

Some of the datasets we include (6dFGS, SDSS MGS and eBOSS QSO) measure $D_V(z)/r_d$, which is a combination of the BAO peak coordinates above. $D_V(z)$ is defined as:

$$D_V(z) \equiv [zD_H(z)D_M^2(z)]^{1/3}. \tag{5.2}$$

The Friedman equation in flat ΛCDM completes our model:

$$\frac{H(z)^2}{H_0^2} = \Omega_r(1+z)^4 + \Omega_m(1+z)^3 + \Omega_\Lambda, \tag{5.3}$$

where Ω_r, Ω_m and Ω_Λ are the fractional densities of radiation, matter and dark energy today (at redshift $z = 0$). Furthermore, in flat ΛCDM, the dark energy fraction can be computed as: $\Omega_\Lambda = 1 - \Omega_m - \Omega_r$. In the late universe, at the redshifts probed by BAO, the radiation fraction is very small. Nevertheless, we model it assuming a fixed neutrino sector with $N_{\rm eff} = 3.046$ and 2 massless species (the third one is massive with $m_\nu = 0.06$ eV and contributes to Ω_m), and a CMB temperature of $T_{CMB} = 2.7255$K. This has been measured by COBE/FIRAS (Fixsen et al., 1996; Fixsen, 2009), and we consider this measurement independent of *Planck*. Therefore, the only free parameters in $H(z)$ are H_0 and Ω_m.

As previously mentioned, when we measure BAO we are measuring a combination of H_0 and r_d, which means the two parameters are fully degenerate. As such, we sample their product: $H_0 r_d$. We will discuss ways to break this degeneracy in the next section, but for the purpose of investigating possible internal tensions in BAO measurements we will work in the $\Omega_m - H_0 r_d$ plane.

5.2.2 Quantifying tension

The aim of this section is to quantify the tension between the different Lyα BAO measurements and Galaxy BAO measurements. This tension is clear when looking at the posteriors (see bottom panel of Figure 5.1), but quantifying it is a non-trivial problem, due to the non-Gaussianity of the posteriors. There is a large number of available approaches in the literature to quantify tension between datasets (e.g. Inman and Jr, 1989; Charnock, Battye, and Moss, 2017; Nicola, Amara, and Refregier, 2019; Adhikari and Huterer, 2019; Raveri and Hu, 2019). One of the most widely used methods is the evidence ratio R (Marshall, Rajguru, and Slosar, 2006; Trotta, 2008; Verde, Protopapas, and Jimenez, 2013):

$$R \equiv \frac{\mathcal{Z}_{AB}}{\mathcal{Z}_A \mathcal{Z}_B}, \tag{5.4}$$

where $\mathcal{Z}$ are evidences, A and B denote the two datasets on their own, and AB denotes the joint results. The Bayesian evidence (the probability of the data D given a model M: $P(D|M)$) is the normalization term in Bayes' theorem, and is usually ignored if one is only interested in the shape of the posterior. However, it has useful applications, e.g. in Bayesian Model Selection (e.g. Sivia and Skilling, 2006), and as mentioned in quantifying concordance between datasets.

As highlighted in Handley and Lemos (2019b), the R-statistic can hide tension when the priors

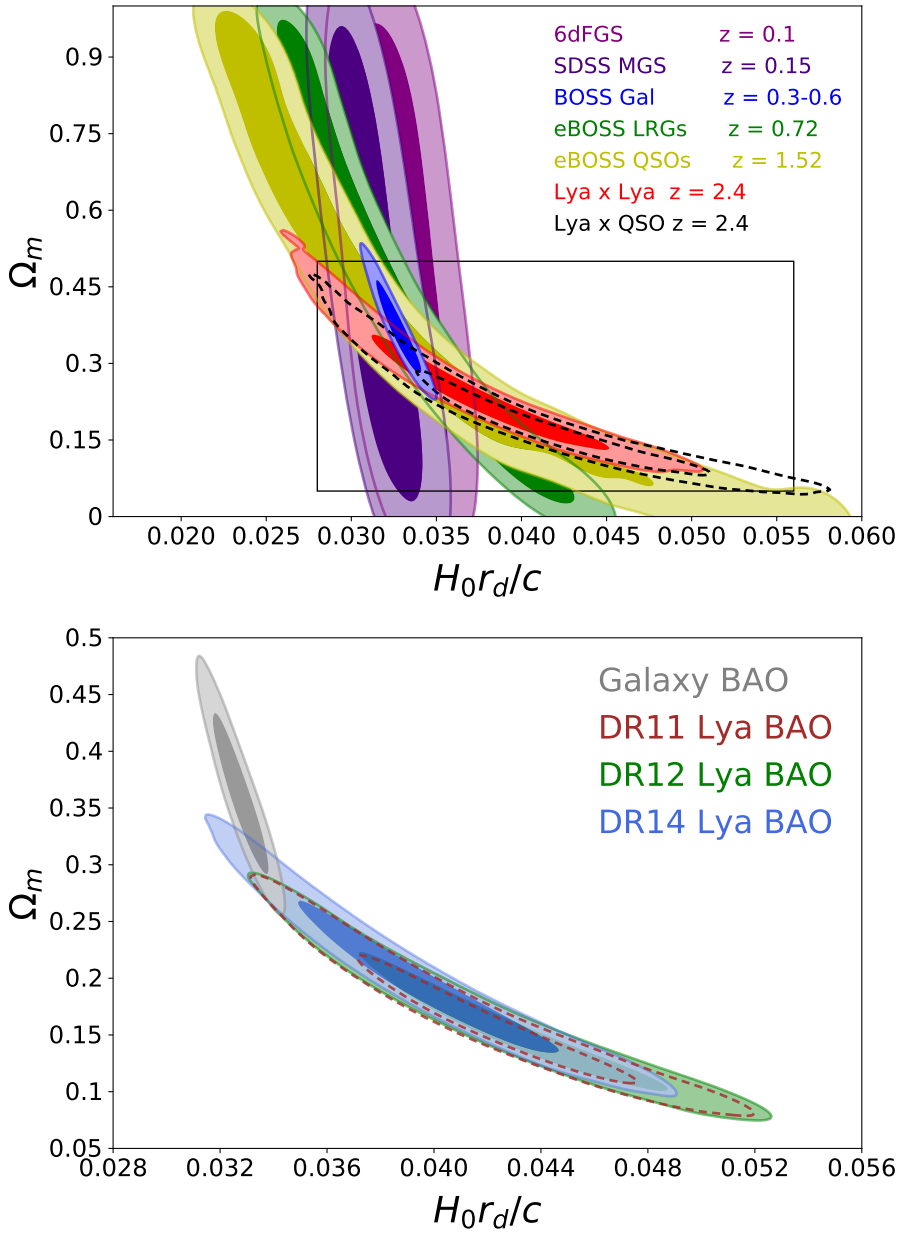

Figure 5.1. (Top) Parameter constraints in a flat ΛCDM cosmology from each BAO dataset individually. The different contour orientations are due to the different redshifts of separate datasets. The box represents the boundaries of the plot on the right with the combined BAO measurements. (Bottom) Comparison of BAO constraints from galaxy clustering and different Lyα forest measurements. The recently released eBOSS DR14 Lyα BAO measurements are visibly more consistent with galaxy BAO than previous results from DR11 and DR12. This is quantified in Table 5.2.

are arbitrarily chosen, since it is proportional to the prior volume shared by both datasets. In this work, we will use the method introduced in Handley and Lemos (2019b): We calculate the 'suspiciousness' S as the ratio between the evidence ratio R, and the information ratio I: $S \equiv R/I$. The information ratio is defined as:

$$\log I \equiv \mathcal{D}_A + \mathcal{D}_B - \mathcal{D}_{AB}, \tag{5.5}$$

where $\mathcal{D}$ is the Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951):

$$\mathcal{D} \equiv \int \mathrm{d}\theta \; \mathcal{P}(\theta) \log \frac{\mathcal{P}(\theta)}{\pi(\theta)}, \tag{5.6}$$

with $\mathcal{P}$ the posterior, π the prior, and θ the parameters. The KL divergence is also called the relative entropy because it quantifies the information difference between two distributions (in this case, the information gain between the prior and the posterior).

The suspiciousness S can be seen as an evidence ratio R from which the dependence on prior volume has been subtracted, in the form of the information ratio I. Therefore, it preserves the qualities that make R a desirable statistic for dataset comparison (such as its Bayesian interpretation and its independence in the choice of parameters), but it is no longer proportional to the prior volume, and therefore it does not hide tension when wider priors are chosen.

As described in Handley and Lemos (2019b), the suspiciousness can be calibrated using the fact that, for Gaussian posteriors, it follows a χ^2_d distribution, where d is the number of parameters simultaneously constrained by the combination of the datasets. From this distribution, a tension probability p of two datasets being discordant by chance can be assigned as the 'p-value' of the distribution. The remaining problem is the calculation of the number of dimensions simultaneously constrained by both datasets. This is done using the Bayesian model dimensionality (BMD) introduced in Handley and Lemos (2019a). It is worth mentioning that the BMD can be smaller or larger than the number of constrained parameters in our model if the posterior is significantly non-Gaussian. While it is clear by looking at the bottom panel of Figure 5.1 that the posteriors are non-Gaussian in the case of Lyα BAO, this method will give us an estimate of the tension between the datasets.[1] We use `Polychord` (Handley, Hobson, and Lasenby, 2015a; Handley, Hobson, and Lasenby, 2015b) to sample our posteriors and compute evidences.

We use the three Lyα BAO measurements published by the BOSS and eBOSS collaborations

[1] In addition, as discussed in Handley and Lemos (2019b), these posteriors can be 'Gaussianised' using Box-Cox transformations (Schuhmann, Joachimi, and Peiris, 2016), which preserve the value of $\log S$.

datasets	$\log R$	$\log I$	$\log S$	d	$p(\%)$	σ
Gal - DR11 Lyα	0.35 ± 0.19	4.04 ± 0.18	-3.68 ± 0.05	2.43 ± 0.15	1.20 ± 0.15	$\simeq 2.5$
Gal - DR12 Lyα	0.26 ± 0.19	3.79 ± 0.18	-3.53 ± 0.05	2.34 ± 0.15	1.31 ± 0.16	$\simeq 2.5$
Gal - DR14 Lyα	1.93 ± 0.19	3.78 ± 0.19	-1.85 ± 0.05	2.19 ± 0.14	6.30 ± 0.61	$\simeq 1.9$

Table 5.2. Tension statistics for combining Galaxy BAO and different Lyα BAO measurements. We show results for the R-statistic, the Bayesian information and the 'suspiciousness'. The Bayesian model dimensionality (d) introduced by Handley and Lemos (2019a) is used to compute a p-value for the suspiciousness, and we use this to compute the approximate number of standard deviations for this tension. The older DR11 and DR12 Lyα results give small p-values, indicating a small probability of this tension being statistical in nature. On the other hand, the recent DR14 results show better agreement with the Galaxy BAO results.

using SDSS data releases 11, 12 and 14. We compare each of these with the combined Galaxy BAO sample within a flat ΛCDM cosmology, and present the tension statistics in Table 5.2. We compute probability values of $\simeq 1.2\%$ and $\simeq 1.3\%$ for the consistency between the Galaxy BAO sample and the DR11 and DR12 Lyα results respectively, indicating that there is a very small probability that this tension appears purely by chance. On the other hand, using the latest DR14 results we compute $p \simeq 6.3\%$, consistent with the tension being statistical in nature.

5.3 BAO and the Hubble Constant

BAO data must be combined with other measurements in order to break the $H_0 - r_d$ degeneracy and obtain a constraint on H_0. The sound horizon at the drag epoch is given by:

$$r_d = \int_{z_d}^{\infty} \frac{c_s(z)}{H(z)} dz, \tag{5.7}$$

where $c_s(z) = c[3 + \frac{9}{4}\rho_b(z)/\rho_\gamma(z)]^{-1/2}$ is the speed of sound in the baryon-photon fluid (Aubourg et al., 2015), $\rho_b(z), \rho_\gamma(z)$ are the baryon and photon densities respectively, and z_d is the redshift of the drag epoch. Precise computations of r_d require a full Boltzmann code, however, following Aubourg et al. (2015), we use a numerically calibrated approximation to avoid the additional computational cost:

$$r_d \approx \frac{55.154 \exp[-72.3(\omega_\nu + 0.0006)^2]}{\omega_m^{0.25351}\omega_b^{0.12807}} \text{ Mpc}, \tag{5.8}$$

where $\omega_X = \Omega_X h^2$, and $X = m, \nu, b$ are matter, neutrinos and baryons respectively, and $h = H_0/100$ with H_0 in [km s^{-1} Mpc^{-1}]. This approximation is accurate to 0.021% (Aubourg et al., 2015) for a fixed neutrino sector with $N_{\rm eff} = 3.046$ and $\sum m_\nu < 0.6$ eV. Our main results are also benchmarked against independent runs using CosmoMC (Lewis and Bridle, 2002), which uses the Boltzmann

solver CAMB (Lewis, Challinor, and Lasenby, 2000), to validate the approximation.

BAO measurements provide a good constraint on Ω_m, and, as discussed, the neutrino sector is fixed to the minimal mass[2]. Therefore, to compute r_d, only a measurement of the baryon density, $\Omega_b h^2$, is still needed. *Planck* results currently provide the best constraints on $\Omega_b h^2$, however, the goal of this work is to constrain H_0 without using CMB anisotropy information. As such, we instead use primordial deuterium abundance measurements and BBN to put a constraint on the baryon density.

5.3.1 BBN MEASUREMENTS

Deuterium is one of the most widely used primordial elements for constraining cosmology because of its strong dependence on the baryon density (Cyburt et al., 2016). Most of the primordial deuterium was processed into helium during BBN, however, trace amounts remain because the reaction is not completely efficient (Dodelson, 2003). If there are more baryons, this reaction is more efficient, and more deuterium gets processed into helium before freeze-out. Therefore, the deuterium abundance depends strongly on the baryon density. An upper bound can easily be placed on the primordial deuterium abundance because there are no known astrophysical sources that can produce significant quantities of deuterium (Epstein, Lattimer, and Schramm, 1976; Prodanović and Fields, 2003). Deuterium can, however, be destroyed, and as such a lower bound on the abundance requires finding pristine systems with the lowest possible metallicities. These systems have undergone only modest chemical evolution, so they provide the best available environments for measuring the primordial deuterium abundance (see Cyburt et al., 2016 for a review). Recently, Cooke, Pettini, and Steidel (2018) reported a one percent measurement of the primordial deuterium abundance using 7 near-pristine damped Lyα systems (DLAs). However, the sample size should be greatly improved upon with the next generation of 30m telescopes (Grohs et al., 2019).

To obtain a constraint on $\Omega_b h^2$, the deuterium abundance must first be converted to the baryon to photon ratio, η (Cooke et al., 2016). The required calculations (Cooke et al., 2016) need precise measurements of the cross-sections of reactions happening in BBN (see Adelberger et al., 2011 for a review of measurements of these reaction rates). The radiative capture of protons on deuterium to produce ^{3}He: $d(p,\gamma)^3$He, is one reaction whose cross-section is proving difficult to determine in the energy range relevant to BBN. Current laboratory measurements have an uncertainty of $\gtrsim 7\%$, and as such theoretical estimates are mostly used as they provide about $\sim 1\%$ precision

[2]Small deviations from the minimal neutrino mass, within the range allowed by current CMB constraints, would not have a large impact on our results.

Datasets	$\Omega_b h^2$ prior	Ω_m	r_d [Mpc]	H_0 [km s^{-1} Mpc^{-1}]
DR14 BAO + BBN	theoretical	$0.302^{+0.017}_{-0.020}$	149.0 ± 3.2	67.6 ± 1.1
DR14 BAO + BBN	empirical	0.300 ± 0.018	148.0 ± 3.1	68.1 ± 1.1
DR12 BAO + BBN	empirical	0.290 ± 0.018	150.0 ± 3.5	67.5 ± 1.2
DR11 BAO + BBN	empirical	$0.289^{+0.016}_{-0.021}$	$150.3^{+3.7}_{-3.3}$	67.4 ± 1.2
Planck 2018	-	0.3153 ± 0.007	147.09 ± 0.26	67.4 ± 0.5
SH0ES	-	-	-	74.0 ± 1.4

Table 5.3. Latest DR14 BAO + BBN constraints using either theoretical or empirical $d(p,\gamma)^3$He reaction rate. We add results using the Lyα DR11 and DR12 measurements to show the consistency in H_0 results. Results from the *Planck* Collaboration (Planck Collaboration et al., 2018) and the SH0ES collaboration (Riess et al., 2019) are included for comparison.

(Cooke et al., 2016). We will use both theoretical and empirical results and compare them. The best theoretical estimates of the $d(p,\gamma)^3$He reaction rate come from Marcucci et al. (2016), and lead Cooke, Pettini, and Steidel (2018) to compute:

$$100\Omega_b h^2 = 2.166 \pm 0.015 \pm 0.011 \qquad \text{(BBN theoretical)}, \tag{5.9}$$

where the first error comes from the deuterium abundance measurement, and the second from the BBN calculations. Using the empirical value for the reaction rate computed by Adelberger et al. (2011), the baryon density is:

$$100\Omega_b h^2 = 2.235 \pm 0.016 \pm 0.033 \qquad \text{(BBN empirical)}. \tag{5.10}$$

These two results are in mild $\sim 1.7\sigma$ tension with each other, but more importantly, the first measurement (using the theoretical rate) is in $\sim 2.9\sigma$ tension with the latest CMB results from the *Planck* Collaboration[3]:

$$100\Omega_b h^2 = 2.237 \pm 0.015 \qquad (\textit{Planck}\ 2018[4]). \tag{5.11}$$

There are some prospects for solving this tension by allowing the effective number of neutrinos $N_{\rm eff}$ to vary (see Figure 7 of Cooke, Pettini, and Steidel, 2018). A slightly larger value of $N_{\rm eff}$ would reconcile BBN and CMB measurements of $\Omega_b h^2$ (Cooke, Pettini, and Steidel, 2018). However, for the purposes of the present work, we use both values $\Omega_b h^2$ from BBN with the standard $N_{\rm eff} = 3.046$ in order to study the impact of this tension on H_0 measurements.

[3] We use the results from *Planck* 2018 TT,TE,EE + lowE + lensing likelihoods

5.3.2 Results

We combine the BAO data presented in Section 5.2 with the two $\Omega_b h^2$ measurements from BBN deuterium abundance. Using equation 5.8, we compute the size of the sound horizon at the drag epoch r_d and obtain constraints on H_0. The left panel of Figure 5.2 shows results using Lyα BAO + BBN and Gal BAO + BBN, as well as their combination. Individually they are both consistent with higher values of H_0 (latest SH0ES results are also plotted), however once we combine Lyα and Gal BAO, the joint constraint prefers lower, *Planck*-like values of the Hubble constant.

Our results using both the theoretical and empirical $d(p,\gamma)^3$He reaction rates are shown in Table 5.3 and in the right panel of Figure 5.2, together with *Planck* 2018 CMB results (Planck Collaboration et al., 2018) and the SH0ES H_0 measurement from the distance ladder (Riess et al., 2019) for comparison. We also add results using past Lyα measurements (DR11 and DR12) to show the consistency in H_0 constraints. Both our H_0 measurements are consistent with the results of the *Planck* Collaboration. On the other hand, we find that our Hubble constant measurements are in strong tension with local distance ladder results of H_0 from the SH0ES Collaboration. Our results are in approximately $\sim 3.6\sigma$ tension using the theoretical $d(p,\gamma)^3$He reaction rate, and $\sim 3.3\sigma$ tension using the empirical $d(p,\gamma)^3$He reaction rate.

An interesting result can be obtained by reframing this tension in terms of primordial deuterium abundance. If we assume the H_0 constraint from SH0ES (Riess et al., 2019) is true, and we combine it with BAO data, we obtain a constraint on the baryon density of $\Omega_b h^2 = 0.0310 \pm 0.003$. Using BBN (Cooke et al., 2016), we obtain a value for the primordial deuterium abundance of $10^5 (D/H)_P = 1.38 \pm 0.25$ (this assumes ΛCDM and standard BBN). This value is $\sim 4.5\sigma$ below that measured by Cooke, Pettini, and Steidel (2018), and well below the value derived from the interstellar medium of the Milky Way (Linsky et al., 2006). As we discussed, there are currently no known astrophysical sources that can produce significant quantities of deuterium (Epstein, Lattimer, and Schramm, 1976; Prodanović and Fields, 2003). This means D/H measurements have a robust lower limit, which renders such a low value of the primordial deuterium abundance very unlikely.

We find that the relatively large difference between the two $\Omega_b h^2$ measurements from BBN has a small impact on the Hubble constant measurement from current BAO measurements, causing a shift on the best fit value of H_0 of about $\sim 0.5\sigma$. However, with improving BAO data from the next generation of LSS experiments such as DESI (DESI Collaboration et al., 2016) or Euclid (Laureijs et al., 2011), this might change. In the next section, we investigate the advances that DESI data will allow in measuring the Hubble constant independent of CMB data, and the potential impact of

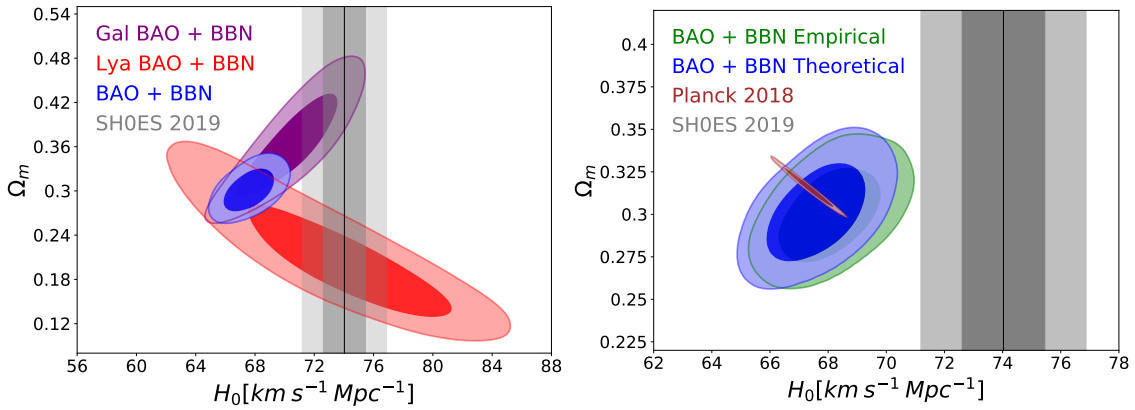

Figure 5.2. (Left) Current state-of-the-art results for H_0 versus Ω_m, independent of CMB anisotropy data. BAO data was combined with a prior on $\Omega_b h^2$ from BBN deuterium measurements (using the theoretical reaction rate). (Right) Our main results using all the BAO samples in Table 5.1, combined with BBN using both reaction rates.

BBN tensions on future results.

5.4 Implications for DESI

The next generation of LSS experiments will be spearheaded by the Dark Energy Spectroscopic Instrument (DESI), starting in 2020. This spectroscopic galaxy survey will cover 14000 square degrees, and measure BAO using both galaxy clustering and the Lyα forest (DESI Collaboration et al., 2016). It will target Luminous Red Galaxies (LRGs) at redshifts $0.4 < z < 1.0$, Emission Line Galaxies (ELGs) at redshifts $0.6 < z < 1.6$, quasars at redshifts $0.6 < z < 2.1$ for clustering only, and quasars at redshifts $2.1 < z < 3.5$ for both clustering and Lyα forest measurements (DESI Collaboration et al., 2016). DESI will also target bright galaxies at redshifts $0 < z < 0.5$ in order to take advantage of the times when moonlight prevents efficient observation of faint targets. This wide redshift coverage means that DESI will be able to precisely constrain the evolution of the Universe up to redshift ~ 3.5. Forecasts for future H_0 constraints from DESI combined with baryon density measurements from the CMB were presented in Wang, Xu, and Zhao (2017). Our objective in this section is to forecast future DESI BAO + BBN constraints on the Hubble constant, and to discuss the role of the discrepant values of the $d(p,\gamma)^3$He reaction rate.

In order to study the impact of BBN tensions on future BAO + BBN measurements of the Hubble constant, we perform a forecast of the future DESI results using the uncertainties presented by DESI Collaboration et al. (2016). For our fiducial cosmology, we use the BAO + BBN empirical results from Section 5.3. We plot results using different components of DESI as well as the

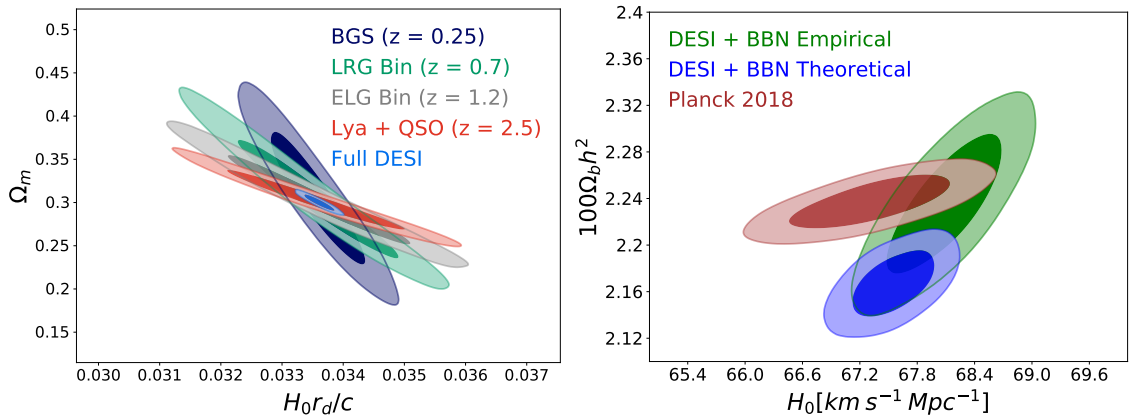

Figure 5.3. (Left) Forecast for future BAO results within flat ΛCDM using different components of DESI. (Right) Forecast for Hubble constant results using the full DESI results combined with the two $\Omega_b h^2$ priors from BBN, and the *Planck* 2018 results (Planck Collaboration et al., 2018) for comparison. The tension in the baryon density between the BBN theoretical constraint (in blue) and the CMB (in red) can clearly be seen in this plot. This shows the importance of solving the BBN tension for the future of BAO + BBN H_0 measurements.

combined results in the left panel of Figure 5.3. For illustration purposes, we only plot one LRG bin at $z = 0.7$ and one ELG bin at $z = 1.2$. With the big improvement in BAO measurements at each redshift, DESI also has the potential to give rise to inconsistent results. If this happens, it will provide a big opportunity to discover unaccounted systematic errors, unforeseen problems with our methods or potentially new physics.

Finally, we turn our attention to the main goal of this section: to quantify the impact of discrepant BBN measurements on future Hubble constant results from BAO + BBN. We perform the forecast described above for DESI and combine it with $\Omega_b h^2$ measurements from BBN, using both the theoretical and empirical $d(p,\gamma)^3$He reaction rates to measure H_0. We plot the results in the right panel of Figure 5.3. The two H_0 constraints are $\sim 1.2\sigma$ apart. This means that solving the BBN $\Omega_b h^2$ discrepancy will play an important role in next generation measurements of H_0 using BAO + BBN. There is hope of better laboratory measurements of the $d(p,\gamma)^3$He reaction rate from the Laboratory for Underground Nuclear Astrophysics (LUNA Gustavino, 2014; Kochanek, 2016).

5.5 Conclusions

We use the suspiciousness statistic proposed by Handley and Lemos (2019b) to investigate the tension between galaxy BAO and the different Lyα BAO measurements. When using the DR11 and DR12 Lyα results, we find probabilities of $\simeq 1.2\%$ and $\simeq 1.3\%$ for the tension being statistical

in nature. On the other hand, the DR14 results show better agreement, with probability of $\simeq 6.3\%$.

We put an independent constraint on H_0 using BAO results with the sound horizon calibrated by baryon density measurements from BBN deuterium abundance studies. One of the BBN reaction rates has very poor laboratory constraints, so we have to rely on either theoretical or empirical estimates (Cooke et al., 2016; Cooke, Pettini, and Steidel, 2018). We obtain two H_0 constraints: $H_0 = 67.6 \pm 1.1$ km s^{-1} Mpc^{-1} using the theoretical reaction rate and $H_0 = 68.1 \pm 1.1$ km s^{-1} Mpc^{-1} using the empirical one. These results are consistent with each other and with CMB results, as can be seen in Figure 5.2. They are also consistent with past BAO + BBN results (Aubourg et al., 2015; Addison et al., 2018), showing that the tension in DR11 and DR12 did not have a large impact on the H_0 constraint. However, they are both in strong ($> 3\sigma$) tension with H_0 results from the distance ladder. Our results again highlight that the tension is not caused by systematic errors in the *Planck* analysis.

Starting in 2020, DESI will accurately measure BAO over a wide redshift range. We use the two BBN $\Omega_b h^2$ measurements and forecast future DESI BAO + BBN results. As can be seen in the right panel of figure 5.3, the choice of BBN reaction rate estimate will have a significant impact on the H_0 constraints. Improved measurements of the $d(p,\gamma)^3$He reaction rate (e.g. from LUNA) will be required in order to obtain accurate constraints of the Hubble constant using BAO + BBN.

5.6 Appendix: Lyα BAO Modules

For Lyα forest datasets, we use the provided χ^2 tables[5]. These tables give the value of the χ^2 as a function of the two BAO peak coordinates scaled using a fiducial cosmology:

$$\alpha_\perp = \frac{[D_M(z_{\rm eff})/r_d]}{[D_M(z_{\rm eff})/r_d]_{fid}} \quad \text{and} \quad \alpha_\parallel = \frac{[D_H(z_{\rm eff})/r_d]}{[D_H(z_{\rm eff})/r_d]_{fid}}. \tag{5.12}$$

We use these tables to interpolate the value of the χ^2 at the points we need for our analysis. For all the other measurements, we use Gaussian likelihoods with the measured means and standard deviations (6x6 covariance matrix for BOSS). We used the Astropy[6] package (Astropy Collaboration et al., 2013; Astropy Collaboration et al., 2018) for the theoretical modelling of the BAO peak coordinates.

Methods to interpolate the χ^2 tables are now available as part of the popular MCMC packages CosmoMC (Lewis and Bridle, 2002) and MontePython (Audren et al., 2013). As such, Lyα BAO

[5] https://github.com/igmhub/picca/tree/master/data
[6] http://www.astropy.org

results can now be easily included by the community in cosmological analyses.

CHAPTER 6

BAYESIAN METHODS FOR FITTING BAO IN THE LYMAN-α FOREST

We study and compare fitting methods for the Lyman-α (Lyα) forest 3D correlation function. We use the nested sampler `PolyChord` and the community code `picca` to perform a Bayesian analysis, which we compare with previous frequentist analyses. By studying synthetic correlation functions, we find that the frequentist profile likelihood produces results in good agreement with a full Bayesian analysis. On the other hand, Maximum Likelihood Estimation with the Gaussian approximation for the uncertainties is inadequate for current data sets. We compute for the first time the full posterior distribution from the Lyα forest correlation functions measured by the extended Baryon Oscillation Spectroscopic Survey (eBOSS). We highlight the benefits of sampling the full posterior distribution by expanding the baseline analysis to better understand the contamination by Damped Lyα systems (DLAs). We make our improvements and results publicly available as part of the `picca` package.

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

6.1 INTRODUCTION

Over the last few decades, cosmology has entered a data driven era, with large surveys providing rich data sets. These data sets encode a vast amount of information, which most of the time is

non-trivial to extract. Performing such large surveys is very expensive and time-consuming, which puts even more emphasis on efficient and accurate extraction of meaningful information. Over the last two decades, the Λ Cold Dark Matter (ΛCDM) model has become widely accepted as the standard cosmological model, however, it has 6 free parameters with extensions adding even more. On top of this, most analyses need extra nuisance parameters to create good models of their data, leading to very high-dimensional parameter spaces (e.g. Planck Collaboration et al., 2018; Abbott et al., 2018). It has become essential to have reliable analysis tools, and this has lead to an increased focus on statistical methods and interpretation.

The efficient and reliable extraction of Baryon Acoustic Oscillation (BAO) information from large scale structure (LSS) data has been a very active topic of research over the last 15 years. Since the first detections of BAO, using the distribution of galaxies (Eisenstein et al., 2005; Cole et al., 2005), there has been much attention given to optimizing statistical methods used on data of discrete tracers (e.g. Vargas Magaña et al., 2013; Anderson et al., 2014; Chan et al., 2018; Hinton, Howlett, and Davis, 2020). However, the newer method of measuring BAO using the Lyman-α (Lyα) forest has received comparatively less attention.

The first detection of the BAO scale in the Lyα forest auto-correlation (Ly$\alpha\times$Lyα) function was done by the Baryon Oscillation Spectroscopic Survey (BOSS) using Data Release 9 (DR9) of the Sloan Digital Sky Survey (SDSS; Busca et al., 2013; Slosar et al., 2013; Kirkby et al., 2013). It was also detected in the cross-correlation of Lyα absorbers with quasar positions (Ly$\alpha\times$QSO) using BOSS DR11 (Font-Ribera et al., 2014). A physical model for the contaminations was first used by Bautista et al. (2017) and du Mas des Bourboux et al. (2017), however, it has a large number of parameters that model the contamination by Damped Lyα systems (DLAs) and different metal absorption lines. Recent analyses have introduced yet more effects (de Sainte Agathe et al., 2019; Blomqvist et al., 2019), and these have led to a large parameter space and potentially complex behaviour. However, as the main aim has always been the measurement of the BAO peak position, these astrophysical parameters have received little attention.

The BAO scale can be measured using just two parameters: $\alpha_{||}$ and $\alpha_{\perp}$. These measure the size of the BAO scale relative to a fiducial cosmology along and across the line of sight, respectively. Because of the focus on measuring BAO, the high-dimensional parameter spaces have so far been investigated only using a frequentist methodology. In particular, the profile likelihood (e.g. Planck Collaboration et al., 2014b; Chan et al., 2018; Algeri et al., 2019) has been used to extract the relevant information. This method approximates the probability at some value of $\alpha_{||}$ and $\alpha_{\perp}$ by the maximum likelihood over the nuisance parameters at that point. In this work we use a Bayesian

approach to fit the Lyα forest 3D correlation function, and we provide a tool for studying this large parameter space and for extracting all the relevant information. While Bayesian methods have been used in Lyα forest analyses before (e.g. Pichon et al., 2001; Kitaura, Gallerani, and Ferrara, 2012; Horowitz et al., 2019; Porqueres et al., 2019), they were never used in BAO analyses of the Lyα correlation function.

The principal difference between the Bayesian (Bayes, 1763; Laplace, 1820) and frequentist (Neyman, 1937) methodologies is their interpretation of the concept of probability. In the Bayesian framework, probability is a degree of belief in an event, while in the frequentist framework, the probability of an event is the limit of its relative frequency in many trials. In the limit of infinite data, the two approaches produce the same results. Our focus, however, is not on philosophical interpretations, but on the practical consequences of the two frameworks when working with real data. One of the reasons for the widespread use of Bayesian methods is the availability of samplers such as Monte Carlo Markov Chain (MCMC; Metropolis et al., 1953; Hastings, 1970) which facilitate the efficient exploration of the complex high-dimensional posterior distributions that often appear in cosmology. It is the efficiency and accuracy of such tools that we want to compare with equivalent frequentist approaches within the context of fitting the Lyα forest correlation function.

The purpose of this work is to investigate the methods used so far and compare them with a Bayesian framework. In particular, we use for the first time a sampler to obtain the full posterior distribution of all parameters. We begin in Section 7.2, where we discuss the frequentist methods used so far, and compare them with the Bayesian methodology. In Section 6.3, we use synthetic correlation functions to showcase the similarities and differences of the two approaches when fitting the BAO parameters. Finally, in Section 6.4, we use a Bayesian framework to analyse the latest extended BOSS (eBOSS) DR14 correlation functions (de Sainte Agathe et al., 2019; Blomqvist et al., 2019), and highlight some potential uses and advantages of the full posterior distribution.

6.2 Bayesian vs. frequentist methods

We first discuss some of the theoretical differences between the Bayesian and frequentist approaches to statistical data analysis. We include this discussion for completeness and also because, as we will see, the two frameworks answer fundamentally different questions when it comes to the quantification of uncertainties. As such, we believe that these theoretical considerations are important to the interpretation of our results. We also discuss some practical differences and their implications for cosmology, and we conclude with a simple toy example.

6.2.1 THE BEST FIT MODEL

The difference in the interpretation of probability between the Bayesian and frequentist approaches leads to a difference in their principal object of study. The frequentist approach treats parameters as fixed quantities and the data as the random variable. The object of study is the probability of obtaining the data, D, given a model, $\mathcal{M}$, and some parameters, $\vec{\theta} = (\theta_1, ..., \theta_n)$; this is also known as the likelihood: $P(D|\vec{\theta}, \mathcal{M})$. For normally distributed data, the likelihood takes the form:

$$P(D|\vec{\theta}, \mathcal{M}) = \frac{\exp\left[-\frac{1}{2}\left(D - \mathcal{M}(\vec{\theta})\right)^T \Sigma^{-1} \left(D - \mathcal{M}(\vec{\theta})\right)\right]}{\sqrt{(2\pi)^n |\Sigma|}}, \tag{6.1}$$

where Σ is the covariance matrix of the data. We will also refer to the logarithm of the likelihood, which we denote by $\mathcal{L} \equiv \log P(D|\vec{\theta}, \mathcal{M})$.

The object of interest for a Bayesian is the posterior distribution of the parameters $\vec{\theta}$, given the data, and a model: $P(\vec{\theta}|D, \mathcal{M})$. This fully encapsulates our knowledge of the probability of possible values of the parameters of interest, by treating these parameters as random variables. The posterior distribution can be computed through Bayes' Theorem (Bayes, 1763; Laplace, 1820):

$$P(\vec{\theta}|D, \mathcal{M}) = \frac{P(D|\vec{\theta}, \mathcal{M}) P(\vec{\theta}|\mathcal{M})}{P(D|\mathcal{M})}, \tag{6.2}$$

where $P(\vec{\theta}|\mathcal{M})$ is the prior probability, and $P(D|\mathcal{M})$ is a constant (for a model $\mathcal{M}$) known as the Bayesian evidence. The evidence is the normalization of the posterior and requires an n-dimensional integral to be computed. It can be used to perform Bayesian model selection (see e.g. Liddle, Mukherjee, and Parkinson, 2006; Trotta, 2007; Trotta, 2008 for applications in cosmology), however, when the only goal is inference this quantity is not necessary.

If we work with wide flat priors, we can deduce from Equations 6.1 and 6.2 that the two frameworks will produce the same best fit $\vec{\theta}_{\text{best}}$, given by the frequentist maximum likelihood $P_{max}(D|\vec{\theta}, \mathcal{M})$, and by the Bayesian maximum posterior probability $P_{max}(\vec{\theta}|D, \mathcal{M})$.

6.2.2 QUANTIFYING UNCERTAINTIES

The two approaches diverge again when it comes to finding the uncertainty on $\vec{\theta}_{\text{best}}$. In this case, it is not only a difference in methodology, but also a fundamental difference in the object of interest. Frequentists quantify uncertainty through confidence intervals[1] (CI), which are defined by the

[1] For n-dimensional distributions these are referred to as confidence regions, however we chose to only talk about intervals to clearly distinguish them from Bayesian credible regions

proportion (frequency) of intervals, measured from the ensemble of possible data sets, that contain the true values of the parameters ($\vec{\theta}_{\rm true}$). On the other hand, Bayesian uncertainty is quantified through credible regions (CR), defined as the smallest region of the posterior that encompasses a certain probability (most often 68% and 95% CRs are quoted).

The two questions asked by Bayesians and frequentists are very different. A Bayesian CR is telling us that, given our data and prior, we are e.g. 95% confident that the true values $\vec{\theta}_{\rm true}$ are within that region. Meanwhile, the frequentist CI is telling us that if we repeat our experiment many times, the confidence intervals we obtain will contain the true values $\vec{\theta}_{\rm true}$ in e.g. 95% of the cases. Note that the frequentist CI that we obtain from our data does not state anything about the probability that it contains $\vec{\theta}_{\rm true}$ (this is a common misconception). In fact, there are extreme cases in the literature where a frequentist CI has 0% probability of containing the truth[2] (Jaynes and Kempthorne, 1976; Welch, 1939). This behaviour is due to the fact that in the frequentist methodology, one never conditions a result on the actually observed data D, but instead on its distribution of possible realizations, which we do not always fully understand given limited data. This issue is fundamental in cosmology, because we only have one Universe to observe. Bayesian CRs can also cause problems because of prior choices. For example, flat priors are only flat for the specific parametrization they are defined on.

We now turn to the computation of these uncertainty intervals and regions. In a frequentist framework, we usually start by computing the best fit parameters $\vec{\theta}_{\rm best}$ by maximizing the likelihood. This is called Maximum Likelihood Estimation (MLE). After that, we can compute a covariance matrix for the parameters by taking the second derivative of the likelihood in parameter space around this peak. However, this covariance matrix is only accurate in general if the errors on the data are normally distributed and the model is linear in all parameters (which would correspond to a Gaussian posterior in a Bayesian framework). If this is not the case, it can still be applied around the peak, but a better approach is to compute the likelihood on a grid in parameter space (we will call this a scan).

The confidence intervals can be computed using regions of equal likelihood around $\vec{\theta}_{\rm best}$. To find these regions, we need $\Delta\mathcal{L}_p \equiv \mathcal{L}_p - \mathcal{L}_{min}$ values such that the region defined by $\mathcal{L}_p$ corresponds to a certain CI of probability p. In the Gaussian case, these values can be computed analytically. In the general case, a large number of Monte Carlo simulations of the data are needed to compute $\Delta\mathcal{L}_p$ values. This procedure is outlined in Section 6.6. Using these $\Delta\mathcal{L}_p$ values, frequentists can

[2] We must stress, however, that even in such extreme cases the frequentist CIs are not wrong. They are just answering a different question.

draw constant likelihood contours using a scan of the parameters and obtain the correct confidence intervals.

We must stress that from a Bayesian perspective, the best fit and the credible regions are just special values computed from the posterior. It is this full posterior distribution that is the real object of interest because it contains all the information about the probable values of the parameters. The computation of the posterior distribution can however be a very demanding task. When dealing with low-dimensional spaces, Bayesians can compute the posterior on a grid, similar to the frequentist method. The scan can be used to compute the Bayesian credible regions by finding the smallest region of the scan that contains a certain probability p. This is usually done by ordering the scan in decreasing order of probability and computing their running sum until the result is a fraction p of the total probability of the grid:

$$\sum_{j=1}^{M} P(\vec{\theta}_j|D,\mathcal{M}) = p \sum_{j=1}^{N} P(\vec{\theta}_j|D,\mathcal{M}), \tag{6.3}$$

where N is the total number of points on the grid, and the M points obtained through this method cover the CR of probability p. Note that the integral of the posterior over a region is normally required to do this, however, this can be approximated by a sum if the grid is equally spaced because the probability density at each point is proportional to the probability mass for that region. As we will show in the next section, these two interpretations of the same grid results often produce identical results. However, the scan quickly becomes infeasible with increasing number of parameters, and Bayesians move on to using more effective methods, such as MCMC.

Bayesians can deal with high-dimensional spaces by efficient and accurate sampling, using tools such as MCMC (Metropolis et al., 1953; Hastings, 1970; see e.g. Press et al., 2007; Gelman et al., 2013 for detailed introductions). Furthermore, the efficient computation of the Bayesian evidence has also become possible with the introduction of Nested Sampling (Skilling, 2004). The underlying principle of such methods is the creation of samples from continuous random variables with probability density proportional to a known function, in our case the unnormalized posterior distribution. Once a sufficient number of samples have been generated, they can be used to compute summary statistics. For example, computing the credible regions amounts to finding the smallest region containing a certain fraction of samples. This fraction corresponds to the probability that, given the data and the prior, the true values of the parameters are within that region.

6.2.3 Nuisance parameters

The handling of nuisance parameters is also an intense topic of debate. In a Bayesian framework, the answer is marginalization. If the parameter vector contains two sets: interesting parameters $\vec{\theta}_{\rm i}$ and nuisance parameters $\vec{\theta}_{\rm n}$, then the posterior distribution of the interesting parameters is given by:

$$P(\vec{\theta}_{\rm i}|D,\mathcal{M}) = \int_{\vec{\theta}_{\rm n}} P(\vec{\theta}_{\rm i},\vec{\theta}_{\rm n}|D,\mathcal{M})\, d\vec{\theta}_{\rm n}. \tag{6.4}$$

We partition the full posterior by integrating over the nuisance parameter space. This ensures that all the probability mass contained in the nuisance parameters is accounted for. On the other hand, nuisance parameters are a major problem of non-Bayesian statistical theories (Ghosh, 1988). There is no consensus frequentist way of addressing this problem (see Basu, 1977 for a discussion of some of the methods).

One of the most common frequentist methods is the profile likelihood (e.g. Planck Collaboration et al., 2014b; Chan et al., 2018; Algeri et al., 2019), which involves computing the likelihood of $\vec{\theta}_{\rm i}$ on a grid, while conditioning on special values of the nuisance parameters, e.g. the best fit values $\vec{\theta}_{{\rm n},best}$. In practice, this is done by maximizing the likelihood over all nuisance parameters at every point on the grid of interesting parameters:

$$P(D|\vec{\theta}_{\rm i},\mathcal{M}) \propto \max_{\vec{\theta}_{\rm n}} P(D|\vec{\theta}_{\rm i},\vec{\theta}_{\rm n},\mathcal{M}). \tag{6.5}$$

For a Gaussian distribution this is equivalent to marginalization because low dimensional cuts from a high dimensional Gaussian are also Gaussian and their volume scales exactly the same as their peak. However, this method can go wrong when a lower dimensional cut changes shape depending on the position of the cut (we show an example of this in Section 6.2.4), because it effectively ignores any possibility that the nuisance parameters can have other values.

The profile likelihood is also useful for computing confidence levels when there are many interesting parameters. The scanning method presented in Section 6.2.2 becomes infeasible if we have too many parameters, and as such, the recursive application of the profile likelihood can be a useful approximation. We can scan a small subset of the parameters (usually one or two at a time) and treat the other parameters as nuisance by applying the profile likelihood at every point on the grid. If we apply this method recursively, we can compute an approximation of the full CIs. The profile likelihood is the method used so far in analyses of the Lyα forest correlation function by BOSS and eBOSS (Slosar et al., 2013; Delubac et al., 2015; Font-Ribera et al., 2014; Bautista et al.,

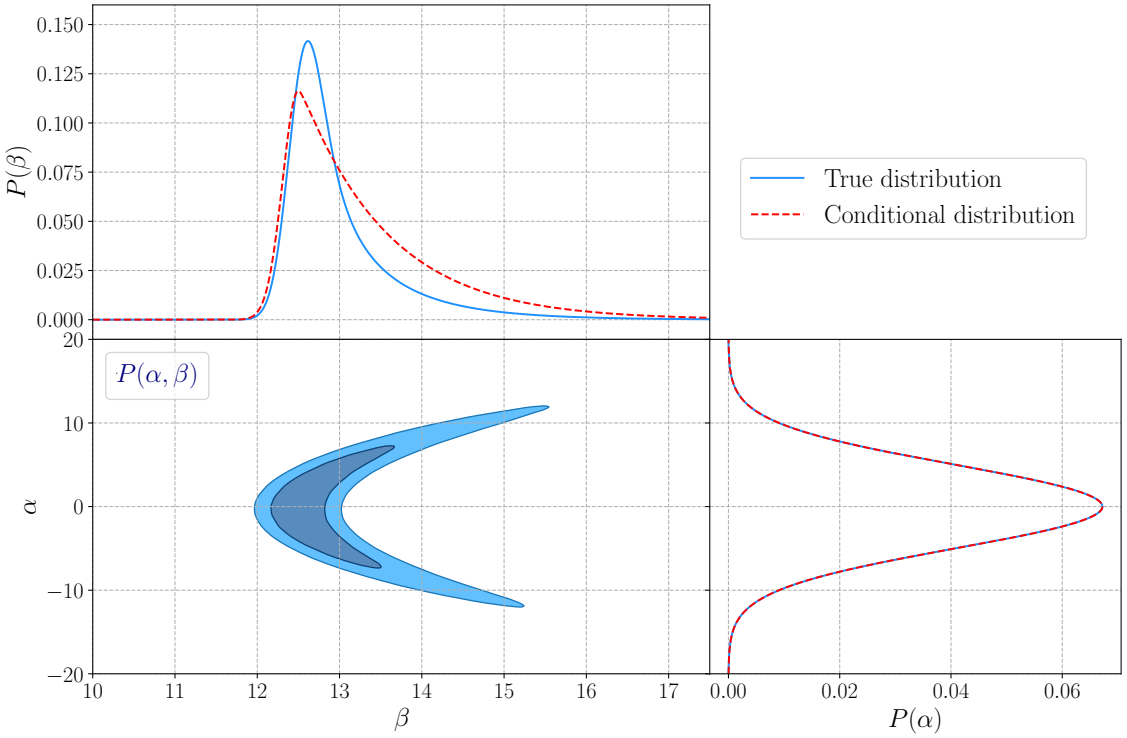

Figure 6.1. Toy model of a bivariate distribution $P(\alpha, \beta)$ used to illustrate the profile likelihood, a common frequentist approximation for dealing with high-dimensional problems. A scan over α conditional on the best fit value of β at each point gives identical results to the marginal distribution (right panel). On the other hand, an equivalent scan over β gives the wrong result (top panel), because the conditional distribution $P(\alpha|\beta)$ changes shape depending on the value of β.

2017; du Mas des Bourboux et al., 2017; Blomqvist et al., 2019; de Sainte Agathe et al., 2019).

6.2.4 Toy Example

We illustrate the profile likelihood using a toy example in Figure 6.1. We chose a bivariate distribution $P(\alpha, \beta)$ used in Wang, Broccardo, and Song (2017), which, depending on the parameter of interest, shows where the profile likelihood works perfectly as well as where it fails. If the interesting parameter is $\theta_i = \alpha$, a scan over α conditional on the best fit value of β at each point gives a result identical to the marginal distribution (bottom right panel of Figure 6.1). On the other hand, if $\theta_i = \beta$, the equivalent procedure gives the wrong result because on the right-hand side of the β grid, the conditional distribution $P(\alpha|\beta)$ is multimodal. This means the shape of $P(\alpha|\beta)$ changes depending on the value of β, resulting in a failure of the profile likelihood approximation.

This behaviour is not due to the difference in interpretation between a Bayesian and a frequentist. In fact, for this bivariate case, a simultaneous scan of both parameters would produce frequentist

results identical to the Bayesian ones, but this is infeasible in high dimensional spaces. If such pathological cases are correctly identified, frequentists can rely on reparametrizations to Gaussianize the problem. However, such cases could appear over arbitrarily many dimensions and be easily missed because usually we can only investigate 1 and 2 dimensional projections.

6.3 Testing BAO measurements on mock data sets

We begin our investigation by applying the different methods introduced above to the problem of efficient and accurate extraction of BAO information from the Lyα forest correlation function. To this end, we use a baseline model of the correlation function to create synthetic data sets. The baseline model is based on the Lyα correlation functions measured by the extended Baryon Oscillation Spectroscopic Survey (eBOSS) using SDSS DR14 data (de Sainte Agathe et al., 2019; Blomqvist et al., 2019), and the publicly available modelling package `picca`[3]. Their main analysis follows the frequentist methods introduced in Section 6.2. We perform a Bayesian analysis of the synthetic data sets and compare the different methodologies, with a focus on accurate BAO measurement.

6.3.1 Synthetic Correlation Functions

In order to compare different fitting methods, we produce 100 Monte Carlo simulations of the Lyα forest flux 3D correlation function. This allows us to investigate the differences over the entire population of possible correlation functions given a data set such as SDSS DR14. For simplicity, in this section we focus only on the Ly$\alpha\times$Lyα auto-correlation function using Lyα absorbers only in the Lyα region, and leave the analysis using the full DR14 data for the next section. We use the measured Lyα correlation function from SDSS DR14 (shown in Figure 8 of de Sainte Agathe et al., 2019). We fit this using the full model from de Sainte Agathe et al. (2019), including metal contaminations.

The mock data sets are drawn randomly from a multivariate normal distribution with mean $\xi(\vec{\theta}_{\rm best})$ and covariance C, where $\xi(\vec{\theta})$ is the best fit model of the correlation function measured in the DR14 analysis (de Sainte Agathe et al., 2019) and C is the covariance matrix of ξ. A new simulated correlation function is then given by:

$$\tilde{\xi} = \xi(\vec{\theta}_{\rm best}) + A\vec{y}, \tag{6.6}$$

[3] Available at https://github.com/igmhub/picca

where the matrix A comes from the Cholesky decomposition $C = AA^T$, and $\vec{y}$ is a vector of N (the size of the C) independent standard normal variates.

The model used to fit the Ly$\alpha\times$Lyα correlation function in the DR14 analyses has 11 free parameters, of which nine are considered nuisance parameters and two ($\alpha_{||}$, $\alpha_{\perp}$) are the parameters of interest (de Sainte Agathe et al., 2019). We follow this distinction and leave the analysis and description of the nuisance parameters for the next section. We fit our mocks using this model and the four methods introduced above. We use large flat priors for most parameters, but we follow de Sainte Agathe et al. (2019) and set tight Gaussian priors for two of the nuisance parameters that are less constrained by the data (see Table 6.1). The choice of uninformative priors means that the shape of the posterior is given only by the likelihood. This means that any differences between Bayesian and frequentist results will be produced by the differences discussed in Section 6.2, and will not be influenced by prior choices.

6.3.2 Fitting Methods

Following the discussion in Section 6.2, we choose four fitting methods to compare over the population of synthetic data sets:

1. **Frequentist MLE**: The likelihood is maximized over all parameters and the uncertainties are given by the covariance around $\vec{\theta}_{\text{best}}$, using the Gaussian $\Delta\mathcal{L}_p$ values.

2. **Frequentist scan**: A scan over $\alpha_{||}$ and $\alpha_{\perp}$ using the profile likelihood, with the uncertainties given by confidence intervals which are set using MC simulations.

3. **Bayesian scan**: A scan over $\alpha_{||}$ and $\alpha_{\perp}$ using the profile likelihood, with the uncertainties given by credible regions, computed using Equation 6.3.

4. **Bayesian sampler**: The full posterior distribution is computed, and the uncertainties on ($\alpha_{||}$, $\alpha_{\perp}$) are given by credible regions after marginalization over the nuisance parameters.

Lyα BAO analyses have so far been frequentist, and used the first two methods. The Bayesian interpretation of the scan was also used whenever scan results were combined with other cosmological probes (e.g. Cuceu et al., 2019) as part of popular packages such as `CosmoMC` or `MontePython` (Lewis and Bridle, 2002; Audren et al., 2013; Brinckmann and Lesgourgues, 2018). However, the two interpretations of the scan were never tested together. This method is also a good middle ground between the frequentist scan and the Bayesian sampler, because the contrast with the first is

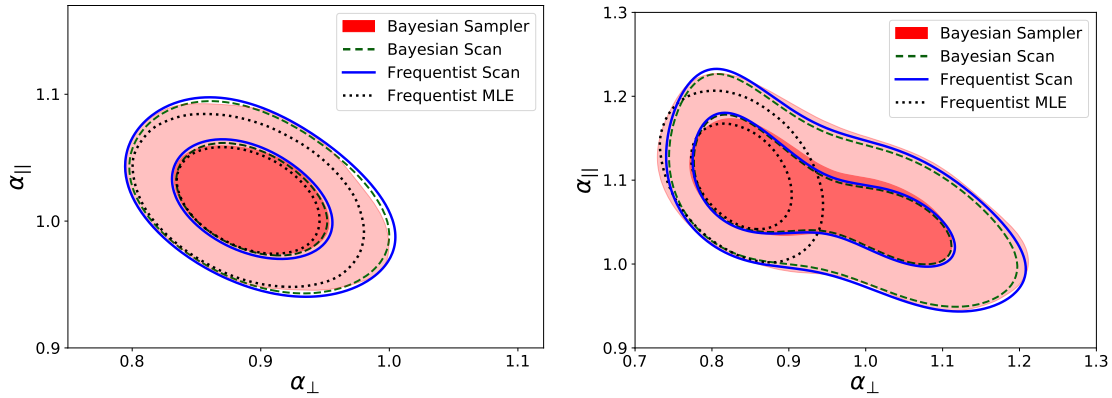

Figure 6.2. Comparison of BAO parameter constraints on mock correlation functions using the frequentist MLE and frequentist scan confidence intervals, and the Bayesian scan and Bayesian sampler credible regions. We showcase two mocks, one where the constraints are very close to Gaussian (left), and one where the constraints are strongly non-Gaussian (right). Note the different scaling. The last three methods produce remarkably similar results, considering they use different methods and very different quantification of uncertainties. On the other hand, the frequentist MLE can fail to properly capture the uncertainty in the results, especially in non-Gaussian cases.

only in the interpretation of uncertainty, while a comparison with the second allows us to directly test the profile likelihood.

The scan requires maximizing the likelihood over all nuisance parameters at each point on the grid. As this operation can be performed independently for each point, we implemented a parallel version of the scan code in `picca` to speed up our analysis. The frequentist analysis also requires a large number of MC simulations in order to compute the $\Delta\mathcal{L}_p^2$ as described in Section 6.2. We also implemented a parallel version of this step, which is now available in `picca`.

In order to sample the full posterior distribution, we implemented an interface to the popular Nested Sampler `PolyChord`[4] (Handley, Hobson, and Lasenby, 2015a; Handley, Hobson, and Lasenby, 2015b) in `picca`. Nested samplers were designed to efficiently compute the Bayesian evidence (Skilling, 2004), but they also provide accurate sampling of the posterior. In particular, nested samplers are very good at dealing with multimodal and highly degenerate posteriors.

6.3.3 Results

Noise on the correlation function can conspire to improve or smear the BAO peak. This means that a population of simulated correlation functions produced from the same covariance gives rise to a range of possible BAO constraints, from very tight, closely Gaussian fits, to very degenerate,

[4] Available at https://github.com/PolyChord/PolyChordLite

multimodal ones where the peak is barely detected. In Figure 6.2, we showcase two examples of results from this population. The one on the left is one of the best BAO constraints in the population, and is very close to Gaussian. On the other hand, the one on the right gives one of the worst constraints and is very non-Gaussian. We note that in the case of DR14-like errors on the correlation function, roughly 70% of constraints (computed using visual inspection of the mock fits) are close to Gaussian, similar to the one on the left. On the other hand, 4% are catastrophic failures, where the constraints are multimodal and strongly non-Gaussian in both parameters (BAO not detected).

Results of the four methods described above are presented in Figure 6.2. The last three methods agree very well with each other, even in the non-Gaussian case. This agreement is remarkable, because as discussed in Section 6.2, they answer fundamentally different questions. The frequentist scan uses confidence intervals to quantify uncertainty, while the two Bayesian methods use credible regions. The agreement between the frequentist scan and the Bayesian scan shows us that BAO results are robust to different interpretations of uncertainty. On the other hand, the agreement between the Bayesian scan and the Bayesian sampler shows us that the profile likelihood is a good approximation for fitting the BAO peak from the Lyα forest correlation function.

The Frequentist MLE, which assumes Gaussianity, does not fare as well. It produces results that are close to the ones obtained using the other three methods in the Gaussian mock, although slightly smaller. On the other hand, it completely fails to capture the uncertainty in the results for the non-Gaussian mock. As these non-Gaussian posteriors make up roughly 30% of the population of possible results, this approximation is inadequate for dealing with current data sets.

The comparison of specific posterior results such as in Figure 6.2 is Bayesian in nature. However, we can also compare them using a frequentist approach by computing the fraction of the interval population that contains the truth $\vec{\theta}_{\rm true}$ within the 68% and 95% regions. We find that the frequentist scan, the Bayesian scan and the Bayesian sampler are again in remarkable agreement, with roughly $55-60$ mocks containing the truth within their 68% regions, and roughly 92 of them containing the truth within their 95% regions. On the other hand, the MLE fails this test as well, with only 44 of the 68% intervals, and 74 of the 95% intervals containing the truth. These numbers are affected by sample noise because we only have 100 simulations, and as such are very rough. In particular, if the profile likelihood works perfectly, by construction the fractions for the frequentist scan will tend to 68% and 95% for a large number of samples.

These results do, of course, rely on the large uniform priors assumed for most parameters. Non-uniform or overly restrictive priors would have an impact on our conclusions, as some of the

differences could come from the choice of priors. In a study similar to our work, Chan et al. (2018) arrive at a different conclusion when tight priors are used on the BAO parameters. We discuss the similarities and differences between our study and Chan et al. (2018) in Section 6.7.

Another relevant question that distinguishes these methods is that of the computational cost. MLE is by far the fastest method, with computation times of order $10^{-1} - 10^{0}$ CPU hours. However, as we just showed, MLE alone is inadequate for current Lyα forest BAO analyses. For the scan, we find computation times of order $10^{2} - 10^{4}$ CPU hours using typical grid sizes of 30×30 up to 50×50. Furthermore, the frequentist interpretation of the scan requires a large set of MC simulations to compute the $\Delta\mathcal{L}_p$ necessary for setting the right CIs. In the best case, these require an extra 10^{3} CPU hours for 10000 MC mocks. `PolyChord` with a typical setup ($nlive = 25\times$ number of parameters, $numrepeats = 3\times$ number of parameters) performs over a timescale similar to the scan ($10^{2} - 10^{4}$ CPU hours), and not only computes the full posterior distribution, but also its integral (the Bayesian evidence).

6.4 The Full Posterior of eBOSS DR14

We now turn our attention to the full posterior distribution. As we have shown above, when the only interest is measuring the BAO peak position, both the scan and the sampler produce very similar results. However, the sampler also computes accurate distributions for all the other parameters. This wealth of information is generally ignored. Some of these are astrophysical parameters, such as the bias of high column density (HCD) absorbers ($b_{\rm HCD}$). These parameters tend to be very correlated with each other and have proved to be very sensitive to modelling choices, which makes their measurement less robust compared to that of the BAO peak.

The computation of the full posterior distribution allows us to access this previously ignored information. It allows us to study the complex high-dimensional distribution of these parameters, and to better understand how our modelling choices affect their measurement. Furthermore, our use of the Nested Sampler `PolyChord` means that we can deal with very degenerate and even multimodal distributions. This sampler also computes the Bayesian evidence, thus providing an accurate tool for model comparison even for strongly non-Gaussian posteriors. As such, it allows us to test possible extensions to the baseline model that may have previously appeared daunting due to their complex interaction with other parameters.

Parameter	Description	Prior
$\alpha_{\|\|}$, $\alpha_{\perp}$	BAO peak position	$\Pi[0.1, 2]$
$b_{\eta Ly\alpha}$	Lyα velocity bias	$\Pi[-0.5, 0]$
$\beta_{Ly\alpha}$	Lyα RSD parameter	$\Pi[0.1, 5.0]$
β_{QSO}	Bias of HCDs	$\Pi[-0.2, 0]$
$\Delta r_{\|\|}[h^{-1}Mpc]$	Shift due to QSO redshift errors	$\Pi[-10, 10]$
$\sigma_v[h^{-1}Mpc]$	Smoothing parameter for QSO non-linear velocities and redshift precision	$\Pi[2, 15]$
ξ_0^{TP}	Amplitude parameter of quasar radiation	$\Pi[0, 2]$
b_{HCD}	Bias of HCDs	$\Pi[-0.2, 0]$
β_{HCD}	RSD parameter of HCDs	$\mathcal{N}(0.5, 0.2^2)$
$b_{\eta CIV(\mathrm{eff})}$	Velocity bias of metal absorber	$\mathcal{N}(-0.005, 0.0026^2)$
$b_{\eta SiII(1190)}$	Velocity bias of metal absorber	$\Pi[-0.2, 0]$
$b_{\eta SiII(1193)}$	Velocity bias of metal absorber	$\Pi[-0.2, 0]$
$b_{\eta SiIII(1207)}$	Velocity bias of metal absorber	$\Pi[-0.2, 0]$
$b_{\eta SiII(1260)}$	Velocity bias of metal absorber	$\Pi[-0.2, 0]$

Table 6.1. Sampled parameters and their prior limits. We use flat priors $\Pi[a, b]$ for most parameters, with limits a and b chosen such that the prior is uninformative. Following de Sainte Agathe et al. (2019) and Blomqvist et al. (2019), we use Gaussian priors $\mathcal{N}(\mu, \sigma^2)$ with mean μ and standard deviation σ for β_{HCD} and $b_{\eta CIV(\mathrm{eff})}$,.

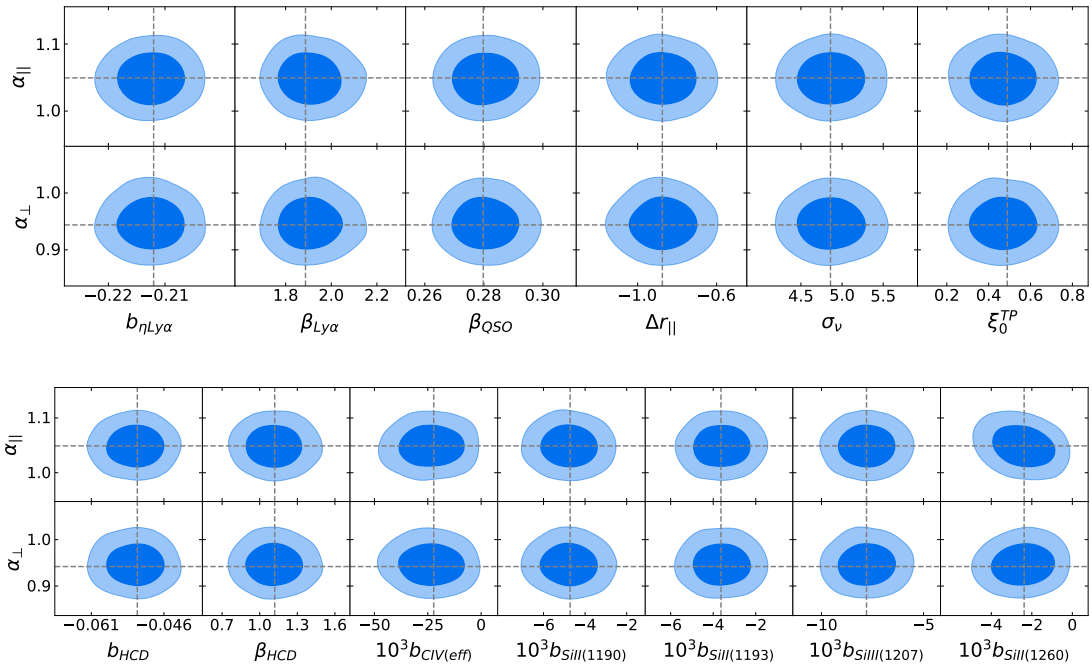

Figure 6.3. Projected posterior distributions of the BAO parameters ($\alpha_{||}$, $\alpha_{\perp}$), versus the other model parameters, using the full eBOSS DR14 Lyα forest data. There are no major correlations between $\alpha_{||}/\alpha_{\perp}$ and any of the other parameters, which shows the robustness of BAO parameters to different modelling choices.

6.4.1 The Baseline Analysis

We start by using `PolyChord` to analyse all Lyα forest eBOSS DR14 correlation functions. In contrast to the last section, we include the correlation of Lyα absorbers in the Lyα region with Lyα absorbers in the Lyβ region, Lyα(Lyα)$\times$Lyα(Lyβ), and the cross-correlation with quasars Ly$\alpha\times$QSO. We use the models presented in de Sainte Agathe et al. (2019) and Blomqvist et al. (2019) to model the auto-correlation and the cross-correlation with quasars, respectively. Our analysis follows the same steps, but we do not include relativistic effects, and we use only one parameter to model the bias of HCDs (instead of three). As in the previous section, we use broad flat priors on all parameters (except β_{HCD} and the bias of foreground CIV absorption (b_{CIV}) which have Gaussian priors). The model parameters used and their priors are presented in Table 6.1.

In Figure 6.3 we show the projected posterior distributions of $\alpha_{||}$ and $\alpha_{\perp}$ versus the other parameters. This shows the robustness of BAO measurements, as none of the parameters are correlated with ($\alpha_{||}$, $\alpha_{\perp}$). The only exception is the bias of the $SiII(1260)$ absorbers ($b_{SiII(1260)}$) which has a small correlation with $\alpha_{||}$. This is due to the line causing an increased correlation along the line of sight at a separation of $r_{||} \approx 105\mathrm{h}^{-1}\mathrm{Mpc}$, which is very close to the BAO peak. However, this metal contamination is barely detected (at $\sim 2\sigma$).

6.4.2 A simple extension

We now use `PolyChord` to illustrate the advantages of sampling the full posterior. We do this by sampling L_{HCD} as an extension to the baseline analysis. This is a parameter used to describe the contamination by HCDs, and it corresponds to a typical length scale. HCDs with a length scale above $\sim 14\mathrm{h}^{-1}\mathrm{Mpc}$ are efficiently identified and masked before computing the correlation function. This means unidentified systems are expected to have a typical length scale below this value, but above the bin width of the correlation function ($4\mathrm{h}^{-1}\mathrm{Mpc}$). This parameter has so far been fixed to a value of $L_{\mathrm{HCD}} = 10\mathrm{h}^{-1}\mathrm{Mpc}$ following the study by Rogers et al. (2018). Different values for this parameter were tested by de Sainte Agathe et al. (2019) and Blomqvist et al. (2019) to confirm there are no biases when measuring the BAO position, however, completely freeing this parameter proved challenging for the minimizer to deal with. We set a flat prior given by: $\Pi[2, 30]\mathrm{h}^{-1}\mathrm{Mpc}$.

As the full projected posterior is too large to show, we chose a subset of parameters and plot their posteriors in Figure 6.4. We show $\alpha_{||}$ and $\alpha_{\perp}$ to study the impact of L_{HCD} on measurements of the BAO peak position. From the other parameters, we chose those whose posterior is affected by L_{HCD}. We also plot the posterior using the basic model, where L_{HCD} is fixed for comparison. Plots of the full posterior are available as part of a Jupyter Notebook at `https://github.com/`

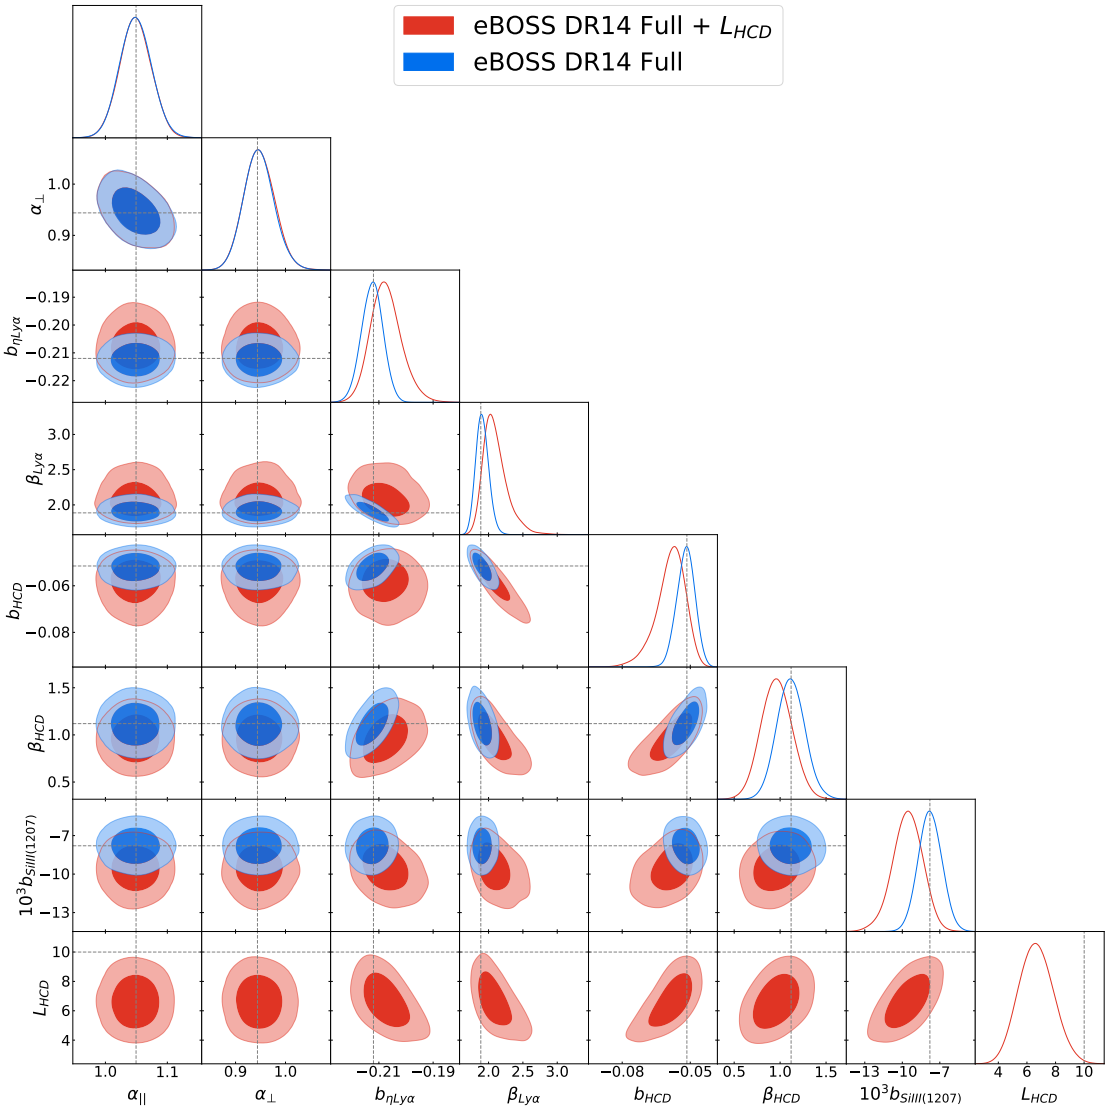

Figure 6.4. Triangle plot showing projected posteriors of DR14 Lyα forest results using the baseline analysis versus an extension where the typical scale of high column density absorbers, $L_{\rm HCD}$, is sampled. The best fit results for the baseline analysis are given by the dashed lines. The first two columns show that the measurement of the BAO scale is very robust to this change. However, the other parameters plotted are correlated with $L_{\rm HCD}$, and as such, their posteriors are significantly affected by this parameter.

andreicuceu/eBOSS-Lya-Posteriors.

The first two columns in Figure 6.4 show that even though the posterior is significantly affected by the choice of $L_{\rm HCD}$, the BAO peak position is very robust to it. In particular, the $\alpha_{||} - \alpha_{\perp}$ posterior of the extended model is in excellent agreement with the results from the baseline analysis. $L_{\rm HCD}$ is constrained within the expected region: $L_{\rm HCD} = 6.6 \pm 1.2,\ ^{+2.5}_{-2.3}$ h^{-1}Mpc $(68\%, 95\%)$. As can be seen in the bottom row, the other parameters shown in Figure 6.4 are correlated with $L_{\rm HCD}$, and as such their posterior is significantly affected. The constraints on the biases and RSD parameters of the Lyα forest and HCDs are all wider than in the baseline analysis, showing the fragility of these parameters to modelling choices. Interestingly, the bias of the SiIII(1207) absorption is also very correlated with $L_{\rm HCD}$.

Comparisons such as the one presented in Figure 6.4 are beneficial because they visually show us how model parameters behave, and how our modelling choices influence our results.

6.5 Conclusions

Analyses of the Lyα forest 3D correlation function have so far focused on measuring the position of the BAO peak, at the expense of analyses of other parameters and their interaction. A frequentist methodology has been used so far, despite the BAO results being subsequently used as part of Bayesian analyses when combined with other probes. In this work, we performed for the first time a Bayesian analysis of the Lyα forest correlation function, and we computed the full posterior using the Nested Sampler `PolyChord`.

We started by discussing the different approaches to fitting the correlation function in Section 7.2, and in particular we focused on the methods used and the difference in the quantification of uncertainty. Frequentists use Confidence Intervals computed from the population of possible data sets, while Bayesians use credible regions computed from the posterior distribution of the parameters given the data and the prior. Furthermore, when dealing with high-dimensional model spaces, Bayesians have access to tools such as MCMC. On the frequentist side, approximations such as the profile likelihood are used. A scan is performed over a few parameters (usually anything above 2 or 3 is computationally infeasible), and the likelihood is maximized over other parameters at each point. This is generally a good approximation, however, it can fail in some cases, such as in the toy example presented in Section 7.2.

We compared the different methodologies on synthetic correlation functions of eBOSS DR14 in Section 6.3. We used both a Bayesian and a frequentist interpretation of scan results. This allows

us to test potential differences from the quantification of uncertainty by comparing the two. We can also test the profile likelihood approximation by comparing the Bayesian scan with Bayesian sampler results. We showed that the three methods agree remarkably well on both a mock with a tight BAO constraint, and one with a large non-Gaussian one. We also plotted results using the frequentist Maximum Likelihood Estimation, where the uncertainties are set using the second derivative around the peak likelihood. This method fails to capture the uncertainty in non-Gaussian cases. Considering these make up roughly $\sim 30\%$ of the population, the MLE is inadequate for current data sets.

By computing the Bayesian interpretation of the scan and showing it agrees very well both with the frequentist scan and the Bayesian sampler, we have confirmed that scan results can safely be combined with other probes as part of Bayesian packages commonly used in cosmology.

In Section 6.4 we turned our attention to the full posterior distribution of eBOSS DR14 Lyα forest correlation functions. We showed that the BAO peak position parameters, $\alpha_{||}$ and $\alpha_{\perp}$, do not have any strong correlation with any of the other parameters, and as such, they are very robust to modelling choices. We extended the baseline analysis by sampling the typical length scale of HCD systems, which was previously fixed to $L_{\rm HCD} = 10 h^{-1}$Mpc. We compared the projected posteriors of the two models, and showed that it has no impact on the measured BAO position. We measured $L_{\rm HCD} = 6.6 \pm 1.2\, h^{-1}$Mpc (68%), and find that this parameter is correlated with many other astrophysical parameters and as such has a significant impact on their posteriors. Plots of the full posterior are available at `https://github.com/andreicuceu/eBOSS-Lya-Posteriors`.

Comparisons such as the one presented in Section 6.4.2 are now easy and fast to perform. Furthermore, this benefit does not come at a major computational cost. As discussed in Section 6.3, `PolyChord` requires similar computational time compared to a 2D scan using the profile likelihood. Beyond constraining the BAO peak parameters, `PolyChord` also computes the full posterior distribution of all parameters and the Bayesian evidence.

Our improvements are freely available as part of the community code `picca`. We hope that these tools will be used to improve future Lyα forest analyses by using the full posterior distribution to study complex parameter distributions and inform modelling choices.

6.6 Appendix A: Computation of CIs

If the model is not a linear function of the parameters, then the $\Delta\mathcal{L}_p$ values, necessary for setting confidence intervals from grid results, usually cannot be computed analytically. In general, these

values must be computed from the population of the possible realizations of the data. However, this population is inaccessible if we can only perform the experiment once, and we must rely on simulated populations of data sets instead of the real one (Press et al., 2007). In cosmology this could be achieved by running the full analysis many times on simulations, however as these are extremely expensive computationally, we usually rely on Monte Carlo (MC) simulations instead. The idea behind MC is to use the measured correlation functions and their covariance matrices together with a model of the data to produce random realizations of our data. See Chapter 15.6 of Press et al. (2007) for a detailed discussion of this algorithm. Note that the χ^2 is generally used instead of the log-likelihood, and $\Delta\chi^2$ values are computed, however, we use the log-likelihood to maintain consistency and because the two are proportional if the likelihood is Gaussian: $\chi^2 \propto -2\log\mathcal{L}$.

We want to produce samples from the distribution of possible data sets given the true model and the true parameter values: $P(D|\vec{\theta}_{\rm true}, \mathcal{M}_{\rm true})$. However, finding this distribution is the goal of our experiment, and as such it is inaccessible at this point. Therefore, we start by assuming that the shape of $P(D|\vec{\theta}_{\rm true}, \mathcal{M}_{\rm true})$ is similar to the shape of $P(D|\vec{\theta}_0, \mathcal{M}_0)$, where we are using our particular model $\mathcal{M}_0$, and the best fit parameters given that model and our data: $\vec{\theta}_0$. This means that we can use $\vec{\theta}_0$ as a surrogate for $\vec{\theta}_{\rm true}$ and create a population of synthetic data sets from a multivariate normal with mean $\mathcal{M}_0(\vec{\theta}_0)$ and covariance matrix given by the measurement errors of the data (again assuming our data is normally distributed).

For each mock data set created using this method, we find the best fit parameter values and call them $\vec{\theta}_\alpha$. If we compute the quantity $\Delta\mathcal{L} \equiv \mathcal{L}(\vec{\theta}_0) - \mathcal{L}(\vec{\theta}_\alpha)$ for each mock data set, we obtain a chi-square distribution with n degrees of freedom, where n is the number of parameters. By analysing a large population of mocks and computing this quantity, the $\Delta\mathcal{L}_p$ is easily obtained by looking at the $\Delta\mathcal{L}$ value that contains the fraction of mocks corresponding to p for that chi-square distribution.

If we are only interested in a subset of ν parameters, then the quantity of interest is $\Delta\mathcal{L}_\nu \equiv \mathcal{L}_\nu - \mathcal{L}(\vec{\theta}_\alpha)$, where $\mathcal{L}_\nu$ is found by fixing the parameters of interest to their best fit values from $\vec{\theta}_0$ and maximizing over the other parameters for each mock. $\Delta\mathcal{L}_\nu$ is a chi-square distribution with ν degrees of freedom, and the same procedure as above is used to find $\Delta\mathcal{L}_p$.

6.7 Appendix B: Comparison with Chan et al. 2018

The Dark Energy Survey (DES) collaboration performed a study similar to the one presented in this work, testing different fitting methods with the aim of measuring BAO using angular correlation

functions in tomographic bins (Chan et al., 2018). They also compared the performance using an MLE approach, the profile likelihood and an MCMC. When they computed the mean and uncertainty from the profile likelihood, they used the mean and variance across the scan weighted by the likelihood value at each point. This is in contrast with our method of finding constant χ^2 surfaces, and as such it may lead to different results.

The main difference is that they worked with isotropic BAO which means they only need to measure one parameter, α. Furthermore, they also pruned their population of synthetic data sets to those where the 68% α constraints are entirely contained within the prior range $[0.8, 1.2]$ (where $\alpha = 1$ is the true value). The result is that the remaining mocks are the ones where the BAO peak has a $\Delta\chi^2 = 1$ region within their prior range on α, which ensures the posterior distribution of α can be approximated as Gaussian. They found that all three methods perform very well, and after investigating the population statistics, they conclude that the MLE is the best tool in these conditions.

In the particular case of isotropic BAO we find similar results. As the model has less freedom, most constraints are Gaussian and all three methods work very well. However, they found that the MCMC has a larger bias in estimating α compared to the MLE (although both are very small). We believe that this is because they used the median to quantify the MCMC result, while we are using the maximum posterior point. The best fit point given by an MCMC (assuming flat priors) and MLE should be the same if both are run appropriately, and as such there should not be any difference in the bias of the peak for an unimodal posterior.

In Appendix A, Chan et al. (2018) worked with a larger range ($[0.6, 1.4]$) and showed that in this case the MCMC is the best approach, with both the MLE and profile likelihood showing small biases. This is in line with our results, considering that in this case there are some non-Gaussian results that are better fit using an MCMC approach. However, we also found the profile likelihood still works very well in these cases. This difference in results could be either caused by the difference in quantifying uncertainty as discussed above, or by a failure of the profile likelihood approximation when dealing with the nuisance parameters.

CHAPTER 7

COSMOLOGY BEYOND BAO FROM THE 3D DISTRIBUTION OF THE LYMAN-α FOREST

We propose a new method for fitting the full-shape of the Lyman-α (Lyα) forest three-dimensional (3D) correlation function in order to measure the Alcock-Paczynski (AP) effect. Our method preserves the robustness of baryon acoustic oscillations (BAO) analyses, while also providing extra cosmological information from a broader range of scales. We compute idealized forecasts for the Dark Energy Spectroscopic Instrument (DESI) using the Lyα auto-correlation and its cross-correlation with quasars, and show how this type of analysis improves cosmological constraints. The DESI Lyα BAO analysis is expected to measure $H(z_{\rm eff})r_{\rm d}$ and $D_{\rm M}(z_{\rm eff})/r_{\rm d}$ with a precision of $\sim 0.9\%$, where H is the Hubble parameter, $r_{\rm d}$ is the comoving BAO scale, $D_{\rm M}$ is the comoving angular diameter distance and the effective redshift of the measurement is $z_{\rm eff} \simeq 2.3$. By fitting the AP parameter from the full shape of the two correlations, we show that we can obtain a precision of $\sim 0.5-0.6\%$ on each of $H(z_{\rm eff})r_{\rm d}$ and $D_{\rm M}(z_{\rm eff})/r_{\rm d}$. Furthermore, we show that a joint full-shape analysis of the Lyα auto and cross-correlation with quasars can measure the linear growth rate times the amplitude of matter fluctuations in spheres of 8 h^{-1}Mpc, $f\sigma_8(z_{\rm eff})$. Such an analysis could provide the first ever measurement of $f\sigma_8(z_{\rm eff})$ at redshift $z_{\rm eff} > 2$. By combining this with the quasar auto-correlation in a joint analysis of the three high-redshift two-point correlation functions,

we show that DESI could be able to measure $f\sigma_8(z_{\rm eff} \simeq 2.3)$ with a precision of $5-12\%$, depending on the smallest scale fitted.

The work shown in this chapter is presented in Cuceu et al. (2021).

7.1 Introduction

The vast amount of cosmological data from spectroscopic surveys is usually compressed into summary statistics such as the correlation function or power spectrum. These statistics can be directly used to measure cosmological parameters; however, it is common to split the inference into two steps (e.g. Beutler et al., 2011; Ross et al., 2015; Alam et al., 2017; eBOSS Collaboration et al., 2021). A template is first used to model the power spectrum or correlation function, in order to measure a few relevant quantities that contain most of the cosmological information. These measurements are then used to fit cosmological parameters for some model, for example flat Λ Cold Dark Matter (ΛCDM), in combination with other probes, usually the cosmic microwave background (e.g. from Planck Collaboration et al., 2018). This approach is used because it contains minimal assumptions, and the full two-point statistic is compressed into a few well understood physical quantities.

Measuring the scale of the acoustic peak from the baryon acoustic oscillation (BAO) signal is one of the most widely used compression methods. This is usually done by splitting the template into a peak and a smooth component for the correlation function, or wiggles and no-wiggles components for the power spectrum. The coordinates of the peak (or wiggles) component are then re-scaled in order to fit the BAO scale from the data. This method has been used to measure the BAO scale using the galaxy distribution at redshifts $z \lesssim 1$ (e.g. Eisenstein et al., 2005; Cole et al., 2005), the quasar (QSO) distribution at redshifts $0.8 < z < 2.2$ (e.g. Ata et al., 2018), and the Lyman-α (Lyα) forest at redshifts $2 < z < 3$ (e.g. Busca et al., 2013; Slosar et al., 2013; Kirkby et al., 2013; Font-Ribera et al., 2014).

The Lyα forest consists of a series of absorption lines blueward of the Lyα emission peak in spectra of high-redshift quasars (e.g. Lynds, 1971; Rauch, 1998). The forest appears due to absorption by neutral hydrogen between the quasar and us, which means it traces the intergalactic medium. This makes it a great tool for cosmology, as it probes the distribution of matter at redshifts ($z \gtrsim 2$) that are generally hard to access with other probes (see e.g. Croft et al., 1999; McDonald et al., 2000; Croft et al., 2002; Viel, Haehnelt, and Springel, 2004, for early cosmological applications).

A common way to extract more information from the two-point statistics of discrete tracers is to fit the full shape (instead of just the peak component) in order to measure the growth rate of structure through redshift space distortions (RSD; e.g Blake et al., 2011; Reid et al., 2012; Beutler et al., 2012; Samushia et al., 2014). This approach is not possible for the Lyα forest because we have to marginalize over an unknown velocity gradient bias, which is degenerate with the growth rate. This bias appears because we work with the two-point statistics of flux, which has a non-linear mapping to the directly distorted field of optical depth (see e.g. Slosar et al., 2011; McDonald, 2003; Givans and Hirata, 2020; Chen, Vlah, and White, 2021). Therefore, an RSD analysis using the Lyα forest three-dimensional (3D) correlation function has so far been out of reach.

The analysis of the Lyα 3D auto-correlation function (Ly$\alpha\times$Lyα) and its cross-correlation with the quasar distribution (Ly$\alpha\times$QSO) has evolved considerably since they were first used to measure the BAO peak from Baryon Oscillation Spectroscopic Survey (BOSS) data (Busca et al., 2013; Slosar et al., 2013; Kirkby et al., 2013; Font-Ribera et al., 2014). A physical model for the correlations was introduced by Bautista et al. (2017) and du Mas des Bourboux et al. (2017). This model includes the effect of metal line contamination and that of high column density (HCD) systems. With the extended BOSS (eBOSS) analyses, the Lyα signal from the Lyβ section of the forest was also used, first through its correlation with Lyα signal in the Lyα section (de Sainte Agathe et al., 2019), and then through its correlation with the QSO distribution (Blomqvist et al., 2019; du Mas des Bourboux et al., 2020). Even though major advancements have been made in modelling and understanding the 3D Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO statistics, so far they have only been used to measure BAO.

In this work, we investigate the potential for extracting more cosmological information from the 3D distribution of the Lyα forest through the Alcock-Paczynski (AP) effect (Alcock and Paczynski, 1979; Hui, Stebbins, and Burles, 1999; McDonald and Miralda-Escudé, 1999; McDonald, 2003). This appears due to the choice of fiducial cosmology, which is used to transform the measured angles and redshifts into comoving coordinates. If this fiducial cosmology is different from the true cosmology, the measured correlation will have an extra anisotropy. Thus isolating this anisotropic AP contribution allows us to determine the true background cosmological model. Some of this AP signal is measured through anisotropic BAO analyses, by measuring two distinct scales along versus across the line of sight. However, this distortion affects the whole correlation function. Therefore, the first objective in this chapter is to complement standard Lyα BAO analyses with AP constraints from a broader range of scales.

The two Lyα forest correlation functions (Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO) are some of our best

probes of the Universe at redshifts $1.8 < z < 4$. However, there is big potential for a third correlation function in this redshift range: the quasar auto-correlation (QSO×QSO). As mentioned above, this has already been used to measure both BAO and the growth rate of structure at effective redshifts $z_{\rm eff} \simeq 1.6$. With the start of the Dark Energy Spectroscopic Instrument (DESI) survey, we will have new quasar catalogues with about 0.7 million expected to be at redshifts $z > 2.1$ (DESI Collaboration et al., 2016). This opens up the potential of performing a joint analysis of the three correlation functions (Lyα×Lyα, Lyα×QSO and QSO×QSO) for the first time. Jointly fitting the full shape of all three correlations would allow us to take full advantage of the synergies between them, and lead to more precise and robust constraints. Our second goal in this work is to investigate how such an analysis could be performed and study its benefits, including the potential for measuring RSD.

We start by introducing our methodology for template-fitting the full shape of the Lyα forest correlation function in Section 7.2. We also compare our approach with that used in past analyses of discrete tracers. After that, in Section 7.3 we perform a forecast analysis to demonstrate how the AP effect can be measured from the full shape of the correlation while preserving the robust BAO measurement. We also demonstrate the usefulness of such a measurement in constraining cosmological parameters in a flat ΛCDM model. Finally, in Section 7.4 we forecast a joint analysis of the three high-redshift two-point (high$-z$ 3×2pt) correlation functions (Lyα×Lyα, Lyα×QSO and QSO×QSO) in order to study their synergies and showcase the potential benefits of such an analysis.

7.2 METHOD

Our model of the 3D correlation function is based on the framework introduced by Kirkby et al. (2013) and used in all Lyα forest BAO analyses. Our approach is meant to extend these analyses to also include information from the broadband. We use a template power spectrum and introduce parameters that re-scale its coordinates. A fit to the data allows us to place constraints on these scale parameters. The resulting measurements can be transformed into constraints on cosmological parameters. We start by introducing these scale parameters in Section 7.2.1. After that, we introduce the components of the template in Section 7.2.2, and compare our approach to BAO analyses and previous full-shape analyses. In Section 7.2.3, we introduce our models for the Lyα forest auto-correlation, its cross-correlation with quasars, and the quasar auto-correlation. Finally, in Section 7.2.4 we showcase the effects of our scale parameters on the Lyα forest correlation function.

7.2.1 Scale parameters

When computing the 3D correlation function, we transform the observed redshift and angular separations $(\Delta z, \Delta\theta)$ into comoving coordinates $(r_{||}, r_{\perp})$. For positions i and j, at redshifts z_i and z_j and separated by an angle $\Delta\theta$, we define the radial coordinates as (du Mas des Bourboux et al., 2020)

$$\begin{aligned} r_{||} &= [D_{\mathrm{C,fid}}(z_i) - D_{\mathrm{C,fid}}(z_j)] \cos\frac{\Delta\theta}{2}, \\ r_{\perp} &= [D_{\mathrm{M,fid}}(z_i) + D_{\mathrm{M,fid}}(z_j)] \sin\frac{\Delta\theta}{2}, \end{aligned} \tag{7.1}$$

where $D_{\mathrm{M}}(z)$ is the comoving angular diameter distance and $D_{\mathrm{C}}(z) = c\int_0^z \mathrm{d}z/H(z)$ is the radial comoving distance, with c as the speed of light and $H(z)$ as the Hubble parameter. The *fid* term indicates that we use a fiducial cosmology to compute these distances. If the true cosmology is different from the fiducial one, the ratio between the inferred line of sight and transverse distances will be different from the true ratio. This means we will observe an apparent anisotropy in the measured correlation, which is the Alcock-Paczynski effect we wish to measure (Alcock and Paczynski, 1979). Note however, that there are other sources of anisotropy, such as RSD. In order to measure the AP effect, we have to correctly model and marginalize over all other anisotropies.

When building a model for the correlation function, we follow past Lyα forest BAO analyses and use a template power spectrum computed using a fixed cosmology. Following Kirkby et al. (2013), we allow for small differences between the template and measured cosmologies by using general coordinate transformations of the form $r_{||} \to r'_{||}(r_{||}, r_{\perp}, z)$ and $r_{\perp} \to r'_{\perp}(r_{||}, r_{\perp}, z)$.

The most commonly used parametrization for anisotropic re-scalings is given by:

$$r'_{||} = q_{||} r_{||}, \quad r'_{\perp} = q_{\perp} r_{\perp}, \tag{7.2}$$

where $(q_{||}, q_{\perp})$ re-scale the coordinates along and across the line of sight, respectively. However, we wish to isolate the AP effect which changes the ratio $r_{\perp}/r_{||}$. Therefore, we define the parameters:

$$\phi(z) \equiv \frac{q_{\perp}(z)}{q_{||}(z)} \text{ and } \alpha(z) \equiv \sqrt{q_{\perp}(z) q_{||}(z)}, \tag{7.3}$$

where $\phi(z)$ re-scales the ratio: $r'_{\perp}/r'_{||} = \phi\, r_{\perp}/r_{||}$, and is meant to measure the AP effect. On the other hand, $\alpha(z)$ re-scales the product $r'_{\perp} r'_{||} = \alpha^2\, r_{\perp} r_{||}$, which translates into an isotropic re-scaling of ξ. The effect of these parameters becomes clearer when we consider their impact on the radial

and transverse coordinates through small deviations around $\phi = 1, \alpha = 1$:

$$\begin{aligned} r'_{||} &= \alpha r_{||} - \frac{\phi - 1}{2} r_{||} + \mathcal{O}[(\phi-1)^2, (\alpha-1)^2], \\ r'_{\perp} &= \alpha r_{\perp} + \frac{\phi - 1}{2} r_{\perp} + \mathcal{O}[(\phi-1)^2, (\alpha-1)^2]. \end{aligned} \tag{7.4}$$

The α parameter produces the same effect on both $r_{||}$ and $r_{\perp}$, which corresponds to isotropic re-scaling. On the other hand, ϕ produces small changes that are directly opposite in $r_{||}$ versus $r_{\perp}$, which corresponds to anisotropy in ξ. We study the effect of these parameters on the correlation function in Section 7.2.4.

These quantities are an intermediate step between fitting the correlation function and constraining cosmological parameters. Having defined the scale parameters we will use, we turn our attention to the template and the application of these parameters.

7.2.2 Two-component full-shape parametrisation

We construct our model based on the separation of the BAO feature from the rest of the correlation. This is achieved by starting with a template isotropic linear power spectrum for an assumed fiducial cosmology, computed using CAMB (Lewis, Challinor, and Lasenby, 2000). This template power spectrum is decomposed into a peak (or wiggles) component, $P_{\rm peak}(k, z_{\rm eff})$, and a smooth (or no-wiggles) component, $P_{\rm smooth}(k, z_{\rm eff})$, using the method described in Kirkby et al. (2013). See Figure 2 of Kirkby et al. (2013) for a visualization of the peak and smooth components. The reason for this separation is that BAO is a clear feature that can be used as a standard ruler; it has been studied extensively, and we know that for the Lyα forest it is very robust when it comes to contaminants (e.g. Cuceu, Font-Ribera, and Joachimi, 2020). Therefore, we consider it advantageous to separate this feature, because it will make it easier to study and understand the information contained in the rest of the correlation (i.e. in the broadband), and how it is affected by contaminants (e.g. HCDs and continuum fitting).

The full transformed correlation in the original coordinates $(r_{||}, r_{\perp}, z)$ is given by:

$$\xi_{\rm full}(r_{||}, r_{\perp}, z) = \xi_{\rm peak}(r'_{||}, r'_{\perp}, z) + \xi_{\rm smooth}(r''_{||}, r''_{\perp}, z), \tag{7.5}$$

where the transformed coordinates of the peak component $(r'_{||}, r'_{\perp})$ are allowed to be different from the transformed coordinates of the smooth component $(r''_{||}, r''_{\perp})$. For comparison, in BAO analyses we would fix the smooth component: $(r''_{||}, r''_{\perp}) = (r_{||}, r_{\perp})$, whereas past full-shape analyses

did not use the peak-smooth decomposition, which would be equivalent to fixing the two sets of transformations to be the same: $(r''_{||}, r''_{\perp}) = (r'_{||}, r'_{\perp})$.

As we have two sets of coordinate transformations, we will need two sets of (ϕ, α) parameters. The AP effect distorts the entire correlation, and ϕ is meant to measure this anisotropy. Therefore, both the smooth and peak components are affected by ϕ in the same way. This means that we would ideally sample only one ϕ parameter that re-scales both components. However, as we wish to understand the cosmological value added by re-scaling the broadband, and also study how each parameter is affected by contaminants, we will keep them separate. Going forward, we will use the notation $\phi_{\rm s}$ for the smooth component and $\phi_{\rm p}$ for the BAO peak component. A measurement of ϕ corresponds to a measurement of:

$$\textbf{AP:}\ \ \phi(z) = \frac{F_{\rm AP}(z)}{F^{\rm fid}_{\rm AP}(z)} = \frac{D_{\rm M}(z)H(z)}{[D_{\rm M}(z)H(z)]_{\rm fid}}, \tag{7.6}$$

where the AP parameter is defined as the ratio of two distances $F_{\rm AP}(z) = D_{\rm M}(z)/D_{\rm H}(z)$, with $D_{\rm H}(z) = c/H(z)$.

On the other hand, the α parameter has different interpretations for the peak and smooth components. We not only need to account for the different expansion histories between the template and the data, but also for the features that set the scale we measure. We denote the parameter that isotropically re-scales the peak as $\alpha_{\rm p}$, and the equivalent parameter for the smooth component as $\alpha_{\rm s}$. In the case of the BAO peak, the relevant scale is the size of the sound horizon at the end of the drag epoch, $r_{\rm d}$. The isotropic scale of the peak component, $\alpha_{\rm p}$, corresponds to a measurement of:

$$\textbf{BAO:}\ \ \alpha_{\rm p}(z) = \sqrt{\frac{D_{\rm M}(z)D_{\rm H}(z)/r_{\rm d}^2}{[D_{\rm M}(z)D_{\rm H}(z)/r_{\rm d}^2]_{\rm fid}}}. \tag{7.7}$$

On the other hand, $\alpha_{\rm s}$ is harder to identify with one clear feature. The scale of matter-radiation equality ($k_{\rm eq}$) is a feature that contributes to $\alpha_{\rm s}$, and has successfully been used to constrain cosmology from the power spectrum (Baxter and Sherwin, 2021; Philcox et al., 2021). However, it is not clear that it is the only feature that contributes to the isotropic scale of the broadband. Furthermore, the effect produced by $\alpha_{\rm s}$ is very similar to that of the Lyα flux bias, which could lead to the two parameters being hard to disentangle. Therefore, we will not focus on the cosmological interpretation of $\alpha_{\rm s}$ in this work, and leave it to future studies to determine if this parameter could be useful.

In past galaxy full-shape analyses, there was no smooth/peak decomposition, and the isotropic

scale parameter was interpreted using $r_{\rm d}$ (e.g. Beutler et al., 2017). This is based on the approximation that most of the signal comes from the BAO peak. This means we can measure $\alpha_{\rm p}$ very precisely, but not $\alpha_{\rm s}$, so the measurement of a parameter $\alpha = \alpha_{\rm p} = \alpha_{\rm s}$ would be dominated by signal from the peak. By fitting two different parameters, we will be able to test this assumption.

Finally, for clarity, we show how our new set of parameters $(\phi_{\rm s}, \alpha_{\rm s}, \phi_{\rm p}, \alpha_{\rm p})$ would be treated in BAO and galaxy full-shape analyses:

$$\begin{aligned} \text{Standard BAO analyses: } & (\phi_{\rm s}, \alpha_{\rm s}) \text{ fixed to } (1,1), \\ \text{Galaxy full-shape: } & (\phi_{\rm s}, \alpha_{\rm s}) \text{ fixed to } (\phi_{\rm p}, \alpha_{\rm p}), \\ \text{Two-component Ly}\alpha & \\ \text{full-shape (this work): } & (\phi_{\rm s}, \alpha_{\rm s}, \phi_{\rm p}, \alpha_{\rm p}) \text{ all free.} \end{aligned} \tag{7.8}$$

In the rest of this work, we show the effects of $\phi_{\rm s}$ and $\alpha_{\rm s}$ on the correlation function, study the potential for measuring them using the Lyα forest and its cross-correlation with quasars, and show their usefulness for constraining cosmology. However, we leave it to future work to investigate how they interact with contaminants and potential systematic errors that may affect them.

7.2.3 Correlation function model

Our models for the Lyα forest auto-correlation and its cross-correlation with quasars follow du Mas des Bourboux et al. (2020), however we use simplified versions with no contaminants or distortion due to the effect of continuum fitting. The Lyα forest analyses of the auto and cross-correlation have so far only been done using models with linear-order perturbations. For Ly$\alpha\times$Lyα, a small scale non-linear correction term is also used, with the parameter values calibrated using simulations (Arinyo-i-Prats et al., 2015). On the other hand, full-shape analyses of QSO$\times$QSO typically use higher-order perturbation theory (e.g. Taruya, Nishimichi, and Saito, 2010). In this work, we restrict ourselves to linear-order perturbation theory. Therefore, the full anisotropic power spectra of Ly$\alpha\times$Lyα ($P_{\rm Ly\alpha}$), Ly$\alpha\times$QSO ($P_\times$) and QSO$\times$QSO ($P_{\rm QSO}$) are given by:

$$P_{\rm Ly\alpha}(k,\mu_{\rm k},z) = b_{\rm Ly\alpha}^2(1+\beta_{\rm Ly\alpha}\mu_{\rm k}^2)^2\, F_{\rm nl,Ly\alpha}^2 P_{\rm fid}(k,z), \tag{7.9}$$

$$\begin{aligned} P_\times(k,\mu_{\rm k},z) = {} & b_{\rm Ly\alpha}(1+\beta_{\rm Ly\alpha}\mu_{\rm k}^2)\times \\ & \times (b_{\rm QSO}+f(z)\mu_{\rm k}^2)\; F_{\rm nl,QSO}\, P_{\rm fid}(k,z), \end{aligned} \tag{7.10}$$

$$P_{\rm QSO}(k,\mu_{\rm k},z) = (b_{\rm QSO}+f(z)\mu_{\rm k}^2)^2\, F_{\rm nl,QSO}^2\, P_{\rm fid}(k,z), \tag{7.11}$$

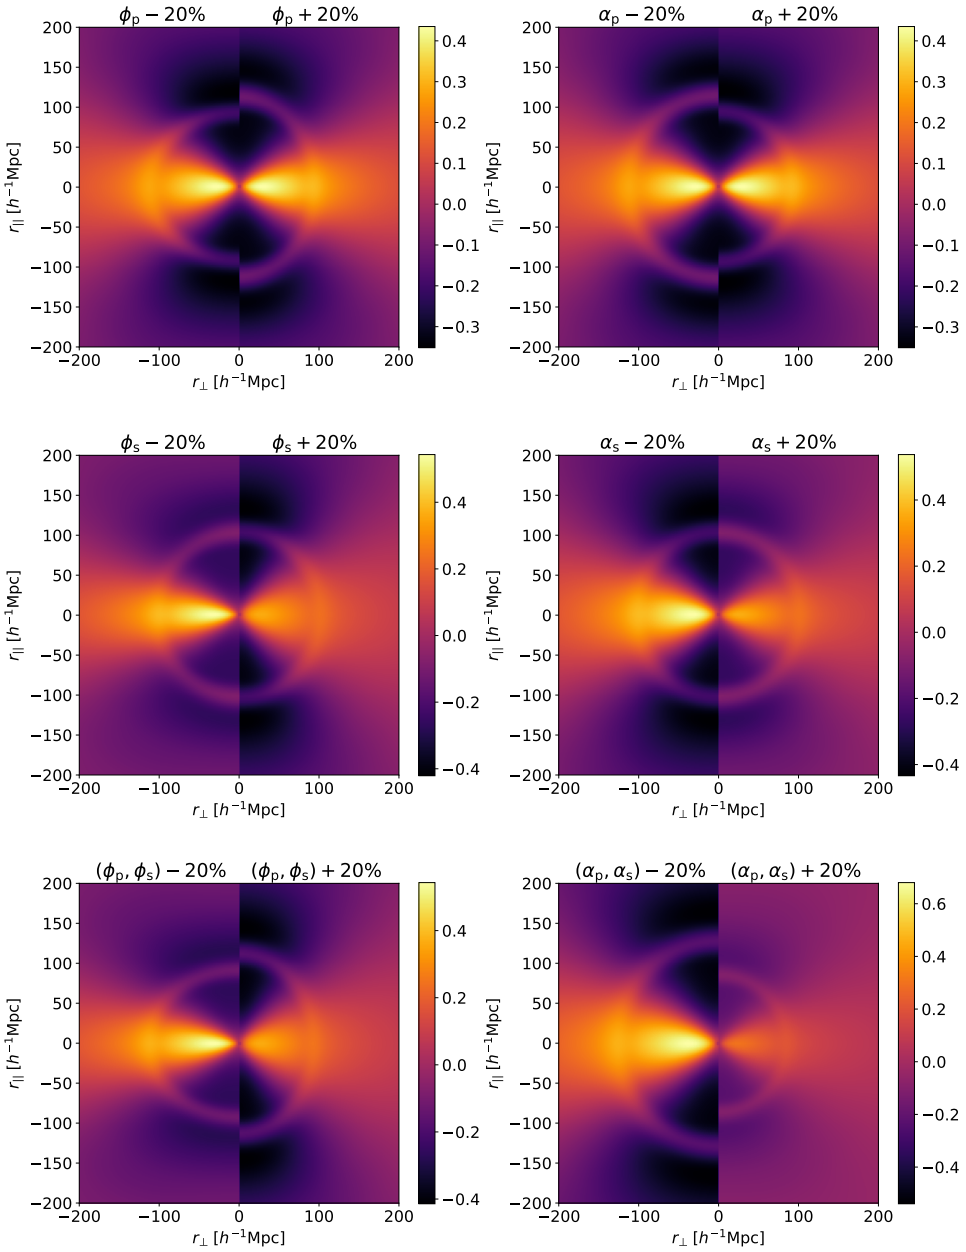

Figure 7.1. Contour plots of the Lyα forest auto-correlation function $(r^2\xi)$ in terms of the radial coordinates along and across the line of sight $(r_{||}, r_\perp)$. Each plot shows the correlation function computed using a smaller value of the given parameter on the left and a higher value on the right. The models here use linear theory (Eq. 7.9) with $\beta_{\mathrm{Ly}\alpha} = 1.669$. The left column shows the effect of the ϕ parameters which change the anisotropy of the correlation (the AP effect). The right column shows the effect of the α parameters which change the isotropic scale of the correlation. The top row shows the scale parameters for the peak component (ϕ_p and α_p). This is what BAO analyses measure. The middle row shows the parameters that re-scale the smooth component (ϕ_s and α_s). We aim to measure both the BAO parameters and the broadband parameters. Finally, the bottom row shows the effect of changing the parameters for the peak and smooth components at the same time ($\phi_\mathrm{s} = \phi_\mathrm{p}$ and $\alpha_\mathrm{s} = \alpha_\mathrm{p}$). This is what past spectroscopic galaxy clustering analyses (e.g. BOSS and eBOSS) measure when fitting the full shape of the correlation.

where $b_{\mathrm{Ly}\alpha}$ and b_{QSO} are the linear biases of the Lyα forest and quasars respectively, $f(z)$ is the logarithmic growth rate, and $\mu_{\mathrm{k}} = k_{||}/k$, with the wavenumber k, and its projection along the line of sight, $k_{||}$.

The RSD parameter of the Lyα forest is given by:

$$\beta_{\mathrm{Ly}\alpha} = \frac{b_{\eta,\mathrm{Ly}\alpha} f(z)}{b_{\mathrm{Ly}\alpha}}, \tag{7.12}$$

where $b_{\eta,\mathrm{Ly}\alpha}$ is the velocity divergence bias. As $b_{\eta,\mathrm{Ly}\alpha}$ and $f(z)$ always appear together in the Lyα forest RSD term, they are completely degenerate. Therefore, we use the parameter $\beta_{\mathrm{Ly}\alpha}$ to define the RSD term of the Lyα forest, and we marginalize over it instead of $b_{\eta,\mathrm{Ly}\alpha}$. This is meant to separate the degeneracies of different parameter combinations, and to clearly differentiate between nuisance parameters $(b_{\mathrm{Ly}\alpha}, b_{\mathrm{QSO}}, \beta_{\mathrm{Ly}\alpha})$ and the parameters of interest (ϕ, f). We also note that the symmetries of Equation 7.10 mean that when fitting only the cross-correlation, $b_{\mathrm{Ly}\alpha}$ is fully degenerate with b_{QSO}, and similarly $\beta_{\mathrm{Ly}\alpha}$ with $f(z)/b_{\mathrm{QSO}}$.

The small-scale non-linear correction for Lyα, $F^2_{\mathrm{nl,Ly}\alpha}$, is given by the model introduced by Arinyo-i-Prats et al. (2015). However, this has only been tested and applied to the Lyα forest auto-correlation, and not for the cross-correlation. Therefore, we only apply this term for Lyα×Lyα. On the other hand, the term $F_{\mathrm{nl,QSO}}$, which models the quasar non-linear velocities, is used for both the cross-correlation and the quasar auto-correlation. Following Percival and White (2009), this is given by:

$$F_{\mathrm{nl,QSO}}(k_{||}) = \sqrt{\frac{1}{1 + (k_{||}\sigma_{\mathrm{v}})^2}}, \tag{7.13}$$

where σ_{v} is a free parameter representing the rms velocity dispersion.

We also model the non-linear broadening of the BAO peak by applying the term $P_{\mathrm{nl,peak}}$ to the peak component of the power spectrum, $P_{\mathrm{peak}}(k, z_{\mathrm{eff}})$, following Eisenstein, Seo, and White (2007). This term is given by:

$$P_{\mathrm{nl,peak}} = \exp\left[-k^2_{||}\Sigma^2_{||}/2 - k^2_{\perp}\Sigma^2_{\perp}/2\right], \tag{7.14}$$

where $k_{\perp}$ is the projection of the wavenumber k across the line of sight, and the smoothing scales $(\Sigma_{||}, \Sigma_{\perp})$ are fixed to the values $(6.42, 3.26)h^{-1}$Mpc (Kirkby et al., 2013).

We use the `Vega` library[1] to compute model correlation functions using the same template power spectrum (and fiducial cosmology) as in du Mas des Bourboux et al. (2020). `Vega` is a

[1] https://github.com/andreicuceu/vega

new, improved version of the BAO fitter in the `picca`[2] library that was used in eBOSS Lyα BAO analyses.

7.2.4 Impact on the correlation function

We investigate how the parameters we introduced $(\phi_s, \alpha_s, \phi_p, \alpha_p)$ change the Lyα forest auto-correlation function using the model presented above. We show the effect produced by these parameters in Figure 7.1, using contour plots of the correlation function. For each plot, we show a model correlation computed using a 20% smaller value of a given parameter on the left and 20% higher value on the right, while the other parameters are kept fixed to one (scale parameters) or their best fit (nuisance parameters) from du Mas des Bourboux et al. (2020). Note that such changes are extreme, and chosen only to clearly showcase the effect of varying the parameters. The coordinate re-scalings we use are only approximations that work for values close to the template cosmology (parameter values around 1).

The first two rows of Figure 7.1 show the effect of the four parameters we introduced, with the ϕ parameters on the left and the isotropic scale parameters on the right. The top row shows the parameters that only affect the BAO peak, ϕ_p and α_p, while leaving the broadband component mostly unchanged. The former produces an anisotropy in the BAO scale (top left), leading to a different position of the peak along versus across the line of sight. The latter isotropically re-scales the BAO peak (top right). On the other hand, the two plots in the middle row show the parameters that only affect the smooth component, ϕ_s and α_s, while leaving the BAO peak unchanged. ϕ_s changes the anisotropy of the smooth component (middle left). Note that ξ is anisotropic even for $\phi_s = 1$ due to RSD. We will need to marginalize over this effect if we want to measure ϕ_s. The α_s parameter isotropically re-scales the smooth component (middle right) without affecting the position of the BAO peak.

Finally, the bottom row of Figure 7.1 shows the effect of re-scaling the smooth and peak components at the same time by fixing $\phi_s = \phi_p$ (bottom left) and $\alpha_s = \alpha_p$ (bottom right). This means that the peak and broadband are entangled, leaving a measurement of α harder to interpret. This is what past full-shape analyses of discrete tracers have measured, but using different parametrizations.

[2] https://github.com/igmhub/picca

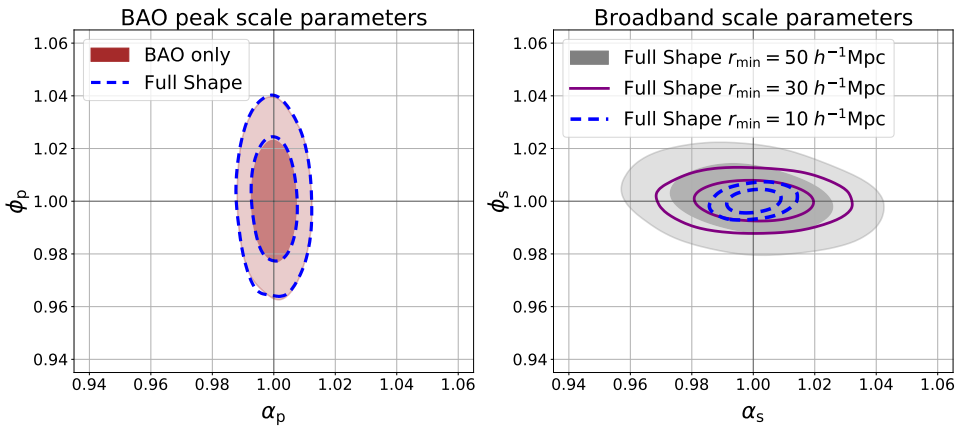

Figure 7.2. Forecast constraints on the four scale parameters in our two-component full-shape analysis from Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO. In the left panel, we compare the posterior distributions of the BAO peak parameters α_p and ϕ_p obtained using a BAO only analysis and our full-shape method. The very good agreement between the two shows that we can isolate the robust BAO information when performing a full-shape analysis. In the right panel, we show posterior distributions of the broadband scale parameters ϕ_s and α_s for different minimum separations used for the fits. This shows that we can obtain much better constraints on the Alcock-Paczynski parameter (ϕ) from the smooth component compared to those from the BAO peak.

7.3 AP Forecasts for the Lyα forest

We start our investigation of a potential full-shape analysis from the Lyα forest by analysing simulated correlation functions. We use these mock correlations to test our proposed two-component full-shape analysis, and forecast how well DESI will be able to measure the four scale parameters we introduced.

7.3.1 Mock data

We compute model correlation functions for Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO as described in Section 7.2.3. The models are computed using the best fit parameter values from eBOSS DR16 ($b_{\mathrm{Ly}\alpha}$=−0.117, $\beta_{\mathrm{Ly}\alpha}$=1.669, b_{QSO}=3.73, f=0.97, σ_v=6.86h^{-1}Mpc), except for the scale parameters, which are all set to equal one. We use these models as our simulated data. For the purposes of this work, we wish to perform a forecast analysis, and therefore we do not add noise to the fiducial data vector. Note that what we refer to as a forecast here is different from the usual Fisher forecasts commonly used in the literature (e.g. Tegmark, 1997; Tegmark, Taylor, and Heavens, 1997; Heavens, Jimenez, and Verde, 2014). We only analyse a single noiseless realization instead of

computing the Fisher information matrix.[3] This is because our main goal is to study the behaviour of our model, and its ability to extract more information from the Lyα correlations.

We use covariance matrices computed from mock data sets by Farr et al. (2020). These mocks were created using the `LyaCoLoRe` package,[4] which uses an initial Gaussian random field to simulate Lyα forest transmitted flux skewers and adds the relevant small scale power and RSD. The mocks were used to create full sky quasar catalogues containing ~ 3.7 million QSOs above redshift $z = 1.8$. These simulated data products were used to compute the Lyα forest auto-correlation function, its cross-correlations with QSOs, the QSO auto-correlation, and the relevant covariance matrices. In order to compute a covariance matrix relevant for DESI, we use the expected DESI survey area of 14000 square degrees, and assume it will measure roughly ~ 1.1 million QSOs above redshift $z = 1.8$ (DESI Collaboration et al., 2016). We then compute a factor that re-scales the covariance matrix ($f_{\rm cov}$) to match the expected DESI number density ($n_{\rm DESI}$) and area ($A_{\rm DESI}$):

$$f_{\rm cov} = \left(\frac{n_{\rm Mock}}{n_{\rm DESI}}\right)^2 \frac{A_{\rm Mock}}{A_{\rm DESI}}, \tag{7.15}$$

where $n_{\rm Mock}$ and $A_{\rm Mock}$ are the number density and area of the mock correlation computed by Farr et al. (2020), and the factor we compute is $f_{\rm cov} \simeq 4$. This factor is based on the fact that Lyα forest measurements are still limited by shot noise, and therefore the number density needs to be accounted for alongside the area, which accounts for cosmic variance. We also validate it by comparing our cosmological constraints with the forecasts from DESI Collaboration et al. (2016) in Section 7.3.3. The DESI simulated covariance matrix is then given by $C_{\rm DESI} = f_{\rm cov} C_{\rm Mock}$, based on the mock covariance, $C_{\rm Mock}$.

We assume that there is no cross-covariance between Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO (du Mas des Bourboux et al., 2017), as has been standard with Lyα BAO analyses so far. We use a Gaussian likelihood, and compute posterior distributions using the Nested Sampler `PolyChord`[5] (Handley, Hobson, and Lasenby, 2015a; Handley, Hobson, and Lasenby, 2015b). We use the recommended setup (live_points = 25$\times$ number of parameters, num_repeats = 3$\times$ number of parameters) when running `PolyChord`. When fitting each correlation independently, we sample the parameters: $\{\phi_{\rm p}, \alpha_{\rm p}, \phi_{\rm s}, \alpha_{\rm s}, b_{\rm Ly\alpha}, \beta_{\rm Ly\alpha}\}$ for the auto-correlation, while the cross-correlation has one extra parameter ($\sigma_{\rm v}$). For the cross, we do not sample the QSO bias and RSD parameters due to the degeneracies with the Lyα parameters (see Section 7.2.3). When performing joint fits, we

[3] Defined as the expected value of the Hessian matrix of the likelihood with respect to the parameters.
[4] https://github.com/igmhub/LyaCoLoRe
[5] https://github.com/PolyChord/PolyChordLite

also sample $b_{\rm QSO}$ and $f(z)$; however, we treat $f(z)$ as a nuisance parameter in this section, and only focus on measuring ϕ_s.

7.3.2 Scale parameters from a two-component full-shape analysis

When creating our method for a two-component full-shape analysis, our first goal was to preserve the robust BAO information that we normally measure by re-scaling only the peak component. In order to check if our method succeeded in isolating this information, we fit the mock data using a BAO type model where we fix the smooth component, and only re-scale the peak. We then compare the posterior distributions of the BAO peak scale parameters ($\phi_{\rm p}$ and $\alpha_{\rm p}$) to the posteriors obtained from the full-shape analysis. The results are shown in the left panel of Figure 7.2. We show the constraints for a joint analysis of Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO. We find that our method arrives at BAO measurements in very good agreement with classic BAO analyses, which means that by re-scaling the smooth component we do not influence the measurement of the position of the acoustic peak.

Our next goal for these forecasts is to understand the constraining power we have on the smooth component scale parameters, $\phi_{\rm s}$ and $\alpha_{\rm s}$. To this end, we consider a few different fitting strategies. As discussed above, before an actual measurement of these parameters, a full analysis of potential systematic errors needs to be performed. This study would inform the different analysis choices that need to be made in order to obtain robust measurements. One of these choices is the smallest scale that we fit. For past Lyα forest BAO analyses, this has been chosen to be $r_{\rm min} = 10\,h^{-1}$Mpc. This choice is not as important for BAO analyses because the BAO peak is a large-scale feature, and so, is not affected by small-scale contaminants. However, when attempting to measure scale parameters using the broadband component, these small scales have the potential to provide a lot of information. This is both because of the extra data points, and also because these data points at small separations have higher signal-to-noise. Therefore, we test a few different values of $r_{\rm min}$ that represent the range of possible options. We showcase the best case scenario where we are not affected by systematic errors all the way down to $10\,h^{-1}$Mpc, a worst case scenario where we have to cut the small scales and $r_{\rm min} = 50\,h^{-1}$Mpc, and an intermediate case where we cut to $r_{\rm min} = 30\,h^{-1}$Mpc. The lower value was chosen based on the value used by Lyα BAO analyses, however, it might be too optimistic given current understanding of the Lyα correlation functions (see Section 7.5 for discussion). On the other hand, the choice of the upper value was made because we might start to lose BAO information when removing scales above $50\,h^{-1}$Mpc (Kirkby et al., 2013).

The forecast broadband scale parameter results (again for Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO) are shown in the right panel of Figure 7.2 for different $r_{\rm min}$. Figure 7.2 also highlights the difference in constraining power between the AP parameter and the isotropic scale parameter. It shows that, using the BAO peak, we can obtain very good measurements of the isotropic scale parameter, $\alpha_{\rm p}$ (the 68% confidence region is at a precision of $\sim 0.5\%$). However, we do not have very good constraining power when it comes to the AP parameter, $\phi_{\rm p}$, for which the 68% confidence region is at a precision of $\sim 1.6\%$. This is in contrast to the AP measurement from the smooth component, where even in the worst case scenario the 68% confidence region is at a precision of $\sim 0.9\%$, and in the best case scenario it is at $\sim 0.3\%$. This shows the large potential gain in cosmological information from adding this AP measurement from the broadband.

7.3.3 Cosmological forecasts

We show the benefits of extracting more information from the Lyα forest 3D correlation functions by performing a simple cosmological analysis using the forecast measurements obtained above. We use a flat ΛCDM model, and we first model each of the measured parameters individually, in order to understand how each of them constrains cosmology. The cosmological interpretations of the scale parameters in terms of distances are given by Equations 7.6 and 7.7. Therefore, in order to complete our model, we just need the expressions for $D_{\rm M}$ and $H(z)$ in a flat ΛCDM cosmology. The comoving angular diameter distance is given by:

$$D_{\rm M}(z) = c\int_0^z \frac{{\rm d}z'}{H(z')}, \tag{7.16}$$

and the Hubble parameter is given by the Friedman equation:

$$\frac{H(z)^2}{H_0^2} = \Omega_{\rm m}(1+z)^3 + \Omega_\Lambda + \Omega_{\rm r}(1+z)^4. \tag{7.17}$$

In flat ΛCDM, the dark energy fraction can be computed from the matter and radiation fractional densities: $\Omega_\Lambda = 1-\Omega_{\rm m}-\Omega_{\rm r}$. We also model the radiation fraction assuming a CMB temperature $T_{\rm CMB} = 2.7255K$ (Fixsen et al., 1996; Fixsen, 2009), and a fixed neutrino sector [6]. This means the only free parameters in $H(z)$ are H_0 and $\Omega_{\rm m}$.

For the AP parameter, we have a ratio of distances: $D_{\rm M}/D_{\rm H}$ (Equation 7.6), which means the Hubble constant cancels out. Therefore, in flat ΛCDM, ϕ corresponds to a measurement of $\Omega_{\rm m}$.

[6] We use $N_{\rm eff} = 3.046$, with 2 massless species and one massive with $m_\nu = 0.06$ eV that contributes to $\Omega_{\rm m}$.

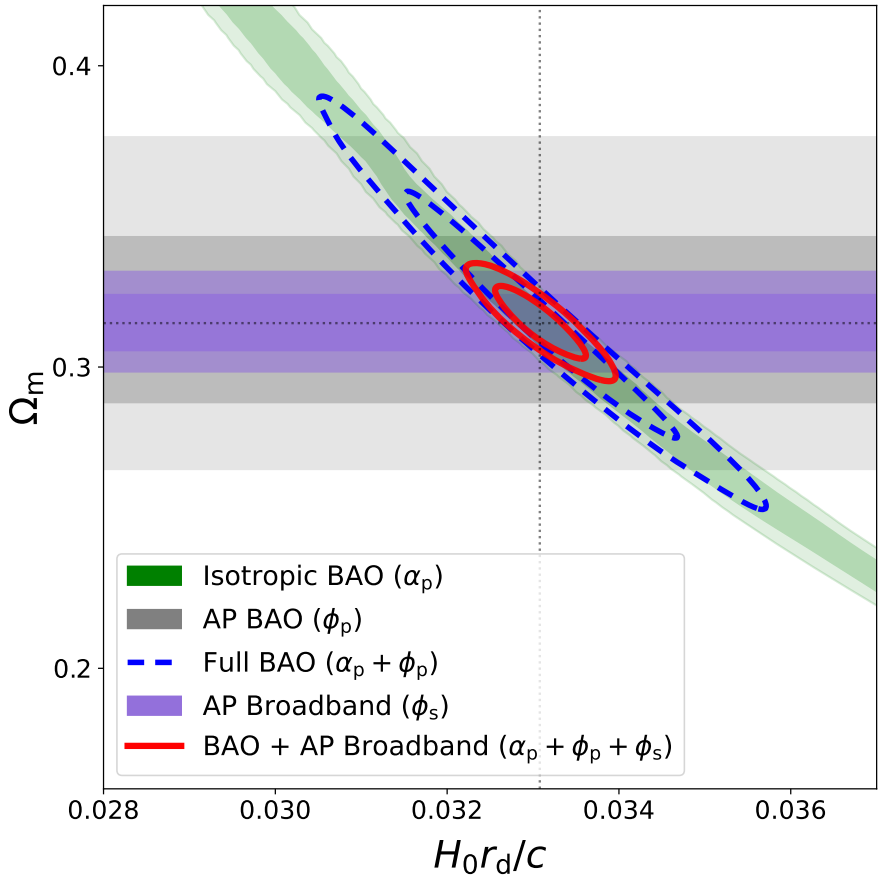

Figure 7.3. Forecast posterior distributions on cosmological parameters in flat ΛCDM using different scale parameter measurements from the Lyα forest correlation function. Measurements of the AP parameter (ϕ) only constrain the matter fraction Ω_m, while the isotropic BAO scale measures Ω_m and the combination $H_0 r_d$. The AP measurement from the broadband (ϕ_s) is significantly better compared to the one from the BAO peak (ϕ_p). Therefore, the improved Ω_m measurement leads to much tighter constraints when combined with the BAO measurement.

On the other hand, for α_p we have a product of distances divided by the scale of the sound horizon squared: $D_M D_H / r_d^2$. As each of the two distances has a factor of $1/H_0$, we are left with the product $H_0^2 r_d^2$, which means the two parameters are fully degenerate. Therefore, with α_p, we measure a combination of Ω_m and the product $H_0 r_d$.

We use the α_p, ϕ_p and ϕ_s measurements presented above to constrain the relevant cosmological parameters. For ϕ_s we use the result from the fit with $r_{min} = 30\ h^{-1}$Mpc, and we again use `PolyChord` to compute the posterior distributions. The constraints on α_p, ϕ_p and ϕ_s translate into measuring $H(z_{eff}) r_d$ and $D_M(z_{eff})/r_d$ with a precision of $\sim 0.5 - 0.6\%$ each. In contrast, the DESI Lyα BAO analysis is expected to measure $H(z_{eff}) r_d$ and $D_M(z_{eff})/r_d$ with a precision of

$\sim 0.9\%$[7] (DESI Collaboration et al., 2016). The cosmological parameter results using the individual measurements and their combinations are shown in Figure 7.3.

The constraint from the isotropic BAO measurement ($\alpha_{\rm p}$) leads to an elongated posterior with a strong degeneracy in the $\Omega_{\rm m} - H_0 r_{\rm d}$ space. This degeneracy is broken when combining with the $\Omega_{\rm m}$ constraint from $\phi_{\rm p}$ to obtain the usual anisotropic BAO measurement. However, as noted above, the AP measurement from the broadband is much better than the one measured from the peak. Therefore, by adding the $\phi_{\rm s}$ measurement to the BAO constraint, we can break the long correlation and obtain much better joint constraints. While the BAO measurements constrain $\Omega_{\rm m}$ and $H_0 r_{\rm d}$ with a precision of 8.9% and 3.3% (68% credible regions) respectively, adding the AP measurement from the broadband improves these constraints to 2.5% and 1.0%.

7.4 A JOINT ANALYSIS OF THE HIGH-z 3×2PT

In Sections 7.2 and 7.3 we focused on extracting more information from the full shapes of Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO through the AP parameter. We now turn our attention to the other source of cosmological information commonly used in full-shape analyses: redshift space distortions. In particular, we focus on the ability of joint analyses of the two Lyα correlations to obtain meaningful measurements from RSD, and on the potential of a joint analysis of the three high redshift two point (high-z 3×2pt) correlation functions: Ly$\alpha\times$Lyα, Ly$\alpha\times$QSO and QSO$\times$QSO.

7.4.1 CONTEXT

As the Lyα forest velocity divergence bias, $b_{\eta,{\rm Ly}\alpha}$, is fully degenerate with the logarithmic growth rate, $f(z)$, we have so far treated RSD as a nuisance that we need to marginalize over. In practice, RSD analyses are sensitive to the combination $f\sigma_8$, where σ_8 is the amplitude of matter perturbations in spheres of 8 Mpc/h. This means that Ly$\alpha\times$Lyα effectively measures the combinations $b_{{\rm Ly}\alpha}\sigma_8$ and $b_{\eta,{\rm Ly}\alpha}f\sigma_8$.

The Lyα-QSO cross-correlation could in theory be used to measure $f\sigma_8$. However, on its own it cannot constrain all the biases $(b_{{\rm Ly}\alpha}, b_{\eta,{\rm Ly}\alpha}, b_{\rm QSO})$ even for BAO analyses where we fix $f\sigma_8$ (see Section 7.2.3 and du Mas des Bourboux et al., 2020). On the other hand, a joint full-shape analysis of Ly$\alpha\times$QSO and Ly$\alpha\times$Lyα could help break these degeneracies and produce an $f\sigma_8$ constraint.

Another option for measuring $f\sigma_8$ at high redshift $(1.8 < z < 4)$ is to use the quasar auto-correlation, QSO$\times$QSO. The growth rate of structure was first measured from the quasar distribution

[7] We recover this precision by translating the measurements of $\alpha_{\rm p}$ and $\phi_{\rm p}$ to $H(z_{\rm eff})r_{\rm d}$ and $D_{\rm M}(z_{\rm eff})/r_{\rm d}$, which validates our approach of re-scaling the covariance matrix presented in Section 7.3.1

by the eBOSS collaboration using the SDSS DR14 data (Gil-Marín et al., 2018; Zarrouk et al., 2018; Hou et al., 2018). They performed full-shape analyses on both the 3D power spectrum and the 3D correlation function. With the last eBOSS analysis using SDSS DR16, these measurements have been updated and now provide a $\sim 10\%$ constraint on the growth rate at an effective redshift $z_{\rm eff} = 1.48$ (Hou et al., 2020; Neveux et al., 2020). The QSO sample contained $343,708$ quasars and spanned a redshift range of $0.8 < z < 2.2$. For comparison, DESI will measure about 1.7 million QSOs at $z < 2.1$ to be used as tracers only, and another 0.7 million at $z > 2.1$ to be used both as tracers and to measure the Lyα forest (DESI Collaboration et al., 2016).

The high redshift QSO×QSO measurement could be combined with the two Lyα forest correlations in a joint analysis. This could lead to improved $f\sigma_8$ constraints because of the information from the cross-correlation, and also due to the potential of the three correlations helping break parameter degeneracies. Therefore, our goal in this section is to study the potential of a high-redshift joint analysis of the three two-point (high-z 3×2pt) correlation functions: Lyα×Lyα, Lyα×QSO and QSO×QSO.

7.4.2 Methods

We use a template linear power spectrum with a fixed normalization, which is proportional to σ_8. The logarithmic growth rate, $f(z)$, and σ_8 are completely degenerate in linear theory (Percival and White, 2009), and therefore we are sensitive to the product $f(z)\sigma_8(z)$ for quasars and $b_{\eta,{\rm Ly}\alpha} f(z)\sigma_8(z)$ for the Lyα forest. As $b_{\eta,{\rm Ly}\alpha}$ is unknown, we will continue sampling over the $\beta_{{\rm Ly}\alpha}$ parameter, effectively treating the Lyα forest RSD term as a nuisance to be marginalized over.

We perform our analysis of the high-z 3×2pt using the two-component full-shape method we introduced in Section 7.2. For the Lyα forest auto and cross correlation, we use the same simulated data and covariance matrices as described in Section 7.3. For the QSO auto-correlation, we also use a covariance matrix computed by Farr et al. (2020), re-scaled to the DESI area and number density as described in 7.3.1. The QSO×QSO simulated correlation function is given by a fiducial model (no noise) following the best fit parameter values from du Mas des Bourboux et al. (2020), again with the scale parameters set to unity.

Our effective parameter vector for joint fits is given by: $\{\phi_{\rm p}, \alpha_{\rm p}, \phi_{\rm s}, \alpha_{\rm s}, f\sigma_8, b_{{\rm Ly}\alpha}\sigma_8, b_{\rm QSO}\sigma_8, \beta_{{\rm Ly}\alpha}, \sigma_{\rm v}\}$. When fitting individual correlations, we follow the approach we took in Section 7.3, of fixing the QSO bias and RSD terms for the cross-correlation.

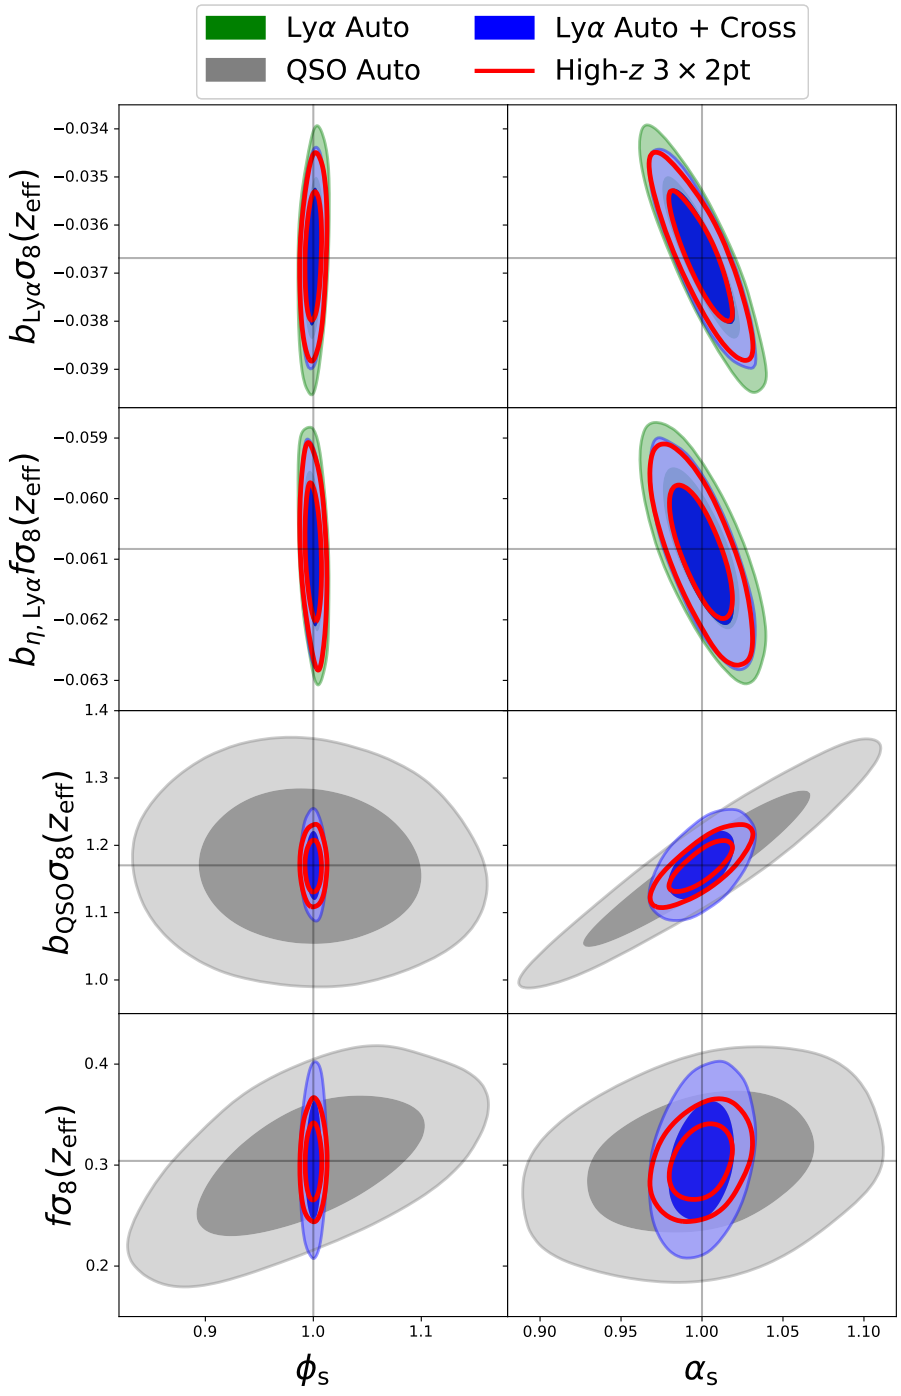

Figure 7.4. Posterior distributions of the Lyα forest auto-correlation (green), the Lyα auto + cross (blue), the QSO auto-correlation (gray), and the joint high-z 3×2pt analysis of: Ly$\alpha\times$Lyα, Ly$\alpha\times$QSO and QSO$\times$QSO (red). We use a minimum separation $r_{\mathrm{min}} = 30\,h^{-1}$Mpc. The first two rows show the parameters measured only by the Lyα forest, while the bottom two rows show parameters constrained only by the quasar distribution. Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO cannot constrain RSD individually, however, a joint full-shape analysis of both gives us an $f\sigma_8(z_{\mathrm{eff}} \simeq 2.3)$ constraint that rivals the one from the quasar auto-correlation.

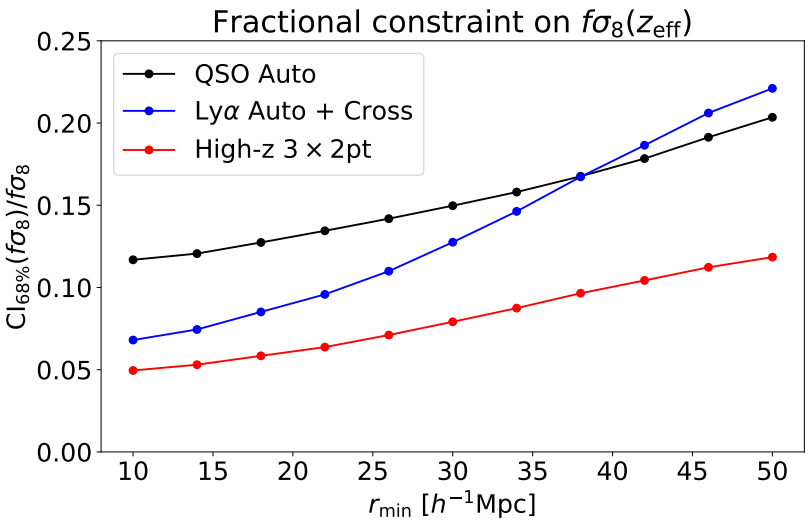

Figure 7.5. Forecast fractional constraints of the growth rate times the amplitude of fluctuations ($f\sigma_8$), as a function of the minimum separation ($r_{\rm min}$) used for the fits. The black line shows the precision for the quasar auto-correlation, while the blue line shows the precision for a joint full-shape analysis of Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO. The most precise and robust $f\sigma_8(z_{\rm eff}\simeq 2.3)$ measurement is obtained by jointly fitting all three correlation functions (red line).

7.4.3 Breaking parameter degeneracies

In our parametrization, the BAO parameters $(\phi_{\rm p},\alpha_{\rm p})$ are decoupled from the rest of the analysis. Therefore, as long as there is negligible cross-covariance between the different correlations, there is no benefit to $(\phi_{\rm p},\alpha_{\rm p})$ constraints from performing a joint analysis (i.e. fitting the correlations as one data vector). This has been the case so far with Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO in BOSS and eBOSS (e.g. Bautista et al., 2017; du Mas des Bourboux et al., 2017; du Mas des Bourboux et al., 2020), but the cross-covariance for DESI remains to be studied.

The benefits of performing the joint analysis should be most pronounced when it comes to the parameters we measure from the full-shape analysis: $\phi_{\rm s}$, $\alpha_{\rm s}$ and $f\sigma_8$. This is firstly due to the fact that these parameters are correlated with some of the nuisance parameters, and therefore, a joint analysis would allow us to disentangle their effects and lead to improved constraints. This is illustrated in Figure 7.4 (where we use a minimum separation $r_{\rm min}=30\,h^{-1}$Mpc). The top two rows show parameters that are only measured by the Lyα auto-correlation, while the bottom two rows show parameters that are only measured by the QSO auto-correlation. Note that the fact that $\phi_{\rm s}$ does not seem to be correlated with $b_{\eta,{\rm Ly}\alpha}f\sigma_8$ in Figure 7.4 for Ly$\alpha\times$Lyα is just due to the scale of the axes, which is set to display the weak QSO$\times$QSO constraint. AP and RSD are correlated, however, we do not expect these correlations to be the same for galaxies and the forest

because the two tracers cluster differently (e.g. $\beta_{\mathrm{Ly}\alpha}\sim1.67$ while $\beta_{\mathrm{QSO}}\equiv f/b_{\mathrm{QSO}}\sim0.26$).

The cross-correlation requires all four parameters: $(b_{\mathrm{Ly}\alpha}\sigma_8, b_{\eta,\mathrm{Ly}\alpha}f\sigma_8, b_{\mathrm{QSO}}\sigma_8, f\sigma_8)$, however, the system is degenerate. On the other hand, when we run a joint analysis of the cross-correlation with the Lyα auto-correlation (blue) we are able to constrain this system, because of the tight measurements of $b_{\mathrm{Ly}\alpha}\sigma_8$ and $b_{\eta,\mathrm{Ly}\alpha}f\sigma_8$ from Ly$\alpha\times$Lyα. This leads to a constraint on $f\sigma_8$ of 12.8% (68% confidence region), which is tighter than the one from the QSO auto-correlation of 15.4% (bottom left panel of Figure 7.4).

The second benefit of performing this joint analysis is due to the correlation between RSD and the AP effect. When measuring $f\sigma_8$, we have to marginalize over the AP parameter. If we knew the true background cosmology, i.e. for fixed AP, we would obtain much better measurements of the growth rate.[8] Even though the Lyα auto cannot directly measure the growth rate, it constrains the AP parameters (especially ϕ_{s}) very precisely. Therefore, including Ly$\alpha\times$Lyα in a joint analysis with QSO$\times$QSO can help break the correlation between RSD and AP, and improve the $f\sigma_8$ constraint. This is illustrated in the bottom left panel of Figure 7.4.

The joint high-z 3×2pt analysis appears to work well in breaking parameter correlations when it comes to $f\sigma_8$ and ϕ_{s}. However, that is not the case with α_{s}. While performing a joint analysis does lead to better constraints on this parameter, the posterior remains very correlated with all three biases (right column of Figure 7.4). This leaves a measurement of α_{s} prone to systematic errors, and therefore supports our decision from Section 7.2 not to focus on its cosmological interpretation.

So far in this section we used a minimum separation of $r_{\mathrm{min}}=30\,h^{-1}$Mpc to show how joint analyses help us break parameter degeneracies. However, we also want to test how these potential $f\sigma_8$ measurements would be affected if we could go to smaller scales (e.g. by having better models for non-linearities), or we had to cut even more data due to systematic effects on these scales. We show this in Figure 7.5, where we plot the marginalized fractional 68% credible regions on $f\sigma_8$ from the QSO auto, the Lyα auto + cross, and the joint analysis of all three correlations. We find that when we can include data at small scales ($r_{\mathrm{min}}\lesssim35\,h^{-1}$Mpc), the Ly$\alpha$ auto + cross combination (blue line) gives us better constraints compared to the QSO auto (black line). On the other hand, if we have to cut the small scales ($r_{\mathrm{min}}\gtrsim35\,h^{-1}$Mpc), the Ly$\alpha$ measurement degrades very fast, and the QSO auto becomes comparable and even slightly better at constraining $f\sigma_8$. This gives another advantage for performing a joint high-z 3×2pt analysis, because it leads to much more stable and robust measurements (red line). While the quasar auto-correlation can constrain

[8] Note that if we mainly wanted to test gravity, adding a tight prior on AP (e.g. from CMB measurements) would be justified in order to obtain the best possible $f\sigma_8$ measurement.

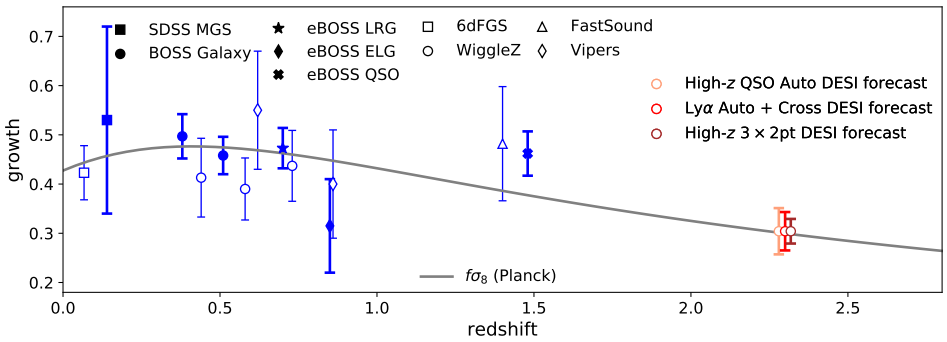

Figure 7.6. The growth rate times the amplitude of fluctuations ($f\sigma_8$) as a function of redshift. The gray line is the best-fit of CMB measurements from Planck. The blue points are some of the existing $f\sigma_8$ measurements. All of these measurements are at redshifts $z < 2$, with most of them at $z < 1$. The three points at high redshift are forecast constraints from DESI for the (high redshift) quasar auto-correlation, Lyα auto and cross-correlations, and the joint high-z 3×2pt analysis. Note that all three measurements are at the same effective redshift (given by the middle point), but are plotted at slightly different redshifts for visualization purposes.

$f\sigma_8(z_{\rm eff} \simeq 2.3)$ with a precision of $12-20\%$ depending on $r_{\rm min}$, the high-z 3×2pt analysis can achieve a precision of $5-12\%$.

We have also checked how the two-component full-shape approach affects our results by comparing it with the approach usually taken in galaxy full-shape analyses of fitting the full correlation as one component (no peak/smooth decomposition). The $f\sigma_8$ constraints are larger when sampling four parameters (our two-component approach) versus two parameters (the one-component approach used in galaxy full-shape analyses). This is to be expected as the model has more degrees of freedom. However, the effect is very small when it comes to the high-z 3×2pt constraints. We found that using a value of $r_{\rm min} = 30\,h^{-1}$Mpc, we obtain a precision of 7.9% on $f\sigma_8$ with the two-component approach, while with the one-component approach we obtain a precision of 7.6%. This does not significantly affect our conclusions in this work, but the two-component approach might be more advantageous when the effects of contaminants are studied, as it decouples the peak from the broadband (see Section 7.5).

Finally, in Figure 7.6 we emphasize how useful a full-shape high-z 3×2pt analysis would be. We show in blue some of the current $f\sigma_8$ measurements from different surveys (Ross et al., 2015; Alam et al., 2017; eBOSS Collaboration et al., 2021; Beutler et al., 2011; Blake et al., 2012; Okumura et al., 2016; Pezzotta et al., 2017). All of these measurements are at redshifts $z < 2$, with most of them at $z < 1$. The three points on the right show our DESI forecasts of $f\sigma_8(z_{\rm eff})$ at an effective redshift $z_{\rm eff} \simeq 2.3$. We use a conservative $r_{\rm min} = 30\,h^{-1}$Mpc. This analysis would allow

us to study the growth rate of cosmic structures at higher redshifts than ever before.

The results in this section show the potential of a joint full-shape analysis of the three correlation functions: Ly$\alpha\times$Lyα, Ly$\alpha\times$QSO and QSO$\times$QSO, when it comes to measuring RSD and the AP effect. The next steps required for such an analysis are to improve the model by adding contaminants and better non-linear models, and to study the potential systematic errors that would affect this measurement, especially on the Lyα forest side where a full-shape analysis of the 3D correlation function has never been done. We discuss these in more detail in the next section.

7.5 Discussion and next steps

In this work, we have shown the potential for extracting more cosmological information from the Lyα forest 3D auto-correlation function and its cross-correlation with quasars. We took the template fitting approach, where we use a template power spectrum to measure a few physically meaningful quantities that are easy to interpret and translate to cosmological constraints given some model. In our case, these quantities are the anisotropic scale parameter (ϕ), the isotropic scale of the BAO peak ($\alpha_{\rm p}$) and the growth rate times the amplitude of fluctuations in spheres of 8 h^{-1}Mpc ($f\sigma_8$). This approach should simplify the study of the impact of contaminants because we only have to deal with a few parameters whose effect we understand very well.[9] Such a study is required before a full-shape analysis of the Lyα forest correlation functions is performed on real data, however, it is outside the scope of this work. Here we wish to briefly go over the most important contaminants, and mention what we can do to minimize their impact. In particular, the most relevant contaminants for the measurement of AP and RSD are those that introduce anisotropies.

High column density (HCD) systems are a significant contaminant for the Lyα forest due to their broad absorption profile and long damping wings (Font-Ribera and Miralda-Escudé, 2012; Rogers et al., 2018). However, they also trace the underlying density field, which means they can add extra signal if modelled correctly. In past BOSS and eBOSS analyses, large damped Lyα systems (DLA) that could be identified were masked (e.g. Bautista et al., 2017; du Mas des Bourboux et al., 2020). However, clustering measurements could potentially be biased by masking part of the spectrum, as the mask is correlated to the density field. This was not a problem for BAO analyses, but its impact on a full-shape analysis needs to be tested. On the other hand, the small HCDs were left in the data and had to be included in the model. Rogers et al., 2018

[9] This is in contrast to a direct fit of cosmological parameters, where a study of contaminants would be much harder. This is due to the larger parameter space, but also because it is harder to identify and separate the effects of these parameters on the correlation function.

showed that HCDs can be successfully modelled down to the smallest scale considered in this work ($\sim 10\,h^{-1}$Mpc), by using a simple model in linear theory, with a separate bias and RSD parameter, convolved with Voigt profiles for the damping wings.

The Lyα forest auto and cross correlation functions are contaminated by metal transitions with rest-frame wavelength close to that of the Lyα transition, that add correlations between themselves and the Lyα forest or quasars. They are also contaminated by metal lines that are further away in rest-frame wavelength through their own auto-correlation. These metal lines have been successfully modelled for BOSS and eBOSS (Bautista et al., 2017; du Mas des Bourboux et al., 2017; de Sainte Agathe et al., 2019; Blomqvist et al., 2019; du Mas des Bourboux et al., 2020) by adding extra correlations with the same form as Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO, and with their own bias and RSD parameters. However, these still need to be tested at DESI-level precision, and for the full-shape analysis we also have to test how the metal lines affect the measurement of the AP parameter and RSD.

Another important source of contamination are QSO redshift errors, which could introduce a systematic bias if not modelled correctly. Non-linear peculiar velocities also have big impact on the anisotropy because they create fingers of god. For Lyα forest analyses (and in this work), these two effects have been modelled using simple damping terms with a Lorentzian or Gaussian profile based on Percival and White (2009). For a full-shape analysis, we might need to use more complex models, as was done for past quasar auto analyses (e.g. Hou et al., 2020; Neveux et al., 2020). Quasar radiation effects (also known as the transverse proximity effect) are also an important source of contamination for the Lyα-quasar cross-correlation. This is because the quasar radiation increases the ionization fraction in the surrounding gas, leading to less Lyα forest absorption (Font-Ribera et al., 2013). This effect has been modelled analytically and was shown to not have a significant impact on BAO analyses (du Mas des Bourboux et al., 2017; du Mas des Bourboux et al., 2020); however, this needs to be tested for a full-shape analysis as well.

The final effect we consider is the fitting of the quasar continuum, which removes power on scales larger than the size of the forest. This produces a distortion in the measured correlations along the line of sight, and therefore introduces another source of anisotropy. This has been successfully modelled through a distortion matrix (Bautista et al., 2017) for BOSS and eBOSS. A similar approach could be sufficient for a full-shape analysis using DESI, but this needs to be tested.

All the contaminants presented here have been studied before and are modelled in existing Lyα BAO analyses. However, what still needs to be understood is how they interact with the new

parameters we wish to study (ϕ_s and $f\sigma_8$). Additionally, du Mas des Bourboux et al. (2020) found that adding broadband polynomials to the model can improve the fit of the correlations, which could point to contaminants that are not modelled well enough, or new effects that have not been considered. The addition of these polynomials was shown not to have a significant impact on BAO measurements, but they cannot be used for full-shape analyses because we want to extract broadband information, not marginalize over it. Therefore, a careful analysis on the impact of contaminants on AP and RSD measurements needs to be performed in order to determine if and on what scales current models are appropriate for a full-shape analysis of the Lyα forest correlations. Furthermore, an analysis of potential systematic errors would inform the decisions related to which scale-parameters to sample (e.g. whether to have two separate ϕ parameters). We also mention that even in the worst-case scenario where we have to cut the small scales due to some significant systematic bias, we have shown that a full-shape analysis of the high-z 3×2pt could still lead to state of the art cosmological measurements at redshifts $1.8 < z < 4$.

7.6 Conclusions

The Lyman-α (Lyα) forest 3D auto-correlation function (Ly$\alpha\times$Lyα) and its cross-correlation with the quasar (QSO) distribution (Ly$\alpha\times$QSO) are currently some of the best cosmological probes of the Universe at redshifts $1.8 < z < 4$. However, so far they have only been used to measure the BAO scale. In this work, we proposed to expand the cosmological information extracted from these statistics by fitting the full shape of these correlations in order to measure the Alcock-Paczynski (AP) parameter.

In Section 7.2 we introduced our model for fitting the correlation function using a two-component approach, where we decomposed the template power spectrum into a peak component which contains the BAO information, and a smooth component. We then re-scaled the two components independently in order to decouple the measurement of the BAO peak from the rest of the analysis. In Section 7.3 we studied the potential for measuring the AP effect from the broadband of the Lyα forest correlations. We used simulated correlation functions and mock DESI covariance matrices within a simple linear model with no contaminants. We showed that our two-component full-shape method successfully isolates the measurement of the BAO peak by comparing it to a BAO only analysis. Furthermore, we showed that using this idealized approach, a joint full-shape analysis of Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO from DESI could measure the AP parameter at an effective redshift $z_{\rm eff} \simeq 2.3$ with a precision of $0.3\% - 0.9\%$ (68% credible regions). Compared to the

expected DESI Lyα BAO constraint on AP, which is $\sim 1.6\%$, we were able to obtain roughly two to four times better precision. In Section 7.3.3 we showed how this measurement would help constrain cosmological parameters in a flat ΛCDM model. In the conservative case where we fit to a smallest scale of 30 h^{-1}Mpc, the inclusion of the AP measurement from the broadband gives us roughly three times better precision on the relevant cosmological parameters compared to the BAO measurement.

In Section 7.4, we studied the potential for measuring the logarithmic growth rate times the amplitude of fluctuations in regions of 8 h^{-1}Mpc ($f\sigma_8$) at high redshift ($1.8 < z < 4$) using the DESI Lyα forest and quasar position measurements. An $f\sigma_8$ measurement at redshifts $z \gtrsim 1.6$ is unprecedented. Neither the Lyα auto-correlation nor the Lyα-QSO cross-correlation can constrain $f\sigma_8$ independently, due to a degenerate system of parameters. However, in Section 7.4.3 we showed that their combination (Ly$\alpha\times$Lyα + Ly$\alpha\times$QSO) is able to break these parameter degeneracies and obtain a measurement of $f\sigma_8(z_{\rm eff})$ at an effective redshift $z_{\rm eff} \simeq 2.3$. This joint analysis was able to obtain constraints of $7\% - 22\%$ (68% credible regions) depending on the minimum separation used. For comparison, with the high redshift quasar auto-correlation (QSO$\times$QSO) from DESI, we were able to obtain a precision of $12\% - 20\%$. Furthermore, we showed that combining the two Lyα correlations with the quasar auto-correlation in a joint analysis of the three high-redshift two-point correlation functions (high-z 3×2pt) would give us the most precise and robust measurement of $f\sigma_8$ at these redshifts. We found that a high-z 3×2pt analysis of the full DESI data could be able to measure $f\sigma_8(z_{\rm eff} \simeq 2.3)$ with a precision of $5\% - 12\%$, depending on the minimum separation used.

In this work, we have shown how to extract more information from the 3D distribution of the Lyα forest through the AP parameter. We have also shown it is possible to measure $f\sigma_8$ through a joint full-shape analysis of Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO. While the DESI Lyα BAO analysis is expected to measure $H(z_{\rm eff})r_{\rm d}$ and $D_{\rm M}(z_{\rm eff})/r_{\rm d}$ with a precision of $\sim 0.9\%$, adding the AP measurement from the broadband could give us constraints of $\sim 0.5\%$. On the other hand, performing a high-z 3×2pt analysis would allow us for the first time to measure $f\sigma_8$ at high redshift.

PART III

CONCLUSION

CHAPTER 8

CONCLUSIONS AND OUTLOOK

In this final part, I will summarize the work described in Part II and outline plans for ongoing and future work. I will also talk about some of the new results that appeared after the publication of the work described in past chapters.

8.1 THE HUBBLE TENSION AND BAO

In Chapter 5, we used baryon acoustic oscillations (BAO) data and deuterium abundance measurements combined with big bang nucleosynthesis (BBN) to constrain the Hubble constant, H_0. We also studied the apparent tension between past BAO measurements using the Lyman-α (Lyα) forest and those using the distribution of galaxies (we call this Galaxy BAO). Finally, we performed forecasts of this BAO+BBN method for the Dark Energy Spectroscopic Instrument (DESI).

The main results of this work are the two H_0 constraints, obtained using two different BBN priors:

$$H_0 = 67.6 \pm 1.1 \text{ km s}^{-1}\text{Mpc}^{-1} \text{ (theoretical prior)}, \tag{8.1}$$

$$H_0 = 68.1 \pm 1.1 \text{ km s}^{-1}\text{Mpc}^{-1} \text{ (empirical prior)}. \tag{8.2}$$

The two different priors come from one reaction rate in the BBN calculation. The laboratory measurements (on which the empirical estimate is based) have a large uncertainty, so we also use the prior based on a theoretical estimate. We found that the two result in very similar H_0 constraints.

On the other hand, we have also shown that for DESI this will no longer be the case, and the choice of prior could be very important. This problem has since been solved due to new measurements from the Laboratory for Underground Nuclear Astrophysics (LUNA), which greatly improved constraints on the relevant reaction cross-section (Mossa et al., 2020a; Mossa et al., 2020b). The resulting $\Omega_b h^2$ constraint improves on the old empirical prior we used, and is also consistent with Planck cosmic microwave background (CMB) measurements.

Since the publication of the work in Chapter 5 (Cuceu et al., 2019), which was based on eBOSS DR14 data, the eBOSS collaboration published results using the final eBOSS data release (SDSS DR16). These results include BAO measurements using luminous red galaxies (Gil-Marín et al., 2020; Bautista et al., 2020), emission line galaxies (Tamone et al., 2020; de Mattia et al., 2021), quasars (Hou et al., 2020; Neveux et al., 2020) and the Lyα forest (which includes both the auto and cross-correlation with quasars; du Mas des Bourboux et al., 2020). These results along with redshift space distortions (RSD) measurements are used to constrain cosmological parameters across a variety of models in eBOSS Collaboration et al. (2021).

The first result relevant to the work presented here is that the difference between Lyα and galaxy BAO measurements shrank even more. In Section 5.2 we found a 2.5σ difference between the two in DR12, and 1.9σ in DR14. This tension is generally reported in terms of a tension between Lyα BAO and Planck CMB measurements. However, the two ways of framing it are similar because the BOSS DR12 result, which is the biggest source of constraining power in galaxy BAO, is consistent with Planck. du Mas des Bourboux et al. (2020) reports the difference between Lyα BAO and Planck to now be at $\sim 1.5\sigma$ using a Gaussian approximation. Furthermore, in du Mas des Bourboux et al. (2020), we performed a study of the evolution of these BAO results between DR12 and DR16, and found that these shifts are consistent with random fluctuations.

Our BAO+BBN analysis was updated with the new DR16 data in eBOSS Collaboration et al. (2021). The new measurement of the Hubble constant is $H_0 = 67.35 \pm 0.97$ km s^{-1} Mpc^{-1}(using the empirical BBN prior), which is consistent with both of our results. Other new results that are consistent with Planck include CMB measurements from the Atacama Cosmology Telescope (ACT; Aiola et al., 2020) and the South Pole Telescope (SPT; Dutcher et al., 2021), and a full-shape direct fit of the BOSS galaxy power spectra combined with BAO and BBN (Philcox et al., 2021).

On the other hand, the newest distance ladder result gives a value of $H_0 = 73.0 \pm 1.4$ km s^{-1} Mpc^{-1}(Riess et al., 2021), in $\sim 4.2\sigma$ tension with Planck. New local measurements of H_0 using the tip of the red giant branch (TRGB) to build a distance ladder have also appeared from the Carnegie-Chicago Hubble Program (Freedman et al., 2019; Freedman et al., 2020). Interestingly,

this constraint is currently $H_0 = 69.6 \pm 1.9$ km s^{-1} Mpc^{-1}, in between the Planck and SH0ES measurements. However, other TRGB analyses find higher values that are more consistent with the distance ladder using Cepheids (Reid, Pesce, and Riess, 2019; Soltis, Casertano, and Riess, 2021). Many other H_0 measurements have also appeared in the literature in recent years. For more complete recent reviews, see Di Valentino et al. (2021) and Shah, Lemos, and Lahav (2021).

At this point, the Hubble tension has fully captured the attention of the cosmology community and beyond. There have been many new measurements, either independent of Planck and SH0ES or reanalysing the two data sets with new methods. However, the tension still stands. On the distance ladder side, there are still questions about the photometric calibration or potential astrophysical bias (e.g. from dust extinction, metallicity dependence etc.). On the CMB side (and other probes that use early Universe measurements), most discussion now focuses on the reliance on the ΛCDM model. While there have been many proposals for a theoretical solution by either extending or replacing ΛCDM, there is no promising candidate so far. For recent reviews of possible solutions, see Di Valentino et al. (2021) and Schöneberg et al. (2021). Therefore, the hunt for the value of the Hubble constant is still ongoing.

8.2 BAO FROM THE LYα FOREST

In Chapter 6, we performed a Bayesian analysis of the 3D Lyα forest correlation functions measured by eBOSS DR14. Using a sample of 100 Monte Carlo simulations, we studied the differences between the frequentist framework used by past BOSS and eBOSS analyses and the Bayesian approach. We tested maximum likelihood estimation (MLE), with uncertainties derived from a Gaussian approximation, and found that this approach fails to fully capture the uncertainty for current data sets.

We also tested the method of performing a grid scan over the two BAO scale parameters (i.e. Profile Likelihood), and computed both frequentist confidence intervals and Bayesian credible regions. We found that when it comes to measuring BAO, these two agree very well both between themselves and with the posterior distribution computed using the Nested Sampler `PolyChord`. However, `PolyChord` not only measures the BAO scale, but the full posterior distribution and its normalization (the Bayesian evidence) while having a similar computational cost as the scan.

The main outcome of this work is the introduction of Bayesian sampling in the community code `picca`. This has already been used in the eBOSS DR16 Lyα analysis (du Mas des Bourboux et al., 2020), Chapter 7 (Cuceu et al., 2021), and other ongoing projects within the DESI collaboration. A

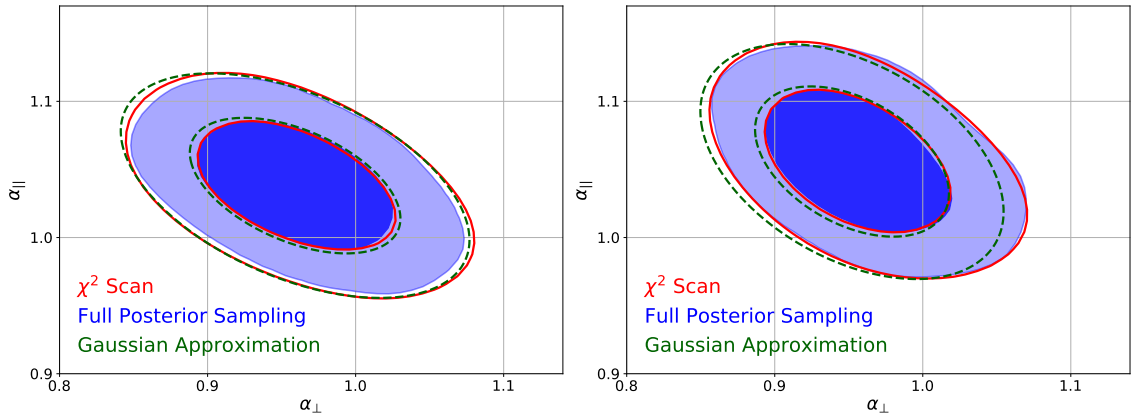

Figure 8.1. Comparison of eBOSS DR16 Lyα BAO results using different methods. Results from the Lyα auto-correlation are on the left, and those from its cross-correlation with quasars are on the right. The frequentist χ^2 scan, also known as the Profile Likelihood, is compared to Bayesian full posterior sampling and the Gaussian approximation. From du Mas des Bourboux et al. (2020).

personal contribution that is relevant to this work, was a test of DR16 Lyα BAO measurements using different methods. This is very similar to the tests on mocks from Section 6.3, and the results from du Mas des Bourboux et al. (2020) are shown in Figure 8.1. We found that in the case of DR16, the three methods agree very well.

After the publication of the work in Chapter 6, I started working on a new library for modelling and fitting Lyα correlation functions. This became the `Vega` package[1] and was used to perform the work from Chapter 7. It also uses the Nested Sampler `PolyChord` to compute posterior distributions, and has all the tools of the old fitting library from `picca`. This package along with the lessons learned in Chapter 6 will be used in the upcoming BAO analyses from DESI.

It may seem that improved data sets will inevitably lead to posteriors that are closer and closer to Gaussianity, thus lessening the importance of the statistical methods used to measure BAO.[2] However, with DESI we will most likely have enough data to perform BAO measurements in different redshift bins (instead of just one). This has the advantage of constraining cosmology at different times that can complement each other. However, it also means that BAO parameters will not be as well constrained, especially in the low and high redshift bins. Therefore, the choice of statistical method will still play an important role.

Beyond BAO, there is also the potential for a full-shape analysis of the Lyα forest correlation

[1] https://github.com/andreicuceu/vega

[2] One caveat for this is that with increased statistical precision comes greater risk of systematic effects from both known unknowns and unknown unknowns. Ways of dealing with these include the use of hyper-parameters that can result in highly non-Gaussian posteriors (see e.g. Lahav et al., 2000; Hobson, Bridle, and Lahav, 2002; Bernal and Peacock, 2018).

functions, which we discussed in Chapter 7. This type of analysis would extend the set of interesting parameters beyond the two BAO scale parameters, meaning that scanning methods may become computationally infeasible. The new parameters are also likely to be correlated with the other nuisance parameters (see Section 7.4), requiring computation of the full posterior in order to study them.

8.3 Beyond BAO with the Lyα forest

In Chapter 7, we studied the potential of performing a full-shape analysis of the Lyα forest 3D correlation functions in order to extract more cosmological information. We started with the Alcock-Paczynski (AP) effect. We already use some of this information when we measure anisotropic BAO. However, the AP effect affects the entire correlation. Our approach is different from the one found in the literature, as we still separate the peak and the smooth components. This allows us to preserve the robustness of BAO measurements, while also gaining extra information from a broader range of scales. By performing an idealized DESI forecast, we showed that even in a conservative case, AP measurements from the full-shape would significantly improve the constraining power of BAO only results.

Besides the AP effect, full-shape analyses are also used to measure redshift space distortions (RSD) through the parameter combination $f\sigma_8$. For the Lyα forest flux, there is an extra velocity divergence bias that is degenerate with the growth rate f (Slosar et al., 2011; McDonald, 2003; Givans and Hirata, 2020; Chen, Vlah, and White, 2021). Therefore, without some knowledge of this bias, we cannot extract the cosmological information from RSD. However, the Lyα-QSO cross-correlation (Ly$\alpha\times$QSO) has a biased term from Lyα and an unbiased term from QSOs. The two terms are degenerate, but we can measure the first one from the auto-correlation (Ly$\alpha\times$Lyα). Therefore, a joint analysis of the Lyα auto and the cross-correlations could constrain $f\sigma_8$. We have shown forecasts of such a measurement in Section 7.4.

A full-shape analysis of the two Lyα forest 3D correlations (Ly$\alpha\times$Lyα and Ly$\alpha\times$QSO) would not only give us the best high-redshift ($2 < z < 4$) cosmological constraints through the AP effect, but could also provide the first $f\sigma_8$ measurements above redshift 2. On the other hand, the quasar auto-correlation (QSO$\times$QSO) has already been used by the eBOSS collaboration to measure both effects (Gil-Marín et al., 2018; Zarrouk et al., 2018; Hou et al., 2018). These measurements are at lower effective redshift compared to the Lyα forest, and include quasars in the range $0.8 < z < 2.2$. However, with DESI the QSO auto-correlation could be measured in the same redshift range as the

forest, in order to perform a full-shape analysis. Therefore, we also studied the potential of a joint high redshift analysis of the three two-point correlation functions (high-z 3×2pt).

By breaking parameter correlations between the parameters of interest (AP, $f\sigma_8$) and nuisance parameters (e.g. biases), a high-z 3×2pt analysis could provide the most precise cosmological measurements at high redshift ($2 < z < 4$). However, such a measurement still requires some work. In Section 7.5, we discussed some of the potential systematic errors that could affect such a full-shape analysis. While we have models for most known relevant contaminants, these models have not been tested in the context of a full-shape analysis, or to the precision required by DESI. In the final section, I will present some of the ongoing and future work that is necessary to achieve this type of analysis.

8.4 Ongoing and future work

The focus of my current and near future work is on cosmology from the 3D Lyα forest correlation functions. This has two components. The first is part of a group effort in the DESI collaboration to perform a BAO analysis using the first year of data (Y1). The second is to perform the full-shape analysis on real data from eBOSS and DESI.

The DESI survey officially began in May 2021, after a short period of survey validation (SV). At the moment of writing, it has already observed about 0.33 million quasars. Out of these, we expect roughly one third to be at redshifts $z > 2.1$, where we can measure the Lyα forest (DESI Collaboration et al., 2016). For comparison, in the final eBOSS analysis, we used Lyα forests from roughly 0.21 million quasars (du Mas des Bourboux et al., 2020). Furthermore, the surveys of the other tracers are also progressing very well, with roughly 0.53 million LRGs, 0.33 million ELGs, and 1.59 million bright galaxies whose redshift has been measured at the time of writing. This means DESI is well on its way towards becoming the largest spectroscopic galaxy survey.

8.4.1 Preparing for the first DESI analysis

Measuring the BAO feature using the Lyα forest correlations is one of the DESI Key Projects for Y1. All key project analyses that will lead to cosmological results in DESI have to be blinded. This means that the collaboration does not know the true results (e.g. BAO position) while the analysis choices are still being decided. A good blinding strategy prevents scientists from introducing their own personal biases, while also interfering as little as possible with the analysis (see e.g. Muir et al., 2020; Brieden et al., 2020).

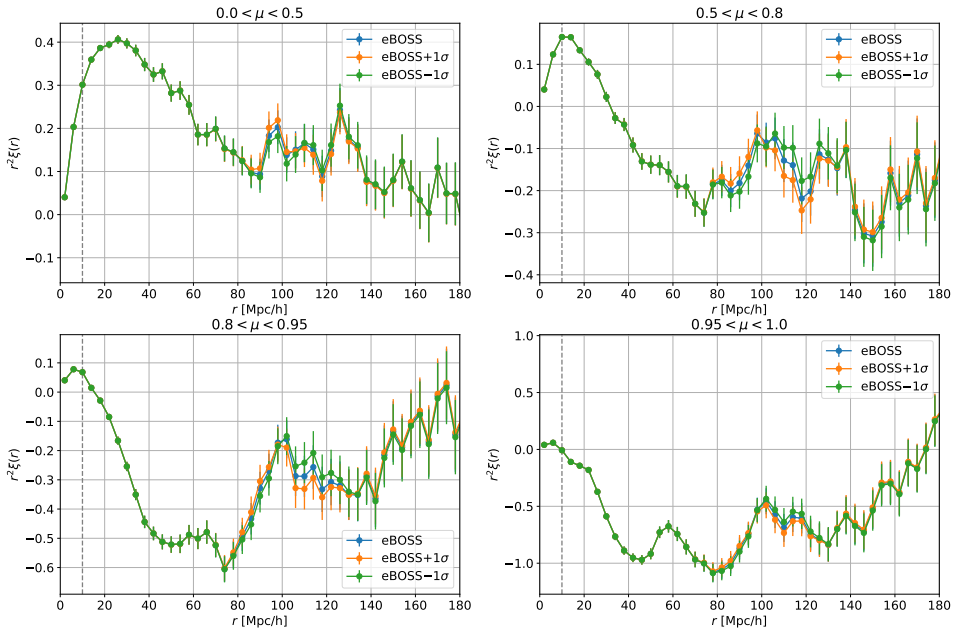

Figure 8.2. The Lyα forest auto correlation function measured by eBOSS DR16, in four μ (the cosine of the line-of-sight angle) wedges. The blue points represent the original correlation. For the orange and green points, the position of the BAO peak has been shifted left and right respectively, using a fiducial model. This illustrates the method of blinding that will be used for the DESI Lyα year one BAO analysis.

For the year one DESI Lyα BAO key project, there were 5 proposed strategies, applied at different stages of the analysis. I was directly involved in developing and testing two of these, including the one that was ultimately chosen by the Lyα working group in DESI. Here I will only present this last strategy, and explain why it was chosen.

This blinding strategy consists of shifting the position of the peak in the measured correlation function. We achieve this by computing a fiducial model based on the best fit parameter values from eBOSS DR16 (du Mas des Bourboux et al., 2020), but with the BAO parameters fixed to $\alpha_{||} = 1$ and $\alpha_{\perp} = 1$. We then create a second model where the BAO parameters take random values[3] around $(1, 1)$. The difference between these two model correlation functions is then automatically applied to every measured correlation function that is using main survey data.[4] An example of this applied to eBOSS DR16 data is shown in Figure 8.2.

This blinding strategy is applied much later in the analysis compared to other strategies found

[3] We use a Gaussian distribution with a standard deviation defined by the constraining power of eBOSS DR16 and a cutoff at 3σ.

[4] The blinding is not applied for SV data.

in the literature (e.g. DES Collaboration et al., 2021; Heymans et al., 2021). While some of the proposed strategies did blind at earlier stages, these are not as desirable for the Lyα forest forest due to the strong metal correlation peaks. The most prominent of these can seen in the bottom two plots of Figure 8.2 at a separation $\sim 60h^{-1}$Mpc. This is due to correlations between Lyα absorption and two metal absorption lines: SiII(1190Å) and SiII(1193Å). Any blinding that changes the position of this peak can be guessed relatively easily because we know where the peak should be[5], and it is very well constrained. Therefore, we need a blinding strategy that only changes the correlation around the position of the BAO peak.[6]

This type of blinding strategy could also be used for a full-shape analysis where we want to measure the AP effect. This is because the metal correlations are added to the model at the last step. This means we can add a random coordinate re-scaling to the full Lyα correlation functions (instead of just the BAO peak), before we add the metal correlations. This brings us to the next topic, which is a full-shape analysis of Lyα correlations using real data.

8.4.2 Towards a full-shape analysis

As we have discussed in Section 7.5, before a full-shape analysis can be performed on real data, we need to study the impact of contaminants on the parameters of interest. Based on the forecast analysis we performed in Chapter 7, we can break down this type of analysis into different steps. The end goal is to measure the BAO peak position, the AP effect from the full-shape and RSD through the $f\sigma_8$ parameter combination. As we have seen, we can only measure $f\sigma_8$ through a joint analysis of the Lyα auto and cross-correlation with quasars. This means we can start by focusing on the AP effect from the Lyα auto and cross individually. Ultimately, this means measuring two parameters; the isotropic BAO scale, $\alpha_{\rm p}$, and the full-shape anisotropic scale parameter, ϕ. Here, we also have the option to split ϕ up into two parameters that re-scale the peak and smooth components, respectively. This could be desirable in order to better understand the correlations with contaminants, because we expect the smooth component to be the one that is affected by them. It would also allow us to quantify the improvement in constraining power over a standard BAO analysis.

A major component of this analysis will be to decide what is the smallest and largest scale we should fit. We can study this by measuring the constraining power (statistical error) on the AP

[5] Its position can be computed from the difference between the Lyα wavelength and the metal wavelength, transformed into comoving coordinates.

[6] A strategy that changes the rest of the correlation in a way that is not coherent between metal correlations and the BAO peak could also work. However, no such strategy has been proposed, and it is hard to imagine one. This is because we cannot tell the difference between pixels that are dominated by metal absorption and those that are dominated by Lyα.

parameter as a function of minimum and maximum separation, and comparing it to the expected systematic bias using mock data sets. To compute the expected systematic bias, we can perform the analysis on a large number of these mocks, and progressively include different contaminants. This will allow us to test how well we can model the contaminants and what their impact on the parameters of interest is.

In terms of performing a full-shape analysis on DESI data, we would have to wait until the year one Lyα BAO analysis is complete. This is because the BAO measurement is a component of the full-shape analysis. Therefore, roughly for the next one and a half years, the best options for this type of analysis are the BOSS and eBOSS data sets. This is also facilitated by the public availability of the eBOSS DR16 correlation functions, and the set of mock data sets that were used for that BAO analysis. For this reason, my current focus is on using these eBOSS mocks to prepare for a full-shape analysis of the eBOSS DR16 Lyα forest correlation functions.

The Lyα forest is now an established cosmological tracer, with the BAO analysis being a key project in DESI. As we have shown, there is still great potential for more types of cosmological analyses to be performed with this tracer. Here we have focused on the incremental steps of measuring AP and RSD. However, it is my hope that by the time we arrive at the final analysis of the full DESI data, we will have progressed to more advanced techniques. Once we are more confident in our understanding of the model and contaminants for the Lyα correlation functions, we could perform a direct fit of cosmological parameters. This would unlock the full potential of this statistic.

BIBLIOGRAPHY

Abbott, T. M. C. et al. (2018). "Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing". In: *Phys.Rev.D* 98.4, 043526, p. 043526. DOI: 10.1103/PhysRevD.98.043526. arXiv: 1708.01530 [astro-ph.CO] (cit. on p. 82).

Abbott, T. M. C. et al. (Feb. 2019). "First Cosmology Results using Type Ia Supernovae from the Dark Energy Survey: Constraints on Cosmological Parameters". In: *ApJ* 872.2, L30, p. L30. DOI: 10.3847/2041-8213/ab04fa. arXiv: 1811.02374 [astro-ph.CO] (cit. on p. 55).

Addison, G. E., G. Hinshaw, and M. Halpern (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 p. 67).

Addison, G. E. et al. (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 pp. 67, 78).

Adelberger, E. G. et al. (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 pp. 73, 74).

Adhikari, Saroj and Dragan Huterer (2019). "A new measure of tension between experiments". In: *Journal of Cosmology and Astroparticle Physics* 2019.01, pp. 036–036. DOI: 10.1088/1475-7516/2019/01/036. URL: https://doi.org/10.1088%2F1475-7516%2F2019%2F01%2F036 (cit. on pp. 67, 69).

Aiola, Simone et al. (Dec. 2020). "The Atacama Cosmology Telescope: DR4 maps and cosmological parameters". In: *JCAP* 2020.12, 047, p. 047. DOI: 10.1088/1475-7516/2020/12/047. arXiv: 2007.07288 [astro-ph.CO] (cit. on p. 130).

Alam, Shadab et al. (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. 51, 66–68, 102, 122).

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 pp. 103, 105).

Algeri, Sara et al. (2019). "Searching for new physics with profile likelihoods: Wilks and beyond". In: *arXiv e-prints*, arXiv:1911.10237, arXiv:1911.10237. arXiv: 1911.10237 [physics.data-an] (cit. on pp. 82, 87).

Allen, TJ, B Grinstein, and Mark B Wise (1987). "Non-gaussian density perturbations in inflationary cosmologies". In: *Physics Letters B* 197.1-2, pp. 66–70 (cit. on p. 38).

Anderson, Lauren et al. (2014). "The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Releases 10 and 11 Galaxy samples". In: *MNRAS* 441.1, pp. 24–62. DOI: 10.1093/mnras/stu523. arXiv: 1312.4877 [astro-ph.CO] (cit. on p. 82).

Arinyo-i-Prats, Andreu et al. (Dec. 2015). "The non-linear power spectrum of the Lyman alpha forest". In: *JCAP* 2015.12, 017, p. 017. DOI: 10.1088/1475-7516/2015/12/017. arXiv: 1506.04519 [astro-ph.CO] (cit. on pp. 54, 108, 110).

Astropy Collaboration et al. (2013). "Astropy: A community Python package for astronomy". In: *A&A* 558, A33, A33. DOI: 10.1051/0004-6361/201322068. arXiv: 1307.6212 [astro-ph.IM] (cit. on pp. 66, 78).

Astropy Collaboration et al. (2018). "The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package". In: *AJ* 156.3, 123, p. 123. DOI: 10.3847/1538-3881/aabc4f. arXiv: 1801.02634 [astro-ph.IM] (cit. on p. 78).

Ata, Metin et al. (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. 67, 68, 102).

Aubourg, Éric et al. (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 pp. 67, 72, 78).

Audren, Benjamin et al. (2013). "Conservative Constraints on Early Cosmology: an illustration of the Monte Python cosmological parameter inference code". In: *JCAP* 1302, p. 001. DOI:

10.1088/1475-7516/2013/02/001. arXiv: 1210.7183 [astro-ph.CO] (cit. on pp. 60, 78, 90).

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. 52).

Bardeen, J. M. et al. (May 1986). "The Statistics of Peaks of Gaussian Random Fields". In: *ApJ* 304, p. 15. DOI: 10.1086/164143 (cit. on p. 49).

Basu, Debabrata (1977). "On the Elimination of Nuisance Parameters". In: *Journal of the American Statistical Association* 72.358, pp. 355–366 (cit. on p. 87).

Bautista, Julian E. et al. (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. 55, 66, 67, 82, 87, 103, 120, 123, 124).

Bautista, Julian E. et al. (Sept. 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 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 p. 130).

Baxter, Eric J. and Blake D. Sherwin (Feb. 2021). "Determining the Hubble constant without the sound horizon scale: measurements from CMB lensing". In: *MNRAS* 501.2, pp. 1823–1835. DOI: 10.1093/mnras/staa3706. arXiv: 2007.04007 [astro-ph.CO] (cit. on p. 107).

Bayes, Thomas (1763). "LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, FRS communicated by Mr. Price, in a letter to John Canton, AMFR S". In: *Philosophical transactions of the Royal Society of London* 53, pp. 370–418 (cit. on pp. 57, 83, 84).

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. 48).

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. 48).

Bernal, José Luis and John A. Peacock (July 2018). "Conservative cosmology: combining data with allowance for unknown systematics". In: *JCAP* 2018.7, 002, p. 002. DOI: 10.1088/1475-7516/2018/07/002. arXiv: 1803.04470 [astro-ph.CO] (cit. on p. 132).

Bernal, José Luis, Licia Verde, and Adam G. Riess (2016). "The trouble with H_0". In: *Journal of Cosmology and Astro-Particle Physics* 2016.10, 019, p. 019. DOI: 10.1088/1475-7516/2016/10/019. arXiv: 1607.05617 [astro-ph.CO] (cit. on p. 66).

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 p. 45).

Betoule, M. et al. (Aug. 2014). "Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples". In: *A&A* 568, A22, A22. DOI: 10.1051/0004-6361/201423413. arXiv: 1401.4064 [astro-ph.CO] (cit. on p. 55).

Beutler, Florian et al. (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. 66–68, 102, 122).

Beutler, Florian et al. (July 2012). "The 6dF Galaxy Survey: z≈ 0 measurements of the growth rate and σ_8". In: *MNRAS* 423.4, pp. 3430–3444. DOI: 10.1111/j.1365-2966.2012.21136.x. arXiv: 1204.4725 [astro-ph.CO] (cit. on p. 103).

Beutler, Florian et al. (Apr. 2017). "The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: anisotropic galaxy clustering in Fourier space". In: *MNRAS* 466.2, pp. 2242–2260. DOI: 10.1093/mnras/stw3298. arXiv: 1607.03150 [astro-ph.CO] (cit. on p. 108).

Blake, Chris et al. (Aug. 2011). "The WiggleZ Dark Energy Survey: the growth rate of cosmic structure since redshift z=0.9". In: *MNRAS* 415.3, pp. 2876–2891. DOI: 10.1111/j.1365-2966.2011.18903.x. arXiv: 1104.2948 [astro-ph.CO] (cit. on p. 103).

Blake, Chris et al. (Sept. 2012). "The WiggleZ Dark Energy Survey: joint measurements of the expansion and growth history at z ¡ 1". In: *MNRAS* 425.1, pp. 405–414. DOI: 10.1111/j.1365-2966.2012.21473.x. arXiv: 1204.3674 [astro-ph.CO] (cit. on pp. 51, 122).

Blas, Diego, Julien Lesgourgues, and Thomas Tram (July 2011). "The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes". In: *JCAP* 2011.7, 034, p. 034. DOI: 10.1088/1475-7516/2011/07/034. arXiv: 1104.2933 [astro-ph.CO] (cit. on p. 50).

Blomqvist, Michael et al. (2019). "Baryon acoustic oscillations from the cross-correlation of Lyα absorption and quasars in eBOSS DR14". In: *arXiv e-prints*, arXiv:1904.03430, arXiv:1904.03430. arXiv: 1904.03430 [astro-ph.CO] (cit. on pp. 55, 66–68, 82, 83, 88, 89, 94, 95, 103, 124).

Brieden, Samuel et al. (Sept. 2020). "Blind Observers of the Sky". In: *JCAP* 2020.9, 052, p. 052. DOI: 10.1088/1475-7516/2020/09/052. arXiv: 2006.10857 [astro-ph.CO] (cit. on p. 134).

Brinckmann, Thejs and Julien Lesgourgues (Apr. 2018). "MontePython 3: boosted MCMC sampler and other features". In: *arXiv e-prints*, arXiv:1804.07261, arXiv:1804.07261. arXiv: 1804.07261 [astro-ph.CO] (cit. on p. 90).

Busca, N. G. et al. (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. 55, 82, 102, 103).

Cardona, Wilmar, Martin Kunz, and Valeria Pettorino (2017). "Determining H_0 with Bayesian hyper-parameters". In: *Journal of Cosmology and Astro-Particle Physics* 2017.3, 056, p. 056. DOI: 10.1088/1475-7516/2017/03/056. arXiv: 1611.06088 [astro-ph.CO] (cit. on p. 66).

Carroll, Sean M. (Feb. 2001). "The Cosmological Constant". In: *Living Reviews in Relativity* 4.1, 1, p. 1. DOI: 10.12942/lrr-2001-1. arXiv: astro-ph/0004075 [astro-ph] (cit. on p. 31).

— (2004). *Spacetime and geometry. An introduction to general relativity* (cit. on p. 25).

Chan, K. C. et al. (2018). "BAO from angular clustering: optimization and mitigation of theoretical systematics". In: *MNRAS* 480.3, pp. 3031–3051. DOI: 10.1093/mnras/sty2036. arXiv: 1801.04390 [astro-ph.CO] (cit. on pp. 82, 87, 93, 100).

Charnock, Tom, Richard A. Battye, and Adam Moss (2017). "Planck data versus large scale structure: Methods to quantify discordance". In: *Phys. Rev. D* 95 (12), p. 123535. DOI: 10.1103/PhysRevD.95.123535. URL: https://link.aps.org/doi/10.1103/PhysRevD.95.123535 (cit. on pp. 67, 69).

Chen, Shi-Fan, Zvonimir Vlah, and Martin White (Mar. 2021). "The Lyα forest flux correlation function: a perturbation theory perspective". In: *arXiv e-prints*, arXiv:2103.13498, arXiv:2103.13498. arXiv: 2103.13498 [astro-ph.CO] (cit. on pp. 103, 133).

Cole, Shaun et al. (2005). "The 2dF Galaxy Redshift Survey: power-spectrum analysis of the final data set and cosmological implications". In: *MNRAS* 362.2, pp. 505–534. DOI: 10.1111/j.1365-2966.2005.09318.x. arXiv: astro-ph/0501174 [astro-ph] (cit. on pp. 51, 82, 102).

Cooke, Ryan J., Max Pettini, and Charles C. Steidel (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. 56, 67, 73–75, 78).

Cooke, Ryan J. et al. (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 pp. 56, 73–75, 78).

Cox, Richard T (1946). "Probability, frequency and reasonable expectation". In: *American journal of physics* 14.1, pp. 1–13 (cit. on p. 58).

Croft, Rupert AC et al. (1999). "The Power spectrum of mass fluctuations measured from the Lyα forest at redshift z= 2.5". In: *The Astrophysical Journal* 520.1, p. 1 (cit. on p. 102).

Croft, Rupert AC et al. (2002). "Toward a precise measurement of matter clustering: Lyα forest data at redshifts 2-4". In: *The Astrophysical Journal* 581.1, p. 20 (cit. on p. 102).

Cuceu, Andrei, Andreu Font-Ribera, and Benjamin Joachimi (July 2020). "Bayesian methods for fitting Baryon Acoustic Oscillations in the Lyman-α forest". In: *JCAP* 2020.7, 035, p. 035. DOI: 10.1088/1475-7516/2020/07/035. arXiv: 2004.02761 [astro-ph.CO] (cit. on pp. 81, 106).

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. 65, 90, 130).

Cuceu, Andrei et al. (Oct. 2021). "Cosmology beyond BAO from the 3D distribution of the Lyman-α forest". In: *MNRAS* 506.4, pp. 5439–5450. DOI: 10.1093/mnras/stab1999. arXiv: 2103.14075 [astro-ph.CO] (cit. on pp. 102, 131).

Cyburt, Richard H. et al. (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 pp. 56, 73).

de Mattia, Arnaud et al. (Mar. 2021). "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* 501.4, pp. 5616–5645. DOI: 10.1093/mnras/staa3891. arXiv: 2007.09008 [astro-ph.CO] (cit. on p. 130).

de Sainte Agathe, Victoria et al. (2019). "Baryon acoustic oscillations at z = 2.34 from the correlations of Lyα absorption in eBOSS DR14". In: *arXiv e-prints*, arXiv:1904.03400, arXiv:1904.03400. arXiv: 1904.03400 [astro-ph.CO] (cit. on pp. 55, 66–68, 82, 83, 88–90, 94, 95, 103, 124).

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. 45).

Dekel, Avishai and Ofer Lahav (July 1999). "Stochastic Nonlinear Galaxy Biasing". In: *ApJ* 520.1, pp. 24–34. DOI: 10.1086/307428. arXiv: astro-ph/9806193 [astro-ph] (cit. on p. 49).

Delubac, Timothée et al. (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. 55, 66, 67, 87).

DES Collaboration et al. (May 2021). "Dark Energy Survey Year 3 Results: Cosmological Constraints from Galaxy Clustering and Weak Lensing". In: *arXiv e-prints*, arXiv:2105.13549, arXiv:2105.13549. arXiv: 2105.13549 [astro-ph.CO] (cit. on pp. 30, 56, 136).

DESI Collaboration et al. (2016). "The DESI Experiment Part I: Science,Targeting, and Survey Design". In: *arXiv e-prints*, arXiv:1611.00036, arXiv:1611.00036. arXiv: 1611.00036 [astro-ph.IM] (cit. on pp. 51, 75, 76, 104, 113, 117, 118, 134).

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 pp. 45, 50).

Di Valentino, Eleonora et al. (July 2021). "In the realm of the Hubble tension-a review of solutions". In: *Classical and Quantum Gravity* 38.15, 153001, p. 153001. DOI: 10.1088/1361-6382/ac086d. arXiv: 2103.01183 [astro-ph.CO] (cit. on p. 131).

Dicke, Robert H et al. (1965). "Cosmic black-body radiation." In: *The Astrophysical Journal* 142, pp. 414–419 (cit. on p. 48).

Dodelson, Scott (2003). *Modern cosmology*. Elsevier (cit. on pp. 37–41, 43, 45, 49, 51, 73).

Doroshkevich, AG, Ya B Zel'dovich, and RA Syunyaev (1978). "Fluctuations of the microwave background radiation in the adiabatic and entropic theories of galaxy formation". In: *Soviet Astronomy* 22, pp. 523–528 (cit. on p. 42).

du Mas des Bourboux, Hélion et al. (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. 55, 66–68, 82, 88, 103, 113, 120, 124).

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. 53–55, 103, 105, 108, 110, 111, 117, 118, 120, 123–125, 130–132, 134, 135).

Dutcher, D. et al. (July 2021). "Measurements of the E -mode polarization and temperature-E -mode correlation of the CMB from SPT-3G 2018 data". In: *Phys.Rev.D* 104.2, 022003, p. 022003. DOI: 10.1103/PhysRevD.104.022003. arXiv: 2101.01684 [astro-ph.CO] (cit. on p. 130).

Dyson, F. W., A. S. Eddington, and C. Davidson (Jan. 1920). "A Determination of the Deflection of Light by the Sun's Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919". In: *Philosophical Transactions of the Royal Society of London Series A* 220, pp. 291–333. DOI: 10.1098/rsta.1920.0009 (cit. on p. 56).

eBOSS Collaboration et al. (Apr. 2021). "Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory". In: *Phys.Rev.D* 103.8, 083533, p. 083533. DOI: 10.1103/PhysRevD.103.083533. arXiv: 2007.08991 [astro-ph.CO] (cit. on pp. 30–32, 51, 102, 122, 130).

Eddington, A. S. (Aug. 1917). "The pulsation theory of Cepheid variables". In: *The Observatory* 40, pp. 290–293 (cit. on p. 56).

Einstein, Albert (1915). "Die feldgleichungen der gravitation". In: *Sitzung der physikalische-mathematischen Klasse* 25, pp. 844–847 (cit. on p. 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. 42, 44, 46, 110).

Eisenstein, Daniel J. et al. (2005). "Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies". In: *ApJ* 633.2, pp. 560–574. DOI: 10.1086/466512. arXiv: astro-ph/0501171 [astro-ph] (cit. on pp. 51, 82, 102).

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. 46).

Epstein, R. I., J. M. Lattimer, and D. N. Schramm (1976). "The origin of deuterium". In: *Nature* 263, pp. 198–202. DOI: 10.1038/263198a0 (cit. on pp. 73, 75).

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 pp. 113, 118).

Feeney, Stephen M., Daniel J. Mortlock, and Niccolò Dalmasso (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. 66).

Feroz, F. and M. P. Hobson (Feb. 2008). "Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses". In: *MNRAS* 384.2, pp. 449–463. DOI: 10.1111/j.1365-2966.2007.12353.x. arXiv: 0704.3704 [astro-ph] (cit. on p. 62).

Feroz, F., M. P. Hobson, and M. Bridges (Oct. 2009). "MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics". In: *MNRAS* 398.4, pp. 1601–1614. DOI: 10.1111/j.1365-2966.2009.14548.x. arXiv: 0809.3437 [astro-ph] (cit. on p. 62).

Feroz, Farhan et al. (Nov. 2019). "Importance Nested Sampling and the MultiNest Algorithm". In: *The Open Journal of Astrophysics* 2.1, 10, p. 10. DOI: 10.21105/astro.1306.2144. arXiv: 1306.2144 [astro-ph.IM] (cit. on p. 62).

Fixsen, D. J. (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. 48, 69, 115).

Fixsen, D. J. et al. (1996). "The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set". In: *ApJ* 473, p. 576. DOI: 10.1086/178173. arXiv: astro-ph/9605054 [astro-ph] (cit. on pp. 48, 69, 115).

Follin, B. and L. Knox (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. 66).

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 p. 123).

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 p. 124).

Font-Ribera, Andreu et al. (2014). "Quasar-Lyman α forest cross-correlation from BOSS DR11: Baryon Acoustic Oscillations". In: *Journal of Cosmology and Astro-Particle Physics* 2014.5,

027, p. 027. DOI: 10.1088/1475-7516/2014/05/027. arXiv: 1311.1767 [astro-ph.CO] (cit. on pp. 55, 66–68, 82, 87, 102, 103).

Freedman, Wendy L. et al. (May 2001). "Final Results from the Hubble Space Telescope Key Project to Measure the Hubble Constant". In: *ApJ* 553.1, pp. 47–72. DOI: 10.1086/320638. arXiv: astro-ph/0012376 [astro-ph] (cit. on p. 56).

Freedman, Wendy L. et al. (Sept. 2019). "The Carnegie-Chicago Hubble Program. VIII. An Independent Determination of the Hubble Constant Based on the Tip of the Red Giant Branch". In: *ApJ* 882.1, 34, p. 34. DOI: 10.3847/1538-4357/ab2f73. arXiv: 1907.05922 [astro-ph.CO] (cit. on p. 130).

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. 130).

Friedman, A. (Jan. 1922). "Über die Krümmung des Raumes". In: *Zeitschrift fur Physik* 10, pp. 377–386. DOI: 10.1007/BF01332580 (cit. on p. 26).

Fukuda, Y. et al. (1998). "Evidence for Oscillation of Atmospheric Neutrinos". In: *Phys. Rev. Lett.* 81 (8), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. URL: https://link.aps.org/doi/10.1103/PhysRevLett.81.1562 (cit. on p. 30).

Gaztanaga, Enrique (1992). "N-point correlation functions in the CfA and SSRS redshift distribution of galaxies". In: *The Astrophysical Journal* 398, pp. L17–L20 (cit. on p. 50).

Gelman, Andrew et al. (2013). *Bayesian data analysis*. Chapman and Hall/CRC (cit. on p. 86).

Ghosh, JK (1988). *Statistical information and likelihood: a collection of critical essays by Dr. D. Basu*. Springer (cit. on p. 87).

Gil-Marín, Héctor et al. (June 2018). "The clustering of the SDSS-IV extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample: structure growth rate measurement from the anisotropic quasar power spectrum in the redshift range 0.8 ¡ z ¡ 2.2". In: *MNRAS* 477.2, pp. 1604–1638. DOI: 10.1093/mnras/sty453. arXiv: 1801.02689 [astro-ph.CO] (cit. on pp. 118, 133).

Gil-Marín, Héctor et al. (Oct. 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 p. 130).

Givans, Jahmour J. and Christopher M. Hirata (July 2020). "Redshift-space streaming velocity effects on the Lyman-α forest baryon acoustic oscillation scale". In: *Phys.Rev.D* 102.2, 023515, p. 023515. DOI: 10.1103/PhysRevD.102.023515. arXiv: 2002.12296 [astro-ph.CO] (cit. on pp. 103, 133).

Grohs, Evan et al. (2019). "Big Bang Nucleosynthesis and Neutrino Cosmology". In: *BAAS*. Vol. 51, p. 412. DOI: https://baas.aas.org/wp-content/uploads/2019/05/412_grohs.pdf (cit. on pp. 56, 73).

Gunn, James E and Bruce A Peterson (1965). "On the Density of Neutral Hydrogen in Intergalactic Space." In: *The Astrophysical Journal* 142, pp. 1633–1636 (cit. on p. 52).

Gustavino, C. (2014). "BBN, Neutrinos and Nuclear Astrophysics". In: *Proceedings of the 52th International Winter Meeting on Nuclear Physics (Bormio2014). 27 - 31 January*, p. 50 (cit. on p. 77).

Handley, W. J., M. P. Hobson, and A. N. Lasenby (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 pp. 62, 71, 91, 113).

— (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 pp. 62, 71, 91, 113).

Handley, Will and Pablo Lemos (2019a). "Quantifying dimensionality: Bayesian cosmological model complexities". In: *arXiv e-prints*, arXiv:1903.06682, arXiv:1903.06682. arXiv: 1903.06682 [astro-ph.CO] (cit. on pp. 71, 72).

— (2019b). "Quantifying tension: interpreting the DES evidence ratio". In: *arXiv e-prints*, arXiv:1902.04029, arXiv:1902.04029. arXiv: 1902.04029 [astro-ph.CO] (cit. on pp. 67, 69, 71, 77).

Hastings, W. K. (Apr. 1970). "Monte Carlo sampling methods using Markov chains and their applications". In: *Biometrika* 57.1, pp. 97–109. ISSN: 0006-3444. DOI: 10.1093/biomet/57.1.97. eprint: https://academic.oup.com/biomet/article-pdf/57/1/97/23940249/57-1-97.pdf. URL: https://doi.org/10.1093/biomet/57.1.97 (cit. on pp. 60, 62, 83, 86).

Heavens, Alan, Raul Jimenez, and Licia Verde (2014). "Standard Rulers, Candles, and Clocks from the Low-Redshift Universe". In: *Phys.Rev.Lett* 113.24, 241302, p. 241302. DOI: 10.1103/PhysRevLett.113.241302. arXiv: 1409.6217 [astro-ph.CO] (cit. on p. 112).

Heymans, Catherine et al. (Feb. 2021). "KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints". In: *A&A* 646, A140, A140. DOI:

10.1051/0004-6361/202039063. arXiv: 2007.15632 [astro-ph.CO] (cit. on pp. 56, 136).

Higson, Edward et al. (Sept. 2019). "Dynamic nested sampling: an improved algorithm for parameter estimation and evidence calculation". In: *Statistics and Computing* 29.5, pp. 891–913. DOI: 10.1007/s11222-018-9844-0. arXiv: 1704.03459 [stat.CO] (cit. on p. 62).

Hikage, Chiaki et al. (Apr. 2019). "Cosmology from cosmic shear power spectra with Subaru Hyper Suprime-Cam first-year data". In: *PASJ* 71.2, 43, p. 43. DOI: 10.1093/pasj/psz010. arXiv: 1809.09148 [astro-ph.CO] (cit. on p. 56).

Hinton, Samuel R., Cullan Howlett, and Tamara M. Davis (2020). "Barry and the BAO Model Comparison". In: *MNRAS*. DOI: 10.1093/mnras/staa361. arXiv: 1912.01175 [astro-ph.CO] (cit. on p. 82).

Hobson, M. P., S. L. Bridle, and O. Lahav (Sept. 2002). "Combining cosmological data sets: hyperparameters and Bayesian evidence". In: *MNRAS* 335.2, pp. 377–388. DOI: 10.1046/j.1365-8711.2002.05614.x. arXiv: astro-ph/0203259 [astro-ph] (cit. on p. 132).

Hobson, Michael P et al. (2010). *Bayesian methods in cosmology*. Cambridge University Press (cit. on p. 60).

Horowitz, Benjamin et al. (Dec. 2019). "TARDIS. I. A Constrained Reconstruction Approach to Modeling the z $\sim$ 2.5 Cosmic Web Probed by Lyα Forest Tomography". In: *ApJ* 887.1, 61, p. 61. DOI: 10.3847/1538-4357/ab4d4c. arXiv: 1903.09049 [astro-ph.CO] (cit. on p. 83).

Hou, Jiamin et al. (Oct. 2018). "The clustering of the SDSS-IV extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample: anisotropic clustering analysis in configuration space". In: *MNRAS* 480.2, pp. 2521–2534. DOI: 10.1093/mnras/sty1984. arXiv: 1801.02656 [astro-ph.CO] (cit. on pp. 118, 133).

Hou, Jiamin et al. (Oct. 2020). "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. 118, 124, 130).

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. 42).

Hui, Lam, Albert Stebbins, and Scott Burles (Jan. 1999). "A Geometrical Test of the Cosmological Energy Contents Using the Lyα Forest". In: *ApJ* 511.1, pp. L5–L8. DOI: 10.1086/311826. arXiv: astro-ph/9807190 [astro-ph] (cit. on p. 103).

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 pp. 67, 69).

Jaynes, Edwin T and Oscar Kempthorne (1976). *Confidence intervals vs Bayesian intervals*. Springer, pp. 175–257 (cit. on p. 85).

Jeffreys, Harold (1946). "An invariant form for the prior probability in estimation problems". In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 186.1007, pp. 453–461 (cit. on p. 59).

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 p. 50).

Kapteyn, J. C. (May 1922). "First Attempt at a Theory of the Arrangement and Motion of the Sidereal System". In: *ApJ* 55, p. 302. DOI: 10.1086/142670 (cit. on p. 30).

Kirkby, David et al. (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. 55, 82, 102–106, 110, 114).

Kitaura, Francisco-Shu, Simona Gallerani, and Andrea Ferrara (Feb. 2012). "Multiscale inference of matter fields and baryon acoustic oscillations from the Lyα forest". In: *MNRAS* 420.1, pp. 61–74. DOI: 10.1111/j.1365-2966.2011.19997.x. arXiv: 1011.6233 [astro-ph.CO] (cit. on p. 83).

Kochanek, Izabela (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, p. 012023. DOI: 10.1088/1742-6596/703/1/012023 (cit. on p. 77).

Kofman, LA and D Yu Pogosyan (1988). "Nonflat perturbations in inflationary cosmology". In: *Physics Letters B* 214.4, pp. 508–514 (cit. on p. 38).

Kullback, S. and R. A. Leibler (Mar. 1951). "On Information and Sufficiency". In: *Ann. Math. Statist.* 22.1, pp. 79–86. DOI: 10.1214/aoms/1177729694. URL: https://doi.org/10.1214/aoms/1177729694 (cit. on p. 71).

Lahav, O. et al. (July 2000). "Bayesian 'hyper-parameters' approach to joint estimation: the Hubble constant from CMB measurements". In: *MNRAS* 315.4, pp. L45–L49. DOI: 10.1046/j.1365-8711.2000.03633.x. arXiv: astro-ph/9912105 [astro-ph] (cit. on p. 132).

Laplace, Pierre Simon (1820). *Théorie analytique des probabilités*. Courcier (cit. on pp. 57, 83, 84).

Laureijs, R. et al. (2011). "Euclid Definition Study Report". In: *arXiv e-prints*, arXiv:1110.3193, arXiv:1110.3193. arXiv: 1110.3193 [astro-ph.CO] (cit. on pp. 52, 75).

Leavitt, Henrietta S. and Edward C. Pickering (Mar. 1912). "Periods of 25 Variable Stars in the Small Magellanic Cloud." In: *Harvard College Observatory Circular* 173, pp. 1–3 (cit. on p. 56).

Lemaître, G. (Mar. 1931). "Expansion of the universe, A homogeneous universe of constant mass and increasing radius accounting for the radial velocity of extra-galactic nebulae". In: *MNRAS* 91, pp. 483–490. DOI: 10.1093/mnras/91.5.483 (cit. on p. 26).

Lesgourgues, Julien (Apr. 2011). "The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview". In: *arXiv e-prints*, arXiv:1104.2932, arXiv:1104.2932. arXiv: 1104.2932 [astro-ph.IM] (cit. on p. 50).

Lewis, Antony and Sarah Bridle (2002). "Cosmological parameters from CMB and other data: A Monte Carlo approach". In: *Phys.Rev.D* 66.10, 103511, p. 103511. DOI: 10.1103/PhysRevD.66.103511. arXiv: astro-ph/0205436 [astro-ph] (cit. on pp. 60, 72, 78, 90).

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. 49).

Lewis, Antony, Anthony Challinor, and Anthony Lasenby (2000). "Efficient computation of CMB anisotropies in closed FRW models". In: *ApJ* 538, pp. 473–476. DOI: 10.1086/309179. arXiv: astro-ph/9911177 [astro-ph] (cit. on pp. 50, 73, 106).

Liddle, Andrew R and David H Lyth (2000). *Cosmological inflation and large-scale structure*. Cambridge university press (cit. on p. 37).

Liddle, Andrew R., Pia Mukherjee, and David Parkinson (2006). "Cosmological model selection". In: *arXiv e-prints*, astro-ph/0608184, astro–ph/0608184. arXiv: astro-ph/0608184 [astro-ph] (cit. on p. 84).

Linsky, Jeffrey L. et al. (2006). "What Is the Total Deuterium Abundance in the Local Galactic Disk?" In: *ApJ* 647.2, pp. 1106–1124. DOI: 10.1086/505556. arXiv: astro-ph/0608308 [astro-ph] (cit. on p. 75).

Lovelock, David (Mar. 1971). "The Einstein Tensor and Its Generalizations". In: *Journal of Mathematical Physics* 12.3, pp. 498–501. DOI: 10.1063/1.1665613 (cit. on p. 28).

— (June 1972). "The Four-Dimensionality of Space and the Einstein Tensor". In: *Journal of Mathematical Physics* 13.6, pp. 874–876. DOI: 10.1063/1.1666069 (cit. on p. 28).

LSST Science Collaboration et al. (Dec. 2009). "LSST Science Book, Version 2.0". In: *arXiv e-prints*, arXiv:0912.0201, arXiv:0912.0201. arXiv: 0912.0201 [astro-ph.IM] (cit. on p. 52).

Lynds, Roger (1971). "The absorption-line spectrum of 4c 05.34". In: *The Astrophysical Journal* 164, p. L73 (cit. on p. 102).

Marcucci, L. E. et al. (2016). "Implication of the Proton-Deuteron Radiative Capture for Big Bang Nucleosynthesis". In: *Phys.Rev.Lett* 116.10, 102501, p. 102501. DOI: 10.1103/PhysRevLett.116.102501. arXiv: 1510.07877 [nucl-th] (cit. on p. 74).

Marshall, Phil, Nutan Rajguru, and An že Slosar (2006). "Bayesian evidence as a tool for comparing datasets". In: *Phys. Rev. D* 73 (6), p. 067302. DOI: 10.1103/PhysRevD.73.067302. URL: https://link.aps.org/doi/10.1103/PhysRevD.73.067302 (cit. on p. 69).

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. 54, 103, 133).

McDonald, Patrick and Jordi Miralda-Escudé (June 1999). "Measuring the Cosmological Geometry from the Lyα Forest along Parallel Lines of Sight". In: *ApJ* 518.1, pp. 24–31. DOI: 10.1086/307264. arXiv: astro-ph/9807137 [astro-ph] (cit. on p. 103).

McDonald, Patrick et al. (2000). "The observed probability distribution function, power spectrum, and correlation function of the transmitted flux in the Lyα forest". In: *The Astrophysical Journal* 543.1, p. 1 (cit. on pp. 54, 102).

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 p. 52).

Meszaros, P. (Dec. 1974). "The behaviour of point masses in an expanding cosmological substratum." In: *A&A* 37.2, pp. 225–228 (cit. on p. 42).

Metropolis, Nicholas et al. (1953). "Equation of State Calculations by Fast Computing Machines". In: *The Journal of Chemical Physics* 21.6, pp. 1087–1092. DOI: 10.1063/1.1699114 (cit. on pp. 60, 62, 83, 86).

Moss, Adam (July 2020). "Accelerated Bayesian inference using deep learning". In: *MNRAS* 496.1, pp. 328–338. DOI: 10.1093/mnras/staa1469. arXiv: 1903.10860 [astro-ph.CO] (cit. on p. 62).

Mossa, V. et al. (May 2020a). "Setup commissioning for an improved measurement of the D(p,γ)^{3}He cross section at Big Bang Nucleosynthesis energies". In: *European Physical Journal A* 56.5, 144, p. 144. DOI: 10.1140/epja/s10050-020-00149-1. arXiv: 2005.00002 [physics.ins-det] (cit. on p. 130).

Mossa, V. et al. (Nov. 2020b). "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. 130).

Muir, J. et al. (May 2020). "Blinding multiprobe cosmological experiments". In: *MNRAS* 494.3, pp. 4454–4470. DOI: 10.1093/mnras/staa965. arXiv: 1911.05929 [astro-ph.CO] (cit. on p. 134).

Neveux, Richard et al. (Nov. 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. 118, 124, 130).

Neyman, Jerzy (1937). "Outline of a theory of statistical estimation based on the classical theory of probability". In: *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 236.767, pp. 333–380 (cit. on p. 83).

Nicola, Andrina, Adam Amara, and Alexandre Refregier (2019). "Consistency tests in cosmology using relative entropy". In: *Journal of Cosmology and Astroparticle Physics* 2019.01, pp. 011–011. DOI: 10.1088/1475-7516/2019/01/011. URL: https://doi.org/10.1088%2F1475-7516%2F2019%2F01%2F011 (cit. on pp. 67, 69).

Okumura, Teppei et al. (June 2016). "The Subaru FMOS galaxy redshift survey (FastSound). IV. New constraint on gravity theory from redshift space distortions at z ˜ 1.4". In: *PASJ* 68.3, 38, p. 38. DOI: 10.1093/pasj/psw029. arXiv: 1511.08083 [astro-ph.CO] (cit. on p. 122).

Peacock, John A (1999). *Cosmological physics*. Cambridge university press (cit. on pp. 37, 40, 41, 49).

Peacock, John A. et al. (2001). "A Measurement of the cosmological mass density from clustering in the 2dF Galaxy Redshift Survey". In: *Nature* 410, pp. 169–173. DOI: 10.1038/35065528. arXiv: astro-ph/0103143 (cit. on p. 51).

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 p. 42).

Penzias, Arno A and Robert Woodrow Wilson (1965). "A measurement of excess antenna temperature at 4080 Mc/s." In: *The Astrophysical Journal* 142, pp. 419–421 (cit. on p. 48).

Percival, Will J. and Martin White (Feb. 2009). "Testing cosmological structure formation using redshift-space distortions". In: *MNRAS* 393.1, pp. 297–308. DOI: 10.1111/j.1365-2966.2008.14211.x. arXiv: 0808.0003 [astro-ph] (cit. on pp. 46, 50, 110, 118, 124).

Percival, Will J. et al. (Nov. 2001). "The 2dF Galaxy Redshift Survey: the power spectrum and the matter content of the Universe". In: *MNRAS* 327.4, pp. 1297–1306. DOI: 10.1046/j.1365-8711.2001.04827.x. arXiv: astro-ph/0105252 [astro-ph] (cit. on p. 51).

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 p. 55).

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 p. 55).

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. 31, 55).

Pezzotta, A. et al. (July 2017). "The VIMOS Public Extragalactic Redshift Survey (VIPERS). The growth of structure at 0.5 ¡ z ¡ 1.2 from redshift-space distortions in the clustering of the PDR-2 final sample". In: *A&A* 604, A33, A33. DOI: 10.1051/0004-6361/201630295. arXiv: 1612.05645 [astro-ph.CO] (cit. on p. 122).

Philcox, Oliver H. E. et al. (Jan. 2021). "Determining the Hubble constant without the sound horizon: Measurements from galaxy surveys". In: *Phys.Rev.D* 103.2, 023538, p. 023538. DOI: 10.1103/PhysRevD.103.023538. arXiv: 2008.08084 [astro-ph.CO] (cit. on pp. 107, 130).

Pichon, C. et al. (Sept. 2001). "Inversion of the Lyman α forest: three-dimensional investigation of the intergalactic medium". In: *Monthly Notices of the Royal Astronomical Society* 326.2, pp. 597–620. ISSN: 0035-8711. DOI: 10.1046/j.1365-8711.2001.04595.x. eprint: https://academic.oup.com/mnras/article-pdf/326/2/597/3249991/326-2-597.pdf. URL: https://doi.org/10.1046/j.1365-8711.2001.04595.x (cit. on p. 83).

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 p. 55).

Planck Collaboration et al. (Nov. 2014a). "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. 48).

Planck Collaboration et al. (2014b). "Planck intermediate results. XVI. Profile likelihoods for cosmological parameters". In: *A&A* 566, A54, A54. DOI: 10.1051/0004-6361/201323003. arXiv: 1311.1657 [astro-ph.CO] (cit. on pp. 56, 82, 87).

Planck Collaboration et al. (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. 56, 66).

Planck Collaboration et al. (2018). "Planck 2018 results. VI. Cosmological parameters". In: *arXiv e-prints*, arXiv:1807.06209, arXiv:1807.06209. arXiv: 1807.06209 [astro-ph.CO] (cit. on pp. 30, 31, 38, 66, 74, 75, 77, 82, 102).

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 p. 48).

— (Sept. 2020b). "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 p. 49).

Planck Collaboration et al. (Sept. 2020c). "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. 49, 56).

Porqueres, Natalia et al. (Oct. 2019). "Inferring high-redshift large-scale structure dynamics from the Lyman-α forest". In: *A&A* 630, A151, A151. DOI: 10.1051/0004-6361/201936245. arXiv: 1907.02973 [astro-ph.CO] (cit. on p. 83).

Press, William H et al. (2007). *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press (cit. on pp. 86, 99).

Prodanović, Tijana and Brian D. Fields (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 pp. 73, 75).

Rauch, Michael (1998). "The Lyman alpha forest in the spectra of quasistellar objects". In: *Annual Review of Astronomy and Astrophysics* 36.1, pp. 267–316 (cit. on p. 102).

Raveri, Marco and Wayne Hu (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 pp. 67, 69).

Rees, Martin J (1997). *Before the beginning. Our universe and others* (cit. on p. 36).

Rees, Martin J and Dennis William Sciama (1968). "Large-scale density inhomogeneities in the universe". In: *Nature* 217.5128, pp. 511–516 (cit. on p. 49).

Reid, Beth A. et al. (Nov. 2012). "The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: measurements of the growth of structure and expansion rate at z = 0.57 from anisotropic clustering". In: *MNRAS* 426.4, pp. 2719–2737. DOI: 10.1111/j.1365-2966.2012.21779.x. arXiv: 1203.6641 [astro-ph.CO] (cit. on p. 103).

Reid, M. J., D. W. Pesce, and A. G. Riess (Dec. 2019). "An Improved Distance to NGC 4258 and Its Implications for the Hubble Constant". In: *ApJ* 886.2, L27, p. L27. DOI: 10.3847/2041-8213/ab552d. arXiv: 1908.05625 [astro-ph.GA] (cit. on p. 131).

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. 31, 55).

Riess, Adam G. et al. (2009). "A Redetermination of the Hubble Constant with the Hubble Space Telescope from a Differential Distance Ladder". In: *ApJ* 699.1, pp. 539–563. DOI: 10.1088/0004-637X/699/1/539. arXiv: 0905.0695 [astro-ph.CO] (cit. on pp. 56, 66).

Riess, Adam G. et al. (2011). "A 3% Solution: Determination of the Hubble Constant with the Hubble Space Telescope and Wide Field Camera 3". In: *ApJ* 730.2, 119, p. 119. DOI: 10.1088/0004-637X/730/2/119. arXiv: 1103.2976 [astro-ph.CO] (cit. on pp. 56, 66).

Riess, Adam G. et al. (2016). "A 2.4% Determination of the Local Value of the Hubble Constant". In: *ApJ* 826.1, 56, p. 56. DOI: 10.3847/0004-637X/826/1/56. arXiv: 1604.01424 [astro-ph.CO] (cit. on pp. 56, 66).

Riess, Adam G. et al. (2018). "Milky Way Cepheid Standards for Measuring Cosmic Distances and Application to Gaia DR2: Implications for the Hubble Constant". In: *ApJ* 861.2, 126, p. 126. DOI: 10.3847/1538-4357/aac82e. arXiv: 1804.10655 [astro-ph.CO] (cit. on pp. 56, 66).

Riess, Adam G. et al. (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. 56, 66, 74, 75).

Riess, Adam G. et al. (Feb. 2021). "Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Milky Way Cepheids Confirm Tension with ΛCDM". In: *ApJ* 908.1, L6, p. L6. DOI: 10.3847/2041-8213/abdbaf. arXiv: 2012.08534 [astro-ph.CO] (cit. on pp. 56, 130).

Robertson, H. P. (Nov. 1935). "Kinematics and World-Structure". In: *ApJ* 82, p. 284. DOI: 10.1086/143681 (cit. on p. 26).

— (Apr. 1936a). "Kinematics and World-Structure II." In: *ApJ* 83, p. 187. DOI: 10.1086/143716 (cit. on p. 26).

— (May 1936b). "Kinematics and World-Structure III." In: *ApJ* 83, p. 257. DOI: 10.1086/143726 (cit. on p. 26).

Rogers, Keir K. et al. (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 pp. 95, 123).

Ross, Ashley J. et al. (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 pp. 66–68, 102, 122).

Rubin, Vera C. and Jr. Ford W. Kent (Feb. 1970). "Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions". In: *ApJ* 159, p. 379. DOI: 10.1086/150317 (cit. on p. 30).

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. 49).

Samushia, Lado et al. (Apr. 2014). "The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: measuring growth rate and geometry with anisotropic clustering". In: *MNRAS* 439.4, pp. 3504–3519. DOI: 10.1093/mnras/stu197. arXiv: 1312.4899 [astro-ph.CO] (cit. on p. 103).

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. 52).

Schmidt, Maarten (Apr. 1965). "Large Redshifts of Five Quasi-Stellar Sources." In: *ApJ* 141, p. 1295. DOI: 10.1086/148217 (cit. on p. 52).

Schöneberg, Nils et al. (July 2021). "The H_0 Olympics: A fair ranking of proposed models". In: *arXiv e-prints*, arXiv:2107.10291, arXiv:2107.10291. arXiv: 2107.10291 [astro-ph.CO] (cit. on p. 131).

Schuhmann, R. L., B. Joachimi, and H. V. Peiris (June 2016). "Gaussianization for fast and accurate inference from cosmological data". In: *MNRAS* 459, pp. 1916–1928. DOI: 10.1093/mnras/stw738. arXiv: 1510.00019 (cit. on p. 71).

Scolnic, D. M. et al. (June 2018). "The Complete Light-curve Sample of Spectroscopically Confirmed SNe Ia from Pan-STARRS1 and Cosmological Constraints from the Combined Pantheon Sample". In: *ApJ* 859.2, 101, p. 101. DOI: 10.3847/1538-4357/aab9bb. arXiv: 1710.00845 [astro-ph.CO] (cit. on pp. 31, 55).

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 p. 54).

Shah, Paul, Pablo Lemos, and Ofer Lahav (Sept. 2021). "A buyer's guide to the Hubble Constant". In: *arXiv e-prints*, arXiv:2109.01161, arXiv:2109.01161. arXiv: 2109.01161 [astro-ph.CO] (cit. on p. 131).

Sheth, Ravi K. and Giuseppe Tormen (Sept. 1999). "Large-scale bias and the peak background split". In: *MNRAS* 308.1, pp. 119–126. DOI: 10.1046/j.1365-8711.1999.02692.x. arXiv: astro-ph/9901122 [astro-ph] (cit. on p. 49).

Silk, Joseph (1968). "Cosmic black-body radiation and galaxy formation". In: *The Astrophysical Journal* 151, p. 459 (cit. on p. 43).

Sivia, Devinderjit and John Skilling (2006). *Data analysis: a Bayesian tutorial*. OUP Oxford (cit. on pp. 57, 59, 61, 62, 69).

Skilling, John (2004). *Nested Sampling*. Ed. by Rainer Fischer, Roland Preuss, and Udo Von Toussaint. Vol. 735. American Institute of Physics Conference Series, pp. 395–405. DOI: 10.1063/1.1835238 (cit. on pp. 60, 62, 86, 91).

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. 54, 103, 133).

Slosar, Anže et al. (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. 55, 82, 87, 102, 103).

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. 48).

Soltis, John, Stefano Casertano, and Adam G. Riess (Feb. 2021). "The Parallax of ω Centauri Measured from Gaia EDR3 and a Direct, Geometric Calibration of the Tip of the Red Giant Branch and the Hubble Constant". In: *ApJ* 908.1, L5, p. L5. DOI: 10.3847/2041-8213/abdbad. arXiv: 2012.09196 [astro-ph.GA] (cit. on p. 131).

Sunyaev, RA and Ya B Zel'Dovich (1980). "Microwave background radiation as a probe of the contemporary structure and history of the universe". In: *Annual review of astronomy and astrophysics* 18.1, pp. 537–560 (cit. on p. 49).

Sunyaev, Rashid A and Ya B Zeldovich (1970). "Small-scale fluctuations of relic radiation". In: *Astrophysics and Space Science* 7.1, pp. 3–19 (cit. on p. 49).

Tamone, Amélie et al. (Oct. 2020). "The completed SDSS-IV extended baryon oscillation spectroscopic survey: growth rate of structure measurement from anisotropic clustering analysis in configuration space between redshift 0.6 and 1.1 for the emission-line galaxy sample". In: *MNRAS* 499.4, pp. 5527–5546. DOI: 10.1093/mnras/staa3050. arXiv: 2007.09009 [astro-ph.CO] (cit. on p. 130).

Taruya, Atsushi, Takahiro Nishimichi, and Shun Saito (Sept. 2010). "Baryon acoustic oscillations in 2D: Modeling redshift-space power spectrum from perturbation theory". In: *Phys.Rev.D* 82.6, 063522, p. 063522. DOI: 10.1103/PhysRevD.82.063522. arXiv: 1006.0699 [astro-ph.CO] (cit. on p. 108).

Tegmark, Max (Nov. 1997). "Measuring Cosmological Parameters with Galaxy Surveys". In: *Phys.Rev.Lett* 79.20, pp. 3806–3809. DOI: 10.1103/PhysRevLett.79.3806. arXiv: astro-ph/9706198 [astro-ph] (cit. on p. 112).

Tegmark, Max, Andy N. Taylor, and Alan F. Heavens (May 1997). "Karhunen-Loève Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets?" In: *ApJ* 480.1, pp. 22–35. DOI: 10.1086/303939. arXiv: astro-ph/9603021 [astro-ph] (cit. on p. 112).

Trotta, Roberto (2007). "Applications of Bayesian model selection to cosmological parameters". In: *MNRAS* 378.1, pp. 72–82. DOI: 10.1111/j.1365-2966.2007.11738.x. arXiv: astro-ph/0504022 [astro-ph] (cit. on p. 84).

Trotta, Roberto (2008). "Bayes in the sky: Bayesian inference and model selection in cosmology". In: *Contemporary Physics* 49.2, pp. 71–104. DOI: 10.1080/00107510802066753. eprint: https://doi.org/10.1080/00107510802066753. URL: https://doi.org/10.1080/00107510802066753 (cit. on pp. 60, 69, 84).

Vargas Magaña, Mariana et al. (2013). "SDSS-III Baryon Oscillation Spectroscopic Survey: Analysis of Potential Systematics in Fitting of Baryon Acoustic Feature". In: *arXiv e-prints*, arXiv:1312.4996, arXiv:1312.4996. arXiv: 1312.4996 [astro-ph.CO] (cit. on p. 82).

Verde, Licia, Pavlos Protopapas, and Raul Jimenez (2013). "Planck and the local Universe: Quantifying the tension". In: *Physics of the Dark Universe* 2.3, pp. 166 –175. ISSN: 2212-6864. DOI: https://doi.org/10.1016/j.dark.2013.09.002. URL: http://www.sciencedirect.com/science/article/pii/S2212686413000319 (cit. on p. 69).

Viel, Matteo, Martin G Haehnelt, and Volker Springel (2004). "Inferring the dark matter power spectrum from the Lyman α forest in high-resolution QSO absorption spectra". In: *Monthly Notices of the Royal Astronomical Society* 354.3, pp. 684–694 (cit. on p. 102).

Walker, A. G. (Jan. 1937). "On Milne's Theory of World-Structure". In: *Proceedings of the London Mathematical Society* 42, pp. 90–127. DOI: 10.1112/plms/s2-42.1.90 (cit. on p. 26).

Wang, Yuting, Lixin Xu, and Gong-Bo Zhao (2017). "A Measurement of the Hubble Constant Using Galaxy Redshift Surveys". In: *ApJ* 849.2, 84, p. 84. DOI: 10.3847/1538-4357/aa8f48. arXiv: 1706.09149 [astro-ph.CO] (cit. on p. 76).

Wang, Ziqi, Marco Broccardo, and Junho Song (2017). "Hamiltonian Monte Carlo Methods for Subset Simulation in Reliability Analysis". In: *arXiv e-prints*, arXiv:1706.01435, arXiv:1706.01435. arXiv: 1706.01435 [stat.CO] (cit. on p. 88).

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. 56).

Welch, B. L. (1939). "Note on Discriminant Functions". In: *Biometrika* 31.1/2, pp. 218–220. ISSN: 00063444. URL: http://www.jstor.org/stable/2334985 (cit. on p. 85).

Whelan, John and Jr. Iben Icko (Dec. 1973). "Binaries and Supernovae of Type I". In: *ApJ* 186, pp. 1007–1014. DOI: 10.1086/152565 (cit. on p. 55).

Zarrouk, Pauline et al. (June 2018). "The clustering of the SDSS-IV extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample: measurement of the growth rate of structure from the anisotropic correlation function between redshift 0.8 and 2.2". In: *MNRAS* 477.2, pp. 1639–1663. DOI: 10.1093/mnras/sty506. arXiv: 1801.03062 [astro-ph.CO] (cit. on pp. 118, 133).

Zhang, Bonnie R. et al. (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. 66).

Zwicky, F. (Jan. 1933). “Die Rotverschiebung von extragalaktischen Nebeln”. In: *Helvetica Physica Acta* 6, pp. 110–127 (cit. on p. 30).

— (Oct. 1937). “On the Masses of Nebulae and of Clusters of Nebulae”. In: *ApJ* 86, p. 217. DOI: 10.1086/143864 (cit. on p. 30).