The Alcock-Paczyński effect from Lyman-α forest correlations: Analysis validation with synthetic data

Andrei Cuceu,[1,2,3,4]⋆ Andreu Font-Ribera,[5,4] Paul Martini,[2,1] Benjamin Joachimi,[4] Seshadri Nadathur,[6] James Rich,[7] Alma X. González-Morales,[8] Hélion du Mas des Bourboux[9] and James Farr[4]

[1]*Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, Ohio 43210, USA*
[2]*Department of Astronomy, The Ohio State University, Columbus, Ohio 43210, USA*
[3]*Department of Physics, The Ohio State University, Columbus, Ohio 43210, USA*
[4]*Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, United Kingdom*
[5]*Institut de Física d'Altes Energies, The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona), Spain*
[6]*Institute of Cosmology and Gravitation, University of Portsmouth, Burnaby Road, Portsmouth, PO1 3FX, United Kingdom*
[7]*IRFU, CEA, Université Paris-Saclay, F91191 Gif-sur-Yvette, France*
[8]*Departamento de Física, DCI, Campus León, Universidad de Guanajuato, 37150, León, Guanajuato, México*
[9]*Department of Physics and Astronomy, University of Utah, 115 S 1400 E, Salt Lake City, UT 84112, USA*

ABSTRACT

The three-dimensional distribution of the Lyα forest has been extensively used to constrain cosmology through measurements of the baryon acoustic oscillations (BAO) scale. However, more cosmological information could be extracted from the full shapes of the Lyα forest correlations through the Alcock-Paczyński (AP) effect. In this work, we prepare for a cosmological analysis of the full shape of the Lyα forest correlations by studying synthetic data of the extended Baryon Oscillation Spectroscopic Survey (eBOSS). We use a set of one hundred eBOSS synthetic data sets in order to validate such an analysis. These mocks undergo the same analysis process as the real data. We perform a full-shape analysis on the mean of the correlation functions measured from the one hundred eBOSS realizations, and find that our model of the Lyα correlations performs well on current data sets. We show that we are able to obtain an unbiased full-shape measurement of $D_M/D_H(z_{\text{eff}})$, where D_M is the transverse comoving distance, D_H is the Hubble distance, and z_{eff} is the effective redshift of the measurement. We test the fit over a range of scales, and decide to use a minimum separation of $r_{\text{min}} = 25~h^{-1}$Mpc. We also study and discuss the impact of the main contaminants affecting Lyα forest correlations, and give recommendations on how to perform such analysis with real data. While the final eBOSS Lyα BAO analysis measured $D_M/D_H(z_{\text{eff}} = 2.33)$ with 4% statistical precision, a full-shape fit of the same correlations could provide a $\sim 2\%$ measurement.

Key words: large-scale structure – distance scale – methods: data analysis

1 INTRODUCTION

Over the last few decades, probes of the large-scale structure (LSS) of the Universe have become one of the main tools for studying its expansion history and the properties of its constituents. Measuring the scale of the baryon acoustic oscillations (BAO) feature is currently the most widely used method for extracting cosmological information from LSS (e.g. Eisenstein et al. 2005; Cole et al. 2005; Beutler et al. 2011; Ross et al. 2015; Alam et al. 2017; Abbott et al. 2019; Alam et al. 2021; Abbott et al. 2022). BAO produces distinct features in the two-point statistics of LSS probes, which allows us to disentangle its signal from the other cosmological, astrophysical, and instrumental effects present. In BAO analyses, other cosmological information is effectively ignored by marginalizing over the broadband of the two-point statistics.

A common way of measuring this extra cosmological information

is to instead model the full shape of the correlation function or power spectrum in order to measure the Alcock-Paczyński (AP) effect and redshift space distortions (RSD; e.g. Blake et al. 2011; Reid et al. 2012; Beutler et al. 2012; Samushia et al. 2014). However, such analyses are more difficult than BAO analyses, because accurately modelling these effects requires a good understanding of the astrophysical and instrumental contaminants present in the broadband of two-point statistics. Furthermore, linear theory approaches that may work well on large scales start to break down due to non-linear growth on smaller scales. These difficulties have been successfully addressed in the case of two-point statistics of discrete tracers, leading to numerous full-shape analyses in the literature (e.g. Blake et al. 2011; Reid et al. 2012; Beutler et al. 2012; Samushia et al. 2014; Satpathy et al. 2017; Beutler et al. 2017; Gil-Marín et al. 2020; Bautista et al. 2021; Hou et al. 2021; Neveux et al. 2020). On the other hand, in the case of the three-dimensional (3D) distribution of the Lyman-α (Lyα) forest, this type of measurement has not yet been performed (Cuceu et al. 2021).

The Lyα forest is made up of absorption lines blueward of the

⋆ E-mail: cuceu.1@osu.edu

Lyα emission peak in spectra of high-redshift quasars (e.g. Lynds 1971; Rauch 1998). This absorption is caused by neutral hydrogen in the intergalactic medium between the quasar and us. Due to the expansion of the Universe, the Lyα line is progressively redshifted as the photons travel towards us. This means the distribution of Lyα absorption can be used to study the structure and evolution of the Universe. The forest has long been used as a tool for cosmology through its one-dimensional[1] two-point statistics (e.g. McDonald et al. 2006; Seljak et al. 2005; Viel et al. 2010; Palanque-Delabrouille et al. 2015, 2020). With the advent of large spectroscopic surveys such as the Baryon Oscillation Spectroscopic Survey (BOSS) and its successor, extended BOSS (eBOSS), the 3D distribution of the Lyα forest has also become a useful source of cosmological information (Busca et al. 2013; Slosar et al. 2013; Kirkby et al. 2013; Font-Ribera et al. 2014; Delubac et al. 2015; 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). However, as mentioned above, it has only been used to measure BAO.

The BOSS and eBOSS Lyα cosmological analyses have progressively improved our understanding of the different processes that contribute to the auto-correlation function of Lyα transmitted flux, and its cross-correlation with the quasar distribution. The first physical model of the two correlations was introduced with the BOSS data release (DR) 12 analyses (Bautista et al. 2017; du Mas des Bourboux et al. 2017). The main effects that had to be understood were the distortion caused by fitting the quasar continuum (see Bautista et al. 2017; Pérez-Ràfols et al. 2018 for the modern treatment and Slosar et al. 2011; Font-Ribera et al. 2012; Blomqvist et al. 2015 for further discussion of the effect), the contamination due to metal absorption in the forest (Pieri et al. 2010, 2014), and the contamination due to high column density (HCD) absorbers (Font-Ribera et al. 2012; Font-Ribera & Miralda-Escudé 2012). Starting with the eBOSS DR14 analyses, Lyα absorption in the Lyβ region (blueward of the Lyβ peak) was also used, first through its correlation with Lyα absorption in the Lyα region (de Sainte Agathe et al. 2019), and then through its cross-correlation with quasars (Blomqvist et al. 2019). The model of the Lyα correlation functions has been progressively improved through these analyses, and we now may be in a position where we can perform a full-shape analysis. The goal of this article is to test this. In particular, we want to study the performance of this model when it comes to fitting the full-shape of the eBOSS DR16 Lyα forest correlation functions.

As mentioned above, a common way of extracting cosmological information from the full shape of two-point statistics is to measure the AP effect and RSD. The AP effect arises due to the fiducial cosmology used to transform redshifts and angles into comoving coordinates. Any differences between this fiducial cosmology and the true cosmology will introduce an extra anisotropy in the auto- and cross-correlations (Alcock & Paczyński 1979). Assuming we can accurately model other sources of anisotropy (e.g. RSD, HCDs, metals, distortion due to continuum fitting), the remaining anisotropy can be used to constrain the background cosmological model. This method is already used in analyses of the BAO, but these only use a small range of scales around the BAO peak. Cuceu et al. (2021) have shown that measuring the AP effect using the full shape is a promising avenue for extracting more cosmological information from the Lyα forest. RSD measurements themselves are also generally used to measure the growth rate of cosmic structure, but for the Lyα auto-correlation this is degenerate with an unknown RSD bias parameter

(McDonald 2003; Slosar et al. 2011; Givans & Hirata 2020; Chen et al. 2021). Cuceu et al. (2021) did show that it is possible to measure the growth rate from the quasar RSD parameter in a joint analysis of the Lyα auto-correlation and its cross-correlation with quasars. However, this would not be a significant measurement with eBOSS. Therefore, in this work we only focus on the AP information, and leave RSD measurements for future work using larger surveys such as the Dark Energy Spectroscopic Instrument (DESI Collaboration et al. 2016).

Our objective in this work is to analyse the full shape of the Lyα forest correlation functions in simulated data in order to understand how well current models perform relative to the expected precision of the eBOSS DR16 data. This analysis is meant as a first step towards a full-shape analysis of the eBOSS Lyα 3D correlation functions. The full-shape measurement using real eBOSS data is presented in a separate article (Cuceu et al. 2023).

This article is structured as follows. In Section 2, we introduce our framework, which includes a description of the mock data sets, the analysis process for computing the 3D correlation functions, and the modelling of the correlations. We show the results of our analysis in Section 3, and discuss their implications for a full-shape analysis using real data in Section 4. We summarize and conclude in Section 5.

2 FRAMEWORK

We begin with an overview of the framework we use to validate a full-shape analysis of the Lyα forest 3D correlation functions. We use the set of one hundred eBOSS mock data sets produced by du Mas des Bourboux et al. (2020), which were created using the method introduced by Farr et al. (2020). We describe these mock data sets in Section 2.1, and the measurement of the transmitted flux correlations in Section 2.2. The analysis methodology we use is similar to the one used in past BAO analyses of Lyα BOSS and eBOSS data. The only differences appear at the step of fitting the correlation functions. Instead of only fitting for the BAO peak position, we fit the full shape of the correlations in order to extract the AP parameter. This is based on the method introduced in Cuceu et al. (2021). We describe our modelling framework in Section 2.3.

2.1 Synthetic data-sets

In order to test the model of the Lyα correlation functions, we use the set of one hundred mock data sets introduced in du Mas des Bourboux et al. (2020). These mocks were produced for the BAO analysis of the Lyα eBOSS DR16 data set. Each mock is based on a Gaussian random field generated with the CoLoRe[2] package (Ramírez-Pérez et al. 2022). The Gaussian field is transformed into a log-normal density field, which is Poisson sampled based on an input number density and bias in order to obtain a set of quasar sources. CoLoRe then computes line-of-sight skewers by interpolating the initial Gaussian field and the radial velocity field[3] from each quasar to the centre of the box.

The skewers generated by CoLoRe require significant post-processing in order to turn them into realistic simulations of the Lyα forest. This is performed by the LyaCoLoRe[4] package (Farr et al. 2020). In order to create the large boxes needed to simulate

---

[1] The one dimension here is along the line-of-sight.

[2] https://github.com/damonge/CoLoRe

[3] Computed using the gradient of the Newtonian potential.

[4] https://github.com/igmhub/LyaCoLoRe

an all-sky light-cone to $z = 3.7$ ($\sim 10\ h^{-1}$Gpc in this case), the CoLoRe grid is limited by memory and computational constraints to a resolution of $\sim 2.4\ h^{-1}$Mpc (for the 4906^3 grid used here). However, one of the contributions to the covariance of 3D correlations (and consequently of BAO uncertainties) is related to the amount of small-scale fluctuations in the Lyα forest (McDonald & Eisenstein 2007), often characterized by the one-dimensional power spectrum, or P_{1D}. In order to simulate spectra with a realistic P_{1D}, LyaCoLoRe adds an extra 1D Gaussian field to each line-of-sight. After that, the field undergoes a log-normal transformation, and the output is used to compute the optical depth field (τ) using the fluctuating Gunn-Peterson approximation (FGPA; Bi & Davidsen 1997; Croft et al. 1998). Redshift-space distortions (RSD) are introduced by convolving the real-space optical depth field with a kernel based on the peculiar velocity field. Finally, the transmitted flux in redshift-space (s) is given by $F(s) = \exp[-\tau(s)]$. For more details on this process, and how it is tuned to produce realistic forests, see Farr et al. (2020).

Beyond the Lyα forest transmitted flux fraction, LyaCoLoRe also simulates high column density (HCD) absorbers, as well as absorption by other Lyman lines and metals. HCDs are sampled from the Gaussian density field using an input bias and number density based on literature results. After that, each HCD is allocated a column density given a column density distribution from the literature. In contrast, the other absorption lines are produced using a rescaled version of the Lyα optical depth, which is then mapped to a different observed wavelength based on the rest-frame wavelength of the absorber. In the case of higher Lyman lines, the scaling factors are tuned based on the oscillator strengths of each transition (Iršič et al. 2013; Farr et al. 2020), with the scaling factor of Lyβ being 0.1901. For metal absorbers, the scaling factors are tuned in order to reproduce the level of contamination in the data (Farr et al. 2020; du Mas des Bourboux et al. 2020). A list of the main metal absorbers along with their relative strength can be found in Table 2 of Farr et al. (2020). These contaminants (HCDs, metal lines, and higher-order hydrogen lines) are then added to the simulated Lyα forest transmitted flux fraction.

Each quasar is assigned a random magnitude using an input quasar luminosity function based on Palanque-Delabrouille et al. (2016). This is used to generate an unabsorbed continuum for each quasar by adding a set of emission lines on top of a broken power law (du Mas des Bourboux et al. 2020). Random redshift errors are drawn from a Gaussian distribution with a standard deviation $\sigma_z = 400$ km s^{-1}. The specsim package (Kirkby et al. 2016) is then used to simulate the eBOSS spectral resolution, the pixelization of the spectra, and instrumental noise (du Mas des Bourboux et al. 2020).

While these synthetic data sets were created for BAO studies, in this article we use them to study full-shape analyses. As described above, the major contaminants that are known to affect Lyα forest correlations are being modelled in these mocks. This will allow us to test the performance of our model, and how full-shape analyses are affected by these contaminants. However, one of the main concerns when attempting to use the full-shape of correlation functions for cosmology is the ability to accurately model the non-linear effects on small scales. While we will be able to test some of these using the mocks described above, some ingredients are missing. These include the fact that the small-scale quasar clustering is overestimated (Youles et al. 2022), and missing IGM effects such as pressure smoothing and the impact of UV radiation on the ionization fraction. We will discuss these issues and the applicability of our results in more detail in Section 4.

2.2 From spectra to correlation functions

Once the simulated eBOSS spectra are produced, they undergo the same analysis process as the real data, in order to measure the 3D auto-correlation of Lyα transmitted flux, and its cross-correlation with the quasar distribution. This is described in detail in du Mas des Bourboux et al. (2020). Here, we only give a brief overview of this process, as all the results in this article[5] come from the same mock correlation functions computed for the eBOSS Lyα BAO analysis.

The first step in the analysis process is to compute the flux transmission fluctuation, $\delta_q(\lambda)$, of each quasar q, based on the observed flux, $f_q(\lambda)$:

$$\delta_q(\lambda) = \frac{f_q(\lambda)}{\bar{F}(\lambda)C_q(\lambda)} - 1, \tag{1}$$

where $\bar{F}(\lambda)$ is the global mean transmission, and $C_q(\lambda)$ is the unabsorbed quasar continuum. In general, the true quasar continuum is unknown, and therefore, the data is used to jointly fit for the product $\bar{F}(\lambda)C_q(\lambda)$. This fit also necessarily includes density modes of the size of the forest and larger, which biases δ_q towards zero for each forest. Even though pixels from the same forest are not cross-correlated, this will still bias pixel correlations in nearby forests, resulting in a distortion of the correlation functions.

The distortion of the correlation function can be removed by building a projection, η_{ij}^q, such that:

$$\sum_j \eta_{ij}^q \delta_q^m(\lambda_j) = \sum_j \eta_{ij}^q \delta_q^t(\lambda_j), \tag{2}$$

where δ_q^m and δ_q^t are the measured and true flux fluctuations, j indexes forest pixels before the projection, and i indexes the projected forest pixels. Therefore, the left-hand side of Equation (2) is applied to the measured flux fluctuation field, while the right-hand side is forward modelled in the correlations, as described in Section 2.3. A detailed description of this projection and how it is built can be found in Bautista et al. (2017), and Appendix B of Pérez-Ràfols et al. (2018).

Forests that contain HCDs with column densities log $N_{HI} > 20.3$[6] are given special treatment. Firstly, the regions where the HCD reduces the transmission by more than 20% are masked. Secondly, the absorption wings are corrected using a Voigt profile (du Mas des Bourboux et al. 2020; Noterdaeme et al. 2012). HCDs with column densities log $N_{HI} < 20.3$ remain in the data, and their effect has to be included in the model of the correlations (Section 2.3).

The correlation functions were computed on a grid in comoving coordinates. For two tracers, i and j, at redshifts z_i and z_j, and separated by an angle $\Delta\theta$, the comoving separations along and across the line-of-sight are given by (de Sainte Agathe et al. 2019):

$$r_{||} = [D_c(z_i) - D_c(z_j)] \cos\frac{\Delta\theta}{2}, \tag{3}$$

$$r_\perp = [D_M(z_i) + D_M(z_j)] \sin\frac{\Delta\theta}{2}, \tag{4}$$

where D_M is the transverse comoving distance, $D_c(z) = c \int_0^z dz'/H(z')$ is the radial comoving distance, with $H(z)$ as the Hubble parameter and c as the speed of light. In this work, we will also refer to the (r, μ) parametrization, defined as $r^2 = r_{||}^2 + r_\perp^2$ and

---

[5] With the exception of those in Appendix B.

[6] This is close to the detection threshold of the DLA detector used when analysing real data. du Mas des Bourboux et al. (2020) also tested the DLA detector on the mock spectra in order to validate this approach.

Type	Name	Tracer 1	Tracer 2
Auto	Lyα(Lyα)×Lyα(Lyα)	Lyα flux in the Lyα region	Lyα flux in the Lyα region
	Lyα(Lyα)×Lyα(Lyβ)	Lyα flux in the Lyα region	Lyα flux in the Lyβ region
Cross	Lyα(Lyα)×QSO	Lyα flux in the Lyα region	Quasars
	Lyα(Lyβ)×QSO	Lyα flux in the Lyβ region	Quasars

Table 1. The types and names of the four correlation functions we use in this work, along with the tracers used in each of them. Here, we use "flux" to refer to the transmitted flux fraction. The Lyα and Lyβ regions are defined by the rest-frame intervals $\lambda_{RF} \in [104, 120]$ nm (situated between the Lyα and Lyβ peaks) and $\lambda_{RF} \in [92, 102]$ nm (situated blueward of the Lyβ peak), respectively.

$\mu = r_{||}/r$. du Mas des Bourboux et al. (2020) used a fiducial cosmological model[7] to compute the two distances, which will result in an extra anisotropy in the measured correlations if it differs from the true cosmology. This is the Alcock-Paczyński effect (Alcock & Paczyński 1979).

The correlation function is first computed independently in each HEALPix[8] pixel (Górski et al. 2005). For the eBOSS footprint, there are about 880 pixels (nside = 16), with each covering $3.7 \times 3.7 = 13.4$ deg^2. This corresponds to a 250×250 $(h^{-1}\mathrm{Mpc})^2$ patch at $z_{\mathrm{eff}} = 2.33$. The population of correlations can then be used to compute the mean and covariance of the correlation function of the entire survey. When computing the correlation function in one HEALPix pixel, pairs with forests in neighbouring pixels are still counted. However, given that we are most sensitive to small scales when measuring AP, we assume that the correlation function measurements in each HEALPix pixel are independent for the purposes of computing the covariance matrix. This method of computing the covariance matrix of Lyα forest correlation functions has been validated against other methods by du Mas des Bourboux et al. (2020).

The process for computing the cross-correlation with the quasar distribution, and its covariance matrix, is similar to the one used for the auto-correlation. However, in this case we also distinguish between Lyα flux in front of a quasar, which is assigned negative $r_{||}$, and Lyα flux behind a quasar, which is assigned positive $r_{||}$. This is because the cross-correlation is not symmetric under permutations of the two tracers.

Four types of correlation functions were computed by du Mas des Bourboux et al. (2020) for the eBOSS DR16 Lyα BAO analysis. The first two are the auto-correlation of the Lyα flux in the Lyα region, Lyα(Lyα)×Lyα(Lyα), and its cross-correlation with Lyα flux in the Lyβ region, Lyα(Lyα)×Lyα(Lyβ). The other two are cross-correlations between Lyα forest flux and quasars, and include the cross-correlation of quasars with Lyα flux in the Lyα region: Lyα(Lyα)×QSO, and with Lyα flux in the Lyβ region: Lyα(Lyβ)×QSO. The Lyα region is found between the Lyα and Lyβ peaks, and is defined by the rest-frame interval $\lambda_{RF} \in [104, 120]$ nm. The Lyβ region is blueward of the Lyβ peak, and is defined by the rest-frame interval $\lambda_{RF} \in [92, 102]$ nm. This information is summarized in Table 1.

In this work, we will be extensively using the mean and covariance of the one hundred mock eBOSS correlation functions. We will also refer to this mean as the stacked correlation function. In order to compute these, we first collect all individual correlation subsamples (in HEALPix pixels) from each of the one hundred mocks. The weighted mean of all correlation subsamples over all one hundred mocks gives us the stacked correlation, while its covariance is given

by the covariance of all the subsamples. This stacked correlation was computed for each of the four correlation types.

We show the four stacked correlations in Figure 1, compressed into four μ bins each. Besides the BAO peak at ~ 100 h^{-1}Mpc, the other features are due to metal contamination. This is why they are present in the wedges closer to the line-of-sight and absent from the ones across the line-of-sight. The most prominent of these is the blended metal peak due to SiII(1190) and SiII(1193), which are at comoving separations of $r_{||} = 64$ h^{-1}Mpc and $r_{||} = 56$ h^{-1}Mpc, respectively. In the auto-correlation, the SiIII(1207) peak can also be seen at $r_{||} = 21$ h^{-1}Mpc. Finally, a fourth metal peak due to SiII(1260) is present at $r_{||} = 111$ h^{-1}Mpc, however, it is not visible due to its proximity to the BAO peak.

2.3 Modelling the correlations

Our model for the Lyα forest auto-correlation (Lyα×Lyα) and its cross-correlation with quasars (Lyα×QSO) closely follows that used by the eBOSS collaboration for the BAO analysis of the SDSS DR16 data (du Mas des Bourboux et al. 2020). The main difference is that we use scale parameters to model the full shape of the correlation, instead of restricting their application to scales around the BAO peak. Our model is based on an isotropic template power spectrum, $P_{\mathrm{fid}}(k)$, which is split into a peak (or wiggles) component, $P_{\mathrm{fid}}^{\mathrm{p}}(k)$, and a smooth (or no-wiggles) component, $P_{\mathrm{fid}}^{\mathrm{s}}(k)$, following Kirkby et al. (2013). As shown in Cuceu et al. (2021), this allows us to separate the information we obtain from the BAO peak from that obtained from the rest of the correlation. All the following steps in the analysis are done with both $P_{\mathrm{fid}}^{\mathrm{p}}(k)$ and $P_{\mathrm{fid}}^{\mathrm{s}}(k)$, in parallel. We use the Vega package to perform this analysis.[9]

The Lyα auto and cross power spectra are given by:

$$P_{\mathrm{auto}}(k, \mu_k, z) = b_{\mathrm{Ly}\alpha}'^2 (1 + \beta_{\mathrm{Ly}\alpha}' \mu_k^2)^2 \, G(k, \mu_k) F_{\mathrm{sm}}(k, \mu_k) P_{\mathrm{fid}}(k),$$
(5)

$$P_{\mathrm{cross}}(k, \mu_k, z) = b_{\mathrm{Ly}\alpha}' (1 + \beta_{\mathrm{Ly}\alpha}' \mu_k^2) \, G(k, \mu_k) F_{\mathrm{sm}}(k, \mu_k)$$
$$\times \, b_{\mathrm{QSO}}(1 + \beta_{\mathrm{QSO}} \mu_k^2) \, F_{\mathrm{NL}}(k, \mu_k) P_{\mathrm{fid}}(k),$$
(6)

where $b_{\mathrm{Ly}\alpha}'$ and $\beta_{\mathrm{Ly}\alpha}'$ are the effective Lyα bias and RSD parameter, and b_{QSO} and β_{QSO} are the quasar linear bias and RSD parameter. $G(k, \mu_k) = \mathrm{sinc}(k_{||}R_{||}/2) \, \mathrm{sinc}(k_\perp R_\perp/2)$ accounts for the binning of the correlation function, with $R_{||}$ and $R_\perp$ the radial and transverse bin widths, respectively.

The Lyα absorption is given by a combination of contributions from the diffuse IGM and from high column density (HCD) systems. Both trace the underlying density field, but HCDs correspond to regions with significantly higher clustering. Here, we define HCDs as systems with neutral hydrogen column density above $10^{17.2}$cm^{-2}, which means they include both Lyman limit systems and damped

---

[7] The fiducial cosmology is based on fits to cosmic microwave background anisotropy data from the Planck satellite Planck Collaboration et al. (2016); see Table 2 of du Mas des Bourboux et al. (2020).
[8] https://healpix.sourceforge.io

[9] https://github.com/andreicuceu/vega
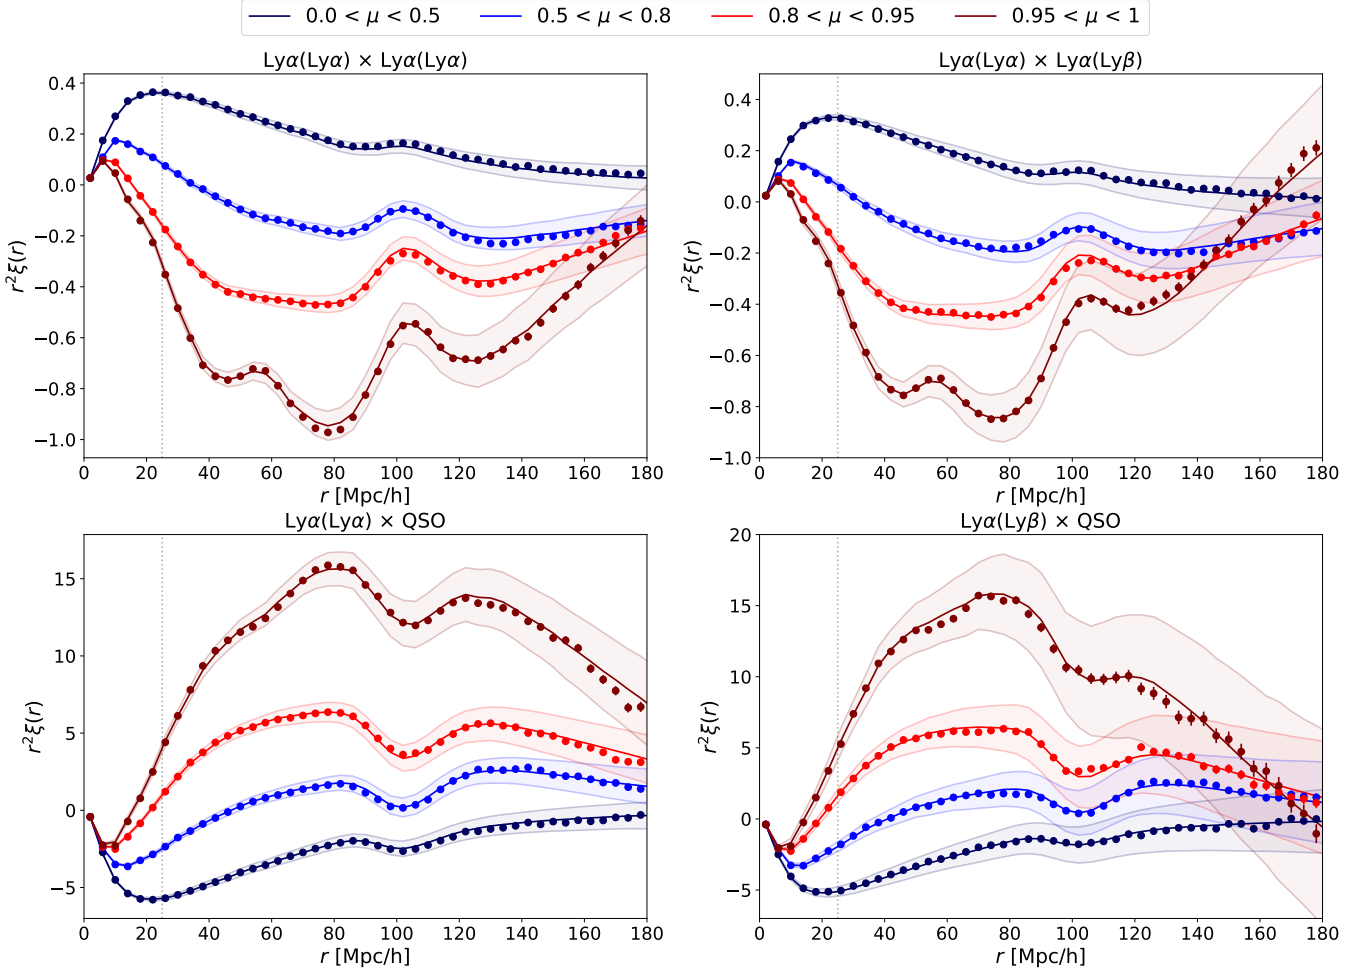
Figure 1. Stacked correlation functions computed from one hundred eBOSS DR16 mocks, compressed into four $\mu = r_{||}/r$ bins. The top panels show the auto-correlation of Lyα forest transmitted flux within the Lyα region on the left and between the Lyα and Lyβ regions on the right. The bottom plots show the cross-correlation between Lyα forest transmitted flux in the Lyα region (left) and Lyβ region (right) and the quasar distribution. The lines represent the joint best fit model of all four correlations. The error-bars are plotted for all points, but except at large separations, they are too small to see. The shaded bands represent the uncertainties from one eBOSS realization.

Lyα (DLA) systems. As described in Section 2.2, large HCDs are detected and masked. However, those with widths smaller than ~ 14 h^{-1}Mpc cannot be detected and remain in the data. These remaining HCDs lead to a broadening effect along the line of sight. Font-Ribera & Miralda-Escudé (2012) showed that this can be modelled through an extra $k_{||}$ dependent term in the effective bias and RSD parameters. This is given by:

$$b'_{Ly\alpha} = b_{Ly\alpha} + b_{HCD}F_{HCD}(k_{||}), \tag{7}$$

$$b'_{Ly\alpha}\beta'_{Ly\alpha} = b_{Ly\alpha}\beta_{Ly\alpha} + b_{HCD}\beta_{HCD}F_{HCD}(k_{||}), \tag{8}$$

where $b_{Ly\alpha}$ and $\beta_{Ly\alpha}$ are the linear bias and RSD parameters associated with the IGM and b_{HCD} and β_{HCD} are the linear bias and RSD parameters associated with HCDs. Following Rogers et al. (2018) and de Sainte Agathe et al. (2019), we model $F_{HCD}(k_{||})$ as an exponential: $F_{HCD} = \exp(-L_{HCD}k_{||})$. The parameter L_{HCD} can be interpreted as the typical length scale of unmasked HCDs (de Sainte Agathe et al. 2019), and is fixed to $L_{HCD} = 10\ h^{-1}$Mpc. We test the impact of this choice in Section 3.3.

The F_{NL} term models the quasar non-linear velocities and statistical quasar redshift error. Following Percival & White (2009), we test two different models for this term, one that introduces a Lorentzian

damping, and one with Gaussian damping. They are given by:

$$F_{NL,Lorentz}^2 = [1 + (k_{||}\sigma_v)^2]^{-1}, \tag{9}$$

$$F_{NL,Gauss}^2 = \exp\left[-\frac{1}{2}(k_{||}\sigma_v)^2\right], \tag{10}$$

where σ_v is a free parameter. Given that the synthetic data is created with Gaussian redshift errors, the Gaussian damping model is more appropriate for modelling the mocks. However, when modelling the real data, the Lorentz damping model is used (du Mas des Bourboux et al. 2017, 2020). This is because the quasar redshift errors generally have long tails that are not well modelled by a simple Gaussian (e.g. Lyke et al. 2020). Therefore, for the purposes of this analysis, we will test both models.

We also introduce Gaussian anisotropic smoothing, $F_{sm}(k, \mu_k)$, to account for the low-resolution of the CoLoRe grid, following Farr et al. (2020). This smoothing model has two free parameters, $\sigma_{||}$ and $\sigma_{\perp}$, which describe the smoothing along and across the line-of-sight, respectively. They are expected to have values of ~ 2.4 h^{-1}Mpc, corresponding to the resolution of the CoLoRe grid.

In order to turn model power spectra into correlation functions, we follow the process described in Kirkby et al. (2013). This involves

a multipole decomposition, followed by a Hankel transform[10] to turn each power spectrum multipole into the corresponding correlation function multipole, and finally the reconstruction of the two-dimensional correlation function from the individual multipoles. Following past Lyα BAO analyses, we use multipole values up to $\ell = 6$. The main components of the theoretical model are the Lyα transmitted flux correlations: $\xi_{\mathrm{Ly}\alpha\times\mathrm{Ly}\alpha}$ for the auto-correlation, and $\xi_{\mathrm{Ly}\alpha\times\mathrm{QSO}}$ for the cross-correlation with quasars. They are computed from the two power spectra in Equations (5) and (6), respectively.

We also model the contamination due to metal absorption, using the same models as for the Lyα auto and cross-correlation. We compute models for the correlations between each metal line and Lyα ($\xi_{\mathrm{Ly}\alpha\times m}$), and among all metal pairs ($\xi_{m_1\times m_2}$) for the auto-correlation, and between each metal line and quasars ($\xi_{\mathrm{QSO}\times m}$) for the cross-correlation. This is described in more detail in Appendix A. Each metal line has its own bias and RSD parameter (b_m, β_m). The same Lyα and QSO (b, β) parameters above are also used for the cross-correlations between metals and Lyα and quasars. In this case, we neglect the HCD effects for the Lyα parameters. Following past Lyα BAO analyses, we fix the metal β parameters to 0.5 (Bautista et al. 2017; du Mas des Bourboux et al. 2020).

The full model correlations are then given by:

$$\xi^t_{\mathrm{auto}} = \xi_{\mathrm{Ly}\alpha\times\mathrm{Ly}\alpha} + \sum_m \xi_{\mathrm{Ly}\alpha\times m} + \sum_{m_1,m_2} \xi_{m_1\times m_2}, \qquad (11)$$

$$\xi^t_{\mathrm{cross}} = \xi_{\mathrm{Ly}\alpha\times\mathrm{QSO}} + \sum_m \xi_{\mathrm{QSO}\times m}, \qquad (12)$$

where the sums are performed over the four metal lines introduced in Section 2.2. The distortion due to quasar continuum errors is modelled using distortion matrices that multiply the results of Equations (11) and (12). See Appendix A and du Mas des Bourboux et al. (2020) for more details.

As described at the start of this section, we perform all these steps using both isotropic power spectrum components, $P^p_{\mathrm{fid}}(k)$ and $P^s_{\mathrm{fid}}(k)$, in parallel. Therefore, we compute the projected correlation function for both the peak component, $\hat{\xi}_p$, and the smooth component, $\hat{\xi}_s$. The two model correlations are then combined into the final theoretical model, allowing for their comoving coordinates to vary:

$$\xi(r_{||}, r_\perp) = \hat{\xi}_s(q^s_{||}r_{||}, q^s_\perp r_\perp) + \hat{\xi}_p(q^p_{||}r_{||}, q^p_\perp r_\perp), \qquad (13)$$

where $r_{||}$ and $r_\perp$ are comoving separations along and across the line-of-sight, respectively. These separations are rescaled using the scale parameters ($q^p_{||}, q^p_\perp$) for the peak, and ($q^s_{||}, q^s_\perp$) for the smooth component. The two peak scale parameters are equivalent to the ($\alpha_{||}, \alpha_\perp$) parameters used in BAO analyses.

In the case of the cross-correlation, we also allow for a systematic shift in the $r_{||}$ coordinate of the cross-correlation through a nuisance parameter, $\Delta r_{||}$. This means the coordinate transform is given by $r_{||} \rightarrow q_{||}(r_{||} + \Delta r_{||})$. We introduce this parameter in order to find if it has any impact on the scale parameters of interest, which could lead to a potential bias when performing the analysis on real data, as it was found to be non-zero in past Lyα BAO analyses (e.g. Blomqvist et al. 2019; du Mas des Bourboux et al. 2020).
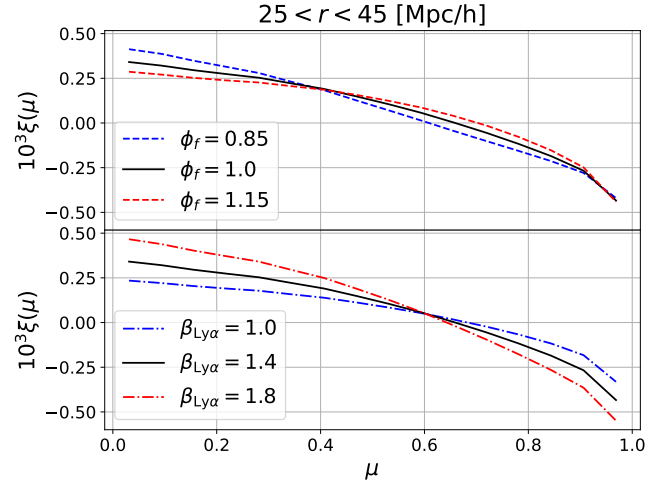
Figure 2. Models of the Lyα forest auto-correlation function in a shell in isotropic separation r, shown as a function of the cosine of the line-of-sight angle, μ. The top panel shows the effect of changing the value of ϕ, the parameter measuring the Alcock-Paczyński effect, while the bottom plot shows the effect of changing the value of the Lyα RSD parameter, $\beta_{\mathrm{Ly}\alpha}$. This shows that the two effects (AP and RSD) lead to different changes in the correlation as a function of μ, which explains why they are not degenerate.

2.4 Modelling the Alcock-Paczyński effect

In order to isolate the Alcock-Paczyński effect, we transform the scale parameter system, following Cuceu et al. (2021):

$$\phi(z) \equiv \frac{q_\perp(z)}{q_{||}(z)} \quad \text{and} \quad \alpha(z) \equiv \sqrt{q_\perp(z)q_{||}(z)}, \qquad (14)$$

where ϕ is the anisotropic parameter, and α is the isotropic scale parameter. Cuceu et al. (2021) showed that a measurement of ϕ is equivalent to a measurement of the AP parameter $D_M(z)H(z)$. As we have two sets of ($q_{||}, q_\perp$) parameters, we also have two sets of (ϕ, α) parameters, for the peak and the smooth component, respectively. The two peak scale parameters have already been tested and measured in BAO analyses. Therefore, our goal here is to use the simulated data sets described above to study measurements of the anisotropic parameter of the smooth component, ϕ_s.

The two main types of analyses we will perform are described in Table 2. In the first one, which we will refer to as the "split AP measurement", we fit two distinct ϕ parameters for the peak and smooth components separately. In the second one, called "full AP measurement", we fit one ϕ parameter to the full shape of the correlations. When performing this analysis on real data, we want to use the full AP measurement, because it coherently extracts the desired information. The split AP measurement can be used as a consistency check of the method or data. However, it is also useful in the current work, because ϕ_s quantifies the gain in information over a standard BAO-only analysis.

In the top panel of Figure 2, we show the impact of changing ϕ on the correlation function compressed in a shell in isotropic separation and plotted as a function of μ. We choose the smallest separation shell, $25 < r < 45 \ h^{-1}$Mpc, in order to best illustrate the anisotropy pattern we aim to recover when fitting the AP effect. The bottom panel of Figure 2 shows the impact of changing the Lyα forest RSD

---

[10] Performed using the FFTLog algorithm (Hamilton 2000). We use the version implemented in the `mcfit` package (https://github.com/eelregit/mcfit).

parameter $\beta_{\text{Ly}\alpha}$, which also affects the anisotropy of the correlation. This figure shows that AP and RSD produce different anisotropy patterns as a function of μ, which means a 2D measurement of the correlation (i.e., in $r_{||}, r_\perp$ or r, μ) can be used to disentangle the two effects.

The isotropic scale parameter from the smooth component, α_s, also contains some cosmological information due to the scale of matter-radiation equality. However, as described in Cuceu et al. (2021), it is not clear how to extract this information given that the effect is very similar to that produced by the Lyα bias, and likely to be heavily impacted by contaminants. Furthermore, Cuceu et al. (2021) found that the constraint on the isotropic scale from the peak, α_p, is much tighter than the one from the smooth component. This is in contrast to ϕ, where the constraint from the broadband is the one that dominates (see Figure 2 of Cuceu et al. 2021). Finally, Gerardi et al. (2023) recently performed simplified forecasts of a direct cosmological analysis of the Lyα forest correlation functions and found that the information extracted can be traced back to the Alcock-Paczyński effect and redshift space distortions. Therefore, for the purposes of this analysis, we will treat α_s as a nuisance parameter, and marginalize over it.

In Table 3, we give a description of all the free parameters in our standard analysis, along with their priors. We follow past BAO analyses, and assign a Gaussian prior to the RSD parameter of HCDs, based on measurements of DLA clustering (Pérez-Ràfols et al. 2018).

We plot the joint best fit model using the split AP analysis configuration in Figure 1. We fit the correlation functions between $r_{\text{min}} = 25$ h^{-1}Mpc and $r_{\text{max}} = 180$ h^{-1}Mpc following the analysis in Section 3.2.[11] This model works surprisingly well given the extreme circumstances of this test. While the fit statistics clearly indicate the model is not good enough for fitting the stack of 100 eBOSS mocks (the fit probability is $\sim 10^{-11}$), the contrast with the eBOSS uncertainties, shown as shaded bands, visually indicates that our model could be good enough for the statistical precision achievable with current data sets. We will show this using the stack of correlations in Section 3.1 and by analysing individual eBOSS mocks in Section 3.4.

3 RESULTS

A large part of the validation of our analysis will focus on studying full-shape fits using a stack of the Lyα forest correlation functions from the eBOSS mocks. Fitting the stack of the mock correlations is much less computationally expensive compared to separately fitting all one hundred mock correlations. Therefore, this method is advantageous when we want to test the performance of our model or study different versions of the analysis. This is because it allows us to study the trends necessary without as much of an emphasis on the effect of noise. Our main goal in this section is to test if the current modelling approach for the Lyα forest correlation functions is appropriate for full-shape analyses of the eBOSS data. In order to achieve this, we want to gain an understanding of the relation between the statistical precision we expect using the eBOSS data and the systematic bias. We also want to investigate modelling choices such as using one versus two anisotropic parameters, the relative contributions of the auto and cross-correlations, and how our modelling of the contaminants affects our measurements.

3.1 Analysis of stacked correlations

As we have four different correlations, we categorise them into the two auto-correlations of Lyα flux, Lyα(Lyα)×Lyα(Lyα) and Lyα(Lyα)×Lyα(Lyβ), and the two cross-correlations with quasars, Lyα(Lyα)×QSO and Lyα(Lyβ)×QSO. This is because the modelling of the two auto-correlations is identical (and similarly for the two cross-correlations) as we treat the Lyα flux in the Lyβ region the same as the Lyα flux in the Lyα region. In the case of the two cross-correlations, the full system of parameters is degenerate. This degeneracy can only be broken when performing a joint analysis of the auto and cross-correlations, because the auto-correlation helps constrain the Lyα bias and RSD parameters. Therefore, we fix the bias and RSD parameter of quasars ($b_{\text{QSO}}, \beta_{\text{QSO}}$) to a fiducial value when fitting the cross-correlation alone.[12] Note that due to correlations between the AP effect and RSD, this may lead the uncertainty in ϕ to be underestimated. However, we will not use these ϕ constraints from the cross-correlation to make any decision or recommendation in this article. We only compute and show them in order to check consistency and for completeness. Also note that, in general, the parameter constraints from the stack of the correlations will not be the same as the mean of the constraints from individual mocks. This is due to the non-linear dependence of the likelihood on the parameters. However, the stacked correlations are useful for testing the performance of our model.

We assume there is no cross-covariance between any of the four correlations, as du Mas des Bourboux et al. (2020) showed that these are negligible. We use a Gaussian likelihood, and compute posterior distributions using the `Vega` package for modelling and the `PolyChord`[13] package for sampling (Handley et al. 2015a,b). When running `PolyChord`, we initialize a number of live points equal to 25 times the number of parameters, and set the length of the slice sampling chains to be of order of the number of parameters.

In Figure 3, we show the results for ϕ_s, ϕ_p, and α_p from the joint fit of the two auto-correlations alone, the two-cross correlations alone, and the combination of all four correlations, using the split AP analysis. These fits use a minimum separation $r > r_{\text{min}} = 25$ h^{-1}Mpc following the analysis in Section 3.2. We find good agreement between the scale parameters measured from the auto- and cross-correlations.

We use the same cosmological model that was used to construct the mocks to also compute comoving coordinates and the power spectrum template.[14] Therefore, we expect to recover unity for all (ϕ, α) parameters. We do recover the expected ϕ values from both the auto and cross-correlation, using both the peak and smooth components independently. The fact that we are able to constrain ϕ_s means it is not degenerate with any of the other effects that introduce their own anisotropies. For a discussion of these effects see Section 3.3, and for plots of correlations between ϕ_s and important nuisance parameters see Appendix E. Finally, Figure 3 shows that we are indeed able to perform an unbiased analysis of the full-shape of Lyα forest correlations for mocks with known contaminants.

For the isotropic scale of the peak, α_p, the result from the cross is consistent with unity. However, the result from the auto-correlation is higher, leading to the combined constraint being $\sim 0.22\%$ larger than unity. While this appears as a $\sim 2\sigma$ systematic bias in the analysis on the stack of 100 eBOSS mocks, it is a small fraction ($\sim 0.2\sigma$) of

---

[11] Note that we still fit for the bias parameter of the SiIII(1207) line, even though that metal peak is outside our fitted range. This is because its contamination is spread all along the line-of-sight by the distortion matrix, and therefore we have to marginalize over it.

[12] The fiducial value is based on the joint constraint of all four correlations.
[13] https://github.com/PolyChord/PolyChordLite
[14] See Appendix B for tests with different fiducial cosmologies.

Analysis type	Parameters	Description	Cosmological Information
Split AP measurement	ϕ_p	Peak anisotropy	AP from BAO
	ϕ_s	Broadband anisotropy	AP from broadband
	α_p	Peak isotropic scale	BAO scale
	α_s	Broadband isotropic scale	Marginalized out
Full AP measurement	ϕ_f	Full-shape anisotropy	AP from full correlation
	α_p	Peak isotropic scale	BAO scale
	α_s	Broadband isotropic scale	Marginalized out

Table 2. The two types of analysis we perform in this work, along with the scale parameters we measure in each of them and their description.
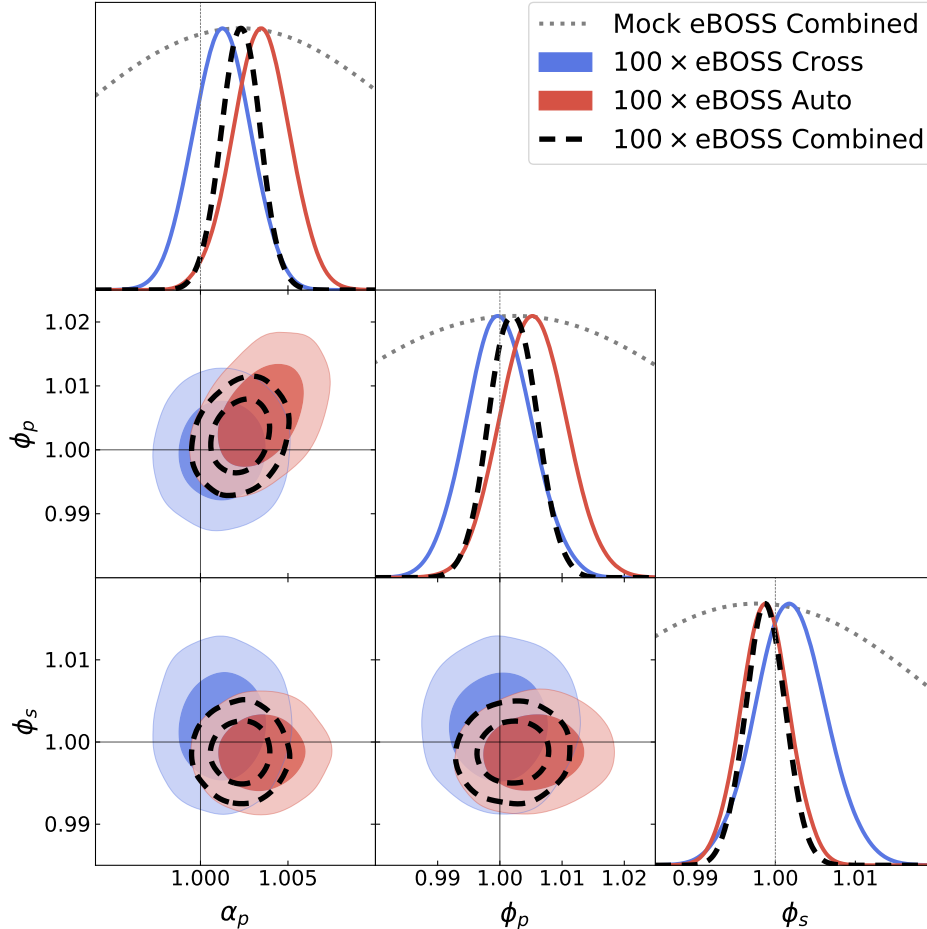
Figure 3. Contour plots of the constraints on scale parameters from the stack of 100 eBOSS mock correlations. The result from the two auto-correlations of Lyα flux, Lyα(Lyα)×Lyα(Lyα) and Lyα(Lyα)×Lyα(Lyβ), is denoted as "Auto", and the one from the two cross-correlations with quasars, Lyα(Lyα)×QSO and Lyα(Lyβ)×QSO, is denoted as "Cross". All scale parameters are expected to be unity, given that the same cosmology was used to construct the mocks and to analyse the data. We also show the expected eBOSS constraint for comparison, but note that the 1σ 2D contours are all larger than the bounds of the plots here. We find that we are able to recover an unbiased measurement of the AP parameter from the full shape of the Lyα forest correlations.

the expected eBOSS statistical uncertainty. We will return to discuss this deviation in Sections 3.4 and 4.

While for the BAO peak parameters, α_p and ϕ_p the constraints obtained from the auto and cross-correlations are very similar, this is not the case for ϕ_s, where the measurement from the auto-correlation is $\sim 30\%$ tighter than that from the cross-correlation. This is in line with the forecasts by Cuceu et al. (2021). We also find that adding the correlations of Lyα flux in the Lyβ forest shrinks the ϕ_s constraint by 16% for the auto-correlation and by 10% for the cross-correlation. In Figure 3 we only focus on the parameters of interest for cosmology. However, unlike BAO parameters, ϕ_s does have correlations with

some nuisance parameters, the most important being $b_{\mathrm{Ly}\alpha}$, $\beta_{\mathrm{Ly}\alpha}$ and b_{HCD}. We show these in Appendix E.

In Figure 4, we compare the split AP measurement with the full AP measurement, based on the joint analysis of all four stacked correlations. We show 1D marginalised posteriors of ϕ measurements from the peak and smooth components (ϕ_p and ϕ_s) and the measurement from the full shape (ϕ_f). We find that the measurements from the peak and smooth components are consistent with each other, with the smooth component providing a constraint that is $\sim 32\%$ tighter than the constraint from the peak component. Note that we do not expect the ϕ_p and ϕ_s measurements to be the same, even though they both

| Parameters	Description	Prior
ϕ, α	Scale parameters (see Table 2)	$U(0.01, 2.0)$
$b_{\mathrm{Ly}\alpha}$	Lyα linear bias	$U(-2.0, 0.0)$
b_{QSO}	QSO linear bias	$U(0.0, 10.0)$
$\beta_{\mathrm{Ly}\alpha}, \beta_{\mathrm{QSO}}$	RSD parameters of Lyα and QSOs	$U(0.0, 5.0)$
b_{HCD}	HCD linear bias	$U(-0.2, 0.0)$
β_{HCD}	RSD parameter for HCDs	$\mathcal{N}(0.5, 0.2^2)$
$\sigma_v \ [\ h^{-1}\mathrm{Mpc}]$	Smoothing for redshift errors and QSO non-linear velocities	$U(0.0, 15.0)$
$\Delta r_\parallel \ [\ h^{-1}\mathrm{Mpc}]$	Shift due to QSO redshift errors	$U(-3.0, 3.0)$
$b_{\mathrm{SiII}(1190)}$	Linear bias of metal absorber	$U(-0.02, 0.0)$
$b_{\mathrm{SiII}(1193)}$	Linear bias of metal absorber	$U(-0.02, 0.0)$
$b_{\mathrm{SiIII}(1207)}$	Linear bias of metal absorber	$U(-0.02, 0.0)$
$b_{\mathrm{SiII}(1260)}$	Linear bias of metal absorber	$U(-0.02, 0.0)$
$\sigma_\parallel \ [\ h^{-1}\mathrm{Mpc}]$	Smoothing along the line-of-sight	$U(0.0, 10.0)$
$\sigma_\perp \ [\ h^{-1}\mathrm{Mpc}]$	Smoothing across the line-of-sight	$U(0.0, 10.0)$

Table 3. List of the free parameters in our model, along with their description and priors. $U(\mathrm{min,max})$ represents a flat prior within that interval, while $\mathcal{N}(\mu, \sigma^2)$ represents a Gaussian prior.
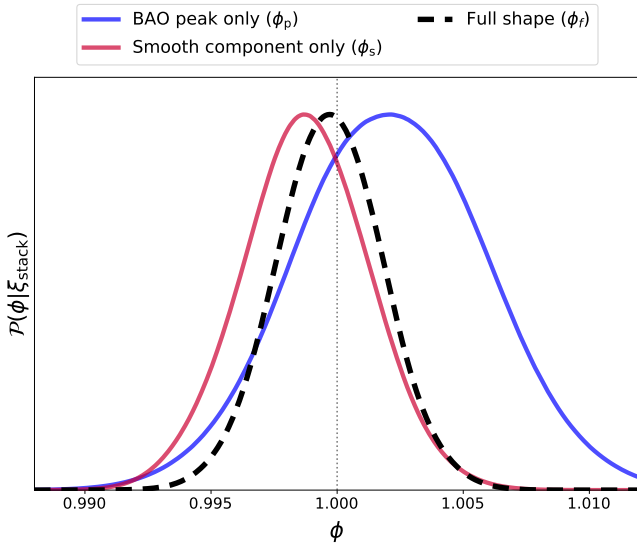
Figure 4. Marginalized posterior distributions of the anisotropic scale parameter ϕ, from a joint analysis of all four Lyα stacked correlation functions. We compare ϕ constraints independently measured from the peak and smooth components (denoted as ϕ_{p} and ϕ_{s}), with constraints measured from the full-shape of the correlations (denoted ϕ_f). All measurements are consistent with the cosmology used to create and analyse the mocks, given by $\phi = 1$.

measure the same effect. This is because they use different parts of the data, resulting in two mostly independent measurements. This is confirmed by the lack of correlation between the two parameters, as shown by their 2D posteriors in Figure 3.

The measurement of ϕ_f is remarkably consistent (to within 0.1σ) with the expected $\phi = 1$. Overall, measuring ϕ from the full-shape of the correlations leads to a $\sim 41\%$ improvement in the 1σ constraints compared to the BAO-only analysis. These numbers will depend on the range of scales we include in our fits, so we next turn our attention to studying the effects of the minimum scale cut we employ.
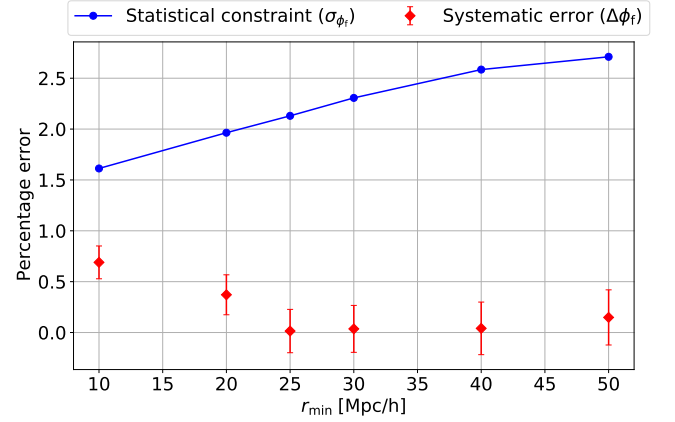
Figure 5. Comparison of the expected statistical constraint on ϕ from eBOSS (σ_ϕ) and the systematic bias ($\Delta\phi$) as a function of minimum separation fitted, r_{min}. The ϕ parameter here is fitted from the full shapes of all four Lyα correlation functions (ϕ_f). The systematic bias is obtained from the difference between the best fit value on the stack of 100 mocks and the true value in the mocks. The statistical constraint is obtained by rescaling the constraint from the stack of mocks to one eBOSS realization. Based on this figure, we choose $r_{\mathrm{min}} = 25 \ h^{-1}\mathrm{Mpc}$ in our baseline analysis.

3.2 Statistical and systematic errors

The results we have shown so far are already quite promising, with ϕ_s and ϕ_f measurements that are consistent with the true cosmology of the mocks ($\phi = 1$). The next step is to quantify this statement. To achieve this, we aim to study how the statistical precision and possible systematic biases depend on the minimum scale included in our analysis. Including more data (especially at small separations) can significantly improve our constraining power (Cuceu et al. 2021), but it may also introduce systematic biases due to imperfect modelling (of e.g. non-linear effects, HCDs, etc.). Therefore, our goal is to find a balance between these two. The key variable in this analysis is the minimum separation scale included in the fit, r_{min}.

In order to quantify the systematic bias of our ϕ measurements, we compute the difference between the best fit value of ϕ from the stacked correlations and the true value in the mocks, i.e., $\Delta\phi = \phi_{\mathrm{best}} - 1$. While we do have the size of the statistical constraints from the stack of the correlations, we want to compare $\Delta\phi$ with the expected constraints from one eBOSS realization. Therefore, we rescale the statistical uncertainty in the fit of the stack by a factor of 10 (given the 100 mocks used to compute the stack) to obtain the expected eBOSS constraint, σ_ϕ.

We compare the statistical constraint and the systematic error on ϕ for different values of r_{min} in Figure 5. We use the full AP setup and show results for ϕ_f from the joint analysis of all four correlation functions. The error-bars on the systematic errors are given by the 1σ statistical constraints of each measurement, which represents the constraining power of the stack of the one hundred mocks.

As expected, σ_ϕ grows significantly as we cut more data (larger r_{min}), with values ranging from 1.6% to 2.7%. The systematic bias starts high for $r_{\mathrm{min}} = 10 \ h^{-1}\mathrm{Mpc}$ (0.7%), but becomes consistent with zero by $r_{\mathrm{min}} = 25 \ h^{-1}\mathrm{Mpc}$. Note that we initially performed this test in r_{min} steps of $10 \ h^{-1}\mathrm{Mpc}$, but chose to compute the $r_{\mathrm{min}} = 25 \ h^{-1}\mathrm{Mpc}$ point after seeing the initial trends. We also performed the same analysis on ϕ_s, using the split AP approach, and found similar results (systematic bias consistent with zero for $r_{\mathrm{min}} \geq 25 \ h^{-1}\mathrm{Mpc}$).

These results confirm that we are able to obtain unbiased measurements of ϕ using the full shapes of the Lyα forest correlation

functions in analyses using $r_{min} \geq 25 \; h^{-1}$Mpc. Therefore, throughout the rest of this work, we will be using $r_{min} = 25 \; h^{-1}$Mpc as our baseline analysis. We also studied the impact of other types of scale cuts, as described in Appendix C. Based on that, we decided to follow BAO analyses and use a maximum scale of $r_{max} = 180$ h^{-1}Mpc, and not impose any anisotropic cuts.

3.3 Modelling of contaminants

The results we have seen so far indicate that our current modelling of the relevant contaminants in the Lyα forest correlations is good enough to recover unbiased full-shape measurements of the AP effect with eBOSS. Before moving to the study of analyses on individual eBOSS mocks, we wish to take a closer look at the individual components of our model. In particular, we want to understand how our measurements are affected by these individual models, and how sensitive we are to different analyses choices. Of particular interest for AP are those effects that introduce their own anisotropies.

The most important effect we have to account for is that of the distortion due to continuum fitting. This effect is modelled by introducing a projection (Equation 2; Bautista et al. 2017) which removes the affected modes in both our data and our models. These distortions are limited to large scales in the power spectrum (low values of $k_\parallel$), but affect the correlation function on all scales (Blomqvist et al. 2015; Bautista et al. 2017). Given the comparison in Figure 1 between the stacked correlations and the best fit model, combined with the fact that we can recover unbiased values of ϕ_s and ϕ_f (Figure 4), we conclude that our modelling of this distortion is accurate enough for current datasets.

While large high column density (HCD) absorbers are masked in the input spectra, those with a typical width smaller than ~ 14 h^{-1}Mpc remain (de Sainte Agathe et al. 2019; du Mas des Bourboux et al. 2020). The current model of this effect (Equations 7 and 8) follows the description introduced by Font-Ribera & Miralda-Escudé (2012), combined with an exponential shape for the $F_{HCD}(k_\parallel)$ function (based on the analysis by Rogers et al. 2018). Note that this shape assumes a model for the column density distribution function of HCDs based on observations (Noterdaeme et al. 2009; Prochaska et al. 2010; Noterdaeme et al. 2012; Zafar et al. 2013), which are very limited for column densities $\log N_{HI} < 20$. This model takes into account the joint contribution of IGM and HCD absorption at the level of two-point functions (which includes the Lyα and HCD auto-correlations and their cross-correlation). However, Font-Ribera & Miralda-Escudé (2012) found that the three-point function, $\langle \delta_{Ly\alpha} \delta_{HCD} \delta_{Ly\alpha} \rangle$, also has a significant contribution. Furthermore, the process of masking the larger HCDs could also bias our correlation measurements, because we are preferentially masking regions with high clustering. These two potential sources of systematic bias have been long known (Slosar et al. 2011; Font-Ribera & Miralda-Escudé 2012). However, so far they have not had a significant impact. This appears to still be the case with full-shape analyses of eBOSS, taking $r_{min} = 25 \; h^{-1}$Mpc. However, they could be contributing to the bias we observe for $r_{min} < 25 \; h^{-1}$Mpc. We have also tested the impact of marginalizing over the L_{HCD} parameter (instead of fixing it to 10 h^{-1}Mpc), and found that it has no impact on ϕ_s measurements.

Contamination due to metal absorption results in clear extra features in the Lyα correlation functions along the line-of-sight. These features correspond to correlations at vanishing separation between metal absorption and either Lyα absorption (auto) or quasars (cross). These correlations get assigned to the wrong bins in comoving coordinates because we interpret everything as Lyα absorption. This is a simple enough process that we can replicate it in the model

through the metal matrix formalism described in Appendix A. However, it means that correlations that used to be at small separations are now spread across the correlation along the line-of-sight. This could result in potential systematic biases, as imperfect modelling of the small scales can now have an impact even on large scales. The results shown above do not indicate this is a problem for eBOSS.

We also note that unlike HCDs, metals were added to these mocks using a re-scaled version of the Lyα optical depth. In reality, metal absorption is associated with the circum- and intra-galactic medium more than the IGM (Pérez-Ràfols et al. 2022). Therefore, these mocks only provide a rough approximation for the clustering of metals. This is especially relevant in the case of metal RSD, as the model has metal RSD parameters fixed to $\beta_m = 0.5$, following du Mas des Bourboux et al. (2020). Synthetic data sets with more realistic metal prescriptions are needed to fully understand their impact on AP measurements.

Quasar redshift errors can be a significant source of uncertainty and also potentially bias both BAO and AP measurements (Youles et al. 2022). For the auto-correlation, quasar redshifts are only used to define the rest-frame wavelength range when computing the continuum. Therefore, redshift errors result in extra spectral diversity, which increases the noise in the measurement of the flux transmission field (du Mas des Bourboux et al. 2020). On the other hand, the cross-correlation is smeared along the line-of-sight, and a potential systematic offset between negative and positive $r_\parallel$ is introduced. We model the first effect using either a Gaussian or a Lorentzian model with a free parameter, as described in Section 2.3. The second effect is modelled through an extra free parameter, $\Delta r_\parallel$, to account for a potential offset. Furthermore, the Gaussian or Lorentzian models are also meant to model the non-linear quasar velocities which give rise to the Finger-of-God effect. The mocks used here do contain statistical redshift errors, modelled as a Gaussian, and the Finger-of-God effect. As analyses on real data have preferred the Lorentzian model, we tested both of them. However, we found no significant shift in any of the scale parameters of interest (α_p, ϕ_p, ϕ_s).

In Figure 6, we compare the baseline results on ϕ with results where we remove the model for each of these contaminants in turn. This means that the contaminants are still present in the measured correlations, but we do not model them. These results are computed from a joint analysis of all four Lyα stacked correlation functions. We also show the baseline result rescaled to correspond to one eBOSS realization. We study three cases: one where we remove the HCD model, one where we do not model metals, and one where we do not model quasar redshift errors. All three introduce some systematic bias in the measurement of ϕ_s, with the largest resulting from the removal of the HCD model, leading to a $\sim 1\%$ shift. On the other hand, for ϕ_p, only the removal of the metal modelling results in a significant bias. This figure shows the relative importance of including each of these contaminants in the model. However, note that even the largest bias (with no HCD modelling at all) represents a shift of only $\sim 0.5\sigma$ of the expected statistical constraint for eBOSS.

3.4 Analysis of individual mock realizations

We now turn our attention to analysing individual eBOSS mocks. The analysis of the stacked correlations has been useful for understanding how well our model performs in the context of a full-shape analysis. We found that, using a minimum separation $r_{min} = 25 \; h^{-1}$Mpc, we are able to obtain unbiased measurements of the anisotropic scale parameter from the full shape of the correlations. However, as mentioned above, parameter constraints from the stack of mocks are not, in general, the same as mean constraints from individual mock
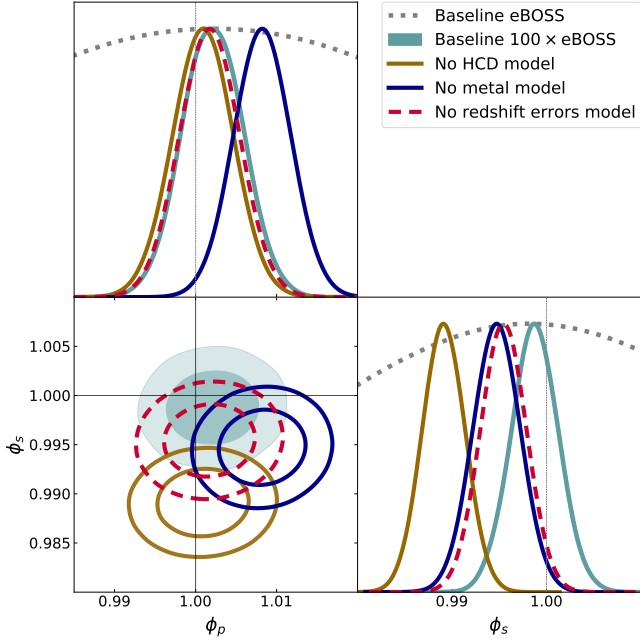
Figure 6. Posterior distributions of the anisotropic scale parameter, ϕ, from the peak and smooth components, respectively. We compare the baseline result, where we model all contaminants following eBOSS (du Mas des Bourboux et al. 2020), with results where we do not model each one of the contaminants in turn. The three results here are obtained by either removing the HCD modelling, the modelling of the metals or the model for quasar non-linear velocities. In each of these cases, not modelling the relevant contaminant introduces a systematic bias in our results. For ϕ_s, the most significant shift appears when removing the model for HCDs.

Figure 7. Histograms of the best fit values of scale parameters obtained by minimizing the χ^2 over all parameters in each of the 100 eBOSS mocks. The red dashed lines represent the mean of the populations, while the dashed-black lines show the expected values of the scale parameters (unity). The top plots show fits of ϕ_p (left) and ϕ_s (right). These show that we are able to obtain unbiased measurement of ϕ from the full shape of the Lyα forest correlations, which is confirmed by the bottom right plot, showing results for ϕ_f. The bottom left plot shows results from fits of the BAO isotropic scale parameter α_p. The mean of the population in this case is slightly higher (by $\sim 0.25\sigma$) than the expected value.

analyses. Therefore, we now wish to test if our conclusions above are also accurate when analysing individual eBOSS mocks.
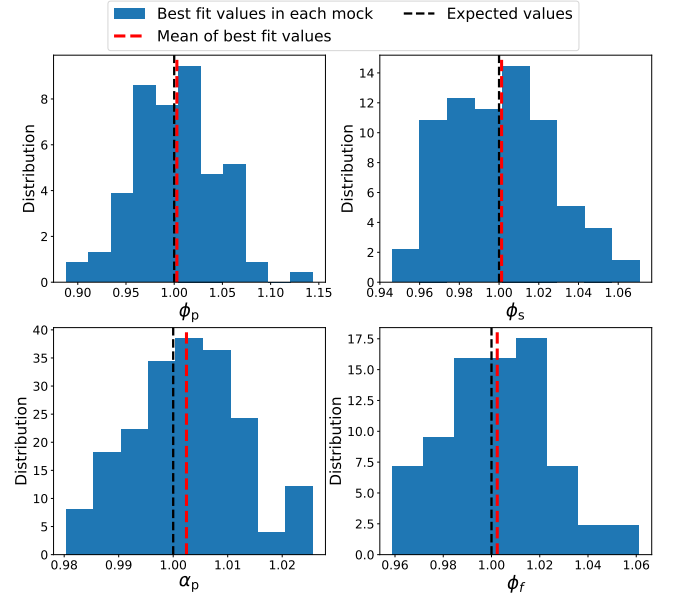
The results shown so far have been obtained by using a sampler to compute full posterior distributions. However, this becomes too computationally expensive for a full set of 100 mocks.[15] Therefore, in this section, we compute the best fit values of the parameters using a minimizer[16]. We estimate the statistical uncertainty in this result using the second derivative in parameter space around the best fit point, assuming a Gaussian likelihood. This allows us to obtain rough estimates of population statistics from the set of mocks. However, note that both the best fits and the constraints obtained using this method are ultimately just an approximation, because we are not properly marginalizing over the nuisance parameters (see Cuceu et al. 2020). We only use this method here because a more complex and accurate method (i.e. sampling) is not practical. We discuss this in more detail in Appendix D.

Our modelling of individual eBOSS mock correlations closely follows the model we used for the stack of correlations, as introduced in Section 2.3. We run the `iminuit` minimizer on each of the one hundred mocks. This provides us with both the best fit value of each parameter and with its Gaussian uncertainty. We compute two versions of the results, using both the split AP analysis with ϕ_p and ϕ_s, and the full AP analysis with ϕ_f.

We show histograms of the best fit values of the scale parameters

in Figure 7. The top plots show results for ϕ_p and ϕ_s, both of which are consistent with the expected value of unity. The means of the populations are $\mu_{\phi_\mathrm{p}} = 1.0028$ and $\mu_{\phi_\mathrm{s}} = 1.0013$, corresponding to shifts from unity of $\sim 0.06\sigma$ and $\sim 0.05\sigma$, respectively.[17] These results reinforce the conclusions from Section 3.1: we are indeed able to recover unbiased measurements of the AP effect using the full shapes of the correlations.

The bottom row of Figure 7 shows results for the isotropic BAO scale (α_p) and the full-shape anisotropic parameter (ϕ_f). The means of the populations in this case are $\mu_{\alpha_\mathrm{p}} = 1.0024$ and $\mu_{\phi_f} = 1.0024$, corresponding with shifts from unity of $\sim 0.23\sigma$ and $\sim 0.11\sigma$, respectively. While these shifts are more significant compared to the ones from ϕ_p and ϕ_s, they are still relatively small, in line with the results from Section 3.1.

4 DISCUSSION

In this work, we have used a set of one hundred eBOSS Lyα forest mocks to study how well we could perform a full-shape analysis of the eBOSS DR16 Lyα correlation functions. We found that we are able to recover unbiased measurements of the anisotropic scale parameter, ϕ, from a joint analysis of the full shapes of the four Lyα correlations when using a minimum separation of $r_\mathrm{min} = 25\ h^{-1}$Mpc. The only significant systematic biases on ϕ were found when either we did not

---

[15] Running the sampler for a joint analysis of all four Lyα correlations with all the contaminants requires $\sim 2 - 3 \times 10^3$ CPU hours, which leads to a rough estimate of $\sim 2 - 3 \times 10^5$ CPU hours to analyse all 100 mocks.
[16] We use the `iminuit` package (Dembinski et al. 2021) for minimization.

[17] The σ values here are based on the expected eBOSS uncertainty given by the standard deviation of these histograms. This is also consistent with the scaled uncertainty from the stack of mocks.

model one of the major contaminants (Section 3.3), or we included scales smaller than 25 h^{-1}Mpc (Section 3.2).

We did find an apparent bias of $0.22\% - 0.24\%$ in the isotropic scale of the peak component, α_p, both when analysing the stack of one hundred mocks (Section 3.1) and when analysing them individually (Section 3.4). This corresponds to a $\sim 0.21\sigma - 0.25\sigma$ bias for eBOSS.[18] du Mas des Bourboux et al. (2020) also found a systematic shift of similar magnitude for the BAO parameters, using the same correlation function measurements with a different parametrization. In the same publication, du Mas des Bourboux et al. (2020) measured BAO using real data, and tested the robustness of the result by adding broadband polynomials meant to marginalize over any unaccounted broadband systematic effects. They found no significant shift in the BAO parameters when performing this test. Therefore, they concluded the BAO results are robust. Such a test cannot be replicated in the context of a full-shape analysis, because the broadband polynomials would erase the information in ϕ_s. However, our model allows the separate fit of BAO (α_p and ϕ_p) and broadband (ϕ_s) information. Therefore, when analysing real data, we can benchmark our BAO result using the results of the eBOSS BAO analysis.

In Section 3.3, we have discussed the impact of the major contaminants that are modelled in the mock data and shown the importance of modelling them accurately. However, we have only touched on the effects that we model and observe in the mock data. While these represent the most important known contaminants for the Lyα forest 3D correlations (which is why they are modelled in the mock data), they are not the only possible sources of systematic uncertainty. Here we wish to discuss some other possible sources of contamination, how they have been treated in the past, and also to give recommendations on which of them warrant attention when performing this analysis on real data.

The mock data used in this work is based on a Gaussian density field with quasars randomly sampled using its log-normal transformation. While this is a good enough approximation on large scales, it does not correctly reproduce the quasar clustering on small scales (Farr et al. 2020). This makes it difficult to accurately test the model on small scales, especially for the cross-correlation. The problem is also exacerbated by the fact that both the distortion due to continuum fitting and the metal contamination take effects that are at small scales and spread them to larger scales along the line-of-sight. Incorrect modelling of the small scales can thus have an impact on all scales, so this analysis would benefit from more realistic mocks to test this effect. However, as these are not currently available, we have to leave such a study for future work.

For the analysis of the Lyα auto-correlation from real data, an empirical model based on the work by Arinyo-i-Prats et al. (2015) is used to fit small-scale non-linearities.[19] This model has five free parameters whose values were measured using hydro-simulations. However, no equivalent model exists for the cross-correlation. A possible approximation would be to use the same model for the cross-correlation as well, but it is not clear if the same parameter values would still work, and this would not account for quasar non-linearities (FoG). Recently, Givans et al. (2022) used high-resolution hydrodynamic simulations to study how well current models perform on small scales in both the Lyα power spectrum and the cross-spectrum of Lyα with massive halos. For Lyα, they found that the model by Arinyo-i-Prats

et al. (2015) performs very well up to wavenumbers much larger than those included here. On the other hand, for the cross-spectrum, they found that the standard Lorentzian damping combined with the Arinyo-i-Prats et al. (2015) model cannot accurately fit the small-scale power. Therefore, for the cross-correlation, they recommend using scales larger than 30 h^{-1}Mpc when fitting it with linear theory. However, they did not include the effect of redshift errors, which suppresses power on these small scales, making deviations from linear theory less important. Therefore, in an analysis of real data, it may be worth testing scale cuts for the cross-correlation of both 25 h^{-1}Mpc, as proposed here, and of 30 h^{-1}Mpc, as proposed by Givans et al. (2022).

Comparing results from the auto and cross-correlations could be a good way of testing for systematic biases due to imperfect modelling of the small scales. This is because the two are driven by different effects, with the cross-correlation affected more by redshift errors and quasars, which cluster very strongly and have large non-linear peculiar velocities, and the auto-correlation driven by the IGM absorption (gas pressure effects, thermal broadening, etc.). Furthermore, for the cross-correlation, testing both the Gaussian and Lorentzian damping models would be a good consistency check for the analysis on real data.

While preparing this manuscript, a new study by Youles et al. (2022) found that large redshift errors can introduce spurious correlations along the line-of-sight in both the Lyα auto and cross-correlation with quasars. Such an effect could lead to systematic biases when performing a full-shape analysis such as the one studied here. Therefore, future work is needed to study its impact on full-shape analyses, and to potentially find ways to account for it (e.g. by forward modelling the effect).

An effect that is not included in the mock data is that of BAO broadening due to non-linear growth. However, this effect is relatively well understood and modelled in BAO analyses (see Kirkby et al. 2013 for a detailed discussion). Therefore, as long as it is correctly modelled when fitting the real data, we do not consider it a potential issue for a full-shape analysis. The effect of quasar radiation, also known as the transverse proximity effect, is another important source of contamination for the cross-correlation, because it increases the ionization fraction in the gas surrounding the quasar, leading to a decrease in Lyα absorption. However, this effect has been successfully modelled (Font-Ribera et al. 2013), and this model has been included in previous BAO analyses. As this only affects the cross-correlation, comparing results from the auto- and cross-correlation should also be useful here.

The metal contamination introduced in the synthetic data used here only includes the effect of the lines that are close in wavelength to the Lyα peak. This type of contamination is mostly given by the cross-correlation between these absorptions lines and either Lyα absorption (for the auto-correlation) or quasars (for the cross-correlation). However, in BOSS and eBOSS BAO analyses, the contamination due to CIV absorption was also modelled. At the relevant comoving coordinates used here, this is dominated by the auto-correlation of CIV absorption, and therefore only affects the Lyα auto-correlation. It is generally modelled in the same way as the other metal lines, as described in Section 2.3. It has its own free bias parameter. However, this parameter has only been marginally detected, with a significance less than 3σ (du Mas des Bourboux et al. 2020). Therefore, the impact of this contamination is minimal. Nevertheless, we recommend that it be modelled and marginalized over in a full-shape analysis on real data, similarly to Lyα BAO analyses.

Fluctuations of the ionizing UV background introduce a scale dependence to the Lyα bias and RSD parameters (Pontzen 2014;

---

[18] Note that the eBOSS uncertainty used for these values are based on estimates from this work, not the real measurement.

[19] We do not use this model here because the small-scale forest-related non-linearities are not included in the mocks.

Gontcho A Gontcho et al. 2014). This scale dependence was modelled in the Lyα BAO analyses using BOSS DR12 and eBOSS DR14 (Bautista et al. 2017; de Sainte Agathe et al. 2019; Blomqvist et al. 2019), following the framework introduced by Gontcho A Gontcho et al. (2014). However, the effect was not detected at a significant level (only $\sim 2\sigma$), and modelling it did not have an impact on the BAO fits. This indicates that it is probably not of significant concern for a full-shape analysis, but its impact on ϕ measurements should be tested when performing this analysis on real data.

In conclusion, the two effects that likely require more attention are the spurious correlations due to redshift errors (identified by Youles et al. 2022), and the small scale non-linear effects in the cross-correlation. The other effects discussed here have been studied before, and modelled at some level in past Lyα BAO analysis. They were not prioritized by the eBOSS collaboration when creating mocks because they were not found to play an important role in past BAO analyses. However, this may not be true for full-shape analyses. Therefore, an important part of the first full-shape analysis on real data will be to determine if the current models for these secondary effects have an impact on measurements of ϕ_s. If they are found to have an impact, it will motivate the development of more accurate mocks in the future.

5 SUMMARY

The Lyman-α (Lyα) forest provides one of the best tracers of large scale structure (LSS) at high redshifts ($2 < z < 4$). However, 3D correlations of the forest have so far only been used to measure the baryon acoustic oscillations (BAO) scale. Recently, Cuceu et al. (2021) showed that a measurement of the Alcock-Paczyński (AP) effect from the full shapes of Lyα forest correlation functions can lead to significant improvements in cosmological constraints compared to BAO-only measurements. In this work, we use synthetic data of the extended Baryon Oscillation Spectroscopic Survey (eBOSS) in order to test such an analysis.

We model the Lyα auto-correlation, and its cross-correlation with quasars, using a similar approach to that used in past Lyα BAO analyses (Kirkby et al. 2013; Bautista et al. 2017; du Mas des Bourboux et al. 2017, 2020), as described in Section 2.3. The only difference appears from the fact that we use scale parameters to fit the full shapes of the correlations. Our parametrization of the scale parameters splits them into an isotropic scale parameter, α, and a parameter for the anisotropy, ϕ, corresponding to a measurement of the AP effect (Cuceu et al. 2021). As we decompose the template power spectrum into a peak component and a smooth component, we have the option of using separate scale parameters for each. We use two setups for the analysis, as described in Table 2. In the first, we decouple the peak and smooth components by fitting two distinct ϕ parameters (ϕ_p and ϕ_s), while in the second we fit only one ϕ parameter (ϕ_f) for the full shapes of the correlations.

We begin our study by analysing stacked correlation functions from the one hundred mock data sets (see Figure 1), in Section 3.1. We perform a joint analysis of all four correlations, as shown in Figures 3 and 4. We find that the results from the two Lyα auto-correlations are consistent with the results from the two cross-correlations. Furthermore, we find that we are able to recover unbiased measurements of ϕ from the full shapes of the correlations, for both types of analysis described above.

In Section 3.2, we study joint analyses of all four stacked correlations using a range of scale cuts. In particular, we focus on the minimum scale used, r_{min}. We run the full AP analysis using r_{min} values from 10 h^{-1}Mpc to 50 h^{-1}Mpc, as shown in Figure 5. We

then compare the systematic bias of ϕ_f measurements, given by the difference between the best fit value and the expected value in the mocks ($\phi = 1$), with an estimate of the expected statistical constraint from eBOSS. We find that we are able to obtain unbiased measurements of ϕ_f for $r_{min} \geq 25$ h^{-1}Mpc. Therefore, we use $r_{min} = 25$ h^{-1}Mpc for the other parts of the analysis.

In Section 3.3, we discuss the main contaminants present in the Lyα forest correlation functions. We also test a few different setups for the analysis, but find that as long as the main contaminants are modelled correctly, there is no significant bias in full-shape ϕ measurements. We show the relative impact of not modelling each main contaminant in Figure 6.

We study full-shape analyses on the individual correlations in Section 3.4. We used approximate results from a minimizer instead of running a sampler on each of the one hundred mocks, due to computational constraints (see Appendix D for a comparison between the two). We show histograms of the best fit parameter values in Figure 7. Our results here reinforce the earlier conclusions. The means of these distributions are consistent with the values expected from the mocks, $\phi = 1$. This again indicates that we are able to obtain unbiased measurements of ϕ, using the full shapes of the correlations.

Finally, in Section 4 we discuss the implications of our results and give recommendations for how such an analysis should be performed on real data from eBOSS. As we have shown, the most important contaminants are well modelled and do not have a significant impact on ϕ measurements for $r_{min} \geq 25$ h^{-1}Mpc. Therefore, we focus on discussing the effects that are not modelled in the synthetic data, and what tests can be done on the real data in order to gauge their impact. Most of these are quite well understood, have already been modelled and tested in past Lyα BAO analyses, and were not found to have a significant impact. However, these tests should also be replicated in the context of a full-shape analysis.

An eBOSS full-shape analysis, as studied here, could improve AP constraints by as much as a factor of two compared to a BAO-only analysis. We have performed this measurement in a follow-up work presented in Cuceu et al. (2023). While our focus here has been on the information extracted through the AP effect, our validation of full-shape measurements could also prove important to other sources of information present in the broadband of Lyα forest correlations. These include growth of structure measurements using redshift space distortions (Cuceu et al. 2021), and fluctuations in the ionising UV background (e.g. Long & Hirata 2023).

DATA AVAILABILITY

Data supporting this research is available on reasonable request from the corresponding author.

REFERENCES

Abbott T. M. C., et al., 2019, MNRAS, 483, 4866
Abbott T. M. C., et al., 2022, *Phys. Rev.* **D**, 105, 043512
Alam S., et al., 2017, MNRAS, 470, 2617
Alam S., et al., 2021, *Phys. Rev.* **D**, 103, 083533
Alcock C., Paczyński B., 1979, Nature, 281, 358
Arinyo-i-Prats A., Miralda-Escudé J., Viel M., Cen R., 2015, J. Cosmology Astropart. Phys., 2015, 017
Bautista J. E., et al., 2017, A&A, 603, A12
Bautista J. E., et al., 2021, MNRAS, 500, 736
Beutler F., et al., 2011, MNRAS, 416, 3017
Beutler F., et al., 2012, MNRAS, 423, 3430
Beutler F., et al., 2017, MNRAS, 466, 2242
Bi H., Davidsen A. F., 1997, ApJ, 479, 523
Blake C., et al., 2011, MNRAS, 415, 2876
Blomqvist M., et al., 2015, J. Cosmology Astropart. Phys., 2015, 034
Blomqvist M., et al., 2019, A&A, 629, A86
Busca N. G., et al., 2013, A&A, 552, A96
Chen S.-F., Vlah Z., White M., 2021, J. Cosmology Astropart. Phys., 2021, 053
Cole S., et al., 2005, MNRAS, 362, 505
Croft R. A. C., Weinberg D. H., Katz N., Hernquist L., 1998, ApJ, 495, 44
Cuceu A., Font-Ribera A., Joachimi B., 2020, J. Cosmology Astropart. Phys., 2020, 035
Cuceu A., Font-Ribera A., Joachimi B., Nadathur S., 2021, MNRAS, 506, 5439
Cuceu A., Font-Ribera A., Nadathur S., Joachimi B., Martini P., 2023, *Phys. Rev.* **L**, 130, 191003
DESI Collaboration et al., 2016, arXiv e-prints, p. arXiv:1611.00036
Delubac T., et al., 2015, A&A, 574, A59
Dembinski H., et al., 2021, iminuit: Jupyter-friendly Python interface for C++ MINUIT2, Astrophysics Source Code Library, record ascl:2108.024 (ascl:2108.024)
Eisenstein D. J., et al., 2005, ApJ, 633, 560
Farr J., et al., 2020, J. Cosmology Astropart. Phys., 2020, 068
Font-Ribera A., Miralda-Escudé J., 2012, J. Cosmology Astropart. Phys., 2012, 028
Font-Ribera A., et al., 2012, J. Cosmology Astropart. Phys., 2012, 059
Font-Ribera A., et al., 2013, J. Cosmology Astropart. Phys., 2013, 018
Font-Ribera A., et al., 2014, Journal of Cosmology and Astro-Particle Physics, 2014, 027
Gerardi F., Cuceu A., Font-Ribera A., Joachimi B., Lemos P., 2023, MNRAS, 518, 2567
Gil-Marín H., et al., 2020, MNRAS, 498, 2492
Givans J. J., Hirata C. M., 2020, *Phys. Rev.* **D**, 102, 023515
Givans J. J., et al., 2022, J. Cosmology Astropart. Phys., 2022, 070
Gontcho A Gontcho S., Miralda-Escudé J., Busca N. G., 2014, MNRAS, 442, 187
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Hamilton A. J. S., 2000, MNRAS, 312, 257
Handley W. J., Hobson M. P., Lasenby A. N., 2015a, MNRAS, 450, L61
Handley W. J., Hobson M. P., Lasenby A. N., 2015b, MNRAS, 453, 4384
Harris C. R., et al., 2020, Nature, 585, 357
Hou J., et al., 2021, MNRAS, 500, 1201
Iršič V., et al., 2013, J. Cosmology Astropart. Phys., 2013, 016
Kirkby D., et al., 2013, J. Cosmology Astropart. Phys., 2013, 024
Kirkby D., Bailey S., Guy J., Weaver B. A., 2016, Quick simulations of fiber spectrograph response v0.5, doi:10.5281/zenodo.154130
Lewis A., 2019, arXiv e-prints, p. arXiv:1910.13970

Long H., Hirata C. M., 2023, MNRAS, 520, 948
Lyke B. W., et al., 2020, ApJS, 250, 8
Lynds R., 1971, The Astrophysical Journal, 164, L73
McDonald P., 2003, ApJ, 585, 34
McDonald P., Eisenstein D. J., 2007, *Phys. Rev.* **D**, 76, 063009
McDonald P., et al., 2006, ApJS, 163, 80
Neveux R., et al., 2020, MNRAS, 499, 210
Noterdaeme P., Petitjean P., Ledoux C., Srianand R., 2009, A&A, 505, 1087
Noterdaeme P., et al., 2012, A&A, 547, L1
Palanque-Delabrouille N., et al., 2015, J. Cosmology Astropart. Phys., 2015, 011
Palanque-Delabrouille N., et al., 2016, A&A, 587, A41
Palanque-Delabrouille N., Yèche C., Schöneberg N., Lesgourgues J., Walther M., Chabanier S., Armengaud E., 2020, J. Cosmology Astropart. Phys., 2020, 038
Percival W. J., White M., 2009, MNRAS, 393, 297
Pérez-Ràfols I., et al., 2018, MNRAS, 473, 3019
Pérez-Ràfols I., Pieri M. M., Blomqvist M., Morrison S., Som D., Cuceu A., 2022, arXiv e-prints, p. arXiv:2210.02973
Pieri M. M., Frank S., Weinberg D. H., Mathur S., York D. G., 2010, ApJ, 724, L69
Pieri M. M., et al., 2014, MNRAS, 441, 1718
Planck Collaboration et al., 2016, A&A, 594, A13
Pontzen A., 2014, *Phys. Rev.* **D**, 89, 083010
Prochaska J. X., O'Meara J. M., Worseck G., 2010, ApJ, 718, 392
Ramírez-Pérez C., Sanchez J., Alonso D., Font-Ribera A., 2022, J. Cosmology Astropart. Phys., 2022, 002
Rauch M., 1998, Annual Review of Astronomy and Astrophysics, 36, 267
Reid B. A., et al., 2012, MNRAS, 426, 2719
Rogers K. K., Bird S., Peiris H. V., Pontzen A., Font-Ribera A., Leistedt B., 2018, MNRAS, 476, 3716
Ross A. J., Samushia L., Howlett C., Percival W. J., Burden A., Manera M., 2015, MNRAS, 449, 835
Samushia L., et al., 2014, MNRAS, 439, 3504
Satpathy S., et al., 2017, MNRAS, 469, 1369
Seljak U., et al., 2005, *Phys. Rev.* **D**, 71, 103515
Slosar A., et al., 2011, J. Cosmology Astropart. Phys., 2011, 001
Slosar A., et al., 2013, J. Cosmology Astropart. Phys., 2013, 026
Viel M., Haehnelt M. G., Springel V., 2010, J. Cosmology Astropart. Phys., 2010, 015
Youles S., et al., 2022, MNRAS, 516, 421
Zafar T., Péroux C., Popping A., Milliard B., Deharveng J. M., Frank S., 2013, A&A, 556, A141
de Sainte Agathe V., et al., 2019, A&A, 629, A85
du Mas des Bourboux H., et al., 2017, A&A, 608, A130
du Mas des Bourboux H., et al., 2020, ApJ, 901, 153

APPENDIX A: MODELLING METALS AND THE DISTORTION

As we assume all forest pixels correspond to Lyα absorption, correlations involving metals are assigned to the wrong bin in comoving coordinates due to one (for Ly$\alpha \times m$ and QSO$\times m$) or both (for $m \times m$) tracers being assigned the wrong redshift. This is modelled through a metal matrix which transforms the model correlations computed using the correct separations, $\tilde{r}_{||}$ and $\tilde{r}_\perp$, to the measured model correlations:

$$\xi_A^{\mathrm{metal}} = \sum_B M_{AB} \xi^{\mathrm{metal}}(\tilde{r}_{||}(B), \tilde{r}_\perp(B)), \qquad (A1)$$

where A and B are comoving separation bins, and the metal matrix is given by:

$$M_{AB} = \frac{1}{W_A} \sum_{(i,j) \in A, (i,j) \in B} w_i w_j, \qquad (A2)$$

where $(i, j) \in A$ refers to separations computed using the assigned redshifts (assuming only Lyα absorption), and $(i, j) \in B$ refers to separations computed using the true redshifts of each tracer. The weights, w_i, are given by the inverse variance of δ_q (see du Mas des Bourboux et al. 2020), and $W_A = \sum_{(i,j) \in A} w_i w_j$. Note that we drop the quasar specification from our notation, as the sums are understood to be over all pixels in each comoving separation bin.

In order to model the distortion due to continuum fitting errors, we need to apply the projection η_{ij}^q to the modelled flux fluctuation δ_q^t (Equation 2). This is done by using η_{ij}^q to construct a distortion matrix that is applied to the theoretical model correlations, ξ_{auto}^t and ξ_{cross}^t to give the projected correlation function:

$$\hat{\xi}_A = \sum_B D_{AB} \xi_B^t, \tag{A3}$$

where A is a bin of the data and B is a bin of the model. The distortion matrices for the auto and cross-correlations are given by:

$$D_{AB}^{\text{auto}} = W_A^{-1} \sum_{ij \in A} w_i w_j \sum_{i'j' \in B} \eta_{ii'} \eta_{jj'}, \tag{A4}$$

$$D_{AB}^{\text{cross}} = W_A^{-1} \sum_{ij \in A} w_i w_j \sum_{i'j \in B} \eta_{ii'}. \tag{A5}$$

See Bautista et al. (2017) and Pérez-Ràfols et al. (2018) for more details.

APPENDIX B: IMPACT OF FIDUCIAL COSMOLOGY

Our method requires the use of a fiducial cosmology, both for defining the comoving coordinate grid on which the correlations are computed, and for the template power spectrum that is used to construct the theoretical model for the Lyα correlations. This fiducial cosmology is based on cosmic microwave background anisotropy data from the Planck satellite (Planck Collaboration et al. 2016). In this appendix, we test the impact of this choice of fiducial cosmology.

We recompute the Lyα auto and cross-correlations of one of the mocks using two different values of the matter density parameter, Ω_m. The two values are $\Omega_m = 0.27$ and $\Omega_m = 0.36$. We then perform our standard analysis and measure α_p and ϕ_f, but we use different power spectrum templates to keep the fiducial Ω_m value consistent in each analysis. In order to compare the results, we compute the relevant combinations of distances as derived parameters. Following Cuceu et al. (2021), measurements of ϕ and α_p correspond to:

$$\textbf{AP:} \quad \phi(z) = \frac{D_M(z)H(z)}{[D_M(z)H(z)]_{\text{fid}}}, \tag{B1}$$

$$\textbf{BAO:} \quad \alpha_p(z) = \sqrt{\frac{D_M(z)D_H(z)/r_d^2}{[D_M(z)D_H(z)/r_d^2]_{\text{fid}}}}, \tag{B2}$$

where $D_H = c/H$ and r_d is the size of the sound horizon at the end of the drag epoch. Based on these equations and the relevant fiducial distances for each value of Ω_m, we transform ϕ_f into D_M/D_H, and α_p into $D_M D_H/r_d^2$.

We compare the results for the large and small values of Ω_m with the one obtained from our standard analysis (using $\Omega_m = 0.315$) in Figure B1. All three results give consistent measurements, however, they are not identical. This can be attributed to the fact that pixel pairs are distributed to different bins in the correlation when using different values of Ω_m, because the comoving distances change. This will have an impact on the measured parameters, especially for BAO, where small changes in the correlation measurements (consistent
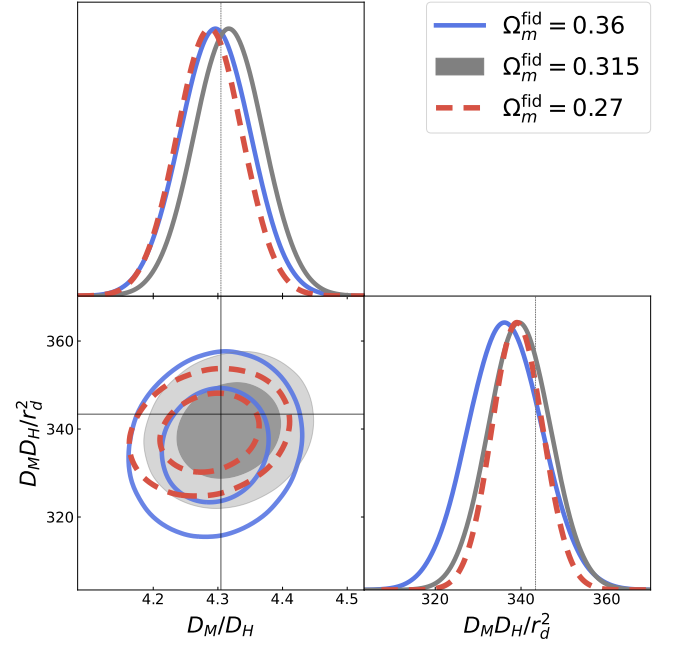
Figure B1. Posterior distributions of the combination of distances we measure for different Ω_m values used in the fiducial cosmology. The combination D_M/D_H corresponds to ϕ_f, while $D_M D_H/r_d^2$ corresponds to α_p. This shows that we measure consistent distances even when using a very different fiducial cosmology.

with the expected noise) can produce large changes in the shape of the posterior (e.g. Cuceu et al. 2020). Furthermore, when changing the fiducial Ω_m, the smallest scale that we fit will be different for the same r_{min}. For example, the smaller value of Ω_m leads to smaller scales being assigned larger comoving separations. This is why the size of the constraints increases as Ω_m increases.[20] Nevertheless, all three results are consistent with each other even for these extreme values of Ω_m.

APPENDIX C: OTHER SCALE CUTS

In Section 3.2 we have focused on the minimum scale fitted, r_{min}. However, we also have to choose the largest scale we fit for, r_{max}. Furthermore, as our measurement is two-dimensional, we could also impose different scale cuts along or across the line-of-sight. We tested different versions of this analysis, which we summarize here.

We show the impact of changing the largest separation included in our fit, r_{max}, in Figure C1. We perform the same type of comparison as we did for r_{min} in Figure 5. The scales we tested range from 140 h^{-1}Mpc to 200 h^{-1}Mpc. We find that measurements of ϕ_f are unbiased for all the r_{max} values we have tried. Furthermore, the improvement in the constraint as we include more data is very weak. Therefore, following eBOSS, we decided to choose $r_{\text{max}} = 180$ h^{-1}Mpc in our standard analysis.

As most of the contaminants of the Lyα forest correlation functions have their largest impact along the line-of-sight, we also decided to

---

[20] This also means that our choice of $r_{\text{min}} = 25$ h^{-1}Mpc is based on the Planck cosmology. Therefore, the analysis should be re-done, and a new minimum scale chosen, if we had reason to believe the real cosmology was very different from the one inferred by Planck.
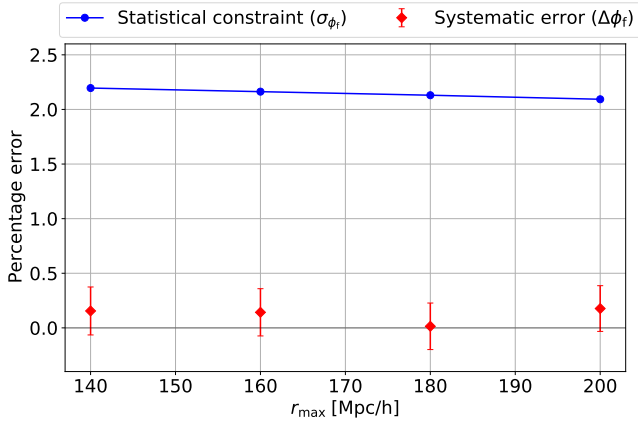
Figure C1. Comparison of the expected statistical constraint on ϕ from eBOSS (σ_ϕ) and the systematic bias ($\Delta\phi$) as a function of maximum scale fitted, r_{max}. The ϕ parameter here is fitted from the full shapes of all four Lyα correlation functions (ϕ_f). The systematic bias is obtained from the difference between the best fit value on the stack of 100 mocks and the true value in the mocks. The statistical constraint is obtained by rescaling the constraint from the stack of mocks to one eBOSS realization.
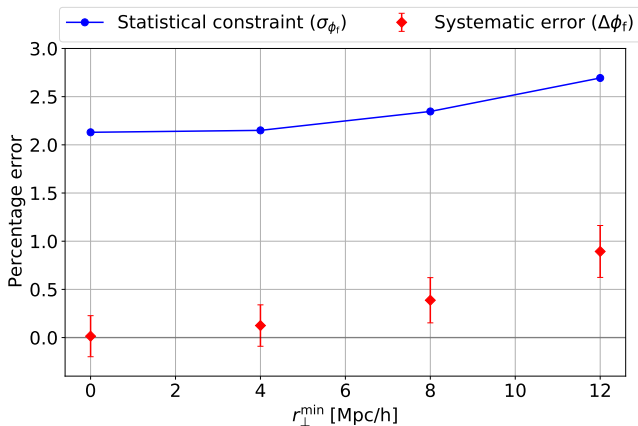
Figure C2. Same as Figure C1, but as a function of the minimum scale across the line-of-sight, $r_\perp^{min}$. We are progressively cutting the columns along the line-of-sight, where the impact of the contaminants is highest. The systematic error on ϕ is increasing as we cut more data, which indicates that we are not able to accurately constrain the nuisance parameters that account for the contaminants when ignoring the information along the line-of-sight.

test a cut in $r_\perp^{min}$. This is the smallest transverse scale that we fit, and values $r_\perp^{min} > 0\ h^{-1}$Mpc mean that we are removing the columns in the correlation that lie along the line-of-sight. The systematic error and expected statistical constraint on ϕ_f for different values of $r_\perp^{min}$ are show in Figure C2. We find that removing the first correlation column along the line-of-sight ($r_\perp^{min} = 4\ h^{-1}$Mpc) still results in an unbiased measurement. However, removing more data after that starts to introduce a significant systematic bias.[21] This could be explained by the fact that removing the columns along the line-of-sight is degrading our ability to fit the parameters that control the contaminants (e.g. HCDs and metals). Even though these contaminants have their

strongest impact along the line-of-sight, they have a small impact on the correlation even at large transverse separations. Therefore, it is still important to model them accurately. Given these results, we chose to not impose any cut in $r_\perp^{min}$.

APPENDIX D: POPULATION CONSTRAINTS

As described in Section 3.4, when running the analysis on each of the one hundred mocks, we are restricted to using a minimizer due to computational constraints. This means that the reported constraints on the measured parameters are just estimates due to two reasons. Firstly, they assume the posterior is Gaussian around the best fitting value. Secondly, the best fit for the parameters of interest (ϕ_f, α_p) is found without properly marginalizing over the nuisance parameters.[22] While these approximations might be correct for the BAO parameters due to their lack of correlations with the nuisance parameters (Cuceu et al. 2020), ϕ_s definitely has such correlations (Appendix E). Therefore, in order to study accurate parameter constraints from the population of mocks, we need to use a sampler.

For the purposes of this appendix, we ran the sampler on correlations computed only using the Lyα region (no Lyβ region) from uncontaminated synthetic data sets. These mocks do not include the effects of HCDs, metal absorption or redshift errors (they are also described in du Mas des Bourboux et al. 2020). Therefore, the parameter space is much smaller, allowing for the computation of the full posterior distributions using a sampler. We performed the joint analysis of the auto and cross-correlations for all one hundred uncontaminated mocks.

In Figure C3, we show histograms of the constraints on α_p, ϕ_p, and ϕ_s from the population of one hundred mocks. The main result is the one given by the sampler, in blue. We also plot the standard deviation of the best fitting values for each of the parameters over the population of mocks. This shows that the variance of the parameter values obtained from the one hundred realizations is consistent with the measured uncertainty in individual realizations, which means the uncertainties are properly accounted for.

We also compare results from the sampler with the estimates from the minimizer in Figure C3. While they roughly agree for the BAO parameters (α_p and ϕ_p), this is not the case for ϕ_s. The two give different constraints, as shown by the difference between the red and blue histograms. This is not surprising given the assumptions involved in the fitter constraints. However, the variance of the ϕ_s best fits is also different between the two (the red and blue vertical lines). This means that marginalization over the nuisance parameters is important in this case.[23]

While this has been a rough test due to computational constraints, most of the parameters that ϕ_s is correlated with are still present in uncontaminated mocks ($b_{Ly\alpha}$, $\beta_{Ly\alpha}$ and β_{QSO}). The main omission is b_{HCD}. Therefore, it would be interesting to perform this analysis with the effect of HCDs modelled in the mocks. Alternatively, future studies should perform the analysis on the fully contaminated mocks if there is access to enough computational resources.

---

[21] Note that we are plotting the absolute value of $\Delta\phi$. Therefore, the trend of increasing systematic bias in Figure C2 does not mean the bias is all in the same direction.

[22] The best fit value of ϕ_s reported by the minimizer is conditional on the best fitting values of all the nuisance parameters. However, in the Bayesian framework, the quantity we are interested in is the value of ϕ_s conditional on the posterior distributions of the nuisance parameters. Only in the second case is the uncertainty due to these parameters properly taken into account.
[23] In a frequentist framework, a possible solution for this would be to use these results to compute non-Gaussian $\Delta\chi^2$ values in order to set more accurate confidence intervals on ϕ_s.
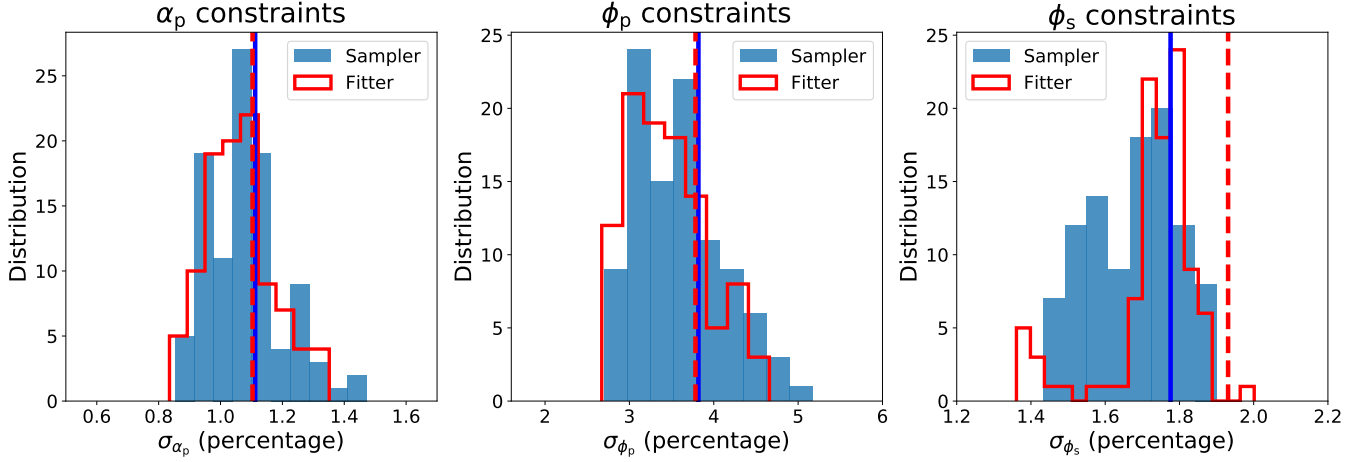
Figure C3. Histograms of the 68% constraints on α_p, ϕ_p and ϕ_s when fitting the correlations measured from one hundred uncontaminated eBOSS mocks. We compare results from a fitter which assume a Gaussian posterior, with results from a sampler where we compute the full posterior distributions. The vertical lines represent the standard deviation of the best fitting results for each parameter over the one hundred mocks. We show that when using the sampler, the variance of the population best fits is consistent with the individual constraints for all three parameters.

APPENDIX E: PARAMETER CORRELATIONS

The advantage of BAO analyses is that they measure a well-defined feature in the correlation function. BAO scale parameters rarely have any significant correlations with the nuisance parameters, which means they are much more robust to different analysis choices when it comes to modelling contaminants. On the other hand, full-shape analyses rely on the ability to accurately model correlation functions over all the scales of interest. We have already shown in Figure 6 what biases are introduced when not modelling certain contaminants. In this appendix, we go into more detail on the sensitivity of ϕ_s measurements to other nuisance parameters.

In Figure E1 we show the posterior distributions of the main parameters that are correlated with ϕ_s. The two posteriors are from the fit to the two Lyα auto-correlations, and from the joint fit of all four correlations. The cross-correlations between Lyα and quasars cannot be included in such a comparison because the full system of parameters is degenerate.[24] The most important correlations are between ϕ_s and the Lyα bias and RSD parameters, $b_{Ly\alpha}$ and $\beta_{Ly\alpha}$. These two are highly correlated between themselves and with the HCD bias, b_{HCD}. The Lyα auto-correlation provides most of the information on these parameters. On the other hand, the information on β_{QSO} comes from the Lyα-QSO cross-correlation. However, a measurement of $\beta_{Ly\alpha}$ (from the auto) is necessary to constrain it due to their degeneracy in the model of the cross-correlation. While we treat β_{QSO} as a nuisance parameter in this work, future studies could also attempt to extract a measurement of the growth rate of structure from it.

This paper has been typeset from a TEX/LATEX file prepared by the author.

---

[24] See Section 3.1 above and Section 4 of Cuceu et al. (2021)
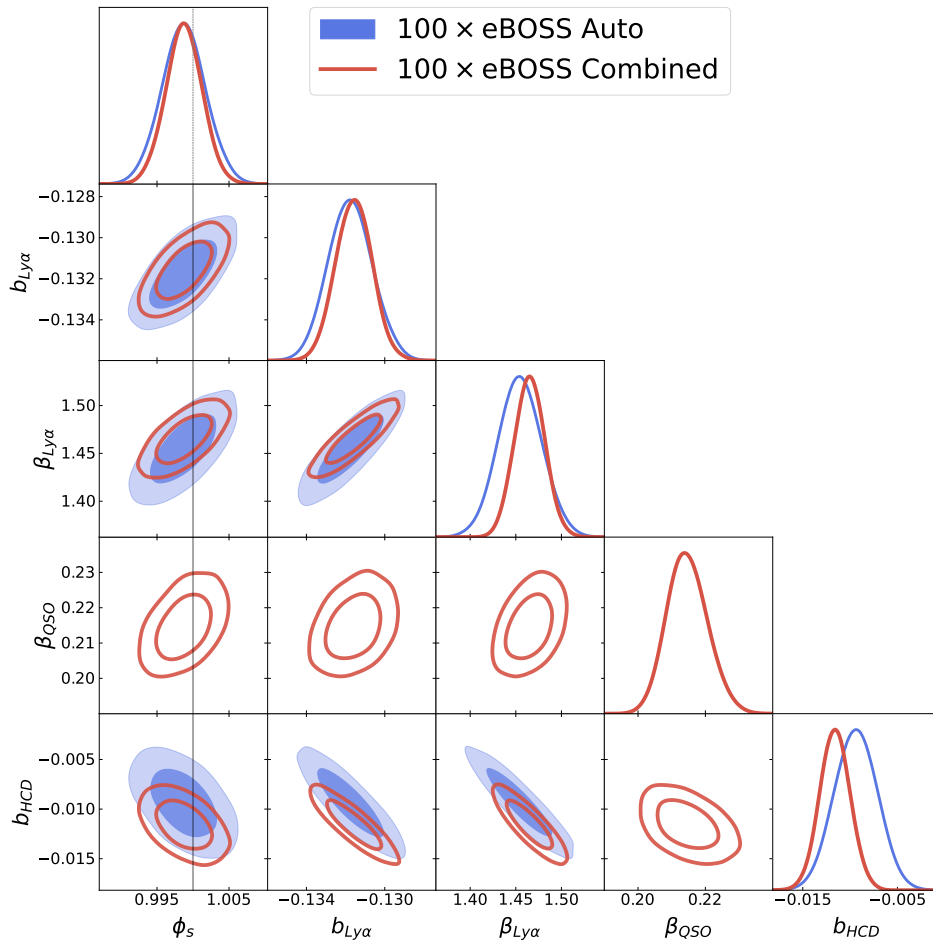
Figure E1. Posterior distributions obtained from fitting the mean of correlation functions computed from one hundred eBOSS mocks. We show results for the joint analysis of the two Lyα auto-correlations, and for the joint analysis of all four correlations. We only plot the parameters that have correlations with ϕ_s, in order to show its sensitivity to the values of these parameters.