Fermilab

DESI DR1 Lyα forest: 3D full-shape analysis and cosmological constraints

FERMILAB-PUB-25-0693-PPD

arXiv:2509.15308

Draft version September 22, 2025
Typeset using LaTeX **twocolumn** style in AASTeX631

DESI DR1 Lyα forest: 3D full-shape analysis and cosmological constraints

Andrei Cuceu [1,2,*] Hiram K. Herrera-Alcantar [3,4] Calum Gordon [5] César Ramírez-Pérez,[5] E. Armengaud [4] A. Font-Ribera [5] J. Guy [1] B. Joachimi,[6] P. Martini [2,7,8] S. Nadathur [9] I. Pérez-Ràfols [10] J. Rich,[4,11] J. Aguilar,[1] S. Ahlen [12] A. Anand [1] S. Bailey [1] A. Bault [1] D. Bianchi [13,14] A. Brodzeller [1] D. Brooks,[6] J. Chaves-Montero [5] T. Claybaugh,[1] K. S. Dawson [15] A. de la Macorra [16] J. Della Costa [17,18] P. Doel,[6] S. Ferraro [1,19] J. E. Forero-Romero [20,21] E. Gaztañaga [22,9,23] S. Gontcho A Gontcho [1,24] A. X. Gonzalez-Morales [25] D. Green [1] G. Gutierrez,[26] C. Hahn [27] M. Herbold [8] K. Honscheid [2,28,8] V. Iršič [29,30,31,32] M. Ishak [33] R. Joyce [18] N. G. Karaçayli [2,7,28,8] D. Kirkby [34] T. Kisner [1] A. Kremin [1] O. Lahav,[6] A. Lambert,[1] C. Lamman [8] M. Landriau [1] J.M. Le Goff,[4] L. Le Guillou [11] M. E. Levi [1] M. Manera [35,5] A. Meisner [18] R. Miquel,[36,5] J. Moustakas [37] A. Muñoz-Gutiérrez,[16] J. A. Newman [38] G. Niz [25,39] N. Palanque-Delabrouille [4,1] W. J. Percival [40,41,42] Matthew M. Pieri [43] C. Poppett,[1,44,19] F. Prada [45] C. Ravoux [46] G. Rossi,[47] E. Sanchez [48] D. Schlegel,[1] M. Schubnell,[49,50] H. Seo [51] J. Silber [1] F. Sinigaglia [52,53] D. Sprayberry,[18] T. Tan [4] G. Tarlé [50] M. Walther [54,55] B. A. Weaver,[18] C. Yèche [4] R. Zhou [1] and H. Zou [56]

[1] *Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA*
[2] *Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA*
[3] *Institut d'Astrophysique de Paris. 98 bis boulevard Arago. 75014 Paris, France*
[4] *IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France*
[5] *Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Edifici Cn, Campus UAB, 08193, Bellaterra (Barcelona), Spain*
[6] *Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, UK*
[7] *Department of Astronomy, The Ohio State University, 4055 McPherson Laboratory, 140 W 18th Avenue, Columbus, OH 43210, USA*
[8] *The Ohio State University, Columbus, 43210 OH, USA*
[9] *Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth, PO1 3FX, UK*
[10] *Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 10, 08930 Barcelona, Spain*
[11] *Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), FR-75005 Paris, France*
[12] *Department of Physics, Boston University, 590 Commonwealth Avenue, Boston, MA 02215 USA*
[13] *Dipartimento di Fisica "Aldo Pontremoli", Università degli Studi di Milano, Via Celoria 16, I-20133 Milano, Italy*
[14] *INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy*
[15] *Department of Physics and Astronomy, The University of Utah, 115 South 1400 East, Salt Lake City, UT 84112, USA*
[16] *Instituto de Física, Universidad Nacional Autónoma de México, Circuito de la Investigación Científica, Ciudad Universitaria, Cd. de México C. P. 04510, México*
[17] *Department of Astronomy, San Diego State University, 5500 Campanile Drive, San Diego, CA 92182, USA*
[18] *NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA*
[19] *University of California, Berkeley, 110 Sproul Hall #5800 Berkeley, CA 94720, USA*
[20] *Departamento de Física, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia*
[21] *Observatorio Astronómico, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia*
[22] *Institut d'Estudis Espacials de Catalunya (IEEC), c/ Esteve Terradas 1, Edifici RDIT, Campus PMT-UPC, 08860 Castelldefels, Spain*
[23] *Institute of Space Sciences, ICE-CSIC, Campus UAB, Carrer de Can Magrans s/n, 08913 Bellaterra, Barcelona, Spain*
[24] *University of Virginia, Department of Astronomy, Charlottesville, VA 22904, USA*
[25] *Departamento de Física, DCI-Campus León, Universidad de Guanajuato, Loma del Bosque 103, León, Guanajuato C. P. 37150, México*
[26] *Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, USA*
[27] *Steward Observatory, University of Arizona, 933 N. Cherry Avenue, Tucson, AZ 85721, USA*
[28] *Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA*
[29] *Institute for Fundamental Physics of the Universe, via Beirut 2, 34151 Trieste, Italy*
[30] *International School for Advanced Studies, Via Bonomea 265, 34136 Trieste, Italy*
[31] *Department of Physics, Astronomy and Mathematics, University of Hertfordshire, College Lane Campus, Hatfield, Hertfordshire, AL10 9AB, UK.*
[32] *Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*
[33] *Department of Physics, The University of Texas at Dallas, 800 W. Campbell Rd., Richardson, TX 75080, USA*
[34] *Department of Physics and Astronomy, University of California, Irvine, 92697, USA*

[35] Departament de Física, Serra Húnter, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain
[36] Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain
[37] Department of Physics and Astronomy, Siena College, 515 Loudon Road, Loudonville, NY 12211, USA
[38] Department of Physics & Astronomy and Pittsburgh Particle Physics, Astrophysics, and Cosmology Center (PITT PACC), University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, PA 15260, USA
[39] Instituto Avanzado de Cosmología A. C., San Marcos 11 - Atenas 202. Magdalena Contreras. Ciudad de México C. P. 10720, México
[40] Department of Physics and Astronomy, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada
[41] Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada
[42] Waterloo Centre for Astrophysics, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada
[43] Aix Marseille Univ, CNRS, CNES, LAM, Marseille, France
[44] Space Sciences Laboratory, University of California, Berkeley, 7 Gauss Way, Berkeley, CA 94720, USA
[45] Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain
[46] Université Clermont-Auvergne, CNRS, LPCA, 63000 Clermont-Ferrand, France
[47] Department of Physics and Astronomy, Sejong University, 209 Neungdong-ro, Gwangjin-gu, Seoul 05006, Republic of Korea
[48] CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain
[49] Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, USA
[50] University of Michigan, 500 S. State Street, Ann Arbor, MI 48109, USA
[51] Department of Physics & Astronomy, Ohio University, 139 University Terrace, Athens, OH 45701, USA
[52] Departamento de Astrofísica, Universidad de La Laguna (ULL), E-38206, La Laguna, Tenerife, Spain
[53] Instituto de Astrofísica de Canarias, C/ Vía Láctea, s/n, E-38205 La Laguna, Tenerife, Spain
[54] Excellence Cluster ORIGINS, Boltzmannstrasse 2, D-85748 Garching, Germany
[55] University Observatory, Faculty of Physics, Ludwig-Maximilians-Universität, Scheinerstr. 1, 81677 München, Germany
[56] National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Road, Chaoyang District, Beijing, 100101, P. R. China

ABSTRACT

We perform an analysis of the full shapes of Lyman-α (Lyα) forest correlation functions measured from the first data release (DR1) of the Dark Energy Spectroscopic Instrument (DESI). Our analysis focuses on measuring the Alcock-Paczynski (AP) effect and the cosmic growth rate times the amplitude of matter fluctuations in spheres of 8 h^{-1}Mpc, $f\sigma_8$. We validate our measurements using two different sets of mocks, a series of data splits, and a large set of analysis variations, which were first performed blinded. Our analysis constrains the ratio $D_M/D_H(z_{\rm eff}) = 4.525 \pm 0.071$, where $D_H = c/H(z)$ is the Hubble distance, D_M is the transverse comoving distance, and the effective redshift is $z_{\rm eff} = 2.33$. This is a factor of 2.4 tighter than the Baryon Acoustic Oscillation (BAO) constraint from the same data. When combining with Lyα BAO constraints from DESI DR2, we obtain the ratios $D_H(z_{\rm eff})/r_d = 8.646 \pm 0.077$ and $D_M(z_{\rm eff})/r_d = 38.90 \pm 0.38$, where r_d is the sound horizon at the drag epoch. We also measure $f\sigma_8(z_{\rm eff}) = 0.37\,^{+0.055}_{-0.065}\,({\rm stat}) \pm 0.033\,({\rm sys})$, but we do not use it for cosmological inference due to difficulties in its validation with mocks. In ΛCDM, our measurements are consistent with both cosmic microwave background (CMB) and galaxy clustering constraints. Using a nucleosynthesis prior but no CMB anisotropy information, we measure the Hubble constant to be $H_0 = 68.3 \pm 1.6$ km s^{-1} Mpc^{-1} within ΛCDM. Finally, we show that Lyα forest AP measurements can help improve constraints on the dark energy equation of state, and are expected to play an important role in upcoming DESI analyses.

1. INTRODUCTION

The physics behind cosmic acceleration represents one of the key unknowns in our current understanding of the nature of the Universe. The simplest explanation, given by the cosmological constant, Λ, is central to the currently preferred model of cosmology, ΛCDM. However, recent Baryon Acoustic Oscillations (BAO) measurements from the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al. 2016, 2022) in combination with cosmic microwave background (CMB) and type Ia supernovae (SNe) have provided strong hints that dark energy might not be described by the cosmological constant (DESI Collaboration et al. 2025a,b). Instead, these datasets prefer more complex explanations involving a dynamic equation of state for dark energy. Furthermore, the ΛCDM model has also come under increasing pressure due to a number of tensions, with perhaps the most discussed being the $> 5\sigma$ discrepancy in the value of the Hubble constant between direct measurements based on the cosmic distance ladder (e.g.

* NASA Einstein Fellow
acuceu@lbl.gov

Breuval et al. 2024; Riess et al. 2022) and indirect probes such as the CMB that rely on ΛCDM (e.g. Planck Collaboration et al. 2020a). One of the primary observational goals that can help us untangle these mysteries is to obtain more precise measurements of the expansion rate of the Universe at different stages of its evolution (Weinberg et al. 2013).

The Lyman-α (Lyα) forest currently provides the tightest constraints on the cosmic expansion rate in the redshift interval $2 < z < 4$ (DESI Collaboration et al. 2025c; Cuceu et al. 2023a). The forest traces the Lyα absorption due to neutral hydrogen in the intergalactic medium (IGM) between us and the background sources, most commonly quasars (e.g. McQuinn 2016). This gives us a unique view into the large-scale structure (LSS) of the Universe, as it allows continuous mapping of the density field along each line-of-sight. For more than a decade, Lyα forest observations from the Baryon Oscillation Spectroscopic Survey (BOSS), extended BOSS (eBOSS), and now DESI, have provided increasingly precise constraints on the cosmic expansion rate through measurements of the BAO scale (e.g. Busca et al. 2013; Slosar et al. 2013; du Mas des Bourboux et al. 2020; DESI Collaboration et al. 2025d,c). These measurements are performed using the three-dimensional (3D) auto-correlation of Lyα flux, as well as the cross-correlation between Lyα flux and quasar positions. While the BAO scale provides a very powerful standard ruler for cosmology, the Lyα forest contains plenty of useful cosmological information beyond BAO.

In this paper, we present an analysis that aims to extract cosmological information using the full shape of the Lyα forest 3D correlation functions measured from the DESI Data Release 1 (DR1). We perform a template-based analysis in which we focus on two well-understood sources of cosmological information: the Alcock-Paczynski (AP) effect and redshift space distortions (RSD). This is in contrast to direct cosmological analyses (also referred to as *full-modelling*) where cosmological parameters are directly measured within a given model (e.g. ΛCDM). In the case of the Lyα forest, Gerardi et al. (2023) used DESI forecasts to show that most of the large-scale 3D information is captured by BAO, AP, and RSD. Therefore, in this work we focus on measuring AP and RSD using the full shapes of the Lyα forest correlation functions.

The Alcock-Paczynski effect introduces an anisotropy in the distribution of LSS when the fiducial cosmology used to transform observed angular and redshift separations into distances is different from the truth (Alcock & Paczynski 1979). In the context of Lyα, AP analyses were first proposed by Hui et al. (1999); McDonald & Miralda-Escudé (1999). Cuceu et al. (2021) and Cuceu et al. (2023a) showed that adding AP information from the full shape of Lyα correlations leads to significantly tighter cosmic expansion rate constraints from BAO. Therefore, the first goal of this work is to extract the full-shape AP information from DESI DR1 correlations measured by DESI Collaboration et al. (2025d, hereafter `DESI24-Ly`α), in order to improve upon their cosmic expansion rate constraints, which are based solely on BAO.

Redshift space distortions are caused by peculiar velocities modifying the observed redshift of quasars, and also the redshift, shape, and amount of Lyα absorption. When interpreting the redshift as purely cosmological in origin, peculiar velocities lead to distortions in the observed distribution of large-scale structure (Kaiser 1987; Hamilton 1998). The RSD signal is sensitive to the growth rate of cosmic structure f, as well as the amplitude of matter fluctuations, conventionally parameterized by σ_8, the fluctuation amplitude in spheres of radius 8 h^{-1}Mpc. In linear theory, these two parameters are degenerate, and therefore we can only measure their product, $f\sigma_8$ (Percival & White 2009). As the Lyα forest transmitted flux is given by a non-linear mapping of the underlying density field, its RSD signal also depends on an unknown velocity divergence bias that multiplies $f\sigma_8$ (e.g. McDonald 2003; Seljak 2012; Chen et al. 2021; Ivanov 2024). This has so far prevented 3D Lyα forest analyses from extracting useful cosmological information from RSD. On the other hand, galaxy clustering analyses, which do not have to deal with this unknown velocity divergence bias, have successfully used RSD to measure $f\sigma_8$ over the last few decades (e.g. Blake et al. 2011; Reid et al. 2012; Alam et al. 2017, 2021; DESI Collaboration et al. 2024a). In this work, we use the quasar RSD signal from large scales in the Lyα-QSO cross-correlation to measure $f\sigma_8$. To obtain this constraint, we perform a joint analysis of the cross-correlation with the Lyα auto-correlation in order to break the degeneracy between the Lyα and QSO RSD signals in the cross-correlation (Cuceu et al. 2021). Therefore, our RSD measurement still depends on constraining the linear Lyα RSD signal, which has not been as thoroughly studied (see Chabanier et al. 2024; de Belsunce et al. 2025, for the current state-of-the-art). This is in contrast to the AP effect, which is geometric in nature and has been used as a cosmological tool for decades.

Our analysis uses the Lyα forest correlation functions measured from DESI DR1 by `DESI24-Ly`α. The DESI survey is in the process of measuring the redshifts of more than 40 million galaxies and quasars over five years (DESI Collaboration et al. 2024b; Miller et al. 2024; Poppett et al. 2024; Guy et al. 2023; Schlafly et al.

2023). The first year of DESI data, which comprises the bulk of DESI DR1, contains more galaxy spectra than any previous spectroscopic survey (DESI Collaboration et al. 2016, 2022, 2025e). The Lyα analysis uses more than 420 000 Lyα forests and 700 000 quasars, which is roughly double what was used in the previous state-of-the-art analyses from eBOSS. The BAO measurements from this dataset were presented in DESI24-Lyα, and they have been recently updated with the DESI Data Release 2 (DR2) dataset in DESI Collaboration et al. (2025c).

We start by introducing the data we use along with our model in Section 2. The Alcock-Paczynski and $f\sigma_8$ measurements are presented in Section 3, and are followed by our validation tests in Section 4, which encompass analyses on mocks, data splits, and a large set of analysis variations. After that, we discuss our results and their validation in Section 5. Finally, we present the cosmological implications of our measurements in Section 6, and conclude in Section 7.

2. ANALYSIS

Our analysis focuses on modelling Lyα forest correlation functions in order to extract cosmological information from their full shape. We use the correlation functions measured by DESI24-Lyα, which are described in Section 2.1. The model we use is based on an extension developed by Cuceu et al. (2021) of the model used for Lyα BAO analyses, and is described in Sections 2.2 and 2.3. Finally, we describe our blinding and unblinding process in Section 2.4.

2.1. *Lyα forest correlation functions*

We use the four 3D Lyα forest correlation functions measured by DESI24-Lyα. The four correlations are computed from three datasets:

1. The Lyα flux overdensity field measured between the Lyα and Lyβ emission lines, in the rest-frame interval $1040 < \lambda_{RF} < 1205$ Å, denoted region A. This dataset encompasses 428,403 Lyα forests, and we refer to it as Lyα(A).

2. The Lyα flux overdensity field measured between the Lyβ emission line and the Lyman limit, in the rest-frame interval $920 < \lambda_{RF} < 1020$ Å, denoted region B. This dataset encompasses 137,432 Lyα forests, and we refer to it as Lyα(B).

3. A catalog of 709,565 DESI quasars at redshift $z > 1.77$, which are used as discrete tracers.

For a detailed description of these datasets, see DESI24-Lyα.

The four correlation functions include two auto-correlations of Lyα flux overdensity, Lyα(A)×Lyα(A) and Lyα(A)×Lyα(B), and two Lyα-QSO cross-correlations, Lyα(A)×QSO and Lyα(B)×QSO. We will refer to the first pair as the auto-correlation, and the second pair as the cross-correlation. These correlation functions are computed on two-dimensional grids in comoving separation along the line-of-sight, $r_{||}$, and across the line-of-sight, $r_\perp$. For a pair of pixels (i, j) at redshifts (z_i, z_j) and separated by an angle $\Delta\theta$, these comoving coordinates are given by:

$$r_{||} = [D_{\rm c}(z_i) - D_{\rm c}(z_j)] \cos\frac{\Delta\theta}{2}, \tag{1}$$

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

where D_c is the comoving distance, and D_M is the comoving angular diameter distance. To compute these distances, a fiducial cosmology based on CMB measurements from the Planck satellite Planck Collaboration et al. (2020a) is used. The parameter values of this fiducial cosmology can be found in Table 2 of DESI24-Lyα.

The correlation functions are computed in 4 h^{-1}Mpc wide bins in both $r_{||}$ and $r_\perp$, and are averaged into a single redshift bin with an effective redshift $z_{\rm eff} = 2.33$. A weighted pair-counting algorithm is used, with the correlation in a comoving coordinate bin M given by:

$$\xi_M = \frac{\sum_{i,j\in M} w_i w_j \delta_i \delta_j}{\sum_{i,j\in M} w_i w_j}, \tag{3}$$

where the sums run over pixel-pixel pairs for the two auto-correlations, and over pixel-quasar pairs for the two cross-correlations. In the case of Lyα forest pixels, δ corresponds to the measured Lyα flux overdensity, while for quasars $\delta = 1$. The weights for Lyα pixels are given by the inverse variance, accounting for both noise and intrinsic large-scale structure variance, and also include the redshift evolution of the Lyα bias. On the other hand, the weights for quasars only account for the redshift evolution of the quasar bias du Mas des Bourboux et al. (2020). For a detailed description of the continuum fitting process used to measure the Lyα δ field, and the weights used in 3D Lyα forest analyses, see Ramírez-Pérez et al. (2024).

Finally, DESI24-Lyα compute one covariance matrix for all four correlation functions, which includes their cross-covariance. This is estimated from 1028 measurements of the correlation function in individual `HEALpix` pixels, corresponding to patches 250×250 $(h^{-1}{\rm Mpc})^2$ in size at $z_{\rm eff} = 2.33$. This noisy estimate is then smoothed by replacing non-diagonal elements of the correlation matrix, ${\rm Corr}_{MN}$, corresponding to the same comoving

separation between bins M and N with their average. This method was tested on DESI DR1 mocks by Cuceu et al. (2025), and was shown to produce accurate uncertainties on BAO parameters measured from the same mocks. We have performed the same tests in the context of the full-shape analysis below in Section 4.1.

2.2. *Modelling of correlations*

Our model for Lyα forest correlation functions closely follows that used in Cuceu et al. (2023b,a), and is implemented in the `Vega`[1] package. It is based on a template approach and at a high level follows these steps:

1. A linear isotropic matter power spectrum, $P_{\rm fid}(k)$, is computed using CAMB (Lewis et al. 2000) at the effective redshift of our measurements ($z_{\rm eff} = 2.33$), with a fiducial cosmology based on Planck CMB results Planck Collaboration et al. (2020a).[2] This is our template.
2. The isotropic power spectrum is split into a wiggles (or peak) component and no-wiggles (or smooth) component following Kirkby et al. (2013). These two components go through the rest of the modelling process independently and are only combined in the final step.
3. Model Lyα power spectra, P_F, and Lyα-QSO cross-spectra, $P_\times$, are computed using the template, the linear Kaiser term Kaiser (1987), along with models for small-scale non-linearities, BAO broadening, high column density (HCD) absorbers, and redshift errors.
4. The now anisotropic model auto- and cross-power spectra are transformed into 2D model correlations, as a function of $r_{||}$ and $r_\perp$, following Kirkby et al. (2013).
5. The coordinates of the model correlations, $r_{||}$ and $r_\perp$, are re-scaled using free scale parameters that we fit for. These are independent for the smooth and peak components.
6. The impact of metal absorption is incorporated by computing theoretical models for all Lyα-Metal, QSO-Metal, and Metal-Metal correlations. A number of smaller effects discussed below are added at this stage, and also the distortion due to continuum fitting errors is modeled by multiplying the model correlation with a distortion matrix.
7. Finally, the two model components (peak and smooth) are summed to produce the final model of the correlation functions.

The model anisotropic auto- and cross-spectra are given by:

$$P_F(\mathbf{k}) = b_F'^2(1+\beta_F'\mu_k^2)^2\, G(\mathbf{k})F_{NL}(\mathbf{k})P_{\rm fid}(k), \quad (4)$$

$$P_{F\times Q}(\mathbf{k}) = b_F'(1+\beta_F'\mu_k^2)(b_Q + f\mu_k^2)\, G(\mathbf{k})X_{NL}(\mathbf{k})P_{\rm fid}(k), \quad (5)$$

where the wavenumber vector $\mathbf{k} = (k, \mu_k) = (k_{||}, k_\perp)$ has modulus k, line-of-sight and transverse components $(k_{||}, k_\perp)$, and $\mu_k = k_{||}/k$. The $G(k, \mu_k)$ term models the binning of the correlation function ($4\ h^{-1}$Mpc bins), and is given by the Fourier transform of a top-hat function. F_{NL} and X_{NL} model small-scale deviations from linear-theory, and are described in detail below. b_F' and β_F' are effective Lyα linear bias and RSD parameters, which account for both Lyα flux and HCD absorption, and are described in detail below. b_Q and f are the linear quasar bias and the cosmic growth rate (i.e., the logarithmic derivative of the linear growth factor), respectively. Note that the Lyα RSD parameter, β_F, also depends on the growth rate through $\beta_F = b_{\eta,F} f / b_F$. However, due to the presence of the unknown velocity divergence bias $b_{\eta,F}$, we cannot use the linear Lyα auto-correlation to measure f, as the two parameters are degenerate. Therefore, we work with the β_F parameter and marginalize over it in our analysis.

High column density (HCD) absorbers are treated in two different ways depending on their strength. Strong HCDs, represented by Damped Lyα Absorbers (DLAs), with column densities $\log N_{\rm HI} > 20.3$, that are found in high signal-to-noise (SNR > 3) spectra are masked (Ramírez-Pérez et al. 2024). This is because strong DLAs have damping wings that extend over a significant part of the spectra they are found in, increasing the noise on large scales in the measured Lyα correlations (Font-Ribera & Miralda-Escudé 2012). The thresholds above were chosen based on the performance of DLA finder algorithms. On the other hand, HCDs that are either too weak to be detected or in spectra with signal-to-noise below 3 are not masked, and their impact must be modelled at the level of the correlation function. We follow Font-Ribera & Miralda-Escudé (2012); du Mas des Bourboux et al. (2020); DESI Collaboration et al. (2025d), and model this effect through the effective bias and RSD parameters given by:

$$b_F' = b_F + b_{\rm HCD}F_{HCD}(k_{||}), \quad (6)$$

$$b_F'\beta_F' = b_F\beta_F + b_{\rm HCD}\beta_{\rm HCD}F_{HCD}(k_{||}), \quad (7)$$

where b_F and $b_{\rm HCD}$ are the linear Lyα and HCD flux biases, β_F and $\beta_{\rm HCD}$ are the RSD parameters for Lyα

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

[2] The template used here is exactly the same as that used for the BAO analysis of the same data set in `DESI24-Ly`α.

and HCDs. The $F_{\rm HCD}$ function accounts for the profile of the HCDs (i.e., the fact that the absorption is not localized along the line-of-sight). We follow de Sainte Agathe et al. (2019) and use the approximate form $F_{\rm HCD} = \exp(-L_{\rm HCD} k_{||})$, in line with DESI24-Ly$\alpha$. $L_{\rm HCD}$ is a free parameter that is allowed to vary in our fits, and can be thought of as the typical scale of undetected HCDs.

For the auto-correlation, we follow previous Lyα BAO and full-shape analyses and model small-scale deviations from linear theory using the empirical model developed by Arinyo-i-Prats et al. (2015). This is given by:

$$F_{NL}(\mathbf{k}) = \exp\left\{ q_1 \Delta^2(k) \left[1 - \left(\frac{k}{k_v} \right)^{a_v} \mu_k^{b_v} \right] - \left(\frac{k}{k_p} \right)^2 \right\}, \tag{8}$$

where $\Delta^2(k) \equiv k^3 P_{\rm fid}(k)/(2\pi^2)$, and we treat $\{q_1, k_v, a_v, b_v, k_p\}$ as unknown nuisance parameters that we marginalize over.[3] This model, which we will refer to as the Arinyo model, represents an empirical fitting formula that was shown to accurately reproduce small-scale Lyα power spectra in hydrodynamical simulations (Arinyo-i-Prats et al. 2015; Givans et al. 2022; Chabanier et al. 2024; Chaves-Montero et al. 2025). In previous Lyα forest analyses, the parameters of this model have been fixed to the values predicted from simulations (DESI Collaboration et al. 2025d; Cuceu et al. 2023a), while in this work we allow them to vary. We use wide flat priors for all 5 parameters, with the edges of the priors informed by the range of possible values Arinyo-i-Prats et al. (2015) measured in simulations (see Appendix B).

For the cross-correlation, the largest impact on small scales is due to quasar redshift errors Font-Ribera et al. (2013). We follow DESI24-Lyα and model this effect along with quasar non-linear velocities using a simple Lorentzian damping, given by:

$$X^2_{NL} = [1 + (k_{||}\sigma_z)^2]^{-1}, \tag{9}$$

where σ_z is a free parameter. As part of our validation in Section 4.3, we will also compare our baseline results to those obtained using a Gaussian damping model instead of the Lorentzian.

In order to account for the broadening of the BAO peak due to non-linear evolution, the model power spectrum for the peak component is also multiplied by $\exp[-(k_{||}^2\Sigma_{||}^2 + k_\perp^2\Sigma_\perp^2)/2]$. The two broadening parameters, are fixed to the theoretical prediction of $\Sigma_\perp = 3.24$ h^{-1}Mpc and $\Sigma_{||} = 6.37$ h^{-1}Mpc, following Eisenstein et al. (2007); du Mas des Bourboux et al. (2020). DESI Collaboration et al. (2025c) has recently tested this choice in the context of the DESI DR2 BAO analysis, and found that these values are consistent with the data, and that allowing these parameters to vary does not impact BAO constraints.

The anisotropic power spectra from Equations (4) and (5) are transformed into 2D model correlation functions through the following process. First, a multipole decomposition is performed up to $\ell = 6$, followed by a Hankel transform, and then the resulting correlation multipoles are assembled into the 2D correlation model as a function of $r_{||}$ and $r_\perp$.

Besides absorption due to hydrogen, the Lyα forest regions we use also contain metal absorption. For metals with rest-frame wavelength close to the rest-frame of Lyα (1215.67Å), the small-scale Lyα-Metal (for the auto) and QSO-Metal (for the cross) cross-correlations will appear as extra peaks along the line-of-sight in our correlation functions. This is because we treat all pixels as Lyα pixels when computing their redshift, and therefore wrongly assign null separation Lyα-Metal and QSO-Metal pairs to correlation function bins at large separations (Pieri et al. 2014; Morrison et al. 2024). The lines found to cause the peaks observed along the line-of-sight in Lyα correlations are: SiIII(1207), SiII(1193), SiII(1190), and SiII(1260) (Bautista et al. 2017). Other metal absorption at longer wavelengths only contribute through their auto-correlation. Guy et al. (2025) used measurements in regions redward of the Lyα line to show that this type of contamination is dominated by CIV absorption. Therefore, we follow DESI24-Lyα, and model the impact of these 5 metal lines.

The model for metal absorption involves computing all Lyα-Metal and Metal-Metal correlations for the Lyα auto, and all QSO-Metal correlations for the Lyα-QSO cross. The model for each of these correlations is very similar to those of the Lyα correlations, with each metal having its own linear bias and RSD parameters. The metal biases are free to vary, while the metal RSD parameters are fixed to 0.5, following DESI24-Lyα.[4] Following Bautista et al. (2017) and du Mas des Bourboux et al. (2017), the model metal correlations are transformed from their correct separations $(\tilde{r}_{||}, \tilde{r}_\perp)$ to the co-

[3] Note that the original model developed by Arinyo-i-Prats et al. (2015), also contained a higher order $\Delta^4(k)$ term with a corresponding q_2 parameter. However, this was found to be very close to zero and has therefore been ignored in subsequent analyses.

[4] This is because we lack the sensitivity to constrain both the bias and RSD parameter for each metal. Given that metal peaks are localized along the line-of-sight, we do not expect this to play a role in our results.

ordinate grid of the measured correlations:

$$\xi_m^M = \sum M_{MN} \xi_m(\tilde{r}_{||}(N), \tilde{r}_{\perp}(N)), \tag{10}$$

where the metal matrices M_{MN} are given by:

$$M_{MN} = \frac{1}{W_M} \sum_{(i,j)\in M,(i,j)\in N} w_i w_j, \tag{11}$$

where $(i,j) \in N$ refers to bins computed using the correct redshifts, $(i,j) \in M$ refers to bins computed using the assumed redshifts, and m refers to each of the 5 metal lines described above.

For the auto-correlation, the following components are added besides the main Ly$\alpha\times$Lyα auto-correlation (ξ_F):

$$\xi_F^t = \xi_F + \sum_m \xi_{F\times m} + \sum_{m_1,m_2} \xi_{m_1\times m_2} + a_{\rm noise}\xi_{\rm noise}, \tag{12}$$

where $F \times m$ refers to Lyα-Metal correlations, $m \times m$ refers to Metal-Metal correlations, and the sums are performed over the 5 metal lines above. The $a_{\rm noise}\xi_{\rm noise}$ term models the correlated sky noise of spectra from fibers of the same spectrograph. Guy et al. (2025) showed that the dominant contribution is the sky background model noise and that it only affects bins with purely transverse separation (i.e. $r_{||} = 0$ h^{-1}Mpc). Following DESI24-Lyα, $\xi_{\rm noise}$ has a fixed shape as a function of $r_\perp$, computed based on the fraction of pairs at each $r_\perp$ that were observed with the same spectrograph, and $a_{\rm noise}$ is a free parameter we marginalize over.

For the cross-correlation, the contaminated model is given by:

$$\xi_{F\times Q}^t = \xi_{F\times Q} + \sum_m \xi_{Q\times m} + \xi_{\rm TP}, \tag{13}$$

where $\xi_{F\times Q}$ is the main Ly$\alpha\times$QSO model, $Q \times m$ refers to QSO-Metal correlations, and ξ_{TP} accounts for the transverse proximity effect. Following Font-Ribera et al. (2013), we use the isotropic form:

$$\xi_{\rm TP} = \xi_0^{\rm TP} \left(\frac{1h^{-1}{\rm Mpc}}{r}\right)^2 \exp\left(-r/\lambda_{\rm UV}\right), \tag{14}$$

where the amplitude $\xi_0^{\rm TP}$ is a free parameter, and we fix $\lambda_{\rm UV} = 300$ h^{-1}Mpc (Rudie et al. 2013).

Once the contaminated models ξ_F^t and ξ_Q^t are built, the only missing ingredient is the distortion due to continuum fitting errors. This arises because we fit the amplitude and slope of the continuum for each individual forest in order to account for Quasar diversity (see Ramírez-Pérez et al. 2024). However, this process also fits for any large-scale structure fluctuations of the size of the forest or larger that are present in each forest. These continuum fitting errors lead to distortions in the measured correlation functions. Following Bautista et al. (2017) and du Mas des Bourboux et al. (2017), this effect is forward modeled through distortion matrices, which for the auto- and cross-correlation are given by:

$$D_{MN}^{\rm auto} = \frac{1}{W_M} \sum_{i,j\in M} w_i w_j \sum_{i',j'\in N} \eta_{ii'}\eta_{jj'}, \tag{15}$$

$$D_{MN}^{\rm cross} = \frac{1}{W_M} \sum_{i,j\in M} w_i w_j \sum_{i',j\in N} \eta_{ii'}, \tag{16}$$

with projection matrices, η_{ij}, given by:

$$\eta_{ij} = \delta_{ij}^K - \frac{w_j}{\sum_k w_k} - \frac{w_j\kappa_i\kappa_j}{\sum_k w_k\kappa_k^2}, \tag{17}$$

where δ_{ij}^K is the Kronecker delta, and $\kappa_k = \log\lambda_k - \overline{\log\lambda_q}$, with λ_i the observed wavelength in pixel i, and $\overline{\log\lambda_q}$ the mean of $\log\lambda$ in the forest of quasar q. The distortion matrices in Equations (15) and (16) are computed using model bins N that are 2 h^{-1}Mpc in size, while the data bins M correspond to the binning size of the measured correlation functions (4 h^{-1}Mpc in size). This means the models computed so far use the smaller 2 h^{-1}Mpc binning for better precision, while the distorted model, $\hat{\xi}_M$, matches the data binning. The distorted model is given by:

$$\hat{\xi}_M = \sum_N D_{MN}\xi_N^t. \tag{18}$$

The computation of the distortion matrices is very intensive, and therefore only a small fraction of pixel pairs are used to measure them. In the baseline analysis, this consists of 1% of pairs, but we also test a variation where that number is doubled in Section 4.3.

Finally, the peak and smooth components, $\hat{\xi}_p$ and $\hat{\xi}_s$, which have so far undergone the modelling process separately, are added to produce the full correlation function model:

$$\xi(r_{||}, r_\perp) = \xi_s(q_{||}^s r_{||}, q_\perp^s r_\perp) + \xi_s(q_{||}^p r_{||}, q_\perp^p r_\perp), \tag{19}$$

where $(q_{||}^s, q_\perp^s)$ rescale the coordinates of the smooth component, while $(q_{||}^p, q_\perp^p)$ rescale the coordinates of the peak component. In a BAO analysis, $q_{||}^s = q_\perp^s = 1$ and the goal is to fit for the BAO scale parameters $(q_{||}^p, q_\perp^p)$. On the other hand, in this work we wish to extract information from the full shape of the correlation functions, and therefore we also fit for $(q_{||}^s, q_\perp^s)$. In the case of the cross-correlation, the line-of-sight coordinates $r_{||}$, are further modified by an additive free parameter, $\Delta r_{||}$, which accounts for any systematic quasar errors.

2.3. *Fitting Alcock-Paczynski and the growth rate*

In this section, we introduce the parametrization used to extract cosmological information from Lyα forest correlation functions. In particular, in this work we focus on the three main sources of information: BAO, the Alcock-Paczynski effect, and redshift space distortions. As BAO has already been measured from this data set by DESI24-Lyα, here we focus on the other two effects.

To isolate the AP effect, we use the parametrization introduced by Cuceu et al. (2021):

$$\phi(z) \equiv \frac{q_\perp(z)}{q_{||}(z)} \quad \text{and} \quad \alpha(z) \equiv \sqrt{q_\perp(z) q_{||}(z)}, \tag{20}$$

where ϕ captures the anisotropy introduced by AP, and is related to the Alcock-Paczynski parameter $D_M H(z)$ through:

$$\phi(z) = \frac{D_M H(z)}{[D_M H(z)]_{\rm fid}}, \tag{21}$$

where $H(z)$ is the Hubble parameter, and the subscript *fid* indicates quantities computed using the fiducial cosmology. As we have two sets of $(q_{||}, q_\perp)$ parameters, we can measure ϕ independently for the peak component and from the broadband. We will use the symbol $\phi_{\rm p}$ to refer to the measurement of ϕ from BAO, while $\phi_{\rm s}$ will be used for the smooth component. As our end goal is to measure AP from the full shape, we will also use the symbol $\phi_{\rm f}$ to refer to measurements of AP where the anisotropy in the two components is fitted through a single parameter (i.e., we impose $\phi_{\rm p} = \phi_{\rm s}$).

On the other hand, α measures an isotropic scale, which for the peak component is related to the comoving acoustic scale r_d, through:

$$\alpha_{\rm p}^2(z) = \frac{D_M(z) D_H(z)/r_d^2}{[D_M(z) D_H(z)/r_d^2]_{\rm fid}}, \tag{22}$$

with $D_H(z) = c/H(z)$, and speed of light, c. As before, we use the symbol $\alpha_{\rm p}$ to refer to the isotropic scale parameter for the peak component, and $\alpha_{\rm s}$ for the isotropic parameter for the smooth component. The $\alpha_{\rm s}$ parameter is sensitive to the scale of matter-radiation equality, which has been used as a stand-alone source of cosmological information with galaxy clustering (e.g. Philcox et al. 2021; Zaborowski et al. 2025). However, in the case of the Lyα forest, its measurement has not been validated in the presence of systematics. Given that $\alpha_{\rm s}$ constraints are significantly weaker than $\alpha_{\rm p}$ constraints, and would need to be independently validated, we follow Cuceu et al. (2021) and treat $\alpha_{\rm s}$ as a nuisance parameter to be marginalized over.

From Equation (5), we can see that measurements of the anisotropy in the Lyα-QSO cross-correlation are sensitive to the growth rate, f. However, both Equations (4) and (5) have an implicit scaling of σ_8^2. This means that instead of $\{b'_F, b_Q, f\}$, we are actually sensitive to the products $\{b'_F\sigma_8, b_Q\sigma_8, f\sigma_8\}$. The Ly$\alpha$ autocorrelation constrains $b'_F\sigma_8$ and β'_F, which allows us to break the degeneracy between the Lyα and the QSO part of Equation (5), and therefore measure $b_Q\sigma_8$ and $f\sigma_8$. Also note that all the parameters mentioned above evolve with redshift, and our measurements are made at the effective redshift of our data, $z_{\rm eff} = 2.33$.

In practice, we fit for $\{b'_F, b_Q, f\}$, and rescale the end result by σ_8. Because σ_8 is defined using spheres of 8 h^{-1}Mpc, there is an additional dependence on the isotropic scale because we are rescaling the coordinates of the model to match the data. As the RSD information primarily comes from the broadband, the relevant scale is the one measured from the smooth component. Following Gil-Marín et al. (2020) and Bautista et al. (2021), we account for this effect by using a rescaling of σ_8 with spheres of $8\alpha_{\rm iso}$ h^{-1}Mpc, where $\alpha_{\rm iso} = q_{||}^{1/3} q_\perp^{2/3}$. This effectively keeps the scale σ_8 is fitted at constant relative to the data. The measurements of $f\sigma_8$ presented in this work are rescaled using the $\alpha_{\rm iso}$ values derived from the best-fit $\alpha_{\rm s}$ and $\phi_{\rm s}$.

2.4. *Blinding*

The analysis presented in this work was first performed blinded in order to avoid human biases. Given that we are using the same measured correlation functions as DESI24-Lyα, we had to work with unblinded correlations and therefore resorted to using the following blinding method. Within the likelihood, we add random values to the parameters of interest, $\phi_{\rm s}$ and f. These random values are kept consistent across all tests we do by fixing the random seed used to generate them. They are generated from Gaussians with null mean, and standard deviations based on the previous eBOSS DR16 measurement in the case of Alcock-Paczynski, $\sigma_{\phi_{\rm s}} = 0.02$, and the forecasted DESI DR1 uncertainty in the case of the growth rate, $\sigma_f = 0.14$[5]. The actual values of the blinding were $\Delta\phi_{\rm s} = -0.020$ and $\Delta f = 0.02$, which corresponds to a significant shift in $\phi_{\rm s}$, and a very small shift in f.

The decisions related to what constitutes the baseline analysis were taken without knowledge of the final results. Also, all the validation tests presented in Section 4 were first performed blinded. In light of the tests on mocks presented in Section 4.1, we decided before un-

[5] This is because our analysis is the first to measure $f\sigma_8$ from the Lyα forest. The forecasted uncertainty is based on Cuceu et al. (2021).

blinding that only the Alcock-Paczynski result has been fully validated, and will therefore be used for the cosmological inference, as discussed in Section 5. We still present our measurement of $f\sigma_8$, along with all of its validation tests, but we do not include it in the cosmological results presented in Section 6.

3. COMPRESSED RESULTS

We model and fit the four correlation functions computed by DESI24-Lyα using the `Vega` package. We use the `iminuit` minimizer (Dembinski & et al. 2020; James & Roos 1975) to find the best fit, and sample posterior distributions using the Nested Sampler `PolyChord` (Handley et al. 2015a,b). We use 4 h^{-1}Mpc bins and compute the auto-correlations up to 200 h^{-1}Mpc in both $r_{||}$ and $r_{\perp}$, while the cross-correlations are computed from -200 h^{-1}Mpc to 200 h^{-1}Mpc in $r_{||}$, and from 0 h^{-1}Mpc to 200 h^{-1}Mpc in $r_{\perp}$. This means the two Lyα auto-correlations have a total of 2500 data points each, while the two Lyα-QSO cross-correlations have 5000 points each, for a total of 15000 data points.

The maximum scale we fit is $r_{\rm max} = 180$ h^{-1}Mpc for all four correlations. On the other hand, the minimum scale is different for the auto vs the cross due to differences in our confidence of the modelling of small scales. For the auto-correlations, we use a minimum scale of $r_{\rm min} = 25$ h^{-1}Mpc, in line with Cuceu et al. (2023b,a), while for the cross-correlation we choose a more conservative $r_{\rm min} = 40$ h^{-1}Mpc. This choice is partly driven by the fact that for the auto-correlation we have an empirical model for small-scale non-linearities (Arinyo-i-Prats et al. 2015), that has been extensively tested on measurements of the Lyα power spectrum from hydro-dynamical simulations. However, for the cross-correlation, we do not have an equivalent model, and therefore, we restrict our analysis to larger scales. Furthermore, Cuceu et al. (2023b) found using eBOSS DR16 mocks that AP measurements are unbiased when using scales larger than 25 h^{-1}Mpc, while a significant bias was detected when using scales smaller than that. We find consistent results using the DESI DR1 mocks here (see Section 4.1). Scale cuts are implemented in isotropic separation (i.e., $\sqrt{r_{||}^2 + r_{\perp}^2}$), and after that we are left with a total of 9188 data points.

We show the best-fit model versus data, along with the residuals, in Figure 1. As the quantities we aim to measure come from the anisotropy in correlation functions, we compress the 2D correlations into shells of isotropic separation and show them as a function of the line-of-sight angle θ.[6] The compression is performed through a weighted average using the covariance matrix. Plots of the same data compressed into wedges instead of shells can be found in Figures 4 and 5 of DESI24-Lyα. Here we focus on two shells at smaller separations ($r < 75$ h^{-1}Mpc for Ly$\alpha\times$Lyα, and $r < 85$ h^{-1}Mpc for Ly$\alpha\times$QSO), as these provide most of the signal for both the AP and RSD measurements. The equivalent plots for the shells at larger separation are in Appendix A. Figure 1 also illustrates the impact of changing $\phi_{\rm f}$ and $f\sigma_8$ through the dashed and dotted lines, respectively. We only show the $f\sigma_8$ variation for the cross-correlation, as this is our only source of information for this effect (see Section 2.3). A similar figure showing the impact of several important nuisance parameters can be found in Appendix B.

The model gives a good fit to the data, with $\chi^2_{\rm min} = 9145.1$ for 9163 degrees of freedom (PTE= 0.55). We have found that changing the minimum separation used does not significantly degrade the fit quality, which is not surprising given that DESI24-Lyα used $r_{\rm min} = 10$ h^{-1}Mpc for both the auto and cross-correlation, and still obtained a good fit (PTE= 0.23). When interpreting the goodness-of-fit from the shell plots shown in Figure 1, it is important to note that there is significant correlation between nearby data points. Directly adjacent points have correlation coefficients of ~ 0.5, and the correlation matrix remains above 0.1 up to the third or fourth off-diagonal.

Our main measurements are presented in Figure 2 and in Table 1, where we focus on the three sources of cosmological information that we use: BAO, the Alcock-Paczynski effect, and redshift space distortions. We compare our full-shape results in blue, with the BAO measurement by DESI24-Lyα from the same dataset in black, and the broadband-only constraint in red (where the BAO parameters $\alpha_{\rm p}$ and $\phi_{\rm p}$ are marginalized). The isotropic BAO parameter ($\alpha_{\rm p}$) constraint is very similar to that of DESI24-Lyα, although not identical due to the small correlation between this parameter and the AP effect (correlation coefficient of ~ 0.16 for $\phi_{\rm p}$, and ~ 0.08 for $\phi_{\rm f}$).

The AP measurement from the broadband ($\phi_{\rm s}$) is also consistent with the AP measurement made by DESI24-Lyα ($\phi_{\rm p}$), but represents a 1.7% constraint compared to the previous 3.8% constraint from the BAO peak in the same dataset. Combining the two, the full-shape of the DESI DR1 Lyα forest correlations gives a

[6] Where $\theta = 0°$ corresponds to $r_{\perp} = 0$, and $\theta = 90°$ corresponds to $r_{||} = 0$.

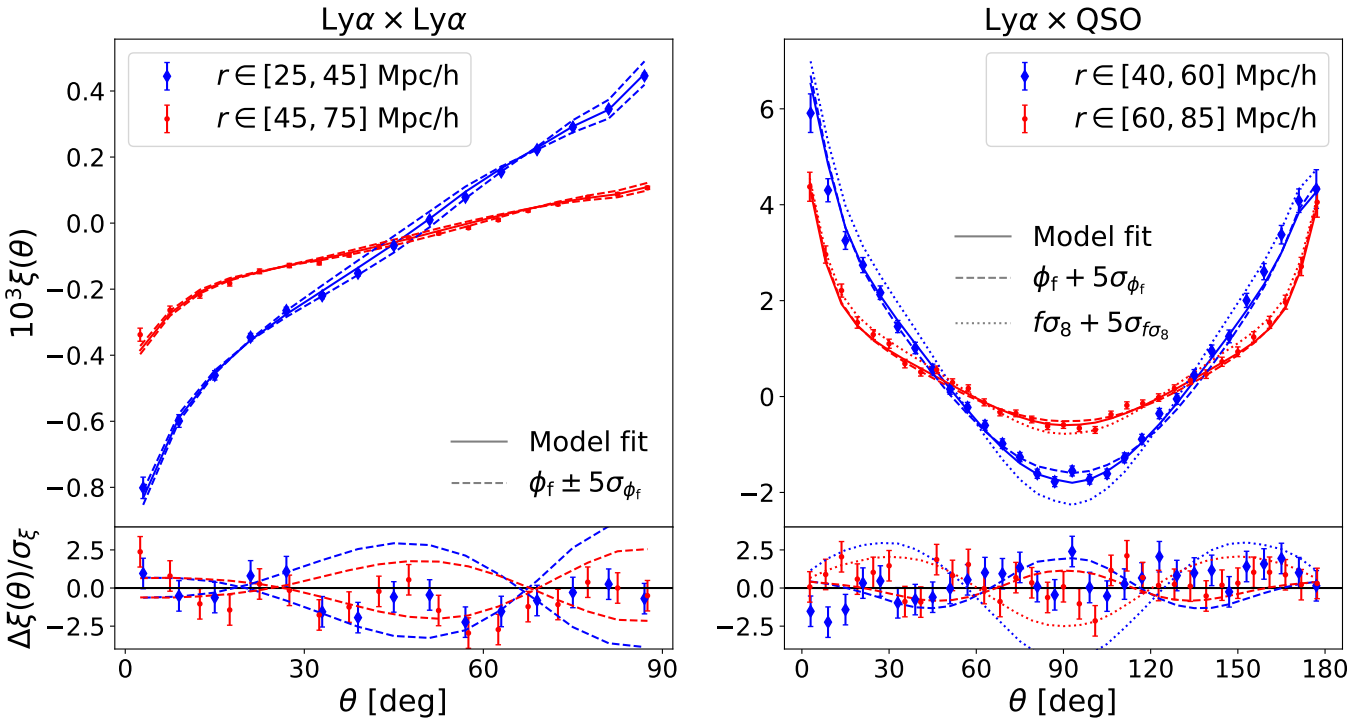

Figure 1. (left) Lyα forest auto-correlation function (points with error-bars) compressed into shells of separation r, and shown as a function of the line-of-sight angle, θ. The top panel shows the shell compression, while the bottom panel shows the residuals of the shells. The best-fit model is also compressed into the continuous lines and captures well the anisotropy present in the data. The dashed lines indicate models where the AP effect is modified by 5σ in either direction, illustrating its impact on these shells. (right) Shell compression of the Lyα-QSO cross-correlation function. θ values below 90° correspond to Lyα pixels in front of the quasar they are correlated with, whereas values above 90° correspond to Lyα pixels behind the quasar (i.e., at larger redshift). Here we also add a model where the value of $f\sigma_8$ is modified by 5σ, shown using dotted lines. Note that the separations used to select the shells are different for the auto- and cross-correlation due to differing scale cuts.

	Isotropic BAO	Alcock-Paczynski	Growth ($z_{\rm eff} = 2.33$)
BAO (DESI24-Lyα)	$\alpha_{\rm p} = 1.000 \pm 0.011$	$\phi_{\rm p} = 1.026 \pm 0.038$	-
Broadband	-	$\phi_{\rm s} = 0.987 \pm 0.017$	$f\sigma_8 = 0.364^{+0.056}_{-0.065}$
Full-shape	$\alpha_{\rm p} = 0.999 \pm 0.010$	$\phi_{\rm f} = 0.995 \pm 0.016$	$f\sigma_8 = 0.371^{+0.055}_{-0.065}$

Table 1. Constraints on BAO, the Alcock-Paczynski effect, and the growth parameter combination $f\sigma_8$, from BAO-only, broadband-only, and full-shape analyses of the DESI DR1 Lyα forest correlation functions. The BAO-only results come from DESI24-Lyα, while the other two analyses are presented here.

1.6% constraint on the Alcock-Paczynski effect ($\phi_{\rm f}$), representing an improvement of about 2.4× in constraining power.

We also obtain the first ever measurement of the growth rate from 3D Lyα forest correlations, quantified through the combined parameter $f\sigma_8$. The constraint is much weaker ($15\% - 17\%$) when compared to the AP result, due to the fact that we are only using the quasar RSD signal in the Lya-QSO cross-correlation for this measurement.

4. ANALYSIS VALIDATION

Our approach for analysis validation closely follows that of DESI24-Lyα. First, we used a set of 150 mock realizations of DESI DR1 to validate our analysis pipeline, by showing that we can recover unbiased constraints and accurate uncertainties (Section 4.1). After that, we performed a large set of data splits (Section 4.2) and analysis variations (Section 4.3) in order to study the robustness of our measurement. These tests were first performed using the blinding method described above, and informed the choice of baseline model.

As our analysis starts from the correlation functions measured by DESI24-Lyα, the main purpose of this section is to validate the robustness of the model for these correlation functions when it comes to measurements of $\phi_{\rm f}$ and $f\sigma_8$. Therefore, we imposed a threshold of 1/3

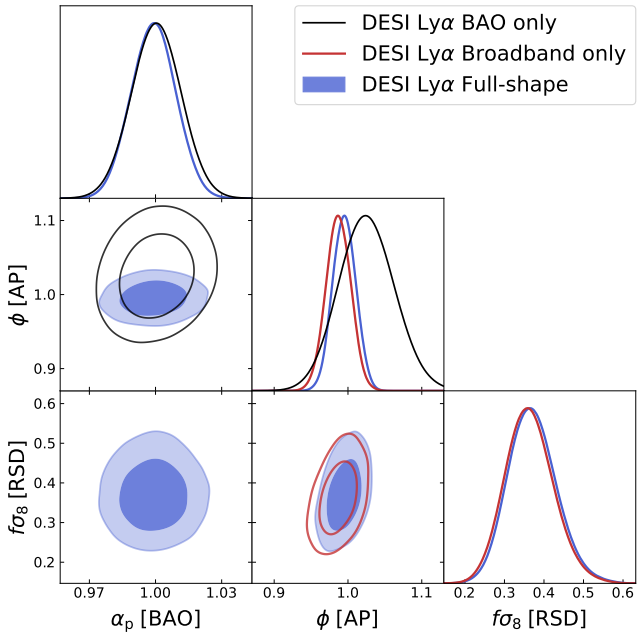

Figure 2. Constraints on the isotropic BAO scale, $\alpha_{\rm p}$, the Alcock-Paczynski effect, ϕ, and the growth parameter combination, $f\sigma_8$. Our main results are shown in blue for the full-shape analysis of the DESI DR1 Lyα forest correlations, where the AP constraint is given by $\phi_{\rm f}$. For comparison, we also show the BAO measurement measured by DESI24-Lyα from the same dataset in black (AP given by $\phi_{\rm p}$), and results from only the broadband in red (where we marginalize over BAO and AP is given by $\phi_{\rm s}$). The Alcock-Paczynski constraint from the broadband is more than a factor of two tighter than the one from BAO.

of the DESI statistical uncertainty, $\sigma_{\rm DESI}$, for tests of the model to be considered successful. This corresponds to ~ 0.005 for $\phi_{\rm f}$ and ~ 0.02 for $f\sigma_8$. We will show this threshold in all figures where it was applied. Note that for model variations, this threshold was applied to $\phi_{\rm s}$ before unblinding, but we did not observe any significant changes when working with $\phi_{\rm f}$ post-unblinding, and therefore present all our results in terms of $\phi_{\rm f}$.

While our main results in Section 3 were based on the full posterior distribution as measured by the Nested Sampler `PolyChord`, for this section we were limited by computational constraints to using the `iminuit` minimizer with Gaussian approximated uncertainties[7]. This was also the case for DESI24-Lyα, but in their case the posterior distribution of the BAO parameters was very well approximated by a Gaussian. On the other hand, in our case the Gaussian approximation can lead to small differences with respect to the full posterior results as described in Appendix C (also see Cuceu et al. 2020). We will discuss the implications of this limitation wherever it is relevant in this section. We will refer to measurements based on the full posterior distribution as sampler results, while those based on the best-fit with Gaussian approximated uncertainties as fitter results.

[7] Computed from the second derivative of the likelihood around the best-fit point.

4.1. *Validation with mocks*

We use the set of mocks created by Cuceu et al. (2025) to validate the DESI DR1 Lyα forest BAO measurement in DESI24-Lyα. For a detailed description of how these mocks are created, see Herrera-Alcantar et al. (2023). Furthermore, the same type of mocks were also used to validate the Lyα full-shape analysis of eBOSS DR16 in Cuceu et al. (2023b), although that analysis only focused on the Alcock-Paczynski effect and did not study RSD. Therefore, here we only give a very brief description of these mocks, focusing on the relevant implications for full-shape analyses.

The mocks were generated using two different sets of fast simulations based on the `Ly`α`CoLoRe` (Farr et al. 2020; Ramírez-Pérez et al. 2022) and `Saclay` packages (Etourneau et al. 2023). Both of these use log-normal density fields and simplified recipes to generate the quasar distribution and paint the Lyα forest absorption. The recipes were calibrated to approximately match the mean flux, 1D power spectrum, and large-scale bias and RSD parameters measured by eBOSS. These simulations are then used to generate mock realizations of the DESI DR1 quasar catalog with matching synthetic DESI spectra containing simulated Lyα forests. This is done using the `desisim` package[8], which simulates the DESI DR1 footprint, magnitude and redshift distributions, spectrograph resolution, noise, and redshift errors. At this stage, astrophysical contaminants are also added, which include Damped Lyα systems (DLAs), Broad Absorption Lines (BALs), and metal absorption (Herrera-Alcantar et al. 2023). We use the set of 100 `Ly`α`CoLoRe` and 50 `Saclay` mocks of DESI DR1 generated in Cuceu et al. (2025) and used to validate the analysis of DESI24-Lyα.

The Lyα forest correlation functions were computed using the same analysis pipeline and methodology that was used to analyze the real data (Section 2). The only difference with respect to the analysis presented in Cuceu et al. (2025) is that we now go beyond BAO in our model and focus on the full-shape information as described in Section 2.2. Following Cuceu et al. (2025), the model applied to mocks differs slightly from the one

[8] https://github.com/desihub/desisim

applied to data in the following aspects. First, we include extra Gaussian smoothing to account for the finite cell size in the log-normal simulations. This includes two new free parameters corresponding to the smoothing amplitude along and across the line-of-sight, respectively. Secondly, we ignore correlated sky residuals, CIV contamination, the transverse proximity effect, and BAO broadening due to non-linear evolution since these effects are not included in the mocks. None of these effects are correlated with either Alcock-Paczynski or RSD given our scale cuts, and we study their impact below using the real data[9]. Finally, log-normal mocks do not have correct small-scale clustering, so we do not use the Arinyo model in our baseline mock fits. We have tested including this model and freeing all 5 extra parameters, and we find no significant impact on our measurements. In fact, similar to the measurements on data, we find no preference for deviations from linear theory (i.e., the q_1 parameter is consistent with zero).

We start by fitting the mean of the correlation functions computed from the two sets of mocks, which we will refer to as stacked correlations. Cuceu et al. (2025) showed that the model described in Section 2.2 provides an excellent fit to these stacked correlations, even when fitting down to 10 h^{-1}Mpc and fixing the broadband scale parameters. Furthermore, Cuceu et al. (2025) also found that BAO measurements are unbiased in these mocks. Therefore, our goal here is to study whether we recover unbiased measurements of AP and RSD. Our measurements of $\phi_{\rm f}$ and $f\sigma_8$ from the stack of mocks are shown in Figure 3. For comparison, we also show the input mock cosmology marked by the gray cross, along with the DESI contours centered on the input cosmology (continuous black), as well as the contour corresponding to 1/3 of the DESI uncertainty (dashed black).

Figure 3 shows two measurements for each of the two sets of mocks. The baseline results are represented by the filled contours, and in the case of the Alcock-Paczynski effect they show excellent agreement with the fiducial cosmology used to generate the mocks for both `Ly$\alpha$CoLoRe` and `Saclay` mocks (vertical gray line). On the other hand, the two sets of mocks give slightly different constraints on $f\sigma_8$. `Ly$\alpha$CoLoRe` mocks are in agreement with the fiducial value (horizontal gray line), while `Saclay` mocks show a small but significant bias. The $f\sigma_8$ measurement from `Saclay` mocks is on the edge of the $1/3\,\sigma_{\rm DESI}$ contour in 2D. However, in 1D it corresponds to a shift of $\sim 0.5\,\sigma_{\rm DESI}$ from the truth, and is detected

[9] See Cuceu et al. (2023b) for a discussion of these effects and their impact on full-shape analyses.

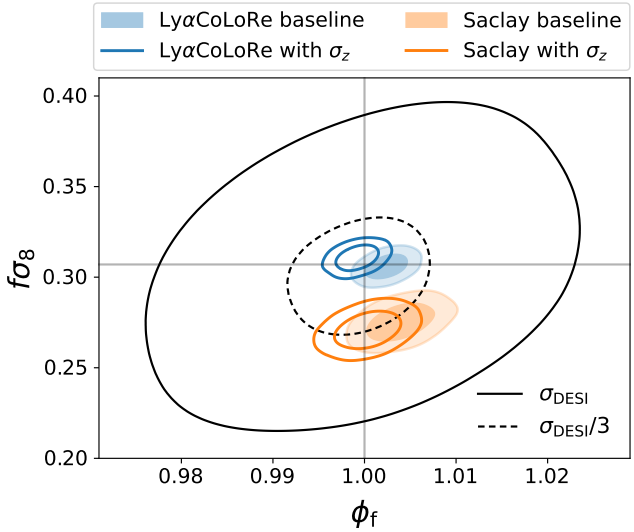

Figure 3. 68% and 95% credible regions of AP ($\phi_{\rm f}$) and RSD ($f\sigma_8$) measurements from the stack of 100 `Ly$\alpha$CoLoRe` (blue) and 50 `Saclay` (orange) mocks. The filled contours show the results for the baseline mock analysis, while the empty contours show results for a variation where redshift errors affect the continuum fitting. The gray cross indicates input mock cosmology, while the dashed black contour shows 1/3 of the DESI DR1 uncertainty. The AP measurement is unbiased in all cases, while the RSD measurement is unbiased for `Ly$\alpha$CoLoRe` mocks, and shows a small but significant bias for `Saclay` mocks.

at a significance of more than 4σ. Given that `Ly$\alpha$CoLoRe` mocks are consistent with the fiducial cosmology, it is difficult to draw strong conclusions from this result, as it may be indicative of an issue with one or both sets of mocks rather than a true bias. The redshift-space mapping is not incorporated in the same way for the two sets of mocks, and the two methods rely on different approximations. We discuss this further in Section 5.

The empty contours in Figure 3 show results from a different type of analysis of the same sets of mocks. This variation simulates Gaussian redshift errors that affect the continuum fitting process. The baseline analysis does include redshift errors, but only in the quasar catalog used to compute cross-correlations. Youles et al. (2022) found that redshift errors can also impact the continuum fitting process, leading to spurious correlations appearing in both Lyα auto and Lyα-QSO cross-correlations. Figure 3 shows that while these spurious correlations do have an impact on our measurements, it is much smaller than the $1/3\,\sigma_{\rm DESI}$ threshold, and does not impact our conclusions. While this variation is useful for testing our sensitivity to this effect, the way redshift errors are simulated in these mocks leads to an extreme version of this contamination (see Youles et al.

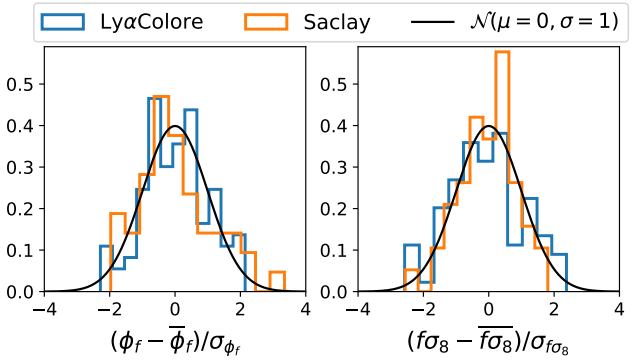

Figure 4. Pull distributions for $\phi_{\rm f}$ (left) and $f\sigma_8$ (right) from the set of 150 mock analyses. The black lines show the expected Gaussian distributions with null mean and unit variance.

2022; Cuceu et al. 2025; Gordon et al. 2025). This is why Cuceu et al. (2025) did not include this effect in their baseline analysis of these mocks, and we take the same approach here.

Using the mock baseline stacked correlation functions, we also studied how scale cuts affect our ability to recover unbiased constraints on AP and growth. Cuceu et al. (2023b) previously performed this test for $\phi_{\rm s}$ using eBOSS mocks (see their Figure 5), and our results are in good agreement with their findings. In particular, we find that using scales below 20 h^{-1}Mpc produces biased constraints on $\phi_{\rm f}$ at more than 2σ significance in both `Ly`α`CoLoRe` and `Saclay` mocks. On the other hand, for $f\sigma_8$ we find biased results when using cross-correlation scales below 35 h^{-1}Mpc in the case of `Ly`α`CoLoRe` mocks. For Saclay mocks, restricting both the auto- and cross-correlation to scales larger than 50 h^{-1}Mpc does produce $f\sigma_8$ constraints consistent with the truth at the 2σ level. However, that is largely due to the significant increase in uncertainty, rather than a shift in the result.

Next, we perform full-shape fits to individual mocks and study the population statistics of the results. The main goal here is to test the robustness of our uncertainties. Similar to Cuceu et al. (2023b), we are limited by computational constraints to using the `iminuit` minimizer instead of sampling the full posterior distributions. Therefore, our individual mock uncertainties are computed under the assumption that the posterior distributions are Gaussian.

Figure 4 shows the pull distributions for $\phi_{\rm f}$ and $f\sigma_8$, obtained from the difference between individual results and the mean of all mocks, divided by the measured uncertainty (i.e. $[X-\overline{X}]/\sigma_X$). The results from `Ly`$\alpha$`CoLoRe` mocks are shown in blue, while those from `Saclay` mocks are shown in orange. We do not find any significant outliers, with the most extreme result being 3.3σ away from the truth, which is not surprising in a sample of 150 mocks. Figure 4 also shows the unit variance Gaussian for comparison. The distribution of pull values for $\phi_{\rm f}$ matches well the expected Gaussian, while for $f\sigma_8$ the histogram shows some deviations. This could be explained by the fact that we are using the Gaussian approximation for the individual mock fits. We present a comparison of sampler and fitter results for the data in Appendix C, where we show that the Gaussian approximation used to estimate uncertainties in the fitter case produces an uncertainty for $f\sigma_8$ of ±0.057, while the uncertainty from the sampler is asymmetric: $^{+0.055}_{-0.065}$. Therefore, we should not expect a perfect match between the histogram of pull values for $f\sigma_8$ and the unit variance Gaussian. However, a full confirmation of the robustness of $f\sigma_8$ uncertainties would require running the sampler on the full set of mocks, which we leave for future work.

We also compute the standard deviations of these distributions, which we find to be 1.04 ± 0.06 for $\Delta\phi_{\rm f}/\sigma_{\phi_{\rm f}}$ and 1.08 ± 0.06 for $\Delta f\sigma_8/\sigma_{f\sigma_8}$, with uncertainties obtained through bootstrap. The 68% credible regions[10] are also very similar, with 1.03 ± 0.10 for $\Delta\phi_{\rm f}/\sigma_{\phi_{\rm f}}$ and 1.07 ± 0.11 for $\Delta f\sigma_8/\sigma_{f\sigma_8}$. These results show that the uncertainties for both $\phi_{\rm f}$ and $f\sigma_8$ are well estimated given the uncertainty from the limited number of mocks.

The tests we have performed in this section show that measurements of the AP effect using the full shape of Lyα forest correlations are unbiased with well-estimated uncertainties. On the other hand, growth measurements are unbiased only when using `Ly`α`CoLoRe` mocks, while `Saclay` mocks show a small but significant bias of about $\sim 0.5\,\sigma_{\rm DESI}$. Based in part on these results, we decided not to use our growth rate measurements for cosmological inference. We discuss this further in Section 5.

4.2. *Data splits*

The next set of tests involves data splits, where we perform full-shape analyses of different subsets of the data in order to check for potential inconsistencies. We start by fitting the two auto-correlations, Lyα(A)×Lyα(A) and Lyα(A)×Lyα(B), separately from the two cross-correlations, Lyα(A)×QSO and Lyα(B)×QSO. As a joint fit of the auto and cross-correlation is needed to break the degeneracy between the Lyα and quasar RSD terms, for this test we can only measure the Alcock-Paczynski effect. We show the constraints on $\phi_{\rm f}$ and $\alpha_{\rm p}$ from this data split in Figure 5. The results from

[10] Computed from half the distance between the 16th and 84th percentiles.

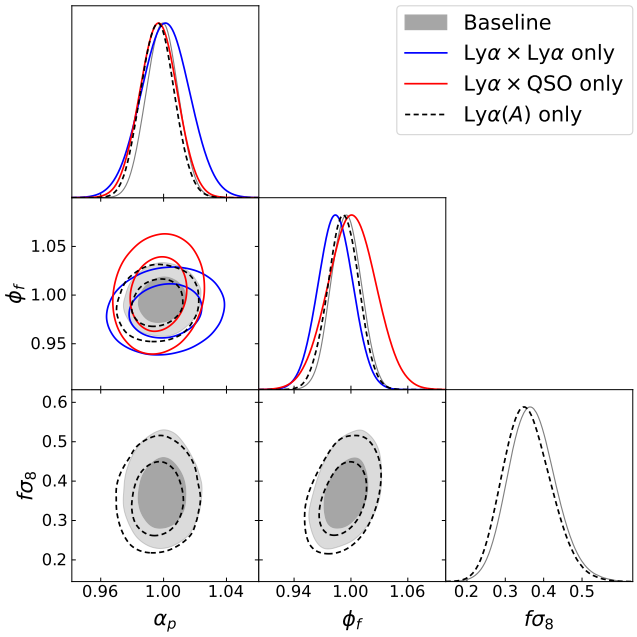

Figure 5. Full-shape constraints from the two Lyα auto-correlations (blue) versus constraints from the two Lyα-QSO cross-correlations (red), as well as results from the two correlations that only use the Lyα(A) region (black dashed). All three data splits are in good agreement with each other and the baseline results that use all 4 correlations (shown in gray).

the auto and cross-correlations are consistent with each other. Figure 5 shows that while the $\alpha_{\rm p}$ constraint from the cross is tighter than the one from the auto, the opposite is true for $\phi_{\rm f}$. In the case of $\alpha_{\rm p}$, this is consistent with what DESI24-Lyα found and also with the variance of the correlation functions, which is larger for the auto compared to the cross in DESI DR1. On the other hand, we find that the broadband AP constraint from the cross is much weaker than the one from the auto, with the $\phi_{\rm f}$ uncertainty from the cross being $\sim 40\%$ larger compared to the auto. This is related to the fact that the Lyα RSD parameter is significantly larger than the QSO one (~ 1.5 versus ~ 0.25), which means that the anisotropy on large scales is stronger in the Lyα auto-correlation compared to the cross-correlation.

Figure 5 also shows the results when we only use the Lyα(A) region. In this case, we are fitting both the auto and cross-correlation (i.e., Lyα(A)×Lyα(A) and Lyα(A)×QSO), and therefore have both AP and RSD constraints. The results are very similar to our baseline constraints using all 4 correlations. This is because the Lyα(B) correlations are much noisier due to a few reasons, including that there are fewer forests that contain this region, and that the region is shorter and noisier due to the presence of other Lyman lines. We find that

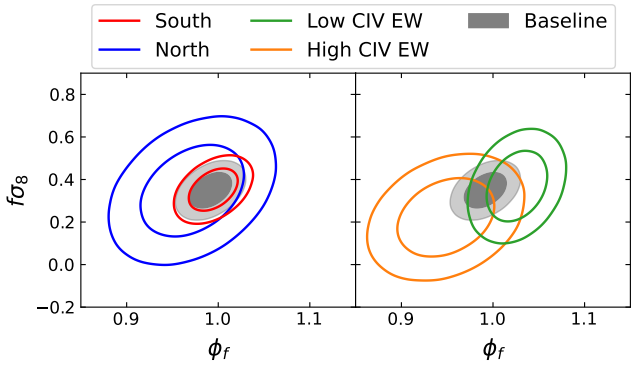

Figure 6. AP and RSD constraints from the baseline analysis (gray) and from two different data splits. In the left panel, we show results from the samples using the South (red) vs North (blue) imaging for quasar target selection. In the right panel, we show results from the samples with Low (green) vs High (orange) CIV equivalent width in the quasar spectra.

adding the two Lyα(B) correlations only improves the constraints by about 3% in both $\phi_{\rm f}$ and $f\sigma_8$.

The final two data splits we test are shown in Figure 6, and are based on measurements by DESI24-Lyα. In the left panel, we show results from a data split based on the imaging surveys used for quasar target selection. The sample designated as "South" is based on imaging from the DECam camera, and includes the entire South Galactic Cap and the southern part of the North Galactic Cap, at $\delta < 32.375°$. The sample designated as "North" is based on imaging from the BASS and MzLS surveys, and only contains 18% of quasars. In the right panel, we show results from a split based on the equivalent width (EW) of the CIV line. This split is done because of the anti-correlation between the quasar continuum luminosity and the EW of the emission lines, known as the Baldwin effect Baldwin (1977). We find consistent results for both of these data splits, as shown in Figure 6.

4.3. *Analysis variations*

The next set of tests we perform includes a large set of variations in the analysis methods used up to the measurement of the correlation functions. As the starting point in this work is from the correlation functions measured by DESI24-Lyα, these tests were all originally performed for the validation of the BAO analysis, and here we only reanalyze those correlation functions in the context of the full-shape analysis. We present the shifts in $\phi_{\rm f}$ and $f\sigma_8$ produced by these variations in Figure 7 (points with errorbars), along with the uncertainty from the baseline analysis for comparison (gray shaded bands).

The first subset of analysis variations consists of changes to the dataset. Therefore, we expect these variations to be affected by statistical fluctuations. The shifts in $\phi_{\rm f}$ and $f\sigma_8$ due to these variations are shown in brown in Figure 7, and consist of the following tests:

- Only quasar targets: Use only quasars that were originally targeted as quasars.
- No sharp lines mask: Do not mask the sharp lines in quasar spectra, as described in Ramírez-Pérez et al. (2024).
- >50 pixels in forest: Use lines-of-sight that have at least 50 valid Lyα pixels, as opposed to the 150 pixel threshold in the baseline analysis.
- Mask-Lyα redshift estimates: Use a different estimator for quasar redshifts, that only takes into accounts wavelengths longer than the Lyα emission line.
- $\lambda_{RF} < 1200$Å: Use only pixels below 1200Å in rest-frame wavelength (1205Å in the baseline analysis).
- $\lambda_{obs} > 3650$Å: Use only pixels above 3650Å in observed wavelength (3600Å in the baseline analysis).
- $\lambda_{obs} < 5500$Å: Use only pixels below 5500Å in observed wavelength (5577Å in the baseline analysis).
- $z_{\rm Q} < 3.78$: Use only quasars at redshifts below 3.78, which corresponds to the highest redshift included in the mocks.
- DLAs SNR > 1: Mask DLAs in spectra that have SNR > 1, as opposed to SNR > 3 for the baseline analysis.

The next subset of analysis variations includes those that affect the continuum fitting process. These variations can also introduce statistical fluctuations, as they impact the weights used both for continuum fitting and for measuring the correlation functions. The shifts in AP and growth due to these changes are shown in purple in Figure 7, and consist of the following tests:

- $\Delta\lambda = 2.4$Å: Perform the pixel re-binning from 0.8Å to 2.4Å before the continuum fitting process, as done in du Mas des Bourboux et al. (2020). For this variation, we set $\sigma^2_{\rm mod} = 3.1$, following Ramírez-Pérez et al. (2024).

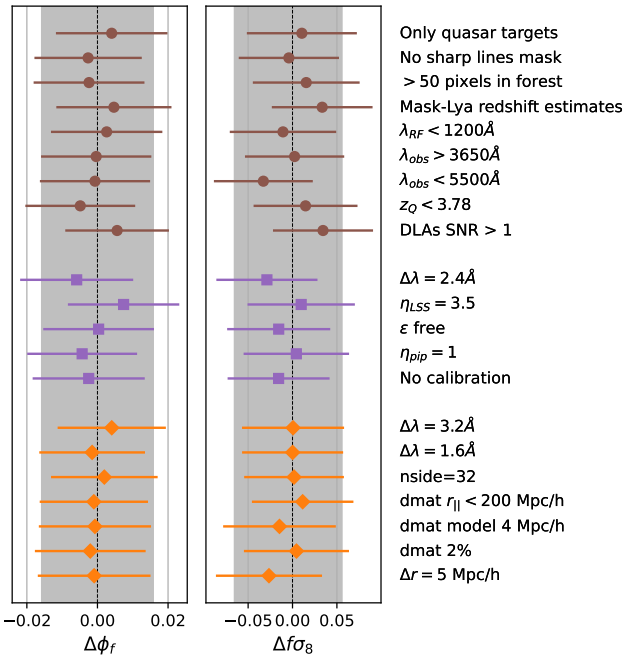

Figure 7. Shifts in the AP and growth measurements from the set of analysis variations performed by DESI24-Lyα. This includes variations where the dataset is different (brown), variations in the method to estimate the Lyα overdensities (purple), and variations in the method used to compute correlations, covariances, and distortion matrices (orange). The gray bands show the uncertainty of the baseline measurement. Most of these variations are subject to statistical fluctuations (especially the ones in brown and purple).

- $\eta_{\rm LSS} = 3.5$: Reduce the factor that scales the intrinsic Lyα forest variance from 7 to 3.5, thereby reducing its contribution to the weights.
- ϵ free: Include an extra term in the Lyα weights meant to capture quasar diversity, following du Mas des Bourboux et al. (2020).
- $\eta_{\rm pip} = 1$: Do not re-calibrate the instrumental noise. See Ramírez-Pérez et al. (2024) for more details.
- No calibration: Do not re-calibrate spectra with the CIII region, as described in Ramírez-Pérez et al. (2024).

The final subset of analysis variations includes changes to how correlation functions, covariance matrices, and distortion matrices are computed. These are shown in orange in Figure 7, and consist of the following tests:

- $\Delta\lambda = 3.2$Å: Rebin Lyα pixels in groups of 4 (3 in the baseline analysis).
- $\Delta\lambda = 1.6$Å: Rebin Lyα pixels in groups of 2 (3 in the baseline analysis).

- nside = 32: Measure the covariance matrix from a set of correlations measured in HEALPix pixels with nside = 32, as opposed to nside = 16 in the baseline analysis. This results in more HEALPix pixels that are smaller in size.

- dmat $r_{||} < 200$ Mpc/h: Compute the distortion matrix up to $r_{||} < 200$ h^{-1}Mpc, as opposed to $r_{||} < 300$ h^{-1}Mpc in the baseline analysis.

- dmat model 4 Mpc/h: Compute the distortion matrix with 4 h^{-1}Mpc bins for the model, matching the binning of the data. In the baseline analysis, 2 h^{-1}Mpc bins are used for the model.

- dmat 2%: Use 2% of pixels to compute the distortion matrix, as opposed to 1% in the baseline analysis.

- $\Delta r = 5$ Mpc/h: Use 5 h^{-1}Mpc bins for the correlation function measurements, as opposed to the 4 h^{-1}Mpc bins used in the baseline analysis.

All these variations produce results consistent with the baseline, with shifts larger than 1/3 of the DESI uncertainty only observed for variations where we expect statistical fluctuations. As noted earlier, we used the fitter with Gaussian approximated uncertainties for the individual tests, while for the baseline analysis, we computed the full posterior distribution. As shown in Appendix C, the fitter results can be shifted with respect to the sampler results due to the posterior being asymmetric, and more importantly, due to the difference between conditioning and marginalizing over nuisance parameters. Therefore, we also expect some extra variance in these results due to these differences.

4.4. *Model variations*

We next turn our attention to variations in the modelling process. Unlike the tests in the previous section, most of these variations were designed specifically to test the robustness of the full-shape analysis. This means we perform a set of tests that is fairly different from the one done for the BAO analysis in DESI24-Lyα.

The first subset of model tests involves changes to the scale cuts, and the resulting shifts in $\phi_{\rm f}$ and $f\sigma_8$ are shown in Figure 8. In the baseline analysis, we use the same maximum fitted scale for all correlations, $r < 180$ h^{-1}Mpc, while for the minimum fitted scale we use $r_{\rm auto} > 25$ h^{-1}Mpc for the auto-correlation, and $r_{\rm cross} > 40$ h^{-1}Mpc for the cross-correlation. Here we test changes to these choices that go in both directions, with choices that are more conservative than the baseline shown using green circles, and less conservative choices shown using red squares. These variations also introduce statistical fluctuations because we are changing which data affects our results. This is especially true of changes to the minimum separation, as our measurements come in large part from small scales (Cuceu et al. 2021).

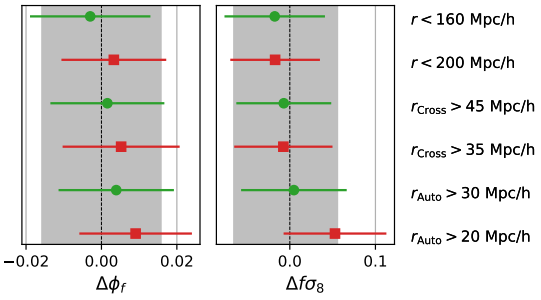

Figure 8. Similar to Figure 7, but showing shifts due to changing the scale cuts used when fitting the model to the measured correlation functions. Changes to the minimum separation are performed independently for the auto and cross-correlation, as they do not use the same limit. Cuts that are more conservative than the baseline are shown using green circles, while less conservative cuts are shown using red squares. These variations are also affected by statistical fluctuations.

The results shown in Figure 8 are consistent with those obtained using the baseline cuts. The only large shifts are observed when decreasing the minimum scale for the auto-correlation. For this particular case, we also observe a significant increase in the magnitude of the HCD bias, indicating that HCDs start to play an important role at separations below 25 h^{-1}Mpc. While in the baseline analysis the HCD bias is consistent with zero (see Appendix B), in the variation with $r_{\rm auto} > 20$ h^{-1}Mpc we strongly detect a non-zero HCD bias at $\sim 12\sigma$.

The final set of variations includes changes to the model used to fit the correlation functions. The shifts in AP and growth produced by these variations are shown in Figure 9, and consist of the following tests:

- Gaussian redshift errors: Use a Gaussian distribution to model redshift errors and peculiar velocities of quasars, as opposed to a Lorentzian distribution in the baseline analysis.

- Weak $\Delta r_{||}$ prior: Use a flat prior for the parameter that models the systematic quasar redshift error, $-3 > \Delta r_{||} > 3$ h^{-1}Mpc. In the baseline analysis, a Gaussian prior is used, with $\mathcal{N}(0, 1)$.

- Weak CIV prior: Use a flat prior for the CIV bias parameter, with $-0.5 < b_{\rm CIV} < 0$. In the baseline analysis, a Gaussian prior is used, with $\mathcal{N}(-0.0243, 0.0015)$.

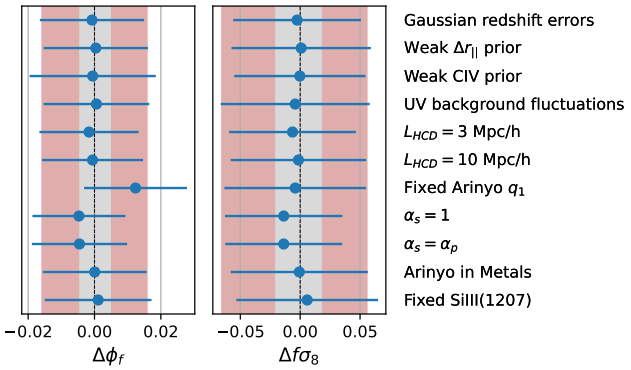

Figure 9. Shifts in the AP and growth measurements when changing the model used to fit the correlation functions. The red bands show the uncertainty from the baseline analysis, while the smaller gray bands show the threshold used for these tests, which represents 1/3 of the uncertainty from the baseline analysis.

- UV background fluctuations: Add the impact of fluctuations in the UV background following Gontcho A Gontcho et al. (2014) and Bautista et al. (2017).
- $L_{\rm HCD}$ = 3 Mpc/h: Fix the value of $L_{\rm HCD}$ to 3 h^{-1}Mpc. This is allowed to vary in the baseline analysis.
- $L_{\rm HCD}$ = 10 Mpc/h: Fix the value of $L_{\rm HCD}$ to 10 h^{-1}Mpc. This is allowed to vary in the baseline analysis.
- Fixed Arinyo q_1: Fix the Arinyo q_1 parameter to the value used in BAO analyses: $q_1 = 0.86$. This is based on the measurements from hydro-dynamical simulations of Arinyo-i-Prats et al. (2015), interpolated to our effective redshift. In the baseline analysis, this is allowed to vary.
- $\alpha_{\rm s} = 1$: Fix the value of the isotropic scale parameter for the smooth component ($\alpha_{\rm s}$) to 1. In the baseline analysis, this is allowed to vary.
- $\alpha_{\rm s} = \alpha_{\rm p}$: Impose the condition that the isotropic scale parameters for the peak and smooth components are the same, therefore fitting for the full-shape isotropic scale. In the baseline analysis, they are varied independently of each other.
- Arinyo in Metals: Use the Arinyo non-linear correction when modelling the Lyα-Metal correlations as well.
- Fixed SiIII(1207): Fix the SiIII(1207) bias parameter to the best-fit value obtained by DESI24-Lyα: $10^3 b_{\rm SiIII(1207)} = -9.79$. This parameter is allowed to vary in the baseline analysis. The Lyα-SiIII(1207) peak is outside our fitting range ($r_{||} \sim 21$ h^{-1}Mpc), but we are still sensitive to its presence due to the distortion matrix (Cuceu et al. 2023b).

The only variation that produces a shift larger than the $1/3\,\sigma_{\rm DESI}$ threshold is the one where the Arinyo q_1 parameter is fixed to the value measured from the reference hydrodynamical simulation of Arinyo-i-Prats et al. (2015), $q_1 = 0.86$. However, the data disfavors such a high value for this parameter with a significance of $\sim 3\sigma$ ($\Delta\chi^2 = 9.4$). See Appendix B for a discussion of this result. This test has informed our choice of allowing all parameters in the Arinyo model to vary freely in the baseline analysis, instead of imposing constraints based on other measurements. We note that the parameters of this model also have a cosmology dependency, so it would be interesting to explore the implications of their constraints, as recently proposed by Chaves-Montero et al. (2025). However, since these parameters are treated as a nuisance here, their measurement has not been validated, so a thorough study into these constraints is needed before any cosmological interpretation can take place.

The remaining tests shown in Figure 9 are all consistent with the baseline measurements to within the $1/3\,\sigma_{\rm DESI}$ threshold. Also note that the two variations that relate to the isotropic scale parameter, $\alpha_{\rm s}$, both produce significantly smaller uncertainties in both $\phi_{\rm f}$ and $f\sigma_8$ compared to the baseline analysis, and produce shifts that are close to our threshold. For these two constraints, we have validated the results by running the sampler as well. We discuss these further in Section 5.

5. DISCUSSION

The measurements presented in this work represent the tightest Alcock-Paczynski constraint at $z > 1$ to date, and the first measurement of $f\sigma_8$ from the Lyα forest cross-correlation with quasars. To validate these results, we performed a large set of tests, including mock analyses, data splits, and analysis and model variations. In this section, we wish to discuss our measurements and the results of the validation tests.

As the starting point of our analysis was from the correlation functions measured by DESI24-Lyα, we largely focused our validation on the modelling of correlations, and all analysis choices we made were at the level of the model. Therefore, we did not investigate in detail the impact of changes to the analysis methodology prior to the measurement of the correlations. The tests shown in Figure 7 are generally consistent with the baseline

result, with no shifts larger than 0.5σ for $\phi_{\rm f}$ and 0.6σ for $f\sigma_8$. Furthermore, for all of the larger shifts, we expect some variation due to statistical fluctuations caused by either changing the dataset or changing the weights used to compute the correlations. Future analyses that also consider changes to the analysis method prior to the measurement of the correlations would also need to study whether the observed shifts are consistent with statistical fluctuations, and if any of them are not, study changes to the analysis method that address those shifts. Also, as the current Lyα forest analysis framework was built over many iterations with the goal of maximizing BAO signal-to-noise, it would be interesting to consider changes that instead minimize the impact of contaminants on full-shape analyses. For example, in-depth tests of the impact of DLA and BAL masking on full-shape measurements (similar to Brodzeller et al. 2025; P. Martini et al. 2025) could inform more conservative analysis choices.

Our validation process was also limited by the use of the fitter with Gaussian uncertainties, instead of full posterior sampling. As discussed in Appendix C, the primary difference is caused by the fitter not fully marginalizing over nuisance parameters. This could introduce extra variance in the set of validation tests, especially for analysis variations where the shape of the posterior changes significantly compared to the baseline. One way to minimize this impact would be to impose Gaussian priors on all parameters that are not well constrained by the data. This is already done with some of the less important parameters ($b_{\rm CIV}$, $\beta_{\rm HCD}$, $\Delta r_{||}$), but those are all informed by direct measurements from other datasets. The parameters that are unconstrained by the data in our analysis are the Arinyo a_v, b_v, k_v, and k_p parameters, as well as the typical scale of un-masked HCDs, $L_{\rm HCD}$. Priors on these parameters could be informed by studies with mocks in the case of $L_{\rm HCD}$, or with simulations in the case of the Arinyo parameters. However, these are beyond the scope of the current work.

Another important aspect of our analysis that differentiates it from galaxy clustering analyses, is that we only measure the isotropic scale from the peak component while marginalizing over the isotropic scale of the broadband, $\alpha_{\rm s}$. This was a central part of the framework introduced by Cuceu et al. (2021), and it allows us to marginalize over potential systematics that isotropically affect the broadband. However, this also introduces a complication, because we still need this parameter to rescale our $f\sigma_8$ measurements as it affects the definition of σ_8 (Section 2.3). Furthermore, this parameter is correlated with both $\phi_{\rm f}$ ($\rho = 0.38$) and with $f\sigma_8$ ($\rho = 0.53$), meaning that systematics that affect it could still spill over into our primary measurements as well. Because $\alpha_{\rm s}$ has so far been treated as a nuisance parameter, not much attention was given to validating its constraint. In our analysis we have found that it is very weakly constrained when compared to $\alpha_{\rm p}$ (3.5% for $\alpha_{\rm s}$ versus 1% for $\alpha_{\rm p}$), but still consistent with the fiducial cosmology ($\alpha_{\rm s} = 1.040 \pm 0.035$). Therefore, the approach of fitting for a single full-shape isotropic scale may be preferred by future analyses. This does mean that the resulting $\alpha_{\rm f}$ measurement would also need to be validated as part of full-shape analyses (similar to what is done for galaxy clustering). However, this measurement would be dominated by the BAO signal given the difference in constraining power on $\alpha_{\rm p}$ and $\alpha_{\rm s}$, and therefore likely to still be robust. Finally, if this approach was taken for future analyses, the resulting constraints would be roughly 9% tighter for $\phi_{\rm f}$ and 18% tighter for $f\sigma_8$ based on our tests.

One of the most important parts of the validation process is showing that we can recover unbiased constraints using mocks. In Section 4.1 we found that $\phi_{\rm f}$ constraints are unbiased for both `Ly`α`CoLoRe` and `Saclay` mocks, while in the case of $f\sigma_8$ only constraints from the `Ly`α`CoLoRe` mocks were consistent with the input cosmology. Both sets of mocks contain all the major contaminants known to affect Lyα forest correlations (HCDs, metals, BALs, redshift errors). These contaminants are injected in the same way for both types of mocks (Herrera-Alcantar et al. 2023; Cuceu et al. 2025), and the analysis and modelling are also the same. Therefore, we consider it unlikely that the observed difference in $f\sigma_8$ is caused by contaminants. However, this is something that could be tested in future analyses by creating and analysing a set of uncontaminated mocks.

The impact of the fiducial cosmology was tested by `DESI24-Ly`α for $\alpha_{\rm p}$ in the context of DESI DR1. For $\phi_{\rm f}$, Cuceu et al. (2023b) studied this using eBOSS DR16 mocks and found no significant impact. Given that our statistical significance is similar to that of eBOSS DR16, we use the same mock and analysis pipelines, and our $\phi_{\rm f}$ result is within 1σ of the fiducial cosmology, we expect this conclusion to hold for DESI DR1 as well. On the other hand, this test has not been done for $f\sigma_8$. As we were not able to fully validate this measurement with mocks, it is unclear whether a test of the impact of the fiducial cosmology on $f\sigma_8$ constraints using the current generation of mocks could be trusted. Therefore, in the case of $f\sigma_8$, we leave these tests for future work.

Arguably, the most important limitation with the current Lyα forest mocks is the fact that they rely on log-normal approximation, which means they cannot be used to test the impact of small-scale deviations from linear theory. On one hand, our analysis focuses on large

scales where linear theory holds well (e.g. Hadzhiyska et al. 2023). Also, the onset of non-linearities happens at smaller scales for Lyα when compared to galaxy clustering, both due to the higher redshift and because it traces significantly less dense environments. On the other hand, a number of effects introduce extra sensitivity to small-scale non-linearities, in particular the distortion matrix and contamination due to metals (Bautista et al. 2017; Busca et al. 2025). The first spreads small-scale information to all scales along the line-of-sight, while the second manifests as small-scale Lyα-metal and QSO-metal correlations appearing as extra peaks along the line-of-sight. This means that incorrect modelling of the small scales in mocks can have an impact on the validation.

In the case of the Lyα auto-correlation, the Arinyo model gives us a way of gauging the importance of deviations from linear theory. We find that the data prefers values of the Arinyo q_1 parameter consistent with zero (Appendix B), and disfavours large deviations from linear theory over the scales we fit. Furthermore, we do not find significant correlations between any of the Arinyo parameters and $\phi_{\rm f}$ and $f\sigma_8$, nor with any of the other nuisance parameters. On the other hand, this type of test is more difficult for the cross-correlation, because our model for small-scale deviations from linear theory in the cross-correlation is very simplistic, and dominated by the impact of redshift errors. Furthermore, the small-scale quasar clustering in the `Ly`α`CoLoRe` mocks is significantly different than the one measured from data and N-body simulations (Ramírez-Pérez et al. 2022; Youles et al. 2022; Casas et al. 2025). These facts led us to take a much more conservative approach in our analysis of the cross-correlation, with the minimum scale used being $r > 40\ \ h^{-1}$Mpc, as opposed to $r > 25\ \ h^{-1}$Mpc for the auto-correlation. Nevertheless, we find that the $f\sigma_8$ measurement is somewhat correlated with both σ_z ($\rho = 0.56$) and a_{TP} ($\rho = -0.39$), indicating it is still fairly sensitive to the modelling of small-scale deviations from linear theory.

When comparing the realism of the mock correlation functions, Cuceu et al. (2025) found that the stacked cross-correlation of the DESI DR1 `Saclay` mocks matches the DESI data better when compared to the `Ly`α`CoLoRe` one. On the other hand, the `Ly`α`CoLoRe` stacked auto-correlation performs better than the one from `Saclay` mocks. However, in the end both sets of mocks are well fit by our model (see Figures 6 and 7 in Cuceu et al. 2025). Given this, it is difficult to draw a clear conclusion from our $f\sigma_8$ results in Section 4.1. While the $f\sigma_8$ signal comes from the cross-correlation, this measurement relies on the effective Lyα RSD term being calibrated by the auto-correlation. If we only used the cross-correlation alone, the two would be degenerate. Based on this, we conclude that at this stage we cannot fully confirm an unbiased $f\sigma_8$ measurement, or provide a well-understood measurement of a systematic uncertainty for this parameter. Further study with more realistic mocks that go beyond the log-normal approximation is needed to fully validate an $f\sigma_8$ measurement from the Lyα forest (e.g. Hadzhiyska et al. 2023; Sinigaglia et al. 2023). Therefore, in the case of the present analysis, we decided before unblinding that we would not use the resulting $f\sigma_8$ constraint for cosmology inference.

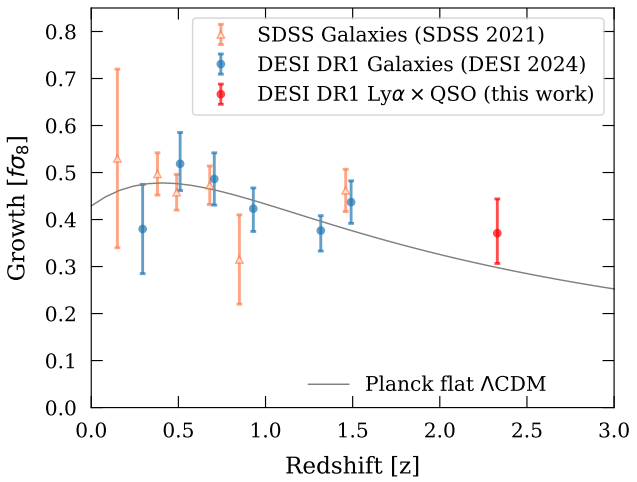

Figure 10. Measurements of the growth rate times the amplitude of fluctuations, $f\sigma_8$, as a function of redshift. The prediction from the Planck CMB data assuming the ΛCDM model is shown as the gray line. Measurements from galaxy clustering are shown in orange for the Sloan Digital Sky Survey (SDSS) and in blue for DESI. Our constraint using the cross-correlation between the Lyα forest and quasars is shown in red, and it includes an additional systematic uncertainty based on results from mocks.

To put our $f\sigma_8$ measurement in context, in Figure 10 we compare it to constraints from galaxy clustering at lower redshifts (Alam et al. 2017, 2021; DESI Collaboration et al. 2024a). We include an extra systematic uncertainty based on the shift measured from the stack of `Saclay` mocks, but caution that this is not a well-understood systematic as discussed above. The result we obtain is:

$$f\sigma_8(z_{\rm eff} = 2.33) = 0.37\ ^{+0.055}_{-0.065}\,({\rm stat}) \pm 0.033\,({\rm sys}). \quad (23)$$

To our knowledge, this is the first direct measurement of $f\sigma_8$ from large-scale structure at $z > 2$.

We now turn our attention back to the validation of the AP measurement. In this case, both sets of mocks show unbiased results. Furthermore, we found that the

impact of spurious correlations due to redshift errors (Youles et al. 2022) is much smaller than the current statistical uncertainties, and well within the threshold of $1/3\,\sigma_{\rm DESI}$ (Figure 3). A model for this contamination was recently introduced by Gordon et al. (2025). This was built by propagating the effect of redshift errors from the smearing of the mean quasar continuum to the resulting spurious correlations. Using the same set of mocks we use here, they showed that this model can correct the small shift observed in AP and RSD. As the statistical uncertainty of our measurements will shrink with the upcoming DESI data releases, these spurious correlations are likely to become a dominant source of systematic uncertainty, and this model will likely play an important role in future such analyses. Indeed, it has already become important for the DESI DR2 BAO analysis (Casas et al. 2025). However, at the level of the DR1 full-shape analysis, this systematic is still small, and therefore, we did not use the model developed by Gordon et al. (2025).

While the mocks we use contain all the major contaminants known to affect Lyα forest correlations, they do not include a number of the less significant effects (see Cuceu et al. 2023b, 2025). These include CIV contamination, the transverse proximity effect, BAO broadening due to non-linear evolution, UV background fluctuations, and correlated sky residuals. In the case of CIV contamination, the test shown in Figure 9, labeled *Weak CIV prior*, shows that the impact of this contamination is negligible. Without the Gaussian prior, we find that the CIV bias parameter is consistent with zero (within 1σ), indicating no detection of CIV contamination in the Lyα forest correlations given current uncertainties. The transverse proximity effect only impacts the cross-correlation, and as discussed in Section 4.2, the cross is less constraining than the auto when it comes to $\phi_{\rm f}$ measurements. The results from the auto and cross are also consistent, as shown in Figure 5, indicating that this contamination is not a cause for concern in case of AP measurements. The BAO broadening due to non-linear evolution only impacts the peak component, while most of our AP information comes from scales smaller than the BAO scale. It does affect the measurement of $\alpha_{\rm p}$ and $\phi_{\rm p}$ but those were validated as part of the BAO analysis by `DESI24-Ly`α.

UV background fluctuations introduce a scale-dependent bias (Pontzen 2014; Gontcho A Gontcho et al. 2014). This is not present in our baseline analysis as it has not been detected at a significant level so far. In Section 4.4, we perform a test where we model this effect following Gontcho A Gontcho et al. (2014) and Bautista et al. (2017). Similar to previous works (e.g. Bautista et al. 2017; Cuceu et al. 2023a; DESI Collaboration et al. 2025d), we do not detect it at a significant level (within 1σ of zero). We also find no impact on the AP measurement (Figure 9).

Correlated sky residuals were studied by Guy et al. (2025) using a synthetic data set created from the DESI DR1 data by replacing the original Lyα δ with realizations of the sky subtraction residuals. The correlation function of this mock data set gives the contamination expected in Lyα forest correlations. Guy et al. (2025) then built a model for this contamination and showed that it fits the measurement from the mock data set very well. This model is used both here and in `DESI24-Ly`α, and it has a free amplitude parameter, $a_{\rm noise}$. This parameter does have a mild correlation with $\phi_{\rm f}$ ($\rho = 0.24$). However, given that this contamination is very well understood and modelled, we do not consider it a significant source of concern. Nevertheless, this effect should be prioritized for inclusion in future mock data sets.

In conclusion, we have found that the measurement of the AP effect is robust, while the measurement of $f\sigma_8$ could not be fully validated due to deficiencies with the mock data sets. Therefore, in the following section we only use our measurements of ϕ to derive cosmological constraints.

6. COSMOLOGICAL INTERPRETATION

In this section, we study the cosmological implications of our Alcock-Paczynski measurement, alone and in combination with other data sets. We begin in Section 6.1 with the constraints on combinations of cosmic distances corresponding to our measurement. After that, we briefly introduce the external data sets we use in Section 6.2, and then present our results using the ΛCDM model in Section 6.3. Finally, we present constraints on the equation of state of dark energy in Section 6.4.

6.1. *Distance measurements*

We first focus on the relevant distance combinations for our AP and isotropic BAO measurements, which are given by Equations (21) and (22), respectively. The measurements corresponding to our full-shape results are:

$$\begin{cases} \textbf{AP}: & D_M/D_H(z_{\rm eff}) = 4.525 \pm 0.071, \\ \textbf{BAO}: & [D_M D_H(z_{\rm eff})]^{1/2}/r_d = 18.36 \pm 0.18, \end{cases} \tag{24}$$

with the effective redshift of our measurements equal to $z_{\rm eff} = 2.33$, and the correlation coefficient between the two distance combinations equal to 0.09. We also compute the distance ratio corresponding to our AP mea-

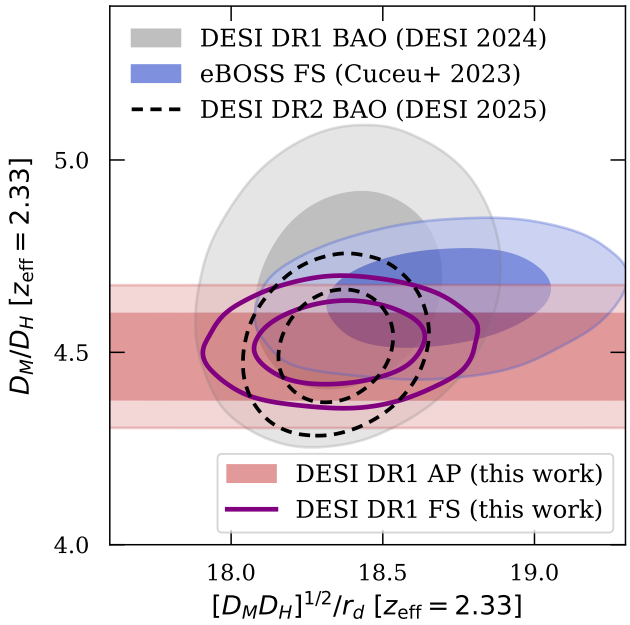

Figure 11. Constraints on the Alcock-Paczynski parameter given by the distance ratio D_M/D_H, and the distance combination corresponding to our isotropic BAO measurement, $[D_M D_H(z_{\rm eff})]^{1/2}/r_d$. DESI Ly$\alpha$ BAO results are shown in grey for DR1 (DESI24-Lyα) and in black dashed for DR2 (DESI Collaboration et al. 2025c), while the eBOSS Lyα full-shape results from Cuceu et al. (2023a) are shown in blue. Our broadband-only AP constraint is given by the horizontal bands in red, while the full-shape result, which includes broadband AP and BAO information, is shown in purple.

surement from the broadband (BB) alone:

$$\textbf{broadband AP}: \quad D_M/D_H(z_{\rm eff}) = 4.489 \pm 0.076. \quad (25)$$

In this case, we marginalize over the scale parameters of the peak component. We will use this constraint when combining our measurement with DESI DR2 Lyα BAO.

In Figure 11, we compare our results with the DR1 and DR2 DESI Lyα BAO measurements (DESI Collaboration et al. 2025d,c), and with eBOSS Lyα full-shape (Cuceu et al. 2023a). Our broadband-only AP constraint is shown as the horizontal shaded area in red. The DESI full-shape constraint is shown in purple, and it contains information from both the broadband through the AP effect and from BAO. Adding the AP information from the broadband improves upon the BAO constraint by an impressive factor of 2.4× (purple versus grey). Furthermore, the full-shape Alcock-Paczynski result from DESI DR1 is 37% tighter than the DESI DR2 BAO constraint (purple versus black), with the two measurements being in excellent agreement. While the BAO information is highly correlated between DR1 and DR2, the broadband AP measurement is independent of both DESI DR1 and DR2 BAO as discussed in Appendix D.

When compared to the results from the eBOSS full-shape analysis (blue in Figure 11), our constraints prefer slightly lower values in both the isotropic and the AP direction. We have computed the correlation coefficient between our results and those from eBOSS using the same method as described in Appendix F of DESI24-Lyα, and in Appendix D below. We find a correlation coefficient of 13%, which is very similar to the 10% correlation DESI24-Lyα found between the DESI DR1 BAO and eBOSS BAO results. Given that the correlation is small, we conclude that the eBOSS and DESI full-shape results are consistent.

We also provide our constraints of the more common distance combinations, D_M/r_d and D_H/r_d. The results from the full-shape $\phi_{\rm f}$ and $\alpha_{\rm p}$ measurements are:

$$\text{DESI DR1 full-shape} \quad \begin{cases} D_H(z_{\rm eff})/r_d & = 8.632 \pm 0.105, \\ D_M(z_{\rm eff})/r_d & = 39.05 \pm 0.52, \\ \rho(D_M/r_d, D_H/r_d) & = -0.244, \end{cases} \quad (26)$$

where ρ gives the correlation coefficient between D_H/r_d and D_M/r_d. For comparison, the DESI DR2 BAO constraints from DESI Collaboration et al. (2025c) are $D_H(z_{\rm eff})/r_d = 8.632 \pm 0.101$ and $D_M(z_{\rm eff})/r_d = 38.99 \pm 0.53$ ($\rho = -0.43$), again showing the excellent agreement between the two. Both our analysis and that of DESI Collaboration et al. (2025c) were performed blinded, and there were no modifications that impact the position of either constraint post-unblinding. Our analysis was unblinded about five months prior to the unblinding of the DESI DR2 BAO measurement. When compared to the distances measured by DESI24-Lyα using DESI DR1 BAO, our uncertainties are 35% smaller for D_H/r_d and 43% smaller for D_M/r_d.

As our broadband Alcock-Paczynski constraint is independent from the DESI DR2 BAO measurement, we also combine the two, and obtain the following distance ratios:

$$\text{DR2 BAO + DR1 AP} \quad \begin{cases} D_H(z_{\rm eff})/r_d & = 8.646 \pm 0.077, \\ D_M(z_{\rm eff})/r_d & = 38.90 \pm 0.38, \\ \rho(D_M/r_d, D_H/r_d) & = -0.016. \end{cases} \quad (27)$$

This corresponds to a 0.9% measurement of D_H/r_d, and a 1% measurement of D_M/r_d.

In the following sections, we will use both the DESI DR1 full-shape constraints from Equation (26), and the combination of the DR1 broadband AP and DR2 BAO constraints from Equation (27). We will refer to constraints that use Equation (26) as *Lyα-FS* , those that use Equation (27) as *Best-Lyα*, and those that only use the Lyα broadband AP constraint as *Lyα-AP*.

6.2. *Other datasets*

For the cosmological analysis, we combine our measurement with results from a number of external probes in order to obtain the tightest constraints. These closely match the probes used in DESI Collaboration et al. (2025b), and they also describe the motivation behind these choices in more detail.

Besides the DESI DR2 Lyα forest BAO measurement, described above, we also use DESI BAO measurements from galaxy clustering at lower redshifts. We use the latest results from DESI DR2 presented in DESI Collaboration et al. (2025b). These consist of D_M/r_d and D_H/r_d measurements in 5 redshift bins that span the range $0.5 < z < 1.5$, and one isotropic BAO constraint at $z = 0.295$.

In order to calibrate the scale of the sound horizon, r_d, BAO data needs to be combined with an external measurement of the baryon density $\Omega_b h^2$. This allows us to constrain the Hubble constant H_0, by breaking its inherent degeneracy with r_d. We use the $\Omega_b h^2$ estimate from Schöneberg (2024), which is based on Big Bang Nucleosynthesis (BBN) and measurements of the primordial deuterium abundance (Cooke et al. 2018), and is given by:

$$\Omega_b h^2 = 0.02218 \pm 0.00055. \tag{28}$$

This is based on the `PRyMordial` code (Burns et al. 2024), and includes marginalisation over uncertainties in nuclear reaction rates.

We use CMB measurements from Planck (Planck Collaboration et al. 2020b,a) and the Atacama Cosmology Telescope (ACT; Madhavacheril et al. 2024; Qu et al. 2024; MacCrann et al. 2024). This includes the Planck temperature (TT), polarization (EE) and cross (TE) power spectra using the `simall`, `Commander` (for $\ell < 30$) and `CamSpec` likelihoods (for $\ell \geq 30$; Efstathiou & Gratton 2021; Rosenberg et al. 2022), and the combination of Planck and ACT DR6 CMB lensing from Madhavacheril et al. (2024). DESI Collaboration et al. (2025b) found that using some of the other CMB likelihoods available in the literature (e.g. Tristram et al. 2024) does not significantly impact dark energy constraints. Therefore, in this work we only focus on combinations with the likelihoods mentioned above, and refer the reader to Appendix A of DESI Collaboration et al. (2025b) and to Garcia-Quintero et al. (2025), for discussions of the impact of different CMB likelihoods on current dark energy constraints.

Finally, for constraints on the equation-of-state of dark energy, we also combine our measurement with results from Type Ia supernovae (SNe) analyses. Following DESI Collaboration et al. (2025a) and DESI Collaboration et al. (2025b), we study combinations with each of the three SNe samples from Pantheon+ (Scolnic et al. 2022; Brout et al. 2022), Union3 (Rubin et al. 2023), and the Dark Energy Survey Year 5 (DESY5; DES Collaboration et al. 2024). The Pantheon+ and Union3 samples consist of spectroscopically-classified SNe drawn from multiple observational programs. They have a significant fraction of objects in common, but use different calibration, modelling and treatment of systematic uncertainties. On the other hand, DESY5 consists of a large sample of photometrically-classified SNe at $z > 0.1$, with uniform calibration based on a single survey, and a small set of SNe at $z < 0.1$, which are drawn from historical sources. We use the publicly available Cobaya likelihoods for the three SNe samples.[11]

6.3. *ΛCDM constraints*

We start with cosmological constraints in the ΛCDM model, focusing on the measurements from Lyα alone (Section 6.1) and in combination with DESI DR2 BAO (Section 6.2). We use the `Cobaya` package (Torrado & Lewis 2019, 2021) and the associated Markov Chain Monte Carlo Metropolis sampler (Lewis & Bridle 2002; Lewis 2013) for cosmological inference. Chains are run until the convergence criterion based on the Gelman-Rubin statistic (Gelman & Rubin 1992) satisfies $R-1 < 0.01$. Theory models are computed using `CAMB` and `Cobaya` likelihoods, and we use the same priors as DESI Collaboration et al. (2025b) for our cosmological analyses.

In ΛCDM, measurements of D_M/r_d and D_H/r_d are sensitive only to two cosmological parameters: Ω_m and the degenerate combination $hr_{\rm d}$, where $h = H_0/(100\ {\rm km\ s^{-1}\ Mpc^{-1}})$. Therefore, we directly sample these two parameters. The constraints from DESI DR1 full-shape are:

$$\left.\begin{aligned} \Omega_{\rm m} &= 0.309^{+0.024}_{-0.028} \\ hr_{\rm d} &= (100.1 \pm 3.4)\ {\rm Mpc} \end{aligned}\right\} \text{DESI DR1 Ly}\alpha\text{-FS}, \tag{29}$$

with a correlation coefficient $\rho = -0.95$. This is consistent with the individual DESI DR2 BAO measurements, as well as the combined measurement presented in DESI Collaboration et al. (2025b).

The constraints from the combination of DESI DR2 Lyα BAO and DR1 Lyα AP, which we denote *Best-Lyα*, are:

$$\left.\begin{aligned} \Omega_{\rm m} &= 0.299^{+0.019}_{-0.023} \\ hr_{\rm d} &= (101.3 \pm 2.8)\ {\rm Mpc} \end{aligned}\right\} \text{DESI Best-Ly}\alpha, \tag{30}$$

[11] https://github.com/CobayaSampler/sn_data

with $\rho = -0.97$. DR1 Lyα full-shape gives an $\sim 8\%$ constraint on the matter fraction, while the Best-Lyα measurement gives a $\sim 7\%$ constraint.

For comparison, the three SNe samples introduced above give Ω_m constraints of $\sim 5\%$ (DESY5), $\sim 7\%$ (Union3), and $\sim 5\%$ (Pantheon+), respectively. This means the Lyα forest measurements from DESI already provide constraints that are competitive with current SNe measurements. However, Lyα prefers smaller Ω_m values when compared to the three SNe samples. Our Best-Lyα measurement is 1.3σ lower than that from Pantheon+, with the difference increasing to 1.7σ for Union3, and 2σ for DESY5. While these differences are not statistically significant, they follow the trend of DESI BAO measurements (DESI Collaboration et al. 2025a,b), which also prefer smaller Ω_m values compared to SNe.

Our measurements are fully consistent with the CMB constraints on Ω_m and $hr_{\rm d}$, and therefore cannot currently provide insight into the small tension reported in DESI Collaboration et al. (2025b) between CMB and DESI BAO constraints. We discuss this later in this section in the context of Figure 12.

We next combine our broadband AP measurement (denoted *Lyα-AP*) with all of the DESI DR2 BAO results (now including galaxy BAO). We will refer to this combination as either "Lyα-AP+BAO" or "Best-Lyα+Galaxy BAO", but the two are equivalent. The result we obtain is:

$$\left.\begin{array}{l}\Omega_{\rm m} = 0.2972 \pm 0.0081, \\ hr_{\rm d} = (101.55 \pm 0.70)\ {\rm Mpc}\end{array}\right\} \begin{array}{c}\text{Best-Ly}\alpha \\ + \text{ Galaxy BAO},\end{array} \quad (31)$$

with $\rho = -0.92$. This corresponds to a 6% tighter constraint on Ω_m, and 4% tighter on $hr_{\rm d}$ when compared to DESI DR2 BAO alone. Given that we are only adding the DR1 Lyα AP constraint on top of the combined DESI DR2 BAO measurement, we consider that even this small improvement is fairly impressive, and shows the potential of Lyα AP measurements with future DESI data releases. However, given there is no significant shift in the result when adding Lyα AP, the tensions between this constraint and SNe Ω_m results remain the same as the ones reported in DESI Collaboration et al. (2025b).

In order to break the degeneracy in $hr_{\rm d}$, we need to calibrate the relative BAO distance measurements. This can be achieved by adding an $\Omega_b h^2$ prior from BBN, which, along with our Ω_m measurement, helps constrain $r_{\rm d}$, and results in a measurement of the Hubble constant H_0. The constraint from DR1 Lyα FS is:

$$H_0 = (67.7 \pm 2.1)\ {\rm km\,s^{-1}\,Mpc^{-1}}\ (\text{Ly}\alpha\text{-FS} + \text{BBN}), \quad (32)$$

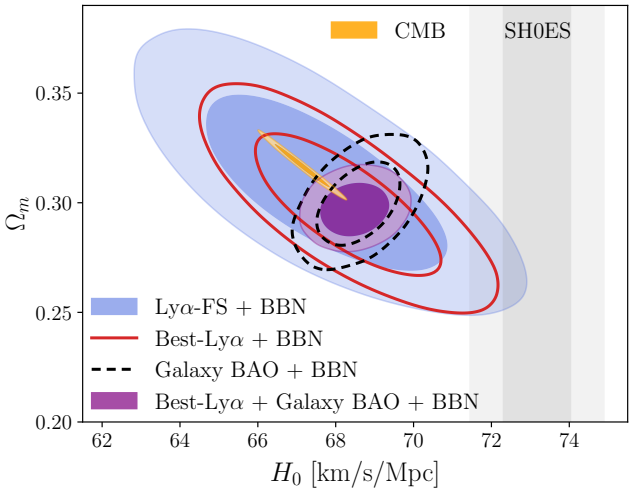

Figure 12. Constraints on the Hubble constant and the matter fraction within the ΛCDM model from DESI measurements combined with an $\Omega_b h^2$ prior from BBN, which helps break the degeneracy in $hr_{\rm d}$. The results from DR1 Lyα full-shape are shown in blue, while the Best-Lyα combination is shown in red, and includes DR1 Lyα AP and DR2 Lyα BAO. The Lyα constraints agree well with both the DESI DR2 galaxy BAO results (black) and the CMB (orange), but are in tension with SH0ES (grey). The combined result in purple includes all DESI DR2 BAO measurements from both galaxies and Lyα and the DR1 Lyα AP constraint.

while the constraint from the Best-Lyα measurement is:

$$H_0 = (68.3 \pm 1.6)\ {\rm km\,s^{-1}\,Mpc^{-1}}\ (\text{Best-Ly}\alpha + \text{BBN}). \quad (33)$$

These represent 3.1% and 2.3% constraints on H_0, respectively.

We compare our measurements with those from the CMB, DESI DR2 galaxy BAO, and SH0ES (Breuval et al. 2024), in Figure 12. The Lyα constraints (blue and red) are in good agreement with both the CMB and DESI galaxy BAO measurements. This plot shows that the 2.3σ tension reported in DESI Collaboration et al. (2025b) between DESI BAO and the CMB is driven by DESI galaxies. Interestingly, the Lyα results align well with the direction of the tension due to the high redshift of the Lyα forest, and because the Lyα isotropic BAO constraint ($\alpha_{\rm p}$) is in excellent agreement with the CMB prediction in ΛCDM. Significant improvements in the Lyα AP constraint (e.g. from future DESI data releases) would shrink the major axis of the Lyα ellipses in Figure 12, and could shed light on this tension. We discuss this further in Section 6.6.

While the Lyα measurements agree with the CMB constraints, they are significantly lower than the cosmic distance ladder results from SH0ES (Breuval et al. 2024), which are shown using the vertical grey bands in

Figure 12. The two Lyα constraints are in 2.4σ (LyαFS) and 2.7σ (Best-Lyα) tension with SH0ES.

Finally, when combining our Lyα AP constraint with all DESI DR2 BAO measurements, we obtain an H_0 value of:

$$H_0 = (68.49 \pm 0.58)\ \mathrm{km\,s^{-1}\,Mpc^{-1}} \quad (\text{DESI DR2 BAO + DR1 Ly}\alpha\text{-AP + BBN}). \tag{34}$$

This is almost identical with the constraint from DESI DR2 BAO alone, which means both the tension with the CMB (2.3σ) and with SH0ES (4.5σ) remain the same as reported in DESI Collaboration et al. (2025b).

Our results show that a weaker Hubble tension of the same sign persists within ΛCDM using only DESI Lyα forest data (2.7σ) or only DESI BAO plus Lyα-AP (4.5σ), with no CMB data. These constraints do still assume standard pre-recombination physics to compute r_d, and they rely on the BBN prior of Equation (28). Therefore, early dark energy or other departures from ΛCDM that affect the sound horizon r_d would change the level of H_0 tension for both DESI results and for the CMB.

6.4. *Dark energy*

Next, we turn our attention to constraints on the nature of dark energy, focusing on measurements of its equation-of-state parameter: $w(z) = P(z)/\rho(z)$, where $P(z)$ is the dark energy pressure, and $\rho(z)$ is its energy density. We use the common Chevallier-Polarksy-Linder parametric model (Chevallier & Polarski 2001; Linder 2003), where the time-evolution of the equation-of-state as a function of the expansion factor $a = (1+z)^{-1}$, is given by:

$$w(a) = w_0 + w_a(1-a), \tag{35}$$

where w_0 is the equation-of-state today and w_a describes its rate of change, with $w \sim w_0 + w_a$ in the distant past. Following DESI Collaboration et al. (2025b), we use wide flat priors with $w_0 \in \mathcal{U}[-3,1]$, $w_a \in \mathcal{U}[-3,2]$, and also impose the condition $w_0 + w_a < 0$ to enforce early matter domination. To quantify the statistical significance of the preference for evolving dark energy, we use the $\Delta\chi^2_{\mathrm{MAP}}$ between the best fit ΛCDM and w_0w_aCDM models for a particular combination of datasets.

As the Lyα measurements are at $z > 2$, combinations with the CMB do not lead to significant constraints on $w_0 - w_a$. This is because both results come from before the epoch of dark energy domination, which leads to similar parameter degeneracies. Instead, we need to combine our results with low-redshift probes, so Lyα can play the role of the high-redshift anchor for the dark energy evolution (similar to the role played by the CMB). Here we consider two low-redshift probes: DESI Galaxy BAO and Type Ia Supernovae. We present $w_0 - w_a$ constraints for various combinations of the four probes (Lyα, Galaxy BAO, Supernovae, CMB), focusing on those where the Lyα measurements play a significant role.

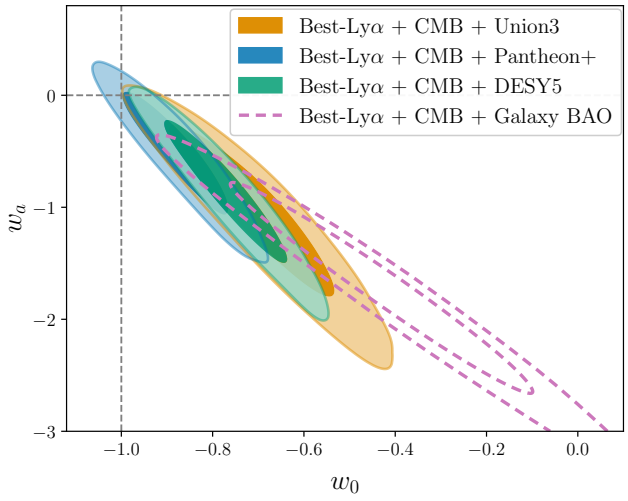

Figure 13. Posterior distributions of the dark energy equation-of-state parameters w_0 and w_a when combining the Best-Lyα measurement from DESI with the CMB and either of the three SNe datasets (filled contours), or with DESI Galaxy BAO constraints (dashed empty contour). The intersection of the gray dashed lines indicates the ΛCDM limit ($w_0 = -1$, $w_a = 0$). The significance of rejection of ΛCDM is not significant for combinations with SNe, ranging between 1.1σ and 2.4σ. On the other hand, the combination with DESI Galaxy BAO shows a more significant 3.1σ deviation from ΛCDM, at the same level found by DESI Collaboration et al. (2025b).

As our best measurement is given by the combination of DR1 Lyα AP and DR2 Lyα BAO (*Best-Lyα*), we will use this result for our dark energy constraints. We have also performed the same analysis using only the DR1 Lyα full-shape constraints and found consistent but less constraining results.

In Table 2, we show $\Delta\chi^2_{\mathrm{MAP}}$ values and the associated frequentist significance of the preference for w_0w_aCDM over ΛCDM given combinations between our measurement and various other data sets (described in Section 6.2). The top half of the table shows results when combining the Best-Lyα constraint with non-DESI probes, while the second half shows results when adding our broadband AP measurement to DESI DR2 BAO measurements (i.e., also including Galaxy BAO). We find that data combinations without DESI Galaxy BAO do not significantly exclude ΛCDM, with the preference for w_0w_aCDM ranging between 0.8σ and 2.1σ when combining Lyα with SNe, and slightly increasing to $1.1\sigma - 2.4\sigma$ when also adding CMB data. The marginal-

Datasets	$\Delta\chi^2_{\rm MAP}$	Significance
Best-Lyα+Pantheon+	-1.8	0.8σ
Best-Lyα+Union3	-5.6	1.9σ
Best-Lyα+DESY5	-6.6	2.1σ
Best-Lyα+CMB+Pantheon+	-2.5	1.1σ
Best-Lyα+CMB+Union3	-6.3	2.0σ
Best-Lyα+CMB+DESY5	-8.4	2.4σ
Lyα-AP+BAO	-4.6	1.6σ
Lyα-AP+BAO+CMB	-12.3	3.1σ
Lyα-AP+BAO+CMB+Pantheon+	-10.8	2.8σ
Lyα-AP+BAO+CMB+Union3	-17.6	3.8σ
Lyα-AP+BAO+CMB+DESY5	-21.2	4.2σ

Table 2. Difference in the effective $\chi^2_{\rm MAP}$ value of the best-fit w_0w_aCDM model and the best-fit ΛCDM model for various combinations of datasets, along with the corresponding significance levels given two extra free parameters. Combinations labeled Lyα-AP+BAO are equivalent to Best-Lyα+Galaxy BAO.

ized posterior distributions of the equation-of-state parameters, w_0 and w_a, from the combination of Lyα, CMB, and SNe are shown as filled contours in Figure 13, and are given by:

$$\left.\begin{aligned} w_0 &= -0.874 \pm 0.078 \\ w_a &= -0.53^{+0.41}_{-0.32} \end{aligned}\right\} \begin{array}{l}\text{Best-Ly}\alpha\text{+CMB} \\ \text{+Pantheon+,}\end{array} \tag{36}$$

when combined with Pantheon+,

$$\left.\begin{aligned} w_0 &= -0.70^{+0.11}_{-0.13} \\ w_a &= -1.06^{+0.58}_{-0.43} \end{aligned}\right\} \begin{array}{l}\text{Best-Ly}\alpha\text{+CMB} \\ \text{+Union3,}\end{array} \tag{37}$$

with Union3, and

$$\left.\begin{aligned} w_0 &= -0.764^{+0.084}_{-0.094} \\ w_a &= -0.88^{+0.47}_{-0.37} \end{aligned}\right\} \begin{array}{l}\text{Best-Ly}\alpha\text{+CMB} \\ \text{+DESY5,}\end{array} \tag{38}$$

for the combination with DESY5.

On the other hand, combinations that include Galaxy BAO (second half of Table 2) show more significant deviations from ΛCDM, in line with the results presented in DESI Collaboration et al. (2025b). In fact the results presented in the second half of Table 2 are almost identical to those based on DR2 BAO alone (DESI Collaboration et al. 2025b). This is not surprising given that DESI DR2 BAO includes the BAO measurement from Lyα, which provides a similar constraining power on AP as our DR1 broadband AP measurement here (see Section 6.1).

The constraint from the combination of Lyα, Galaxy BAO, and CMB is also shown in Figure 13 (empty contour). This constraint is not as tight as the ones using SNe, but shows a more pronounced deviation from ΛCDM, at 3.1σ significance. The marginalized posterior for this combination is:

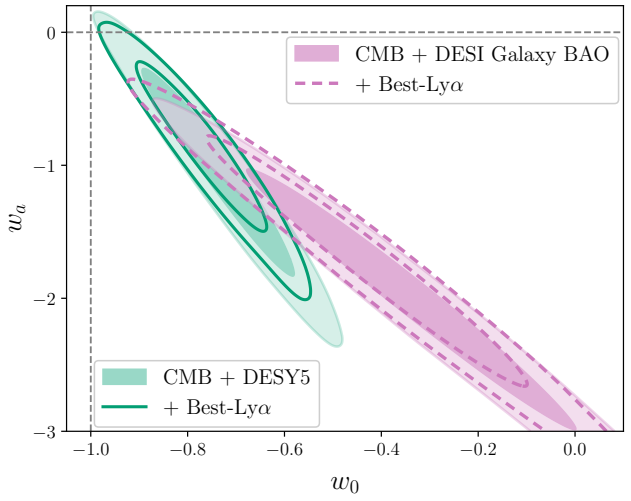

Figure 14. Similar to Figure 13, but showing the impact of adding the DESI Best-Lyα constraint to the combination of CMB and DESY5 SNe (green filled versus empty contours), and to the combination of CMB and DESI Galaxy BAO measurements (pink filled versus empty contours). In the first case, adding Lyα leads to tighter constraints that are closer to ΛCDM. For the second combination, adding Lyα does not shrink the constraints, but does shift the posterior towards ΛCDM, slightly reducing the tension from 3.3σ to 3.1σ.

$$\left.\begin{aligned} w_0 &= -0.43 \pm 0.21 \\ w_a &= -1.72 \pm 0.58 \end{aligned}\right\} \begin{array}{l}\text{Best-Ly}\alpha\text{+CMB} \\ \text{+Galaxy BAO,}\end{array} \tag{39}$$

which is very similar to the constraint from DESI DR2 BAO and CMB presented in DESI Collaboration et al. (2025b).

All the combinations shown in Figure 13 prefer a similar region in the $w_0 - w_a$ space, with $w_0 > -1$ and $w_a < 0$, and have similar degeneracy directions, although not identical and not all pointed exactly at ΛCDM. The regions where the empty contour intersects each of the three filled contours very roughly correspond to the joint posteriors obtained when combining Lyα, Galaxy BAO, CMB, and SNe. These combinations give results very similar to those presented in DESI Collaboration et al. (2025b), with the same significance levels of the preference for w_0w_aCDM over ΛCDM, ranging from 2.8σ to 4.2σ (Table 2).

To better understand the role played by Lyα, in Figure 14 we show results from two of the combinations presented above, with and without the Best-Lyα constraint. First, the green contours show the constraints

from the combination of CMB and DESY5 SNe without Lyα (filled) and with Lyα (empty). In this case, adding Lyα leads to tighter constraints, corresponding to $\sim 22\%$ larger figure-of-merit, and an increase in the tension with ΛCDM from 2.2σ to 2.4σ. On the other hand, adding Lyα to the combination of CMB and DESI Galaxies (pink filled versus empty contours) does not improve the constraints, but it does shift the posterior towards ΛCDM, slightly reducing the tension from 3.3σ to 3.1σ. Figure 14 also illustrates that shifts along the degeneracy direction for either of these combinations cannot fully reconcile the combined DESI, CMB, and DESY5 SNe constraint with ΛCDM.

6.5. *Other constraints*

In this section, we present several other cosmological constraints that are improved by the addition of our Lyα AP measurement. These include neutrino mass and curvature constraints.

Our AP measurement is not directly sensitive to massive neutrinos. However, it provides a geometrical measurement on Ω_m, which can help improve CMB constraints on the neutrino mass. This is because the neutrino constraints from the CMB are highly correlated with several other parameters, including Ω_m (Planck Collaboration et al. 2020a; Loverde & Weiner 2024). This is exactly the same way that BAO measurements contribute to neutrino mass constraints (DESI Collaboration et al. 2025a,b).

We sample the sum of the neutrino masses, $\sum m_\nu$, assuming three degenerate mass eigenstates, and closely follow the analysis setup described in DESI Collaboration et al. (2025b). Within ΛCDM, combining our Best-Lyα constraint with the CMB produces the following upper bound:

$$\sum m_\nu < 0.156 \text{ eV } (95\%, \text{ Best-Ly}\alpha\text{+CMB}). \quad (40)$$

When adding our Lyα-AP measurement to the combination of CMB and DESI DR2 BAO, we obtain:

$$\sum m_\nu < 0.0638 \text{ eV } (95\%, \text{ Ly}\alpha\text{-AP+BAO+CMB}), \quad (41)$$

which is similar but slightly tighter than the constraint presented in DESI Collaboration et al. (2025b). When extending to w_0w_aCDM, our Lyα-AP measurement does not change the constraints from BAO+CMB. We note that these neutrino mass constraints do depend on the likelihood used for the CMB. Here we only use the `CamSpec` likelihood, but DESI Collaboration et al. (2025b) and Elbers et al. (2025) present results using all three likelihoods currently available.

When allowing curvature to vary, Alcock-Paczynski measurements constrain a degenerate combination of Ω_m and the curvature fraction Ω_k (see Cuceu et al. 2023a). Combining our constraint with other AP and BAO measurements at lower redshifts helps break this degeneracy, leading to constraints on curvature. The result from the combination of Lyα-AP and DESI DR2 BAO is given by:

$$\Omega_k = 0.017 \pm 0.036 \text{ (Ly}\alpha\text{-AP+BAO)} \quad (42)$$

This is consistent with a flat Universe, and represents a 14% tighter constraint compared to that from BAO alone. However, when combined with BAO+CMB, the Lyα constraint does not change the result presented in DESI Collaboration et al. (2025b).

6.6. *Cosmology Discussion*

We conclude this part with a discussion of our cosmological constraints, with an eye on the value added by the improved Lyα forest Alcock-Paczynski measurement presented here, and the potential value of future such analyses. Within ΛCDM, our measurement is consistent with both the CMB and DESI Galaxy BAO results, which means it cannot currently provide insight into the tension observed between the two. However, the disagreement is exactly along the direction constrained by the AP effect at high redshift, which means Lyα full-shape analyses with future DESI data releases could shed light on this tension.

Within ΛCDM, adding our DR1 Lyα broadband AP to DESI DR2 BAO results in a 6% reduction in the Ω_m uncertainty. The AP measurement also helps tighten upper bounds on the neutrino mass and DESI constraints on curvature. However, given that here we are combining the DR1 Lyα AP constraint with DR2 BAO, the improvements are fairly small. We expect the improvements to be more significant with the next iteration of this analysis using DESI DR2.

When it comes to constraints on the dark energy equation-of-state parameters, we found that our Lyα AP constraint is particularly useful in combination with Supernovae alone and when also including the CMB. These combinations produce fairly tight constraints on w_0 and w_a, rivaling those from the joint BAO and CMB analysis. On the other hand, adding our Lyα AP constraint to combinations that contain DESI DR2 BAO does not lead to any significant changes in the posterior. This is in part because DESI DR2 BAO contains a Lyα BAO measurement at the same redshift as our DR1 broadband AP constraint, and with a similar uncertainty (1.7% for DR1 broadband AP, and 2.1% for DR2 BAO). However, even when we compared with the results from CMB and Galaxy BAO alone (i.e., without Lyα), we found that adding Lyα does not result in

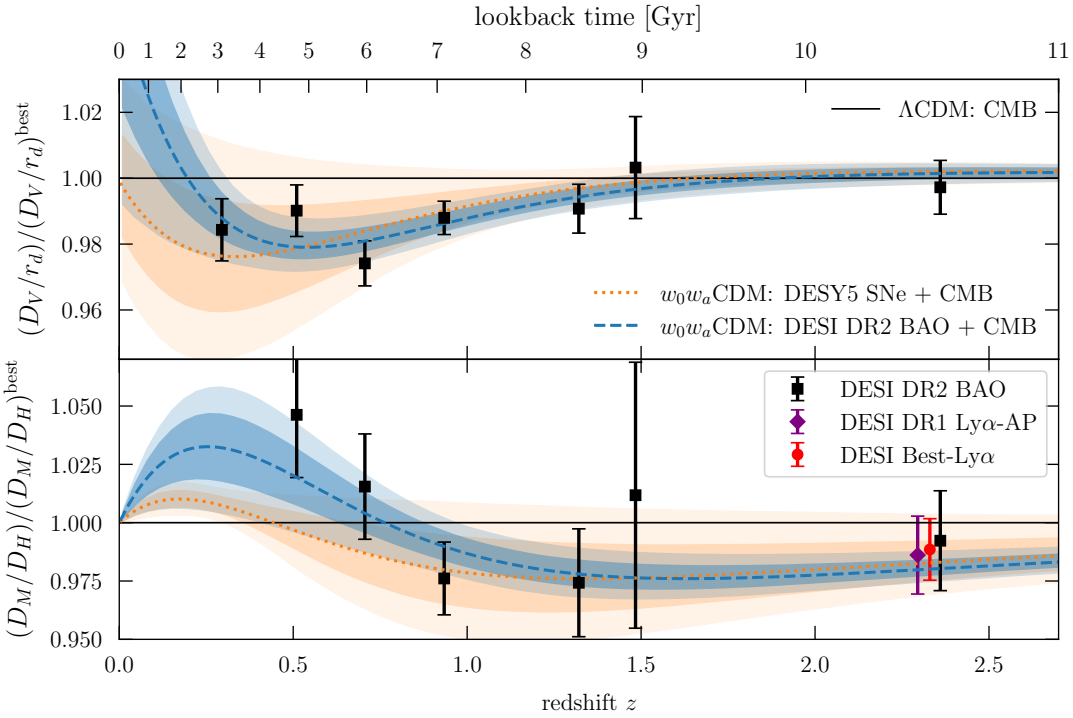

Figure 15. DESI constraints on the isotropic BAO distance D_V/r_d (top), and on the Alcock-Paczynski parameter D_M/D_H (bottom), relative to the CMB best-fit ΛCDM model (black line). The DESI DR2 BAO measurements are shown in black, our DR1 Lyα broadband AP constraint is shown in purple, and the Best-Lyα constraint on AP is shown in red. For visualization purposes, the Lyα constraints are shown at slightly different redshifts. The constraints within the w_0w_aCDM model from the combination of DESI BAO and CMB are shown in blue, with the dashed line indicating the best-fit model. We also show constraints from the combination of DESY5 SNe and CMB within w_0w_aCDM in orange (best-fit represented by the dotted line). This shows that within w_0w_aCDM, current data prefer AP values that are $\sim 2\%$ lower than the CMB ΛCDM predictions in the range $1 < z < 2.5$. Our Best-Lyα combination gives a 1.3% constraint on AP at $z = 2.33$, but this falls right in between the ΛCDM and w_0w_aCDM best-fit models. Cuceu et al. (2021) forecasted Lyα broadband AP constraints of $0.3\% - 0.9\%$ with a 5-year DESI survey, meaning Lyα could play an important role in future dark energy constraints from DESI.

tighter constraints, but does produce a small shift towards ΛCDM.

To gain a better understanding of these results, we project the w_0w_aCDM posterior distributions into the space of our measurements in Figure 15. The figure shows the DESI constraints on the isotropic BAO distance D_V/r_d in the top panel and on the Alcock-Paczynski parameter in the bottom panel, both normalized relative to the best-fit ΛCDM model from the CMB. The actual measurements are represented by points with error bars, while the inferred posteriors within w_0w_aCDM are shown using the colored bands, with the dashed and dotted lines representing the best-fit models. The blue bands show results from DESI DR2 BAO and the CMB, while the orange bands show results from DESY5 SNe and the CMB.

The top panel of Figure 15 shows that DESI Galaxies at $z < 1$ are the primary driver of the tension with the CMB ΛCDM model and of the preference for w_0w_aCDM. This is because the uncertainties of the BAO measurements in this region are similar to the uncertainty in the w_0w_aCDM posterior (black points versus blue bands), and are significantly tighter than the uncertainty in the SNe+CMB posterior. On the other hand, at $z > 1.5$ the w_0w_aCDM posteriors agree well with each other and with the CMB ΛCDM model, as this region is constrained very tightly by the CMB acoustic scale. This means future improvements in the isotropic BAO measurements at high redshift will not be very helpful to dark energy constraints that include CMB information.

When it comes to the Alcock-Paczynski constraints, shown in the lower panel of Figure 15, the conclusions are quite different. At $z > 1$, both w_0w_aCDM posteriors prefer AP values that are $\sim 2\%$ smaller than the CMB prediction in ΛCDM. In this region, our measurement from DESI DR1 Lyα full-shape (purple) gives the tightest constraint to date (1.7%; shown in purple), and its combination with DR2 Lyα BAO produces a 1.3% constraint, which is shown in red in Figure 15. The inferred constraints from SNe+CMB are similar in this

region ($\sim 0.9\%$), while the ones from BAO+CMB are significantly tighter ($\sim 0.3\%$). This explains why adding Lyα to the former helped improve the constraints, while adding it to the latter did not. Furthermore, the Lyα constraints are currently consistent with both the CMB ΛCDM model and the best-fit w_0w_aCDM models, meaning they do not strongly pull in either direction.

The fact that AP uncertainties are significantly larger than the BAO+CMB inferred constraints in w_0w_aCDM (blue bands) means they do not currently play as significant a role as low-redshift isotropic BAO constraints. However, these constraints are expected to improve with full-shape analyses of DESI DR2 on both the Lyα and Galaxy clustering side. Figure 15 shows that future full-shape analyses with emission line galaxies in the range $1 < z < 1.5$, and with Lyα at $z > 2$, could help improve upon current constraints. The Lyα forest is particularly powerful when it comes to AP constraints, as we have seen here. With DESI DR1, it led to a 2.4$\times$ tighter constraint when compared to BAO. Cuceu et al. (2021) forecasted Lyα broadband AP uncertainties in the range $0.3\% - 0.9\%$ for the full 5-year DESI survey, with the main contributing factor being the minimum scale used in the analysis. Given that the difference between ΛCDM and w_0w_aCDM is $\sim 2\%$ in this region, future DESI Lyα forest analyses could play an important role in deciphering the nature of dark energy.

More generally, the DESI DR2 analysis suggests that the dark energy density $\rho_{\rm DE}(z)$ peaks at $z \sim 0.3 - 1$ and declines significantly towards higher redshift. The w_0w_aCDM model provides a compact parameterization that allows this behavior, but if the basic finding of evolving dark energy does hold up, there is no reason to expect that w_0w_aCDM provides a full description. The correct physical model might even involve exchanges of energy between the dark energy, matter, and radiation components, rather than dark energy evolution on its own. In this context, measurements of $D_H(z)$, which directly probe the energy density at $z > 2$ are especially valuable, whereas $D_M(z)$ constrains an integral over $H(z)$ and has limited room to depart from the trend measured by galaxies at lower redshift and θ_* at $z \sim 1100$ (see, e.g., Eqs. 3 and 6 of DESI Collaboration et al. 2025b). Lyα forest BAO is already quite good at measuring D_H because the strong redshift-space distortions amplify the BAO signal in the line-of-sight direction. However, a full-shape Lyα AP measurement in combination with BAO gives greater leverage on the energy density, improving D_H/r_d constraints by a factor of ~ 1.5 relative to BAO alone (Cuceu et al. 2021). Even moderate gains in precision may play a critical role in distinguishing physical models of evolving dark energy.

Finally, while the discussion here focuses on a dynamic dark energy equation-of-state as the way to reconcile BAO and CMB results, several other potential explanations have recently appeared in the literature (e.g., Chen & Zaldarriaga 2025; Sailer et al. 2025). If alternative explanations lead to different predictions for the behaviour of the high-redshift AP parameter compared to w_0w_aCDM, the Lyα forest would also help test these alternatives. A concrete example is non-zero curvature, as proposed by Chen & Zaldarriaga (2025), for which the best-fit BAO+CMB constraints predict AP values roughly consistent with ΛCDM at $z \sim 2$ (see their Figure 6). Therefore, future DESI Lyα forest full-shape analyses could also shed light on whether curvature or a dynamic dark energy equation-of-state offer better explanations for the current tension.

7. SUMMARY

We have presented a full-shape analysis of Lyα forest 3D correlation functions measured from the first data release (DR1) of DESI. We used both the Lyα autocorrelation and its cross-correlation with quasars, which were measured by DESI Collaboration et al. (2025d). The first goal of this analysis was to measure the Alcock-Paczynski (AP) effect from a broad range of scales in order to improve upon the expansion rate constraints based on BAO measurements from the same data set (DESI Collaboration et al. 2025d). The second goal was to use RSD in the Lyα-QSO cross-correlation to measure the product of the growth rate and the amplitude of fluctuations in spheres of 8 h^{-1}Mpc, $f\sigma_8$.

Our analysis, introduced in Section 2, largely follows the methods previously used for the Lyα full-shape analysis of eBOSS (Cuceu et al. 2023b,a), and for the first DESI Lyα BAO measurement (DESI Collaboration et al. 2025d). The DESI Lyα forest analysis leading to the measurement of correlation functions was presented in Ramírez-Pérez et al. (2024); Gordon et al. (2023); DESI Collaboration et al. (2025d). The model used to fit the Lyα correlations and extract full-shape information is described in Sections 2.2 and 2.3. Our analysis was first done blinded as described in Section 2.4.

Our main results are presented in Section 3, and include a 1.6% constraint on the Alcock-Paczynski effect and a $\sim$ 16% constraint on $f\sigma_8$, at an effective redshift $z_{\rm eff} = 2.33$. The AP measurement represents a factor of 2.4$\times$ improvement compared to the BAO analysis of the same dataset, and is almost 40% better than the BAO constraint from the second DESI data release (DR2; DESI Collaboration et al. 2025c). The $f\sigma_8$ measurement is the first of its kind from the Lyα-QSO cross-

correlation, and to our knowledge, the first direct growth rate constraint at $z > 2$.

In Section 4, we performed a large set of tests to validate our measurements, including tests with two different sets of mocks, data splits, and analysis and modelling variations. The tests on mocks, presented in Section 4.1, show that we can recover unbiased AP constraints in the presence of all the major Lyα forest contaminants. Furthermore, we found no evidence of significant systematic shifts in AP with the tests on data, as detailed in Sections 4.2 to 4.4.

On the other hand, for $f\sigma_8$, one of the two sets of mocks shows a $\sim 0.5\sigma$ bias, which we were not able to track down to any particular systematic effect. Therefore, we decided not to use the $f\sigma_8$ constraint for cosmological inference as discussed in Section 5. Our final growth rate measurement, including a systematic uncertainty based on the shift observed in mocks, is: $f\sigma_8(z_{\rm eff}) = 0.37\,^{+0.055}_{-0.065}\,({\rm stat})\,\pm 0.033\,({\rm sys})$.

We present the cosmological interpretation of our AP measurement alone, and in combination with other data sets, in Section 6. The Lyα full-shape analysis from DESI DR1 measures the ratios $D_H(z_{\rm eff})/r_d = 8.632 \pm 0.105$ and $D_M(z_{\rm eff})/r_d = 39.05 \pm 0.52$, where D_H is the Hubble distance, D_M is the transverse comoving distance, and r_d is the sound horizon at the drag epoch. We also combined our DR1 broadband AP constraint with the Lyα BAO measurement from DESI DR2 (DESI Collaboration et al. 2025c), and obtained the ratios $D_H(z_{\rm eff})/r_d = 8.646 \pm 0.077$ and $D_M(z_{\rm eff})/r_d = 38.90 \pm 0.38$.

Within ΛCDM our measurements are consistent with both CMB and galaxy clustering BAO constraints, as discussed in Section 6.3. Using a BBN prior on the baryon density, we measure the Hubble constant to be $H_0 = 68.3 \pm 1.6$ km s^{-1} Mpc^{-1} in ΛCDM. When allowing for evolving dark energy equation-of-state parameters (Section 6.4), we find that Lyα helps improve constraints from the combination of Supernovae and CMB, and slightly increases the tension with ΛCDM. On the other hand, for combinations that include CMB and Galaxy BAO results, adding Lyα leads to a small shift towards ΛCDM, but does not improve constraints.

Finally, in Section 6.6 we discuss our cosmology results and the role future DESI Lyα forest analyses are expected to play when it comes to studying the nature of dark energy. We show that further incremental improvements in isotropic BAO constraints at $z > 1.5$ will not help because this quantity is constrained very precisely by the CMB within w_0w_aCDM. On the other hand, best-fit models using ΛCDM or w_0w_aCDM currently give predictions of the Alcock-Paczynski effect that are $\sim 2\%$ different at $z > 1$. Lyα full-shape analyses using future DESI data releases are expected to produce sub-percent constraints of the AP effect (Cuceu et al. 2021), and could therefore play a key role in deciphering the nature of dark energy.

8. DATA AVAILABILITY

The data used in this analysis is public as part of DESI Data Release 1 (details in https://data.desi.lbl.gov/doc/releases/). The data points corresponding to the figures from this paper will be available in a Zenodo repository.

9. ACKNOWLEDGEMENTS

AC acknowledges support provided by NASA through the NASA Hubble Fellowship grant HST-HF2-51526.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. HKHA acknowledges support from the French National Research Agency (ANR) under grant ANR-22-CE31-0009 (HZ-3D-MAP project) and grant ANR-22-CE92-0037 (DESI-Lya project). CG, CRP, and AFR acknowledge support from the European Union's Horizon Europe research and innovation programme (COSMO-LYA, grant agreement 101044612), as well as support by the Spanish Ministry of Science, Innovation and Universities (MICIU) under grants SEV-2016-0588, PGC-2018-094773-B-C31, RYC-2018-025210, PGC2021-123012NB-C41, PID2024-159420NB-C41 and CEX2024001442-S, some of which include ERDF funds from the European Union. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya.

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE–AC02–05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science and Technology of Mexico (CONAHCYT); the Ministry of Science, Innovation and Universities of Spain (MICIU/AEI/10.13039/501100011033), and by the DESI Member Institutions: https://www.

desi.lbl.gov/collaborating-institutions. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U. S. National Science Foundation, the U. S. Department of Energy, or any of the listed funding agencies.

The authors are honored to be permitted to conduct scientific research on I'oligam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation.

REFERENCES

Alam, S., Ata, M., Bailey, S., et al. 2017, MNRAS, 470, 2617, doi: 10.1093/mnras/stx721

Alam, S., Aubert, M., Avila, S., et al. 2021, *Phys. Rev.* **D** , 103, 083533, doi: 10.1103/PhysRevD.103.083533

Alcock, C., & Paczynski, B. 1979, Nature, 281, 358, doi: 10.1038/281358a0

Arinyo-i-Prats, A., Miralda-Escudé, J., Viel, M., & Cen, R. 2015, JCAP, 2015, 017, doi: 10.1088/1475-7516/2015/12/017

Baldwin, J. A. 1977, ApJ, 214, 679, doi: 10.1086/155294

Bault, A., Kirkby, D., Guy, J., et al. 2025, JCAP, 2025, 130, doi: 10.1088/1475-7516/2025/01/130

Bautista, J. E., Busca, N. G., Guy, J., et al. 2017, A&A, 603, A12, doi: 10.1051/0004-6361/201730533

Bautista, J. E., Paviot, R., Vargas Magaña, M., et al. 2021, MNRAS, 500, 736, doi: 10.1093/mnras/staa2800

Blake, C., Brough, S., Colless, M., et al. 2011, MNRAS, 415, 2876, doi: 10.1111/j.1365-2966.2011.18903.x

Breuval, L., Riess, A. G., Casertano, S., et al. 2024, ApJ, 973, 30, doi: 10.3847/1538-4357/ad630e

Brodzeller, A., Dawson, K., Bailey, S., et al. 2023, AJ, 166, 66, doi: 10.3847/1538-3881/ace35d

Brodzeller, A., Wolfson, M., Santos, D. M., et al. 2025, arXiv e-prints, arXiv:2503.14740, doi: 10.48550/arXiv.2503.14740

Brout, D., Scolnic, D., Popovic, B., et al. 2022, ApJ, 938, 110, doi: 10.3847/1538-4357/ac8e04

Burns, A.-K., Tait, T. M. P., & Valli, M. 2024, European Physical Journal C, 84, 86, doi: 10.1140/epjc/s10052-024-12442-0

Busca, N., Rich, J., Bautista, J., et al. 2025, JCAP, 2025, 020, doi: 10.1088/1475-7516/2025/09/020

Busca, N. G., Delubac, T., Rich, J., et al. 2013, A&A, 552, A96, doi: 10.1051/0004-6361/201220724

Casas, L., Herrera-Alcantar, H. K., Chaves-Montero, J., et al. 2025, arXiv e-prints, arXiv:2503.14741, doi: 10.48550/arXiv.2503.14741

Chabanier, S., Ravoux, C., Latrille, L., et al. 2024, MNRAS, 534, 2674, doi: 10.1093/mnras/stae2255

Chaves-Montero, J., Cabayol-Garcia, L., Lokken, M., et al. 2025, A&A, 694, A187, doi: 10.1051/0004-6361/202452039

Chen, S.-F., Vlah, Z., & White, M. 2021, JCAP, 2021, 053, doi: 10.1088/1475-7516/2021/05/053

Chen, S.-F., & Zaldarriaga, M. 2025. https://arxiv.org/abs/2505.00659

Chevallier, M., & Polarski, D. 2001, International Journal of Modern Physics D, 10, 213, doi: 10.1142/S0218271801000822

Cooke, R. J., Pettini, M., & Steidel, C. C. 2018, ApJ, 855, 102, doi: 10.3847/1538-4357/aaab53

Cuceu, A., Font-Ribera, A., & Joachimi, B. 2020, JCAP, 2020, 035, doi: 10.1088/1475-7516/2020/07/035

Cuceu, A., Font-Ribera, A., Joachimi, B., & Nadathur, S. 2021, MNRAS, 506, 5439, doi: 10.1093/mnras/stab1999

Cuceu, A., Font-Ribera, A., Nadathur, S., Joachimi, B., & Martini, P. 2023a, *Phys. Rev.* **L** , 130, 191003, doi: 10.1103/PhysRevLett.130.191003

Cuceu, A., Font-Ribera, A., Martini, P., et al. 2023b, MNRAS, 523, 3773, doi: 10.1093/mnras/stad1546

Cuceu, A., Herrera-Alcantar, H. K., Gordon, C., et al. 2025, JCAP, 2025, 148, doi: 10.1088/1475-7516/2025/01/148

de Belsunce, R., Chen, S.-F., Ivanov, M. M., et al. 2025, *Phys. Rev.* **D** , 111, 063524, doi: 10.1103/PhysRevD.111.063524

de Sainte Agathe, V., Balland, C., du Mas des Bourboux, H., et al. 2019, A&A, 629, A85, doi: 10.1051/0004-6361/201935638

Dembinski, H., & et al., P. O. 2020, doi: 10.5281/zenodo.3949207

DES Collaboration, Abbott, T. M. C., Adamow, M., et al. 2024, arXiv e-prints, arXiv:2402.10696, doi: 10.48550/arXiv.2402.10696

DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv e-prints, arXiv:1611.00037, doi: 10.48550/arXiv.1611.00037

DESI Collaboration, Abareshi, B., Aguilar, J., et al. 2022, AJ, 164, 207, doi: 10.3847/1538-3881/ac882b

DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024a, arXiv e-prints, arXiv:2411.12021, doi: 10.48550/arXiv.2411.12021
—. 2024b, AJ, 167, 62, doi: 10.3847/1538-3881/ad0b08
—. 2025a, JCAP, 2025, 021, doi: 10.1088/1475-7516/2025/02/021
DESI Collaboration, Abdul-Karim, M., Aguilar, J., et al. 2025b, arXiv e-prints, arXiv:2503.14738, doi: 10.48550/arXiv.2503.14738
—. 2025c, arXiv e-prints, arXiv:2503.14739, doi: 10.48550/arXiv.2503.14739
DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2025d, JCAP, 2025, 124, doi: 10.1088/1475-7516/2025/01/124
DESI Collaboration, Abdul-Karim, M., Adame, A. G., et al. 2025e, arXiv e-prints, arXiv:2503.14745, doi: 10.48550/arXiv.2503.14745
du Mas des Bourboux, H., Le Goff, J.-M., Blomqvist, M., et al. 2017, A&A, 608, A130, doi: 10.1051/0004-6361/201731731
du Mas des Bourboux, H., Rich, J., Font-Ribera, A., et al. 2020, ApJ, 901, 153, doi: 10.3847/1538-4357/abb085
Efstathiou, G., & Gratton, S. 2021, The Open Journal of Astrophysics, 4, 8, doi: 10.21105/astro.1910.00483
Eisenstein, D. J., Seo, H.-J., & White, M. 2007, ApJ, 664, 660, doi: 10.1086/518755
Elbers, W., Aviles, A., Noriega, H. E., et al. 2025, arXiv e-prints, arXiv:2503.14744, doi: 10.48550/arXiv.2503.14744
Etourneau, T., Le Goff, J.-M., Rich, J., et al. 2023, arXiv e-prints, arXiv:2310.18996, doi: 10.48550/arXiv.2310.18996
Farr, J., Font-Ribera, A., du Mas des Bourboux, H., et al. 2020, JCAP, 2020, 068, doi: 10.1088/1475-7516/2020/03/068
Font-Ribera, A., & Miralda-Escudé, J. 2012, JCAP, 2012, 028, doi: 10.1088/1475-7516/2012/07/028
Font-Ribera, A., Arnau, E., Miralda-Escudé, J., et al. 2013, JCAP, 2013, 018, doi: 10.1088/1475-7516/2013/05/018
Garcia-Quintero, C., Noriega, H. E., de Mattia, A., et al. 2025, arXiv e-prints, arXiv:2504.18464. https://arxiv.org/abs/2504.18464
Gelman, A., & Rubin, D. B. 1992, Statistical Science, 7, 457, doi: 10.1214/ss/1177011136
Gerardi, F., Cuceu, A., Font-Ribera, A., Joachimi, B., & Lemos, P. 2023, MNRAS, 518, 2567, doi: 10.1093/mnras/stac3257
Gil-Marín, H., Bautista, J. E., Paviot, R., et al. 2020, MNRAS, 498, 2492, doi: 10.1093/mnras/staa2455
Givans, J. J., Font-Ribera, A., Slosar, A., et al. 2022, JCAP, 2022, 070, doi: 10.1088/1475-7516/2022/09/070
Gontcho A Gontcho, S., Miralda-Escudé, J., & Busca, N. G. 2014, Mon. Not. Roy. Astron. Soc., 442, 187, doi: 10.1093/mnras/stu860
Gordon, C., Cuceu, A., Chaves-Montero, J., et al. 2023, JCAP, 2023, 045, doi: 10.1088/1475-7516/2023/11/045
Gordon, C., Cuceu, A., Font-Ribera, A., et al. 2025, arXiv e-prints, arXiv:2505.08789, doi: 10.48550/arXiv.2505.08789
Guy, J., Bailey, S., Kremin, A., et al. 2023, AJ, 165, 144, doi: 10.3847/1538-3881/acb212
Guy, J., Gontcho, S. G. A., Armengaud, E., et al. 2025, JCAP, 2025, 140, doi: 10.1088/1475-7516/2025/01/140
Hadzhiyska, B., et al. 2023, Mon. Not. Roy. Astron. Soc., 524, 1008, doi: 10.1093/mnras/stad1920
Hamilton, A. J. S. 1998, in Astrophysics and Space Science Library, Vol. 231, The Evolving Universe, ed. D. Hamilton, 185, doi: 10.1007/978-94-011-4960-0_17
Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2015a, MNRAS, 450, L61, doi: 10.1093/mnrasl/slv047
—. 2015b, MNRAS, 453, 4384, doi: 10.1093/mnras/stv1911
Herrera-Alcantar, H. K., Muñoz-Gutiérrez, A., Tan, T., et al. 2023, arXiv e-prints, arXiv:2401.00303, doi: 10.48550/arXiv.2401.00303
Hui, L., Stebbins, A., & Burles, S. 1999, ApJL, 511, L5, doi: 10.1086/311826
Ivanov, M. M. 2024, *Phys. Rev.* **D** , 109, 023507, doi: 10.1103/PhysRevD.109.023507
James, F., & Roos, M. 1975, Comput. Phys. Commun., 10, 343, doi: 10.1016/0010-4655(75)90039-9
Kaiser, N. 1987, MNRAS, 227, 1, doi: 10.1093/mnras/227.1.1
Kirkby, D., Margala, D., Slosar, A., et al. 2013, JCAP, 2013, 024, doi: 10.1088/1475-7516/2013/03/024
Lewis, A. 2013, *Phys. Rev.* **D** , 87, 103529, doi: 10.1103/PhysRevD.87.103529
Lewis, A., & Bridle, S. 2002, *Phys. Rev.* **D** , 66, 103511, doi: 10.1103/PhysRevD.66.103511
Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473, doi: 10.1086/309179
Linder, E. V. 2003, *Phys. Rev.* **L** , 90, 091301, doi: 10.1103/PhysRevLett.90.091301
Loverde, M., & Weiner, Z. J. 2024, JCAP, 2024, 048, doi: 10.1088/1475-7516/2024/12/048
MacCrann, N., Sherwin, B. D., Qu, F. J., et al. 2024, ApJ, 966, 138, doi: 10.3847/1538-4357/ad2610
Madhavacheril, M. S., Qu, F. J., Sherwin, B. D., et al. 2024, ApJ, 962, 113, doi: 10.3847/1538-4357/acff5f
McDonald, P. 2003, ApJ, 585, 34, doi: 10.1086/345945

McDonald, P., & Miralda-Escudé, J. 1999, ApJ, 518, 24, doi: 10.1086/307264

McQuinn, M. 2016, ARA&A, 54, 313, doi: 10.1146/annurev-astro-082214-122355

Miller, T. N., Doel, P., Gutierrez, G., et al. 2024, AJ, 168, 95, doi: 10.3847/1538-3881/ad45fe

Morrison, S., Som, D., Pieri, M. M., Pérez-Ràfols, I., & Blomqvist, M. 2024, MNRAS, 532, 32, doi: 10.1093/mnras/stae1418

P. Martini et al. 2025, in preparation

Percival, W. J., & White, M. 2009, MNRAS, 393, 297, doi: 10.1111/j.1365-2966.2008.14211.x

Philcox, O. H. E., Sherwin, B. D., Farren, G. S., & Baxter, E. J. 2021, *Phys. Rev.* **D** , 103, 023538, doi: 10.1103/PhysRevD.103.023538

Pieri, M. M., Mortonson, M. J., Frank, S., et al. 2014, MNRAS, 441, 1718, doi: 10.1093/mnras/stu577

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, A&A, 641, A6, doi: 10.1051/0004-6361/201833910

—. 2020b, A&A, 641, A1, doi: 10.1051/0004-6361/201833880

Pontzen, A. 2014, Phys. Rev. D, 89, 083010, doi: 10.1103/PhysRevD.89.083010

Poppett, C., Tyas, L., Aguilar, J., et al. 2024, AJ, 168, 245, doi: 10.3847/1538-3881/ad76a4

Qu, F. J., Sherwin, B. D., Madhavacheril, M. S., et al. 2024, ApJ, 962, 112, doi: 10.3847/1538-4357/acfe06

Ramírez-Pérez, C., Sanchez, J., Alonso, D., & Font-Ribera, A. 2022, JCAP, 2022, 002, doi: 10.1088/1475-7516/2022/05/002

Ramírez-Pérez, C., Pérez-Ràfols, I., Font-Ribera, A., et al. 2024, MNRAS, 528, 6666, doi: 10.1093/mnras/stad3781

Reid, B. A., Samushia, L., White, M., et al. 2012, MNRAS, 426, 2719, doi: 10.1111/j.1365-2966.2012.21779.x

Riess, A. G., Yuan, W., Macri, L. M., et al. 2022, ApJL, 934, L7, doi: 10.3847/2041-8213/ac5c5b

Rosenberg, E., Gratton, S., & Efstathiou, G. 2022, MNRAS, 517, 4620, doi: 10.1093/mnras/stac2744

Rubin, D., Aldering, G., Betoule, M., et al. 2023, arXiv e-prints, arXiv:2311.12098, doi: 10.48550/arXiv.2311.12098

Rudie, G. C., Steidel, C. C., Shapley, A. E., & Pettini, M. 2013, ApJ, 769, 146, doi: 10.1088/0004-637X/769/2/146

Sailer, N., Farren, G. S., Ferraro, S., & White, M. 2025. https://arxiv.org/abs/2504.16932

Schlafly, E. F., Kirkby, D., Schlegel, D. J., et al. 2023, AJ, 166, 259, doi: 10.3847/1538-3881/ad0832

Schöneberg, N. 2024, JCAP, 2024, 006, doi: 10.1088/1475-7516/2024/06/006

Scolnic, D., Brout, D., Carr, A., et al. 2022, ApJ, 938, 113, doi: 10.3847/1538-4357/ac8b7a

Seljak, U. 2012, JCAP, 2012, 004, doi: 10.1088/1475-7516/2012/03/004

Sinigaglia, F., Kitaura, F.-S., Nagamine, K., Oku, Y., & Balaguera-Antolínez, A. 2023. https://arxiv.org/abs/2305.10428

Slosar, A., Iršič, V., Kirkby, D., et al. 2013, JCAP, 2013, 026, doi: 10.1088/1475-7516/2013/04/026

Torrado, J., & Lewis, A. 2019, Cobaya: Bayesian analysis in cosmology, Astrophysics Source Code Library, record ascl:1910.019

—. 2021, JCAP, 2021, 057, doi: 10.1088/1475-7516/2021/05/057

Tristram, M., Banday, A. J., Douspis, M., et al. 2024, A&A, 682, A37, doi: 10.1051/0004-6361/202348015

Weinberg, D. H., Mortonson, M. J., Eisenstein, D. J., et al. 2013, PhR, 530, 87, doi: 10.1016/j.physrep.2013.05.001

Youles, S., Bautista, J. E., Font-Ribera, A., et al. 2022, MNRAS, 516, 421, doi: 10.1093/mnras/stac2102

Zaborowski, E. A., Taylor, P., Honscheid, K., et al. 2025, JCAP, 2025, 020, doi: 10.1088/1475-7516/2025/06/020

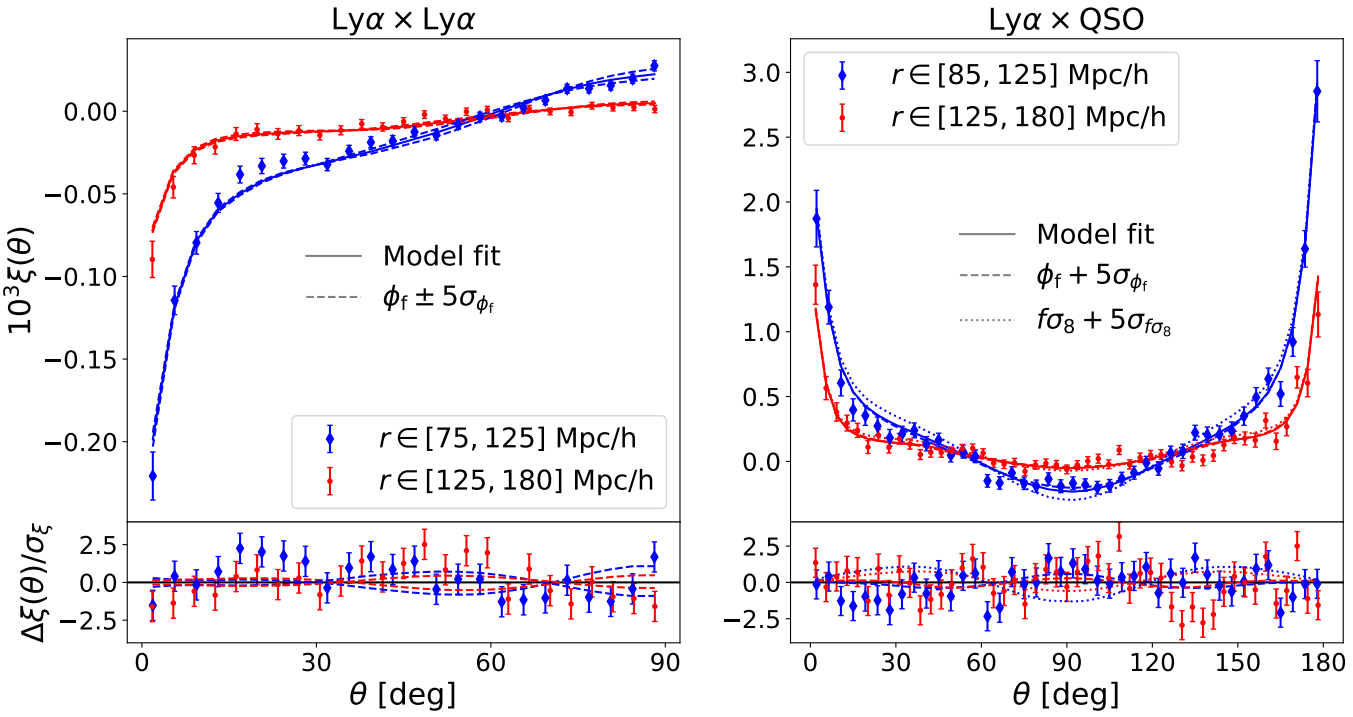

Figure 16. Similar to figure 1, but showing the shell compression for larger scales. The impact of changing the AP and RSD parameters by 5σ is lower than our statistical uncertainty for these shells, indicating that large scales do not contribute very significantly to our constraints.

APPENDIX

A. LARGE-SCALE FITS

In Section 3, we showed a comparison between the best-fit model and the data correlation functions when compressing into shells at smaller separations. Here we show the equivalent plots for the larger separations. These are illustrated in Figure 16, where the blue shells roughly capture the information from the BAO region (left for the auto-correlation, and right for the cross-correlation), while the red shells show the broadband information at scales larger than BAO. Similar to Figure 1, the dashed and dotted lines indicate models where we vary AP and $f\sigma_8$, respectively. In this case, changing ϕ_f and $f\sigma_8$ by 5σ produces a change in the model that is smaller than our uncertainties. This indicates that large scales do not contribute very significantly to our constraints.

B. NUISANCE PARAMETERS

Our analysis has 21 nuisance parameters, introduced in Section 2.2. These include the Lyα bias and RSD parameters, 5 parameters related to the Arinyo model, 13 parameters related to contaminants and systematics, and the isotropic scale of the broadband, α_s. Our priors are presented in Table 3, and largely follow the priors used by DESI24-Lyα.[12] The important exceptions are the systematic quasar redshift error, $\Delta r_{||}$, which had a flat prior in DESI24-Lyα, and the Arinyo parameters, $\{q_1, k_v, a_v, b_v, k_p\}$, which were fixed in DESI24-Lyα. As we use a significantly larger minimum separation for the cross-correlation compared to DESI24-Lyα (40 h^{-1}Mpc versus 10 h^{-1}Mpc), we found that $\Delta r_{||}$ was poorly constrained and decided to impose a wide Gaussian prior to discourage extreme values. $\Delta r_{||}$ has been measured by Bault et al. (2025) for the DESI DR1 quasar sample and found to be significantly smaller than 1 h^{-1}Mpc (also see Brodzeller et al. 2023). The prior has no impact on our results as shown in Section 4.4. This was also recently done in DESI Collaboration et al. (2025c), and we use the same prior. In the case of the Arinyo parameters, we chose wide flat

[12] As mentioned in Section 2.3, we sample the growth rate, f, and treat $f\sigma_8$ as a derived parameter.

Parameter	Priors
$\{\alpha, \phi\}$	$\mathcal{U}[0.01, 2.00]$
f	$\mathcal{U}[0.00, 2.00]$
b_F	$\mathcal{U}[-2.00, 0.00]$
β_F	$\mathcal{U}[0.00, 5.00]$
b_Q	$\mathcal{U}[0.00, 10.00]$
b_{HCD}	$\mathcal{U}[-0.20, 0.00]$
β_{HCD}	$\mathcal{N}(0.500, 0.090)$
$L_{\rm HCD}(h^{-1}{\rm Mpc})$	$\mathcal{U}[0.00, 15.00]$
$10^3 b_{\rm SiII(1190)}$	$\mathcal{U}[-500.00, 0.00]$
$10^3 b_{\rm SiII(1193)}$	$\mathcal{U}[-500.00, 0.00]$
$10^3 b_{\rm SiII(1260)}$	$\mathcal{U}[-500.00, 0.00]$
$10^3 b_{\rm SiIII(1207)}$	$\mathcal{U}[-500.00, 0.00]$
$10^3 b_{\rm CIV(eff)}$	$\mathcal{N}(-24.3, 1.5)$
$\Delta r_\parallel(h^{-1}{\rm Mpc})$	$\mathcal{N}(0.0, 1.0)$
$\sigma_v(h^{-1}{\rm Mpc})$	$\mathcal{U}[0.00, 15.00]$
ξ_0^{TP}	$\mathcal{U}[0.00, 2.00]$
$10^4 a_{\rm noise}$	$\mathcal{U}[0.00, 100.00]$
q_1	$\mathcal{U}[0.00, 2.00]$
k_v	$\mathcal{U}[0.30, 2.00]$
a_v	$\mathcal{U}[0.10, 1.00]$
b_v	$\mathcal{U}[1.00, 2.00]$
k_p	$\mathcal{U}[5.00, 22.00]$

Table 3. Parameter priors used in our analysis. $\mathcal{U}[\min, \max]$ indicates a flat prior while $\mathcal{N}(\mu, \sigma)$ indicates a Gaussian with mean μ and standard deviation σ.

priors based on the fits to hydrodynamical simulations from Arinyo-i-Prats et al. (2015) and Chabanier et al. (2024). Only one of these parameters, q_1, is constrained by the data. The other four parameters, $\{k_v, a_v, b_v, k_p\}$, produce posteriors consistent with the flat priors we imposed and show no significant correlation with our main parameters ($\phi_{\rm f}$ and $f\sigma_8$).

In Figure 17, we show the impact of changing several of the important nuisance parameters on the lowest separation shell for both the auto-correlation (left) and the cross-correlation (right). The magnitude of the variation in each parameter was chosen such that the change in the model can be visualized in the lower panels of Figure 17. For the auto-correlation, we show the impact of changing the Lyα bias and RSD parameters, along with the HCD bias and the main parameter in the Arinyo model (q_1). The other HCD and Arinyo parameters are not included here, as they do not produce significant changes. None of the parameters shown produces changes similar to those produced by AP. Both biases and the RSD parameter produce changes in the slope of the anisotropy, while AP affects the curvature of the anisotropy. On the other hand, for the cross-correlation, the Lyα and QSO RSD terms are degenerate (see Section 2.2), which means the impact of the Lyα RSD parameter is very similar to the impact of $f\sigma_8$. However, β_F is tightly constrained by the auto-correlation, thus allowing us to break this degeneracy. In the case of AP, the impact of the smearing due to redshift errors, modelled through the σ_z parameter, shows some similarity to the impact of $\phi_{\rm f}$. However, as we will see in Figure 18 below, there is no significant correlation between the posteriors of the two parameters. This is partly because a majority of the $\phi_{\rm f}$ information comes from the auto-correlation. Nevertheless, the complete lack of correlation between the two parameters indicates that we are able to distinguish the two effects.

We present the marginalized posterior distribution for the main parameters and a selection of nuisance parameters in Figure 18. We selected the subset of nuisance parameters that show any visible correlation with either of the two main parameters. The $L_{\rm HCD}$ parameter, which represents the typical length scale of unmasked HCDs, is completely unconstrained as we just recover the flat prior. This is not surprising given the scale cuts we use, with both Cuceu et al. (2023a) and DESI Collaboration et al. (2025c) finding similar behaviour. DESI Collaboration et al. (2025c) decided to impose a Gaussian prior on this parameter, as they found this had no impact on their BAO constraints.

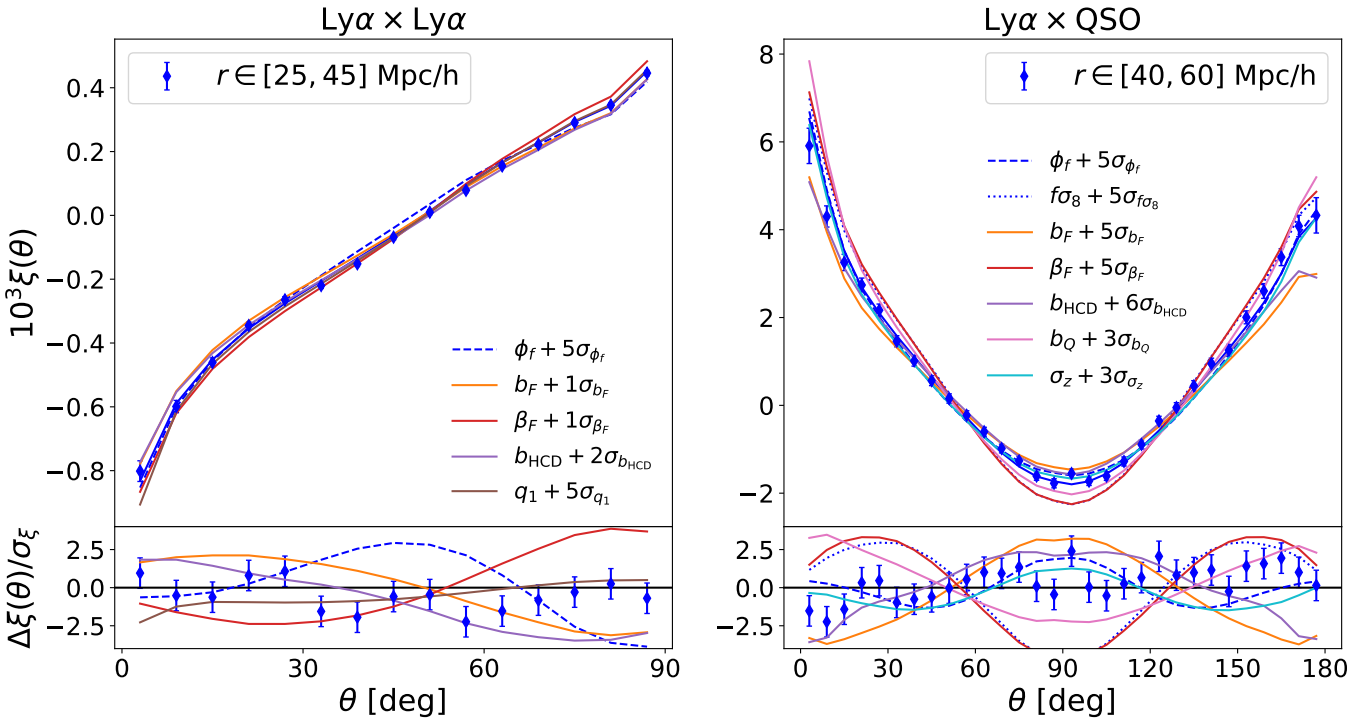

Figure 17. Similar to Figure 1, but showing only one shell at the smallest separations used in our analysis, and including the effect of changing several of the important nuisance parameters. For the auto-correlation (left), none of the nuisance parameters produce changes similar to those produced by changing AP. On the other hand, for the cross-correlation (right), the impact of $f\sigma_8$ is very similar to that of β_F, as the Lyα and QSO RSD terms are degenerate on large scales. This is why the auto-correlation is needed for our $f\sigma_8$ constraints, as it helps constrain β_F. While AP looks quite different from RSD in this compression, its impact has some similarity to that of redshift errors (dashed blue versus cyan).

However, in our case, we found that imposing constraints on this parameter has a small impact on $\phi_{\rm f}$ due to the small correlation between the two parameters (see Figure 18), and decided to keep using the flat prior.

Besides the unconstrained parameters described above, we also have two parameters that only have upper or lower bounds. These are the HCD bias, $b_{\rm HCD}$, which only has a lower bound, and q_1, which only has an upper bound. In the case of $b_{\rm HCD}$, this behaviour was encountered before (Cuceu et al. 2023a; DESI Collaboration et al. 2025c) when using scale cuts $r > 25\ h^{-1}$Mpc, as we do here. For the Arinyo q_1 parameter, this is the first ever measurement on real data, with all previous measurements taking place on hydro-simulations. The constraint we obtain, $q_1 < 0.51$ at 95% CR, is smaller than the value used in previous Lyα BAO analyses, $q_1 = 0.86$, which was based on the result from the reference simulation of Arinyo-i-Prats et al. (2015), interpolated to our effective redshift. However, both Arinyo-i-Prats et al. (2015) and more recent analyses (e.g. Chabanier et al. 2024; Chaves-Montero et al. 2025) have found a wide range of values for this parameter depending on the exact simulation and model setup, some of which are consistent with our results. Nevertheless, we caution against over-interpreting our results in the context of future analyses on hydro-simulations due to two main reasons. First, given that we could not test the Arinyo model in mocks, we do not exclude the possibility that this parameter is also capturing some subdominant systematic effects (e.g., from imperfect modelling of HCDs, metals, or the distortion matrix). Secondly, this parameter was treated as a nuisance parameter in this analysis, and the robustness of its measurement to analysis and modelling variations has not been tested. It would be interesting to perform an analysis that specifically aims to measure the Arinyo parameters for the purpose of comparing with results from hydro-simulations, but that is beyond the scope of this work.

C. COMPARISON OF SAMPLER AND FITTER RESULTS

The main results presented in this paper are given by marginalized posterior distributions over the parameters of interest ($\phi_{\rm f}$, $f\sigma_8$), obtained from the full posterior distribution over all parameters, which is sampled using the `PolyChord` Nested Sampler. However, due to computational constraints, a significant part of our validation results

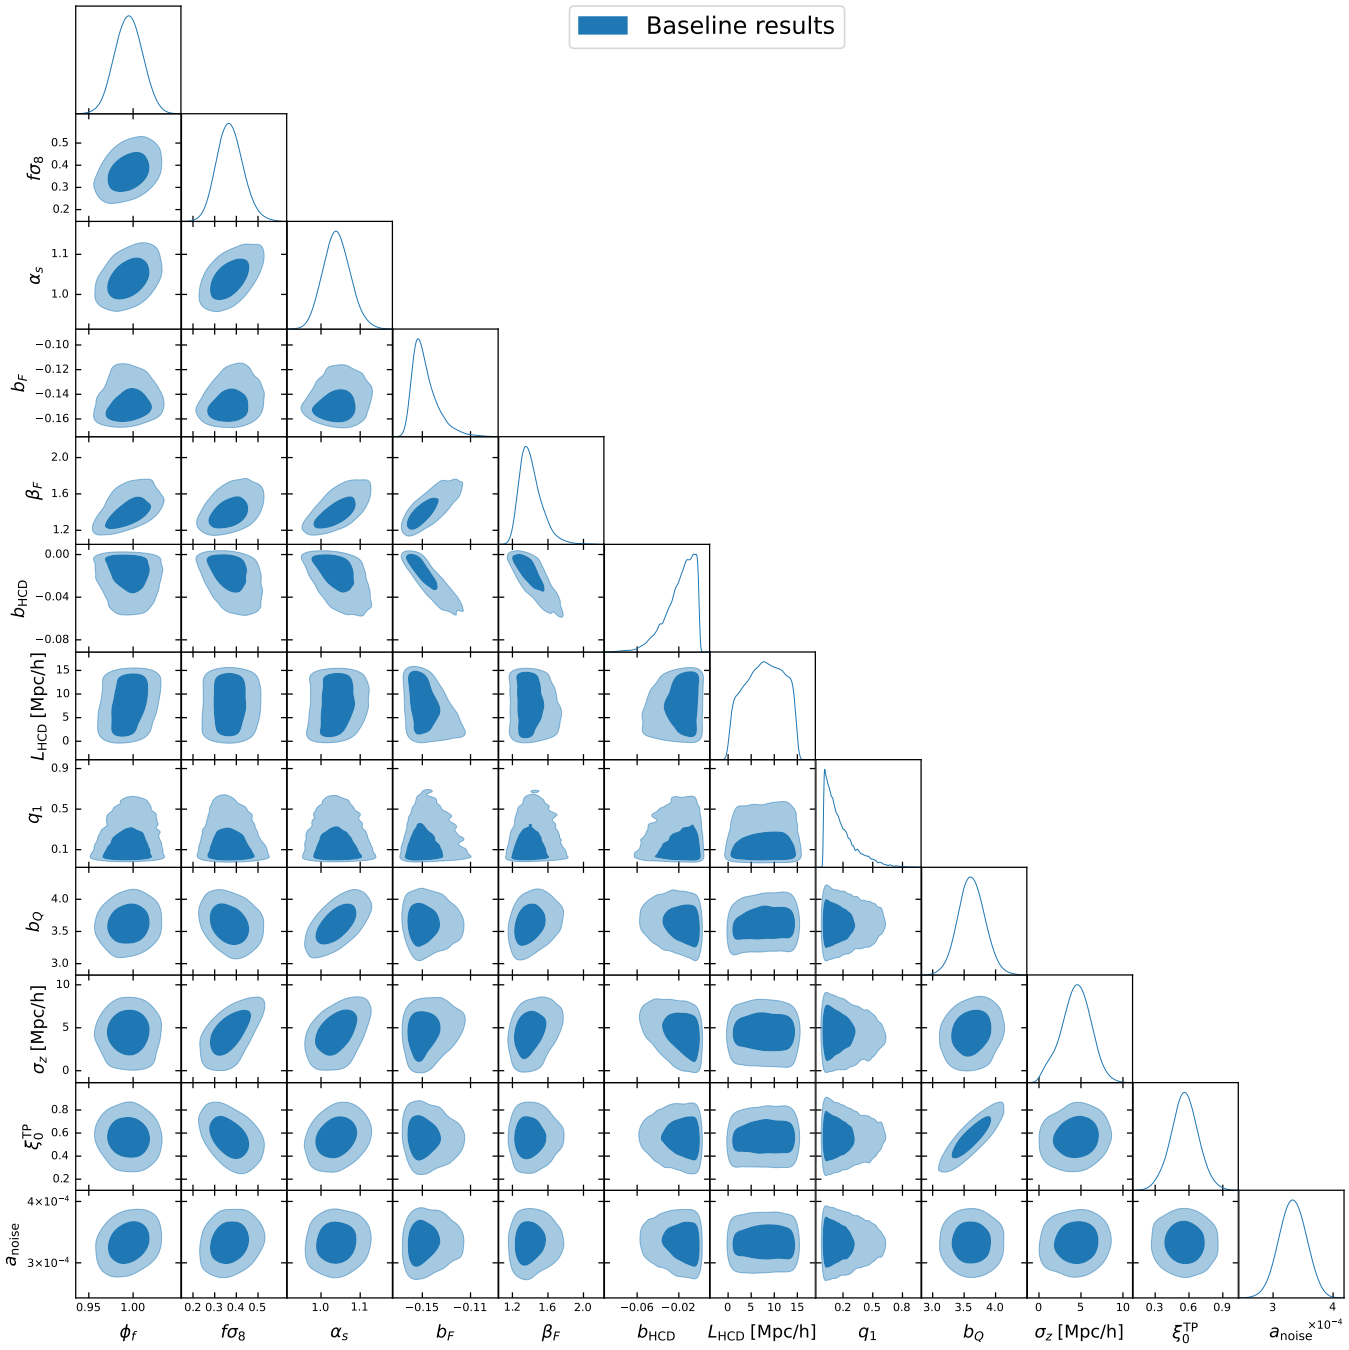

Figure 18. Posterior distribution for $\phi_{\rm f}$ and $f\sigma_8$, along with a selection of nuisance parameters that show some degree of correlation with either of the two main parameters.

were obtained using the `iminuit` minimizer with Gaussian approximated uncertainties. These include the fits to individual mocks, and most of the analysis and modelling variations presented in Section 4. Here, we discuss the differences between the two types of analyses and their impact on our validation.

The first difference relevant to our analysis is due to parameter uncertainties being approximated as Gaussian when using the fitter. The main results affected by this approximation are the tests of the robustness of uncertainties using mocks (Section 4.1), which include the pull distributions presented in Figure 4. This is because the uncertainties in the individual mock fits are approximated as Gaussian, while the distribution of parameter results in mocks is not Gaussian. Nevertheless, we found that the Gaussian approximation works well and produces accurate uncertainties,

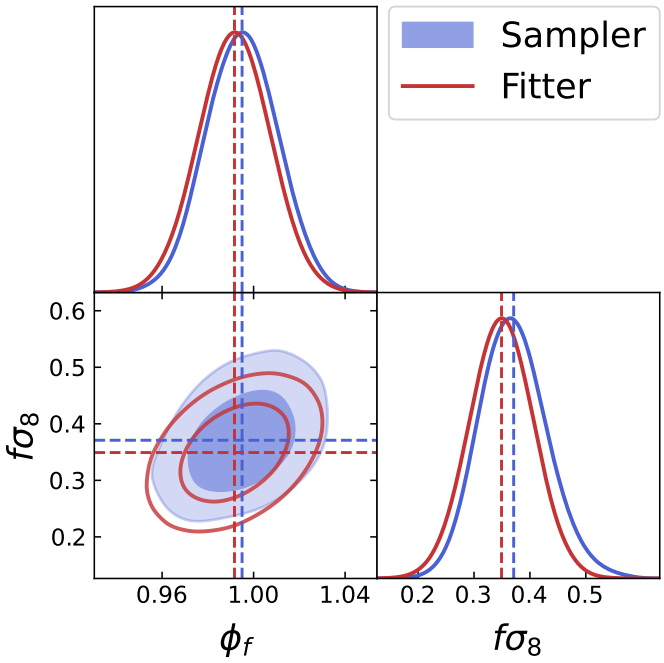

Figure 19. Comparison of AP and growth results using a sampler and marginalization versus a fitter and minimization. Both results use the baseline model described in Section 2.2. The dashed lines indicate the best-fit values obtained by the fitter (red) and the maximum a posteriori values from the sampler (blue). The fitter minimizes the χ^2 over all parameters, and estimates the Gaussian uncertainty using the second derivative of the χ^2 around the best fit. On the other hand, the sampler computes the full posterior distribution, which allows us to marginalize over nuisance parameters. The primary difference arises because the fitter results are conditioned on the best-fit value of each nuisance parameter, whereas the sampler results are marginalized over nuisance parameters.

especially for $\phi_{\rm f}$. Finally, this approximation does not impact our analysis and modelling variations, because in that case, we are only testing for shifts in the best-fit parameter values.

The other main difference relevant to our analysis arises from the fact that when using the sampler, we marginalize over nuisance parameters, whereas the fitter results for $\phi_{\rm f}$ and $f\sigma_8$ are conditioned on the best-fit values of the nuisance parameters. This is relevant because a number of nuisance parameters are either completely or partly unconstrained by the data (see Appendix B). We illustrate this in Figure 19, where we compare sampler and fitter constraints using our baseline model. We also show the best-fit parameter values from the fitter using the dashed red lines, and the maximum a posteriori results from the sampler using the dashed blue lines. We found that the parameter primarily responsible for this difference is $L_{\rm HCD}$.[13] This is because this parameter is completely unconstrained and has a small correlation with $\phi_{\rm f}$ ($\rho \sim 0.2$). Sampler results marginalize over all values of this parameter within the flat prior boundary $0 < L_{\rm HCD} < 15\ h^{-1}$Mpc, while the fitter tries to find a best-fit value, which in most cases ends up being either the upper or lower bound of the flat prior (in the baseline analysis it is the upper bound). Based on tests with mocks, we expect this parameter to have values between 3 h^{-1}Mpc and 7 h^{-1}Mpc, with the larger values obtained when no DLAs are masked (see Appendix B of DESI24-Lyα). The flat prior bounds were chosen to be larger than this range, in order to be conservative. In Section 4.4 we found that fixing this parameter to either 3 h^{-1}Mpc or 10 h^{-1}Mpc does not have a significant impact on $\phi_{\rm f}$ constraints, producing shifts of at most 0.1σ. Nevertheless, because the fitter tends to choose either the upper or lower bound of the flat prior on $L_{\rm HCD}$, our individual mock results and analysis variations likely have some extra variance. Therefore, future Lyα full-shape analyses could choose a more restrictive prior for this parameter, given that it has a small impact on results. For example, a Gaussian prior could be imposed based on the range of results found in mocks, as was recently done by DESI Collaboration et al. (2025c).

[13] This conclusion was reached after performing tests where unconstrained parameters are fixed one-by-one. We found that fixing $L_{\rm HCD}$ brings the sampler and fitter results in very good agreement.

D. COVARIANCE BETWEEN DESI DR1 BROADBAND AP AND DR2 BAO

To quantify the covariance between our DESI DR1 broadband AP measurements and BAO results from DESI DR2, we use the method described in Appendix F of DESI24-Lyα. First, we compute the joint covariance matrix of DESI DR1 and DR2 using only the Lyα(A) region. This is done using the same method for computing the covariance matrix of the four DR1 correlations as described in Section 2.1, but replacing the DR1 Lyα(B) region correlations with the DR2 Lyα(A) region correlations (DESI Collaboration et al. 2025c). This covariance matrix now includes the cross-covariance between the DR1 Lyα(A) correlations and the DR2 Lyα(A) correlations. Similar to DESI24-Lyα, we ignore Lyα(B) correlations, as they contribute a very small amount of information (see Section 3).

We next produce a set of 4096 Monte Carlo realisations of the Lyα(A) auto and QSO cross-correlations for both DESI DR1 and DR2 using the best-fit model from Section 3, with noise generated from the covariance matrix described above. This means the auto and cross-correlation pairs generated for DR1 and DR2 have realistic noise and cross-covariance. After that, we perform full-shape fits on the set of DR1 correlations using the model described in Section 2.2 with independent $\phi_{\rm s}$ and $\phi_{\rm p}$ parameters. In parallel, we also perform BAO fits on the set of DR2 correlations using the model described in DESI Collaboration et al. (2025c). This results in a set of 4096 broadband AP constraints based on DESI DR1, and a set of 4096 BAO constraints based on DESI DR2.

Using these two sets of best fit parameters, we can compute the cross-covariance between DR1 AP and DR2 BAO. The correlation coefficients we obtain between DR1 $\phi_{\rm s}$ and the DR2 BAO parameters, $\alpha_{||}$ and $\alpha_{\perp}$, are:

$$\rho(\phi_{\rm s}^{\rm DR1}, \alpha_{||}^{\rm DR2}) = -0.015 \pm 0.016, \tag{D1}$$

$$\rho(\phi_{\rm s}^{\rm DR1}, \alpha_{\perp}^{\rm DR2}) = 0.001 \pm 0.016, \tag{D2}$$

where the uncertainties are computed through bootstrap. Therefore, we conclude that the cross-covariance between DESI DR1 broadband AP and DR2 BAO is consistent with zero, and the two measurements can be treated as independent.