

UC Irvine

UC Irvine Previously Published Works

Title

DESI DR1 Ly α 1D power spectrum: Validation of estimators

Permalink

<https://escholarship.org/uc/item/9x73w5pk>

Journal

Journal of Cosmology and Astroparticle Physics, 2026(04)

ISSN

1475-7516

Authors

Karaçaylı, Naim Göksel

Ravoux, Corentin

Martini, Paul

et al.

Publication Date

2026-04-01

DOI

10.1088/1475-7516/2026/04/048

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

PAPER • OPEN ACCESS

DESI DR1 Ly α 1D power spectrum: Validation of estimators

To cite this article: Naim Göksel Karaçaylı *et al* JCAP04(2026)048

View the [article online](#) for updates and enhancements.

You may also like

- [Multi-messenger Observations of a Binary Neutron Star Merger](#)
B. P. Abbott, R. Abbott, T. D. Abbott et al.
- [Identification of low-momentum muons in the CMS detector using multivariate techniques in proton-proton collisions at \$\sqrt{s} = 13.6\$ TeV](#)
V. Chekhovsky, A. Hayrapetyan, V. Makarenko et al.
- [Performance of heavy-flavour jet identification in Lorentz-boosted topologies in proton-proton collisions at \$\sqrt{s} = 13\$ TeV](#)
A. Hayrapetyan, A. Tumasyan, W. Adam et al.

RECEIVED: September 19, 2025

REVISED: February 6, 2026

ACCEPTED: February 16, 2026

PUBLISHED: April 16, 2026

DESI DR1 Ly α 1D power spectrum: Validation of estimators

Naim Göksel Karaçaylı et al.

Full author list at the end of the paper

E-mail: karacayli.1@osu.edu, corentin.ravoux@clermont.in2p3.fr

ABSTRACT: The Data Release 1 (DR1) of the Dark Energy Spectroscopic Instrument (DESI) is the largest sample to date for small-scale Ly α forest cosmology, accessed through its one-dimensional power spectrum (P_{1D}). The Ly α forest P_{1D} is extracted from quasar spectra that are highly inhomogeneous (both in wavelength and between quasars) in noise properties due to intrinsic properties of the quasar, atmospheric and astrophysical contamination, and also sensitive to low-level details of the spectral extraction pipeline. We employ two estimators in DR1 analysis to measure P_{1D} : the optimal estimator and the fast Fourier transform (FFT) estimator. To ensure robustness of our DR1 measurements, we validate these two power spectrum and covariance matrix estimation methodologies against the challenging aspects of the data. First, using a set of 20 synthetic 1D realizations of DR1, we derive the masking bias corrections needed for the FFT estimator and the continuum fitting bias needed for both estimators. We demonstrate that both estimators, including their covariances, are unbiased with these corrections using the Kolmogorov-Smirnov test. Second, we substantially extend our previous suite of CCD image simulations to include 675,000 quasars, allowing us to accurately quantify the pipeline's performance. This set of simulations reveals biases at the highest k values, corresponding to a resolution error of a few percent. We base the resolution systematics error budget of DR1 P_{1D} on these values, but do not derive corrections from them since the simulation fidelity is insufficient for precise corrections.

KEYWORDS: Lyman alpha forest, power spectrum

ARXIV EPRINT: [2509.13593](https://arxiv.org/abs/2509.13593)



Contents

1	Introduction	1
2	Synthetic data	3
2.1	Quickquasars mocks	3
2.2	CCD image simulations	5
3	Methods	5
3.1	Continuum fitting	5
3.2	Optimal estimator	6
3.3	FFT estimator	10
3.4	Kolmogorov-Smirnov test	11
4	Validating the optimal estimator	12
4.1	Validation in the absence of contaminants	13
4.2	Validation of masking HCDs, BALs, and atmospheric lines	14
5	Validating the FFT estimator	15
5.1	The power spectrum from transmission files	16
5.2	Validation in the absence of contaminants	17
5.3	Deriving and validating masking corrections	19
6	Accuracy of the spectrograph resolution	22
7	Discussion	23
	Author List	30

1 Introduction

The 1D flux power spectrum of the Ly α forest (P_{1D}) is rich in information on small-scale physics (below one Mpc). It has been used to constrain the thermal state of the intergalactic medium (IGM) [1–3], the amplitude and the slope of the linear matter power spectrum at $k \approx 1 \text{ Mpc}^{-1}$ [4–6], the primordial power spectrum [7], the sum of neutrino masses [8–10], and the nature of dark matter [11–19]. The Dark Energy Spectroscopic Instrument’s (DESI) [20, 21] now-public Data Release 1 (DR1) [22] has over 300,000 Ly α quasars, which is the largest sample to date for P_{1D} measurements by a factor of 1.7 compared to the Extended Baryon Oscillation Spectroscopic Survey (eBOSS) [23]. This substantial increase in data size necessitates a thorough investigation of P_{1D} and covariance matrix estimation techniques and spectroscopic processing to ensure robust cosmological inference.

This work consolidates the validation tests performed for the two estimators used in DESI DR1 P_{1D} measurement — the Fast Fourier Transform (FFT) estimator presented in ref. [24] and the quadratic maximum likelihood estimator (QMLE), also known as the optimal

estimator, presented in ref. [25], which differ at the level of their treatment of the data. The optimal estimator weights pixels using a covariance matrix constructed from the pipeline noise estimates and $\text{Ly}\alpha$ forest correlations, which enables this estimator to be robust against gaps in the spectra that arise from masking contamination from unrelated physical and instrumental effects. The Fourier transform is significantly faster than the inverse covariance weighting of the optimal estimator, but requires simplifying assumptions in the treatment of noise, resolution, and biases due to masking, especially of atmospheric lines.

It is further critical to ensure that the estimated covariance matrix of P_{1D} accurately and conservatively reflects the measurement errors. The FFT technique has a newly developed covariance matrix estimator presented in ref. [24] that we summarize and validate in this paper. The optimal estimator produces a Gaussian covariance matrix that underestimates the errors and needs to be complemented with non-parametric bootstrap techniques [26, 27]. The bootstrap covariance matrix is noisy and numerically unstable for redshift bins with low statistics. We develop and validate a new method for combining Gaussian and noisy bootstrap covariance estimates to ensure smoothness and prevent underestimation of errors.

We produce synthetic 1D $\text{Ly}\alpha$ forest data using a log-normal transformation to validate the statistical robustness and identify any needed bias corrections [28]. These $\text{Ly}\alpha$ transmission fields do not have any instrumental and astrophysical effects, which are added via DESI’s simulation package. We analyze 20 realizations of DR1 and estimate a P_{1D} and a covariance matrix for each realization. Since these mocks are inherently random, the conclusions are interpreted probabilistically using the Kolmogorov-Smirnov test. The corrections derived from these mocks are applied to the measurements in refs. [24, 25].

Since $\text{Ly}\alpha$ forest measurements, and P_{1D} specifically, are highly sensitive to the details of spectroscopic data processing, we also investigate the spectral extraction pipeline. The primary concern with spectral extraction is the accuracy of the resolution matrix, as errors increase quadratically with k values, biasing the small-scale information. We test the accuracy of the resolution matrix by simulating 2D CCD images, which are extracted using DESI’s spectroscopic data processing pipeline [29, 30]. This procedure has been computationally expensive; however, after significant improvements to the pipeline, the images can now be efficiently processed using GPUs at a fraction of the cost, achieving a speedup of over a factor of 20. This enables us to match the sample size of DR1 by expanding our DESI early data release (EDR) simulated CCD images by a factor of 15 [27, 31]. These CCD image simulations help us quantify the systematic error budget in resolution, as we have done in refs. [24, 25]. The simulations can also enable us to quantify any biases introduced by the extraction pipeline and serve as a means to test improvements in pipeline performance. However, simulation fidelity needs to be improved before these results are used as bias corrections.

The outline of this paper is as follows. We describe the two types of simulated data in section 2. Our methodology is presented in section 3, which provides details on the P_{1D} and covariance matrix estimators, as well as the quasar continuum fitting algorithm and the Kolmogorov-Smirnov test. The optimal estimator is validated in section 4, and the FFT estimator is validated in section 5. We estimate the accuracy of the spectroscopic data processing and resolution matrix in section 6. The limitations of our analysis and directions for future work are discussed in section 7.

2 Synthetic data

We follow the same procedure presented in the DESI early data release measurement to generate synthetic data [27]. This procedure consists of two main components: generating 1D transmission fields with forest fluctuations and simulating the DESI instrument.

The 1D transmission fields are created using the lognormal model in ref. [28], which we summarize briefly here. We first generate Gaussian random fields that are correlated along the line of sight, with an input power spectrum tuned to approximately match the desired flux power spectrum (these fields are uncorrelated between different lines of sight). Then, they are multiplied by the appropriate redshift evolution functions and are applied a lognormal transformation to obtain the column density, $n(z)$, and optical depth, $\tau(z)$, fields. The lognormal transform introduces nonlinearity and ensures that $n(z)$ and $\tau(z)$ are non-negative [26, 32]. Finally, the transmission field is defined as $F(z) = e^{-\tau(z)}$. We generate these transmission fields on a linear wavelength grid with a $0.2 \text{ \AA}$ spacing, without any resolution or noise effects.

Simulating the DESI instrument also encapsulates creating a diverse set of quasar spectra, whose continua are generated from a broken power law with emission lines using the `simqso` package [33] as a part of the `desisim` package.¹ We have two methods to simulate the DESI instrument.

The first method uses the `quickquasars` program in the same package. DESI observations are simulated using the `spepsim`² package [34] for quick simulations of the fiber spectrograph response (see ref. [35] for a detailed description of `quickquasars` mocks). This program (after generating random quasar spectra) simulates sky and instrumental noise, and incorporates wavelength-dependent camera resolution, although it does not validate the computationally expensive spectral extraction process.

The second method uses CCD image simulations to validate the spectro-perfectionism resolution matrix [29] and validate the complete spectral extraction process. This method projects mock quasar spectra (generated using `simqso`) onto two-dimensional images that simulate DESI raw data at the CCD pixel level with the `desisim` package. We then process these simulated images with the same spectroscopic reduction pipeline employed for actual DESI observations [30]. This approach is more computationally expensive than the mocks described above.

2.1 Quickquasars mocks

For these mocks, the quasar diversity, DESI instrument, and the sky are simulated through a program called `quickquasars` in the `desisim` package. This program randomly generates quasar spectra from a broken power law with emission lines, convolves the spectra with the wavelength-dependent camera resolution for each spectrograph, adds noise appropriate to a given flux and exposure time, and resamples the resulting simulated spectra onto the output DESI wavelength grid of $\Delta\lambda_{\text{DESI}} = 0.8 \text{ \AA}$ per pixel. The source contribution to noise is smoothed with a Gaussian kernel of $\sigma = 10 \text{ \AA}$ to imitate the DESI pipeline [30].

¹<https://github.com/desihub/desisim>

²<https://github.com/desihub/spepsim>




Figure 1. (*Left*) The average percent error and scatter in the generated mock r -band flux values binned in redshift. We achieve a good match at $z_{\text{qso}} < 4$, which constitutes the majority of the DR1 quasar sample. (*Right*) SNR with respect to observed wavelength. We achieve a good agreement between the data (*blue*) and the mocks (*orange*). Solid lines represent the median, and the area between the 32nd and 68th percentiles is shaded.

Previously, in the DESI EDR analysis, the output resolution matrix was erroneously calculated using a box-car average of the finely sampled camera resolution matrix, both over rows and columns. This smoothed the resolution matrix twice and required a deconvolution as a fix in the estimator. We corrected this double-smoothing issue in the output resolution matrix, so that the deconvolution step described in refs. [27, 36] is no longer applied to these mocks. However, the resolution matrix still needs to be oversampled by a factor of two to improve the accuracy in the quadratic estimator formalism.

`quickquasars` was further improved to match the noise properties of the data [35]. This was achieved by generating quasar continuum templates that match the r -band flux values and adding noise to each spectrum based on the effective exposure time $T_{\text{eff}} = 12.15 \times \text{TSNR}_{\text{LRG}}^2$ seconds, where TSNR_{LRG} is the signal-to-noise ratio of the luminous red galaxy (LRG) template [30]. The input r -band flux values must correspond to the unabsorbed quasar continuum, although the $\text{Ly}\alpha$ forest begins to affect the r -band for quasars at $z > 3.7$. We correct this effect by amplifying r -band flux values based on the average mean flux decrement caused by the forest. The relative error in the generated mock r -band flux values is shown in the left panel of figure 1, where the data points represent the mean error in a redshift bin and error bars represent the scatter in that bin. We achieve a good match for the majority of our sample ($z_{\text{qso}} < 4$) and a maximum of 15% offset on average around $z = 4.7$. The right panel of figure 1 demonstrates that the wavelength dependence of SNR in the synthetic spectra parallels the data. All unique targets identified as quasars are simulated in our mocks at their exact redshifts, approximately at the same r -band flux values and noise properties.

We randomly insert high-column density systems (HCDs), including damped $\text{Ly}\alpha$ absorption systems (DLAs), without any correlations with the underlying matter field. The number of systems per unit redshift is set to be $l(z) = 0.96e^{-7/z^2}$ following ref. [37]. We draw column densities following the column density distribution presented in ref. [38] between $\log N_{\text{HI}} \in [19.5, 22.5]$. In reality, these systems are expected to cluster with the underlying

density field [39, 40]. Therefore, the effect of clustering and incompleteness of real data catalogs still needs to be included in cosmological analyses [41, 42].

We insert broad absorption lines (BALs) into 16% of quasars [43, 44]. Unlike the validation tests for Baryon Acoustic Oscillations studies [45], we do not introduce metal lines or higher-order Lyman lines.

2.2 CCD image simulations

We produced CCD image simulations to validate the implementation of the resolution matrix by the spectroscopic pipeline [30]. This is critical as the accuracy of the resolution matrix is the dominant source of systematic error on small spatial scales. The simulations we produced for this work were constructed in the same manner as those described in ref. [27] for the EDR analysis [27, 36], so here we only provide a short description of how we generate the spectra and then focus on the key differences from our previous work.

We produced the image simulations with the `desisim` package, which was built by the DESI collaboration to develop and validate the spectroscopic data processing pipeline [30]. This package can produce realistic versions of the raw, two-dimensional CCD data that are the starting point for the pipeline processing. The inputs to this code are the intrinsic spectra of targets for a single DESI tile with apparent magnitudes representative of DESI quasar targets. The code models the bias, gain, and readnoise of the detectors, as well as the spectroscopic PSF, trace location, and throughput as a function of wavelength. The code also includes a realistic model of night sky emission and spectroscopic standards.

As in ref. [27], we populate the tiles with only high-redshift quasars (4500 per tile), along with a selection of the standard stars and sky fibers that are required for the operation of the spectroscopic pipeline. As the DESI DR1 sample is substantially larger than EDR, we construct 150 simulated tiles of DESI quasars. Our final dataset is, therefore, 675,000 quasars, which is over an order of magnitude larger than the number we created for EDR and about 50% more than the number of DESI DR1 quasars at $z > 2.09$. Another significant change relative to the EDR simulations is that we changed the redshift distribution from a uniform distribution over $2.6 < z < 3.6$ to one that is a better match to the observations. We also generated mock arc lamp spectra, which are used by the pipeline to measure the PSF and ultimately derive the resolution matrix. Finally, we processed the simulated dataset with the same pipeline we used for the observations.

3 Methods

3.1 Continuum fitting

The power spectrum estimation requires the transmitted flux fluctuations δ_F due to Ly α forest absorption. For a quasar q at redshift z_q , this is given by

$$\delta_{F,q}(\lambda) = \frac{f_q(\lambda)}{\bar{F}(z_q)C_q(\lambda_{\text{RF}})} - 1, \quad (3.1)$$

where λ is the observed wavelength, $f_q(\lambda)$ is the observed flux, and $\lambda_{\text{RF}} = \lambda/(1 + z_q)$ is the wavelength in the quasar's rest frame. Since $\delta_{F,q}$ is relative to the mean transmission

of the IGM ($\bar{F}$), the quasar continuum $C_q(\lambda_{\text{RF}})$ is multiplied by $\bar{F}(z_\alpha)$ at the Ly α forest redshift $z_\alpha = \lambda/\lambda_{\text{Ly}\alpha} - 1$.

In real data, the quasar continuum is unknown and needs to be estimated. We use the same continuum fitting algorithm that was developed and has been applied to both $\xi_{3\text{D}}$ and $P_{1\text{D}}$ measurements to compute $\delta_{F,q}$ [27, 46]. In this formalism, the definition of the quasar continuum absorbs the mean transmission of the IGM, such that the quasar “continuum” $\bar{F}(z_\alpha)C_q(\lambda_{\text{RF}})$ is given by

$$\bar{F}(z_\alpha)C_q(\lambda_{\text{RF}}) = \bar{C}(\lambda_{\text{RF}})(a_q + b_q\Lambda), \quad \text{where} \quad \Lambda \equiv \frac{\log \lambda_{\text{RF}} - \log \lambda_{\text{RF}}^{(1)}}{\log \lambda_{\text{RF}}^{(2)} - \log \lambda_{\text{RF}}^{(1)}}, \quad (3.2)$$

$\lambda_{\text{RF}}^{(1,2)}$ are the minimum and maximum wavelengths considered in the calculation, $\bar{C}(\lambda_{\text{RF}})$ is the global mean continuum, and a_q and b_q are two quasar diversity parameters. These quasar diversity parameters also absorb the mean IGM transmission.

The overview of our method is as follows (see refs. [27, 46, 47] for a detailed description). For each quasar, we find the a_q and b_q values that minimize the following cost function while keeping *all* other parameters fixed:

$$\chi^2 = \sum_j \frac{\left[f_j - (a_q + b_q\Lambda_j)\bar{C}\left(\frac{\lambda_j}{1+z_q}\right) \right]^2}{\sigma_{q,j}^2} + \sum_j \ln \sigma_{q,j}^2, \quad (3.3)$$

where the summation j is over all pixels in the forest region and λ_j is the observed wavelength. The major complication comes from $\sigma_{q,j}^2$, which must take into account the intrinsic large-scale Ly α fluctuations σ_{LSS}^2 :

$$\sigma_{q,j}^2 = \eta(\lambda_j)\sigma_{\text{pipe},j}^2 + \sigma_{\text{LSS}}^2(\lambda_j)(a_q + b_q\Lambda_j)^2\bar{C}^2\left(\frac{\lambda_j}{1+z_q}\right), \quad (3.4)$$

where $\eta(\lambda)$ is the pipeline noise correction term. The FFT estimator uses a specific case in which $\sigma_{q,j}$ is assumed constant across the quasar spectrum. After every quasar is fit, we stack all continua in the rest frame, update the global mean continuum $\bar{C}$, and calculate η and σ_{LSS}^2 as described in the references. This algorithm is implemented in two separate packages: `picca`³ [46] and `qsonic`⁴ [48]. The FFT estimator uses `picca` as it has been part of the same package, whereas the optimal estimator uses `qsonic` since it is significantly faster. We also utilize one of `qsonic`’s differentiating features — employing a fiducial mean IGM transmission to mitigate its coupling to quasar diversity parameters.

3.2 Optimal estimator

The quadratic maximum likelihood estimator (QMLE) formalism has been a highly successful technique in power spectrum measurements from the cosmic microwave background radiation, galaxy surveys, and weak lensing [49–52]. It was also successfully applied to measure Ly α forest $P_{1\text{D}}$ from a variety of datasets such as 3,000 quasars from SDSS [26], 527 high-resolution,

³<https://github.com/igmhub/picca>

⁴<https://qsonic.readthedocs.io/en/stable/>

high-SNR spectra [53], and 54,600 quasars from the DESI early data release [27]. We call this the optimal estimator.

The optimal estimator works in real space (instead of Fourier space) to estimate the power spectrum and therefore allows weighting by the pipeline noise, accounts for intrinsic Ly α large-scale structure correlations, and, most importantly, is not biased by masking unwanted pixels in the spectra. We refer the reader to refs. [28, 53] for more details about our development process and application to high-resolution spectra. In short, the optimal estimator gives a power spectrum $\mathbf{p}_q$ and its covariance matrix $\mathbf{C}_q$ for each quasar q . This covariance matrix $\mathbf{C}_q$ includes contributions from the pipeline noise estimates, the signal contribution based on a fiducial power spectrum, and the “survey window function” due to masking and continuum marginalization. The final P_{1D} estimate is a weighted average: $\mathbf{p} = \mathbf{C}_{\text{tot}}^{-1} \sum_q \mathbf{C}_q^{-1} \mathbf{p}_q$, where $\mathbf{C}_{\text{tot}}^{-1} \equiv \sum_q \mathbf{C}_q^{-1}$. The inverse covariance matrix is also referred to as the Fisher matrix $\mathbf{F} = \mathbf{C}^{-1}$, and as the mode-mixing matrix in the CMB literature. This matrix normalizes the power spectrum in the weighted-average sense and also deconvolves the survey window function, which mixes different (k, z) modes due to masking, continuum marginalization, and the reverse interpolation of pixel pairs into two redshift bins.

We have made large performance improvements to our code since DESI’s early data analysis. These updates include: 1) Using a tile strategy while transpose-copying a matrix, which is ten times faster than the straightforward implementation. 2) Calculating $\mathbf{Q}_\alpha = \mathbf{C}^{-1} \mathbf{C}_{,\alpha}$ instead of $\mathbf{C}^{-1} \mathbf{C}_{,\alpha} \mathbf{C}^{-1}$, which reduces the number matrix multiplications. However, the trace calculation for the Fisher matrix $F_{\alpha\beta} = 0.5 \text{Tr}(\mathbf{Q}_\alpha \mathbf{Q}_\beta^T)$ can only be efficiently performed using the fast transpose update of point 1, leveraging the dot product. 3) Fine-tuning the shared-memory parallelization with `OpenMP`. Thanks to these performance gains, the estimator is now four times faster, and we are now able to estimate P_{1D} at 85 k bins per redshift bin and perform repeated runs on many mock realizations of DR1.

3.2.1 Estimating the covariance matrix

Another major numerical update concerns the bootstrap covariance matrix. Our code was limited to using each MPI task as a subsample in DESI early data analysis. This largely yielded noisy covariance matrix estimates since a few hundred independent samples cannot fully explore all the fluctuations in an approximately 1000×1000 matrix. We had to resort to strong sparsity cuts to suppress this noise, which prevented exploring correlated systematics fluctuations. We now bootstrap over quasars to fully examine the covariance matrix. However, we still have to make a few simplifying assumptions to speed up this procedure. The main approximation concerns the Fisher matrix. Calculating this matrix for every bootstrap realization is computationally expensive. Instead, we use the same Fisher matrix for all realizations. Second, we assign an observation frequency for each quasar based on a Poisson distribution with a mean of one [54, 55]. This means we do not have the same number of quasars across bootstrap realizations, but ref. [54] has shown that this introduces a negligible bias for a large number of independent data points ($n \gtrsim 100$). Finally, we can directly estimate the inverse covariance matrix (i.e. the Fisher matrix), which is what likelihood inferences need, using the constant Fisher matrix approximation and the fact that under the assumption of Gaussianity, $\langle \mathbf{y} \mathbf{y}^T \rangle = \mathbf{C}^{-1}$ where $\mathbf{y} = \mathbf{C}^{-1} \mathbf{p}$. In other words, we calculate the covariance

of the weighted data vectors between bootstrap realizations, which is approximately equal to the inverse of the covariance matrix of P_{1D} .

Unfortunately, the resulting bootstrap is still susceptible to noise problems, mostly at higher redshifts where there are only a few hundred quasars. However, we no longer need strong sparsity cuts thanks to our new regularization scheme that can reliably preserve the fluctuations due to systematics while suppressing the noise in the bootstrap inverse covariance matrix estimates $\mathbf{F}_B$. The optimal estimator formalism calculates the Fisher matrix under the assumption of Gaussianity, which we will refer to as $\mathbf{F}_G$ below. We compute the final bootstrap covariance as follows:

1. Using the diagonals of the Gaussian matrix $v_i = \sqrt{F_{ii}^G}$, normalize both matrices:

$$R_{ij}^G = F_{ij}^G / v_i v_j, \quad R_{ij}^B = F_{ij}^B / v_i v_j. \quad (3.5)$$

2. Subtract the two correlation matrices $\Delta\mathbf{R} = \mathbf{R}_B - \mathbf{R}_G$ and smooth it using a 2D Gaussian filter with $\sigma = 2$. The first half of the redshift bins is smoothed box by box to eliminate edge effects. This step yields a smooth correlation difference matrix $\widetilde{\Delta\mathbf{R}}$.
3. Reconstruct the smoothed bootstrap matrix $\widetilde{\mathbf{F}}_B$:

$$\widetilde{R}_{ij}^B = R_{ij}^G + \widetilde{\Delta R}_{ij} \quad \rightarrow \quad \widetilde{F}_{ij}^B = v_i v_j \widetilde{R}_{ij}^B. \quad (3.6)$$

Theoretically, the Gaussian covariance matrix must yield the smallest errors. In more formal terms, the inequality $\mathbf{x}^T \mathbf{C}_B \mathbf{x} \geq \mathbf{x}^T \mathbf{C}_G \mathbf{x}$ must hold for all vectors $\mathbf{x}$. The smoothing procedure described so far cannot guarantee this theoretical boundary by itself. We now force this lower limit in two more steps.

4. We find the eigenvalues λ_i and eigenvectors $\mathbf{e}_i$ of this smoothed bootstrap inverse covariance matrix $\widetilde{\mathbf{F}}_B$. We calculate the values corresponding to each eigenvector using the Gaussian matrix $\lambda_i^G = \mathbf{e}_i^T \mathbf{F}_G \mathbf{e}_i$. These are the theoretical maxima, therefore we replace $\lambda_i \rightarrow \min(\lambda_i, \lambda_i^G)$ and reconstruct $\widetilde{\mathbf{F}}_B$ with these new eigenvalues. This procedure is the same as refs. [26, 27].
5. The inequality written above means that $\mathbf{M} = \widetilde{\mathbf{C}}_B - \mathbf{C}_G$ must be a positive semi-definite matrix. We calculate $\mathbf{M}_+$ by setting all negative eigenvalues of $\mathbf{M}$ to zero, and obtain the final bootstrap covariance matrix: $\widetilde{\mathbf{C}}_B^{\text{final}} = \mathbf{C}_G + \mathbf{M}_+$. Note that this step cannot be done with inverse covariance matrices.

The left panel in figure 2 shows the initial difference matrix $\Delta\mathbf{R}$, where the redshift increases from left (top) to right (bottom). The first two redshift bins exhibit strong off-diagonal contributions in the bootstrap estimate, whereas higher redshift bins are significantly noisier. The middle panel shows the 2D smoothed difference matrix $\widetilde{\Delta\mathbf{R}}$. The noise is largely suppressed. To preserve the off-diagonal structure of the first few redshift bins at the edges, we smooth each redshift bin separately. The right panel shows the difference for the final bootstrap estimate after forcing the inequality $\mathbf{x}^T \mathbf{C}_B \mathbf{x} \geq \mathbf{x}^T \mathbf{C}_G \mathbf{x}$.




Figure 2. The difference matrices $\Delta \mathbf{R}$ at three stages of the optimal estimator’s regularization scheme. The redshift increases from left (top) to right (bottom) in each panel. The initial estimate shown in the left panel has strong off-diagonal terms in the first few redshift bins and has significant noise at higher redshifts. The middle panel shows the 2D smoothed difference matrix $\widetilde{\Delta \mathbf{R}}$ in which the noise is largely suppressed. The right panel shows the difference for the final bootstrap estimate after it is forced to be larger than the Gaussian covariance matrix.

3.2.2 Analysis choices

We calculate the transmitted flux fluctuations δ_F using `qsonic` with the following settings. We set the rest-frame wavelength range to $1050 - 1180 \text{ \AA}$ following the literature standard, conservatively avoiding the large, hard-to-model profiles of the $\text{Ly}\alpha$ and $\text{Ly}\beta$ emission lines. We limit the observed wavelength range to $3600 - 7000 \text{ \AA}$ based on the minimum wavelength measured and the most distant observed quasar redshift ($z_{\text{qso}} \approx 5$) by DESI. The global mean continuum $\overline{C}(\lambda_{\text{RF}})$ is measured in $0.4 \text{ \AA}$ steps. We select quasars with an average SNR greater than one at wavelengths greater than the $\text{Ly}\alpha$ emission line of the quasar and with an average SNR greater than 0.3 in the forest region. We do not apply any pipeline noise correction, but we still fit for η during continuum fitting.

When estimating $P_{1\text{D}}$ we split spectra into two segments if they have more than 350 pixels, and we ignore segments with fewer than 20 pixels to reduce computation time and improve continuum marginalization. We use 60 linear bins with $\Delta k_{\text{lin}} = 5 \times 10^{-4} \text{ s km}^{-1}$, which is comparable to the fundamental frequency of one segment, followed by 25 log-linear bins with $\Delta k_{\text{log}} = 0.01$, reaching a maximum k that remains smaller than the Nyquist frequency. We measure $P_{1\text{D}}$ with redshift bins of size $\Delta z = 0.2$ from $z = 2.0$ to $z = 4.6$. As noted in the DESI EDR analysis in ref. [27], the $z = 2.0$ bin is susceptible to noise and interpolation errors, so we exclude it from all analyses presented in this work. We use the true power spectrum as a fiducial to calculate the signal matrix [28], and interpolate the signal and derivative matrices using sufficiently high resolution in redshift and velocity spacings.

3.2.3 Smoothing the measured power spectrum

The corrections proposed in this paper for the optimal estimator are scalings of a smooth power spectrum $P_{\text{smooth}}(k, z)$ following ref. [27]. This P_{smooth} is calculated by applying `SmoothBivariateSpline` to the logarithms of $1+z, k$ and measured (noisy) P_{1D} with corresponding statistical weights $\sigma_{\text{boot}}/P_{\text{1D}}$ and a smoothing factor that is five times the number of data points. This function in `scipy` is based on the algorithm presented in refs. [56, 57].

3.3 FFT estimator

The FFT estimator was used in several studies to directly compute P_{1D} from a set of transmitted flux fluctuations [26, 36, 41, 42, 58]. The complete description of the FFT estimator used in this study is given in ref. [36]. We provide a brief description and present the modifications below. The core concept of the FFT estimator is to decompose $\delta_{F,q}(\lambda)$ for each quasar q in the following manner:

$$\delta_{F,q}(\lambda) = (\delta_{\text{Ly}\alpha,q}(\lambda) + \delta_{\text{metals},q}(\lambda)) \otimes W_q(\lambda, \mathbf{R}) + \delta_{\text{noise},q}(\lambda), \quad (3.7)$$

where $\delta_{\text{Ly}\alpha,q}$ is the Ly α contrast we want to measure, $\delta_{\text{metals},q}$ corresponds to the contrast caused by metal absorption in the Ly α forest region, $\delta_{\text{noise},q}$ is the noise contrast, and $W_q(\lambda, \mathbf{R})$ is the effect of the finite resolution of the spectrograph on both astrophysical signals. It is derived from the resolution matrix $\mathbf{R}_q$, which characterizes the spectrograph resolution [29, 30]. Crucially, this estimator is employed in the wavelength configuration, as DESI processes the spectra in an equally spaced wavelength grid.

Following ref. [36], the derived Ly α contrast is cut into segments of the same length, noted as s . We Fourier transform this contrast and account for the different correlations in Fourier space between the signals mentioned above. The FFT estimator for the redshift bin z and the wavenumber bin A is as follows:

$$P_{\text{1D}}(A, z) = \left\langle [P_{\text{raw},s}(k) - P_{\text{noise},s}(k)] \cdot \mathbf{R}_s^{-2}(k) \right\rangle_{s \in z, k \in A} - P_{\text{metals}}(A), \quad (3.8)$$

where the ensemble average is calculated over all segments and wavenumbers that fall into the bin (z, A) , and $P_{\text{raw},s}$ is the direct Fourier transform of the transmitted flux fluctuation in the Ly α forest region, $P_{\text{noise},s}$ is the noise power spectrum estimated directly from the pipeline, and $\mathbf{R}_s$ is the Fourier transform of the resolution matrix for the segment s . Note that the wavenumbers are naturally in $\text{\AA}^{-1}$ units, while the power spectra are in $\text{\AA}$ units. Since our mocks have no metal contamination, the P_{metals} term is ignored in this work. Finally, our objective is to validate and evaluate the astrophysical and instrumental systematics correction terms in the final measurement of P_{1D} presented in ref. [24].

Before computing the FFT estimator, we apply an average signal-to-noise ratio cut $\overline{\text{SNR}} > 1$ per pixel, where the $\overline{\text{SNR}}$ is defined as the average over the considered part of the Ly α forest. To get thinner redshift bins without too much correlation, the δ_F 's are split into three consecutive and non-overlapping segments of equal length. We remove segments shorter than 75 spectrum pixels due to a cut in the UV region or the presence of a DLA. We also do not consider the Ly α segments with more than 120 masked spectrum pixels. We compute the FFT P_{1D} estimator by averaging (as shown with the $\langle \cdot \rangle$ operation in eq. (3.8)) over individual

Ly α segments. For convenience, we often refer to the term inside the averaging operator as an “individual power spectrum”, denoted $p_s(k)$ for the segment s and the wavenumber k , although this is not mathematically a power spectrum.

When averaging individual power spectra within a given redshift and wavenumber bin, we apply weights $w_{A,s} = 1/\sigma^2(A, s)$ based on each segment’s SNR as described in ref. [36]. Here, $\sigma^2(A, s)$ represents the variance associated with each segment s and wavenumber bin A , and is fitted as a function of the segment’s SNR. The eq. (3.8) is then calculated as a weighted average over all individual power spectra’s k wavenumbers that contribute to the wavenumber bin A :

$$\langle P \rangle(A, z) = \frac{\sum_{s \in z, k \in A} w_{A,s} p_s(k)}{\sum_{s \in z, k \in A} w_{A,s}}. \quad (3.9)$$

Contrary to ref. [36], we do not use the error bars obtained from the SNR weighting scheme but choose to estimate the statistical error bars with the following unbiased estimator:

$$\langle \sigma \rangle^2(A, z) = \left(\frac{(\sum_{s \in z} w'_{A,s})^2}{\sum_{s \in z} (w'_{A,s})^2} - 1 \right)^{-1} \left[\frac{\sum_{s \in z} (w'_{A,s})^2 P_{A,s}^2}{\sum_{s \in z} (w'_{A,s})^2} - \langle P_{A,s} \rangle_{s \in z}^2 \right], \quad (3.10)$$

where $P_{A,s} = \sum_{k \in A_s} p_s(k)/N_{A,s}$, A_s is the ensemble of wavenumbers associated with segment s that fall into bin A , $N_{A,s}$ is the number of wavenumbers in A_s , and $P_{A,s}$ is the average of individual power spectra in bin A . The weights used here are defined as $w'_{A,s} = N_{A,s} w_{A,s}$.

3.3.1 Estimating the covariance matrix

We develop a new, unbiased estimator for the full covariance matrix that accounts for the aforementioned SNR weighting scheme:

$$\begin{aligned} \langle C \rangle(A, B, z) = & \left(\frac{\sum_{s \in z} w'_{A,s} \sum_{s \in z} w'_{B,s}}{\sum_{s \in z} w'_{A,s} w'_{B,s}} - 1 \right)^{-1} \\ & \times \left[\frac{\sum_{s \in z} w'_{A,s} w'_{B,s} P_{A,s} P_{B,s}}{\sum_{s \in z} w'_{A,s} w'_{B,s}} - \langle P_{A,s} \rangle_{s \in z} \langle P_{B,s} \rangle_{s \in z} \right]. \end{aligned} \quad (3.11)$$

After ascribing a redshift bin to each segment, the averaging of individual power spectra, along with the computation of the variance and the covariance, is performed independently for each redshift bin.

3.4 Kolmogorov-Smirnov test

Our tests on mocks are based on calculating χ_i^2 for each realization i :

$$\chi_i^2 = (\mathbf{p}_i - \mathbf{p}_{\text{ref}})^T \mathbf{C}_i^{-1} (\mathbf{p}_i - \mathbf{p}_{\text{ref}}), \quad (3.12)$$

where $\mathbf{C}_i$ is the estimated covariance matrix for realization i , and $\mathbf{p}_{\text{ref}}$ is the reference power spectrum. The latter is the analytically calculated power spectrum for the optimal estimator, as both are naturally in velocity configuration [28]. However, the FFT estimator is constructed in wavelength configuration, and the conversion from velocity to wavelength configuration introduces slight biases. So, $\mathbf{p}_{\text{ref}}$ instead needs to be measured from mocks for the FFT measurement.

These χ^2 values are inherently random, drawn from an “empirical” distribution with an underlying cumulative distribution function (CDF): $F(\chi^2)$. We will compare this to a theoretical distribution using the Kolmogorov-Smirnov (KS) test. The KS test is a nonparametric test that does not require any assumptions about $F(\chi^2)$ when comparing it against other probability distributions. It can reject the chosen null hypothesis if the p-value is below a specified threshold. Typically, this threshold is set to 0.05, balancing the trade-off between false negatives and false positives. We show below that this remains a reasonable threshold for detecting a few percent biases with 95% efficiency.

Our theoretical distribution is a χ^2 probability distribution with the given number of degrees of freedom of the data vector, ν . Let us denote the CDF of this distribution as $G(\chi^2)$. We adopt the one-sided KS test with the following:

- *Null hypothesis*: Results are unbiased, and errors are correctly estimated or overestimated, $F(\chi^2) \geq G(\chi^2)$.
- *Alternative hypothesis*: Results are biased, or errors are underestimated, $F(\chi^2) < G(\chi^2)$ for some χ^2 .

We reject the null hypothesis (therefore fail the validation test) if the KS test yields a p-value less than 0.05. This is similar to goodness-of-fit tests, where higher p-values indicate a better fit. Due to computational time limitations in analysis, we generate 20 independent realizations, which limits the power of the KS test.

We first note that our p-value threshold of 0.05 sets an upper bound of 95% on the true-positive rate, even when the null hypothesis holds (p-values are uniformly distributed between zero and one in this case). We run numerical trials to further understand how effective this choice of p-value is against a possible global bias factor on covariance. We run 10,000 numerical trials, where we generate 20 random numbers from the $\chi^2(\nu = 1035)$ distribution in each trial. We find that, with these choices, we can rule out a 2.85% bias in 80% of trials and a 3.75% bias in 95% of trials. In other words, there remains at least 5% chance to falsely pass the validation test even though the distribution was biased by a few percent. A smaller p-value of 0.01 would reduce the chance of falsely rejecting the null hypothesis but would require a bias as large as 3.69% to rule it out in 80% of the trials. In contrast, a larger p-value of 0.1 can rule out a 2.41% bias in 80% of trials, but at the cost of a higher chance of falsely rejecting (failing) the null hypothesis (validation test). The most effective way to improve the KS test is to increase the number of mocks, which we leave to future work.

Lastly, we occasionally quote the two-sided KS test results to understand whether errors are overestimated. The null hypothesis in this case is $F(\chi^2) = G(\chi^2)$, i.e., the errors are correctly estimated. We can logically conclude that the errors are overestimated if the one-sided test passes and the two-sided test rejects the null hypothesis.

4 Validating the optimal estimator

We start our analysis by validating the optimal estimator. Because of its insensitivity to masking and consequently the survey window function, the list of tests simplifies to tests of uncontaminated data and fully contaminated data.




Figure 3. (*Left*) PDF of χ^2 -squared values for the no-systematics case from the optimal estimator. The Gaussian covariance matrix fails the KS test by yielding a p-value of 0.03, while the bootstrap covariance matrix passes it by yielding a p-value of 0.64. (*Right*) The ratio between the estimated P_{1D} from the stack of 20 mock realizations and the true power minus one. The ratios are shifted vertically for clarity. Dashed lines indicate the expected value of zero for each redshift bin. The area within 1% of the expected value is shaded. The fluctuations are between 0.5% and 1% for $z \leq 3.6$. Higher redshifts fluctuate more, so they are not shown for clarity.

Even though the optimal estimator can marginalize over the dominant continuum fitting errors, it cannot completely eliminate them. These errors add additional power to all k scales through higher-order correlations. We derive this bias for the optimal estimator and examine the effect of masked pixels on the induced bias.

4.1 Validation in the absence of contaminants

This section aims to (1) validate the estimators when no contaminants are present and the quasar continuum is perfectly known, (2) validate the covariance matrix, and (3) derive the biases resulting from the continuum fitting procedure. We generate and use 20 `quickquasars` mocks without any contaminants for this goal.

We use the underlying true continuum of quasars and mean flux to derive transmitted flux fluctuations, which bypasses the continuum fitting step. The purpose of this test is to confirm that our pipeline does not introduce any small numerical errors due to factors such as interpolation, weighting, and random noise. We run the quadratic estimator without marginalization, as no continuum fitting errors exist. However, in the case of real data with continuum marginalization, $k < 0.001 \text{ s km}^{-1}$ bins will become unusable due to continuum fitting, as we will discuss shortly. Additionally, high- k bins will be limited by the spectrograph resolution, which we quantify with respect to effective velocity spacing $R_z = c\Delta\lambda_{\text{DESI}}/(1+z)\lambda_{\text{Ly}\alpha}$ of each redshift bin, where $\Delta\lambda_{\text{DESI}} = 0.8 \text{ \AA}$. So, we limit the k modes to the following range: $0.001 \text{ s km}^{-1} \leq k \leq 0.75 k_{\text{Nyq}}(z)$, where the Nyquist frequency is defined as $k_{\text{Nyq}}(z) = \pi/R_z$ and increases with redshift. We perform the KS test described in section 3, and obtain p-values of 0.03 for the Gaussian covariance matrix and 0.64 for the bootstrap covariance matrix (as described in section 3.2.1), which is shown in figure 3. This indicates that the Gaussian covariance matrix does not fully capture the errors. Since the bootstrap matrix's p-value is greater than our 0.05 threshold, we consider this test passed by the optimal estimator.

We performed two additional checks. First, a two-sided KS test yields $p = 0.88$ for the bootstrap covariance matrix, which indicates that it is likely a good description of the true covariance matrix. Second, we stack the results from all 20 mock realizations by averaging P_{1D} estimates. The resulting, more precise P_{1D} follows the underlying true power, which is shown in the right panel of figure 3. This ratio reveals hints of subpercent-level biases in $3.8 \leq z \leq 4.4$, and a 5% bias at $z = 4.6$ due to edge effects, which is not shown in the figure for clarity. Correcting for these biases or inflating the covariance matrices to account for these errors yields p-values greater than 0.9 for both cases, but likely overcorrects or overestimates. We will revisit this point when we include more realistic features in the mock spectra below.

The second goal is to investigate the effect of continuum fitting. We fit each quasar with $a_q + b_q \Lambda$ polynomial and marginalize out polynomials of Λ from δ_F up to and including the second order. Besides contaminating low- k modes, the continuum fitting adds extra power to all scales through three-point correlations. This bias is strongly detected with a p-value of approximately zero. Therefore, we need to quantify and correct this bias in our measurements. We model this as a relation with respect to the underlying true power P_{true} such that $P_{\text{meas}}/P_{\text{true}} - 1 = b(k)$, where P_{meas} is the measured P_{1D} . To eliminate the statistical noise and possible systematic artifacts, we use the estimated P_{1D} with true continuum in the denominator. We fit $b(k)$ from the stack of 20 mocks with a first-order polynomial of $k/(0.009 \text{ s km}^{-1})$ between $0.001 \text{ s km}^{-1} \leq k \leq 0.75 \text{ k}_{\text{Nyq}}$ at each redshift bin. The resulting bias is $b(k) \sim 6 \times 10^{-3}$. Section 4.2 provides more details about the robustness of these bias estimates.

We correct each power spectrum estimate using the smooth power spectrum as $P_{\text{corr}} = P_{\text{meas}} - \frac{b(k)}{1+b(k)} P_{\text{smooth}}$ for each redshift bin. After this correction, the bootstrap covariance matrix passes the KS test with a p-value of 0.1, whereas the Gaussian covariance matrix still fails. We conclude the optimal estimator results with bias correction, and the bootstrap covariance matrix is validated in the uncontaminated case. Since these mocks do not include systematics realistically, this correction is not applied to the data. The correction that is applied to the data is derived in the next subsection.

4.2 Validation of masking HCDs, BALs, and atmospheric lines

The real data is contaminated by BAL features, HCDs, and atmospheric emission lines, and is complicated by the “survey window function” resulting from masking these features. The FFT estimator must correct for this window function (masking bias), whereas the quadratic estimator must confirm its robustness by construction. We generate 20 `quickquasars` mock realizations with BAL features and randomly placed DLAs to focus on the effect of masking specifically.

The optimal estimator method has two goals. The first goal is to confirm that masking these features does not introduce any biases. We calculate the transmitted flux fluctuations using the true continuum of each quasar to separate the effect of continuum fitting and mask all BAL features and HCDs at their exact locations.⁵ We further mask the regions around sky emission lines that will be masked in real data analysis. Then, we estimate P_{1D}

⁵As opposed to searching for their locations using a finder algorithm, which may produce incomplete, impure, and incorrect size estimates.




Figure 4. Fitted linear relation for $b(k)$ for redshifts bins with over 3σ detection significance for the optimal estimator. The solid lines are derived from fully contaminated mocks, and the dashed lines are derived from uncontaminated mocks. The $b(k)$ values for the contaminated mocks are larger than those for the uncontaminated case for all redshifts except $z = 3.8$. The difference arises from the continuum-fitting response to the underlying large-scale density field and the survey window function.

without any marginalization. We obtain p-values greater than 0.7 for both the Gaussian and bootstrap covariance matrices; therefore, the optimal estimator passes this test. The Gaussian matrix passing this test after previously failing may seem unexpected. However, this is due to weakening statistical power, as some pixels are “thrown out”, which results in diagonals of the Gaussian covariance matrix to be 30% larger on average.

The second goal is to investigate the effect of continuum fitting when these features are masked. We follow the same steps outlined in section 4.1. The p-values are again zero without any correction. After the bias correction, the bootstrap covariance matrix passes the KS test with a p-value of 0.83, whereas the Gaussian covariance matrix fails with a p-value of 0.03.

Figure 4 shows the fitted polynomial and its 1σ contours for the first nine redshift bins where the bias is detected at greater than 3σ significance. The solid lines and shaded regions denote $b(k)$ for the fully contaminated mock, whereas the dashed line is $b(k)$ derived from uncontaminated mocks discussed in section 4.1. It is important to note that the uncertainties reflect the statistical power of a 20 times larger dataset. The $b(k)$ values for the contaminated mocks are larger than those for the uncontaminated case for all redshifts except $z = 3.8$. The difference is caused by the continuum-fitting response to the underlying large-scale density field within the unmasked forest region. We conclude that the optimal estimator with this correction term (listed in table 1) and a bootstrap covariance matrix is accurate.

5 Validating the FFT estimator

In this section, we validate the FFT estimator. The KS test is performed over the same wavenumber range used in the EDR and DR1 measurements [24, 36], i.e., $0.0468 \text{ \AA}^{-1} < k < 2.0$

z	$b(k) \times 10^3$	z	$b(k) \times 10^3$	z	$b(k) \times 10^3$
2.2	$6.768 - 0.699x$	2.8	$5.611 - 0.098x$	3.4	$8.913 + 0.093x$
2.4	$5.736 - 0.692x$	3.0	$6.206 - 0.101x$	3.6	$10.825 + 0.172x$
2.6	$5.865 - 0.026x$	3.2	$8.335 - 0.089x$	3.8	$7.599 - 0.059x$

Table 1. Bias corrections $b(k)$ for the optimal estimator method that are derived from the stack of 20 fully contaminated mocks, where $x \equiv k/0.009 \text{ s km}^{-1}$.

$\text{\AA}^{-1}$. Since the small-number statistics in the higher redshift bins result in poor estimation of the covariance matrix and possible outliers, we limit the analysis to the redshift range $2.1 < z < 4.3$, as was done for the DR1 FFT measurement.

The bias correction for FFT is different from the optimal estimator strategy. We fit a second-order polynomial to P_{ref}/P_m obtained from the stack of all realizations, where m denotes the considered case, and correct each realization by directly multiplying it by this fit. Additionally, a systematic uncertainty is added to the covariance matrix used for χ^2 calculation in the KS test. This uncertainty is defined as a percentage of the relative correction on P_{1D} . For the previous EDR measurement [36], the percentage was equal to 30%. We use this percentage as the default for all the corrections tested in this article.⁶ We incorporate this more realistic level of systematic uncertainty into all the cases below as part of the KS test.

5.1 The power spectrum from transmission files

The FFT estimator works in wavelength configuration instead of velocity configuration; that is, it estimates P_{1D} in $\text{\AA}$ units instead of km s^{-1} units. In this section, we test the consistency of the two.

The transformation between the two configurations is non-linear in general: $v = c \ln \lambda / \lambda_p$, such that the velocity separation between two points is $v_2 - v_1 = \Delta v = c \ln \lambda_2 / \lambda_1$. For small separations, this is approximated as $\Delta v = c \frac{\Delta \lambda}{\lambda_c}$, where $\Delta \lambda \equiv \lambda_2 - \lambda_1$ and $\lambda_c = (\lambda_1 + \lambda_2)/2$. For a redshift bin z , the central wavelength becomes $\lambda_c = (1 + z)\lambda_{\text{Ly}\alpha}$. The conversion from $\text{\AA}$ units to km s^{-1} units is then as follows:

$$k_\lambda = \frac{c}{(1+z)\lambda_{\text{Ly}\alpha}} k_v \quad \text{and} \quad P_\lambda = \frac{(1+z)\lambda_{\text{Ly}\alpha}}{c} P_v. \quad (5.1)$$

We compute the mock transmission power spectrum (P_{trans}) and compare it to the input power spectrum (P_{input}) converted from velocity to angstrom units to assess the validity range of this conversion in the absence of other complications. Figure 5 shows the relative difference $P_{\text{trans}}/P_{\text{input}} - 1$. We omit the highest three redshift bins for clarity, since they yield large error bars. We find percent-level biases in wavenumbers below $k \leq 1 \text{ \AA}^{-1}$ across all redshift bins. However, such a difference is not observed in the DR1 data between the optimal and FFT estimators, which may be obscured by other systematics in real data.

This and other biases in the FFT estimator make it difficult to have a settled reference power spectrum for eq. (3.12). In contrast to the optimal estimator analysis, the reference power spectrum is adapted to the considered case to mitigate effects other than the one being investigated.

⁶Except for BAL masking, where a smaller percentage is sufficient to pass the test as discussed later.




Figure 5. The relative difference between the power spectrum from mock transmissions (P_{trans}) and the input power spectrum (P_{input}) for the FFT estimator. The highest three redshift bins yield large error bars and are omitted for clarity. We find that the conversion between wavelength and velocity units introduces small biases at $k \leq 1 \text{ \AA}^{-1}$ across all redshift bins as the small separations assumption gradually breaks down.




Figure 6. Comparison between the power spectrum measured on the transmission files and the one obtained using the true continuum imposed by `quickquasars` for the FFT estimator. Only a representative sample of redshift bins is shown for clarity. This ratio quantifies the misestimation of the resolution damping correction, which affects the smallest scales. We use this ratio as an additional correction for the resolution in the DR1 measurement [24].

5.2 Validation in the absence of contaminants

This section aims to (1) identify biases in the resolution treatment, (2) validate the covariance matrix, and (3) estimate the biases resulting from the continuum fitting procedure. As in section 4.1, we generate and use 20 `quickquasars` mocks without any contaminants for these goals.

We start by using the true continuum of quasars and mean flux as before. Using the above P_{trans} as the reference power spectrum, we investigate if any bias remains in the stack of all 20 mocks. Figure 6 shows the ratio between the two power spectra. We find a bias at small scales, which may be due to the Fourier treatment of the resolution matrix or an issue with noise subtraction. We remind the reader that `quickquasars` mocks do not involve complicated spectro-perfectionism reduction and instead provide an average resolution matrix for simulated cameras. We also note that the bias changes by 0.2% on average across all redshifts (wavelengths), which indicates that this bias is due to the Fourier treatment of the resolution matrix.




Figure 7. (*Left*) PDF of chi-square values for the no systematics case for the FFT estimator. The one-sided KS test indicates that the covariance matrix estimator passes the validation threshold. (*Right*) PDF of χ^2 values before (*orange*) and after (*blue*) continuum bias correction and resulting one-sided KS test p-value.

To achieve our second goal, we set our reference power spectrum (P_{ref} in eq. (3.12)) as the stack of all 20 mocks. Since this reference power has the same resolution bias as the individual P_{1D} estimates, the net contribution of resolution bias is removed from this test. To test the estimated covariance matrix of each realization, we compare the estimated P_{1D} of this realization to the reference power spectrum and calculate χ^2 using the covariance matrix derived for each realization (the errors in P_{ref} are ignored). Using the covariance from eq. (3.11) results in a one-sided p-value of one and a two-sided KS test p-value of zero. This indicates that the covariance is overestimated. We renormalized the covariance by a 10% factor and show the resulting χ^2 distribution in figure 7. We obtain a one-sided p-value of 0.85 and a two-sided p-value of 0.27, which indicates that the χ^2 distribution passes the KS test and is closer to the expected Gaussian distribution. We note that the relatively low value of the two-sided p-value is caused by the shape of the distribution rather than an amplitude offset in the covariance.

To estimate biases arising from continuum fitting, we use the same reference power spectrum, removing other possible causes that could explain the change in P_{1D} while fitting the quasar continuum, such as spectrum pixelization, resolution modeling, and noise, since these affect both true continuum and fitted continuum estimates in the same manner. Figure 8 shows the impact of continuum fitting on P_{1D} . This is a small effect (1% at $z \leq 3.2$ and up to 2% at $z \geq 3.4$) and in agreement with conclusions from the optimal estimator and findings from EDR analysis in ref. [36].

We include a systematic error budget in the covariance matrix, as discussed before, to perform the KS test. The KS test with our correction and the additional error budget passes the p-value threshold, increasing it from zero for the uncorrected case to one, which is shown in the right panel of figure 7. This is expected, considering that we are dealing with a relatively small effect.




Figure 8. Impact of continuum fitting on P_{1D} for the FFT estimator. Comparison between the average of 20 mocks with true continuum (P_{ref}) and fitted continuum (P_{cont}). Only a representative sample of redshift bins is shown for clarity. We fit a second-order polynomial for $k > 0.1 \text{ \AA}^{-1}$ that takes the error bar of each point into account.

5.3 Deriving and validating masking corrections

The FFT estimator requires data points to be regularly spaced. We mask pixels by setting the value of the Ly α contrast δ_F to zero. Since all masked pixels are taken into account when computing P_{1D} , masking introduces a bias that must be corrected. This bias is effectively a survey window function that convolves all Fourier modes. Here, we approximate this effect with a scale-dependent bias in Fourier space as was done in previous Ly α P_{1D} estimates with FFT. Our baseline approach follows refs. [36, 42], i.e. we use the FFT estimator to compute the ratio between the masked and unmasked power spectra of **quickquasars** mocks. For each redshift bin, this ratio is fit with an appropriate polynomial to derive a correction that we will multiplicatively apply to the data in a later step. In this section, the masking corrections for HCD, atmospheric lines, and BAL are computed and tested. Additionally, we test the multiplicativity of these masking corrections when they are combined with the continuum fitting correction defined in previous sections.

5.3.1 HCD masking

The presence of non-IGM components due to HCDs, in addition to the Ly α forest absorption due to neutral hydrogen in the IGM, significantly impacts the measured P_{1D} . These HCDs produce deep and wide absorption in the Ly α forest with extended damping wings. These features occur when the sightline to a quasar passes through the circumgalactic environment of an intervening galaxy. The contamination by these systems is a major source of systematic uncertainty in the estimation and cosmological interpretation of P_{1D} , especially for small values of k (e.g., see ref. [40]).

To assess the impact of these effects, as well as the overall impact on P_{1D} we rely again on 20 **quickquasars** mocks. The effect of pure masking is tested by masking the DLA regions using their exact locations and column densities in *uncontaminated* mocks. This isolates the effect solely to masking by removing the complications that come from a mismatch in the sizes of DLA regions estimated using column densities. We perform continuum fitting correction with and without masking HCDs and compare their ratio. The resulting bias is below the percent level and nearly featureless (constant in k). Nevertheless, when uncorrected, the




Figure 9. Impact of atmospheric emission line masking on the one-dimensional power spectrum for the FFT estimator. The impact is quantified by the ratio of the reference power spectrum (P_{ref}) to the atmospheric emission line masked power spectrum (P_{sky}). The fitted solid lines are used as corrections for the bias induced by atmospheric emission line masking.

estimates yield a p-value of zero, failing the KS test. We pass the KS test with a p-value of one by applying the HCD correction.

5.3.2 Discrete line masking

We mask the atmospheric emission and Galactic absorption lines identified in ref. [36]. These features occur at fixed observed wavelengths across all spectra, so that they introduce a larger bias through these correlations.

In this section, we improve the precision of our bias correction using the new, larger set of mocks tailored for DESI DR1. As before, we calculate the ratios $P_{\text{sky}}/P_{\text{ref}}$ and fit a second-order polynomial to the stack of all realizations. These are shown in figure 9. The largest corrections occur at redshift bins $z = 2.2$ and $z = 2.6$ due to sparse atmospheric line emissions and Galactic absorption due to the Ca II H&K doublet lines at 3968 Å and 3933 Å. For $z \geq 2.8$, the corrections are only caused by atmospheric emission line masking.

The left panel of figure 11 shows the χ^2 distribution for the KS test, which gives a p-value of 0.90. This test indicates that the correction derived is well-suited, even with a significant impact on the power spectrum. It also indicates that the systematic uncertainty associated with this correction (which is set to 30 % of the correction itself [24], included in the covariance matrix) is properly estimated.

5.3.3 BAL masking

The companion paper [24] performs a variation of the FFT analysis that shows that BALs detected with Absorption Index (AI , see ref. [43] for details) cause a significant deviation at large scales and should be accounted for. Here, we derive a correction associated with the masking of BAL systems with $AI > 0$ by comparing P_{1D} when features are masked to the unmasked case. This is shown in figure 10.

For the KS test in this case, we decrease the associated relative systematic uncertainty caused by the BAL AI masking correction to 6%, as the default value (30%) is a significant overestimate according to the χ^2 distribution. With this new uncertainty, we consider the obtained one-sided p-value of 0.45 (shown in the left panel of figure 11) sufficient to validate the correction and its systematic uncertainty.




Figure 10. Impact of BAL *AI* masking on the one-dimensional power spectrum for the FFT estimator. The impact is quantified by the ratio of the reference power spectrum (P_{ref}) to the BAL *AI* masked power spectrum (P_{BAL}). Only a representative sample of redshift bins is shown for clarity. The fitted solid lines are used as corrections for the bias induced by BAL *AI* masking.




Figure 11. (Left) PDFs of χ^2 values for the FFT estimator after atmospheric (sky) line masking correction (blue) and BAL *AI* masking correction (orange). The KS test indicates that the covariance matrix estimator with added systematic error budget passes the validation threshold. (Right) PDF of χ^2 values after the combined effect of DLA, BAL *AI*, and sky masking is corrected by multiplication of individual effects. The KS test shows that these corrections can be multiplied.

5.3.4 Combining corrections

The FFT estimator requires several corrections associated with continuum fitting, DLA, BAL *AI*, and atmospheric line masking. The main measurement assumes the multiplicativity of these corrections, i.e., if all effects are applied simultaneously, the full correction is the product of individual corrections. However, this assumption does not account for potential correlations between the effects. Masked regions of different effects can significantly overlap, resulting in an overestimation when the product of corrections is applied. Additionally, some masking corrections can differ with and without continuum fitting.

In order to test the multiplicativity assumption, we create a set of mocks with continuum fitting for which the three masking effects are applied. We correct the P_{1D} of each mock by the product of individual corrections. Furthermore, we include a systematic uncertainty that is the quadratic sum of the previously described individual uncertainties. The corrected

measurement is compared to the one using the true continuum of quasars, and the results of the KS test are shown in the right panel of figure 11. Without changing the percentage associated with systematics, the one-sided p-value of one indicates that the assumption of correction multiplicativity is verified as correct. The lower χ^2 values compared to expectation indicate that the whole combination of systematic uncertainties is likely overestimated.

6 Accuracy of the spectrograph resolution

Uncertainty in the resolution correction dominates the systematic error budget for P_{1D} at large wavenumbers. When the bias in the spectrograph resolution (b_{res}) is small, the scale-dependent error on P_{1D} can be written as $b_{\text{res}} \times 2k^2 R_z^2 \times P_{\text{true}}(k, z)$, where $R_z \equiv c\Delta\lambda_{\text{DESI}}/(1+z)\lambda_{\text{Ly}\alpha}$ [27]. Our goal is to quantify these systematic errors on P_{1D} using CCD image simulations processed with the DESI spectroscopic pipeline (described in section 2.2). In a previous work, we validated the spectro-perfectionism algorithm and its resolution matrix in P_{1D} estimation when the PSF is perfectly known [27]. We also showed that there is a wavelength-dependent bias of $b_{\text{res}}(z) \sim 1\%$ in the spectrograph resolution when the PSF is empirically measured from arc calibration lamps. Our focus here is on a better quantification of $b_{\text{res}}(z)$.

We simulate 675,000 spectra, which is fifteen times the number used in our EDR analysis, and extract them using empirical PSFs in the same manner as DESI’s nightly operations, using the same version of the pipeline used for DR1 processing. The relative errors of both estimators with respect to the reference power spectrum (analytically calculated for the optimal estimator and measured from the transmission files for the FFT estimator) are shown in figure 12, where the blue points stand for the FFT results and the red triangles for the optimal estimator results. We remind the reader that we generated these simulations with a maximum quasar redshift of $z_{\text{qso}}^{\text{max}} = 3.6$, corresponding to a maximum forest redshift bin of $z_F^{\text{max}} = 3.4$. The two estimators generally display the same trend in error, except for a few key differences. First, they show different error trends at high k for $z = 2.2$ and 2.4 . Such differences are normal, as each estimator has a unique treatment for the resolution correction. Second, the FFT estimator has a larger bias at low k than the optimal estimator. This large-scale bias is unexpected and is unlikely to originate from a purely resolution effect. We do not include this large-scale offset or remove its contribution at high k by keeping the k^2 functional form without additive terms. The $b_{\text{res}}(z)$ are derived in each redshift bin independently and noted in the figure.

Our recommendation is to include these biases in the systematic error budget as correlated fluctuations, rather than correcting for them in the measurement. First, the simulation fidelity needs to be improved, and the large-scale offset must be understood to ensure the precision of such corrections. Second, these biases partially arise from errors in estimating pipeline noise. Later versions of the spectroscopic pipeline include minor changes in this area. We reprocess the simulated spectra with the latest version of the pipeline available and confirm that these changes alter, but do not eliminate, the high k bias. This is an ongoing effort.

The offset at large scales remains in both versions of the spectroscopic pipeline. Various diagnostic quantities do not indicate any problems with the pipeline, such as stacked flux for measuring flux calibration errors and estimating power in spectral regions with no power (null test). We also do not observe a 5% difference in real data between the two estimators




Figure 12. Systematic errors due to the spectral resolution derived from CCD image simulations for both estimators. As we have shown previously in ref. [27] for the optimal estimator, the bias depends on the wavelength, changes direction, and corresponds to 1 – 2% precision. The bias of the FFT estimator seems consistent in direction across wavelengths and agrees with the optimal estimator results except for $z = 3.0$. We do not correct for this bias but instead include it in the systematic error budget as a correlated fluctuation.

at $z = 3.0$, which makes us assume that this is an error in the simulation stage rather than the extraction stage. This will also be investigated in the future.

7 Discussion

We have validated the P_{1D} measurement pipelines for DESI DR1 against major systematics such as masking contaminated pixels (survey window function), biases arising from continuum fitting errors, and the spectroscopic pipeline for both the optimal and FFT estimators. We generated two classes of synthetic spectra.

The first class is based on 1D Ly α transmission fields and includes a diverse set of quasar spectra and simulates the DESI instrument in 1D. We generated 20 realizations of DR1 and used these spectra to (1) validate that the optimal estimator is robust against masked pixels and (2) derive masking bias corrections for the FFT estimator and validate their application. Additionally, we calculated the continuum fitting error bias for both estimators.

The second class of synthetic spectra projects quasar spectra onto realistic, 2D CCD images that simulate raw DESI data. These are then processed with the DESI spectroscopic

reduction pipeline. Using these CCD image simulations, we quantified the accuracy of the spectrograph resolution to be around 1%. The result of all of these tests is that the DESI P_{1D} measurement pipeline is robust and understood at the percent level.

Given the increased number and signal-to-noise ratio of spectra from ongoing DESI observations, it is important to continue improving synthetic data generation to take full advantage of the data’s increased statistical power. First, the mock power spectrum is only tuned to match the observations with 10–20% agreement on average. Improving the fidelity of the mocks to the data will improve the accuracy of the various corrections presented in this article, and enable end-to-end P_{1D} validation analyses that include cosmological inference in future pipelines.

Relatedly, HCDs are randomly placed into our mocks, and their column density distribution is based on measurements from smaller samples than are currently available. The masked regions at HCD locations are correlated with the Ly α forest, which introduces an additional bias for both estimators in real data. In a future study, these systems can be placed with an abundance matching formalism to include forest-HCD correlations and study the cosmological implications of HCDs beyond masked pixels. This 1D formalism would be significantly cheaper than realistic hydrodynamical simulations. Furthermore, in real data analysis, only the densest systems, DLAs with column density $N_{\text{HI}} > 10^{20.3} \text{ cm}^{-2}$, can be identified with 85% completeness and purity depending on the confidence level and SNR cuts [59–61]. This means all weak HCDs ($N_{\text{HI}} < 10^{20.3} \text{ cm}^{-2}$), as well as missed DLAs, are left unmasked in the data. They impact P_{1D} through their intrinsic average line profile (one-halo effect) and their clustering (two-halo effect), since they are biased tracers of cosmic structure. For similar reasons, falsely detected DLAs may also bias the estimated P_{1D} . As the DLA finder is expected to find false positives preferentially in regions of large absorption, the associated mask is correlated with the Ly α forest. We leave the study of these effects to future work.

None of our simulations includes masking due to cosmic rays. However, the statistics of such pixels are minor and are not expected to alter the measurement at levels comparable to those of atmospheric lines or other astrophysical contaminants. Finally, we note that FFT masking can be improved by applying analytic window correction methods, as in ref. [62], thereby drastically reducing the systematic error introduced by masking. This improvement will be performed in future studies.

Software. The optimal estimator⁷ is written in `c++`. It depends on `CBLAS` and `LAPACKE` routines for matrix/vector operations, `GSL`⁸ for certain scientific calculations [63], `FFTW3`⁹ for FFTs [64]; and uses the Message Passing Interface (MPI) standard¹⁰ to parallelize tasks. The FFT estimator and associated plots are generated with `picca`⁴ (9.12.3) and `p1desi`¹¹ (2.0.1) software. The quasar spectra are organized with the `HEALPix` [65] scheme on the sky. We use the following commonly-used software in `python` analysis: `astropy`¹² a community-developed

⁷<https://github.com/p-slash/lyspeq>

⁸<https://www.gnu.org/software/gsl>

⁹<https://fftw.org>

¹⁰<https://www.mpich.org>

¹¹<https://github.com/corentinravoux/p1desi>

¹²<https://www.astropy.org>

core `python` package for Astronomy [66–68], `numpy`¹³ an open source project aiming to enable numerical computing with `python` [69], `scipy`¹⁴ an open-source project with algorithms for scientific computing. `healpy` to interface with `HEALPix` in `python` [70], Finally, we make plots using `matplotlib`¹⁵ a comprehensive library for creating static, animated, and interactive visualizations in `python` [71].

Acknowledgments

NGK and PM acknowledge support from the United States Department of Energy, Office of High Energy Physics under Award Number DE-SC0011726. EA acknowledges support from the Agence Nationale de la Recherche (ANR) through grant ANR-22-CE92-0037.

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.

¹³<https://numpy.org>

¹⁴<https://scipy.org>

¹⁵<https://matplotlib.org>

Data Availability Statement. This article has associated data in a data repository. All data points shown in the figures are available on the following website: <https://doi.org/10.5281/zenodo.18672780>

References

- [1] E. Boera, G.D. Becker, J.S. Bolton and F. Nasir, *Revealing Reionization with the Thermal History of the Intergalactic Medium: New Constraints from the Ly α Flux Power Spectrum*, *Astrophys. J.* **872** (2019) 101 [[arXiv:1809.06980](#)] [[INSPIRE](#)].
- [2] M. Walther, J. Oñorbe, J.F. Hennawi and Z. Lukić, *New Constraints on IGM Thermal Evolution from the Ly α Forest Power Spectrum*, *Astrophys. J.* **872** (2019) 13 [[arXiv:1808.04367](#)] [[INSPIRE](#)].
- [3] B. Villaseñor, B. Robertson, P. Madau and E. Schneider, *Inferring the Thermal History of the Intergalactic Medium from the Properties of the Hydrogen and Helium Ly α Forest*, *Astrophys. J.* **933** (2022) 59 [[arXiv:2111.00019](#)] [[INSPIRE](#)].
- [4] R.A.C. Croft, D.H. Weinberg, N. Katz and L. Hernquist, *Recovery of the power spectrum of mass fluctuations from observations of the Lyman alpha forest*, *Astrophys. J.* **495** (1998) 44 [[astro-ph/9708018](#)] [[INSPIRE](#)].
- [5] P. McDonald et al., *The Observed probability distribution function, power spectrum, and correlation function of the transmitted flux in the Lyman-alpha forest*, *Astrophys. J.* **543** (2000) 1 [[astro-ph/9911196](#)] [[INSPIRE](#)].
- [6] C. Pedersen et al., *An emulator for the Lyman- α forest in beyond- Λ CDM cosmologies*, *JCAP* **05** (2021) 033 [[arXiv:2011.15127](#)] [[INSPIRE](#)].
- [7] M. Viel, J. Weller and M. Haehnelt, *Constraints on the primordial power spectrum from high resolution Lyman-alpha forest spectra and WMAP*, *Mon. Not. Roy. Astron. Soc.* **355** (2004) L23 [[astro-ph/0407294](#)] [[INSPIRE](#)].
- [8] R.A.C. Croft, W. Hu and R. Dave, *Cosmological Limits on the Neutrino Mass from the Ly α Forest*, *Phys. Rev. Lett.* **83** (1999) 1092 [[astro-ph/9903335](#)] [[INSPIRE](#)].
- [9] N. Palanque-Delabrouille et al., *Constraint on neutrino masses from SDSS-III/BOSS Ly α forest and other cosmological probes*, *JCAP* **02** (2015) 045 [[arXiv:1410.7244](#)] [[INSPIRE](#)].
- [10] C. Yèche, N. Palanque-Delabrouille, J. Baur and H. du Mas des Bourboux, *Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100*, *JCAP* **06** (2017) 047 [[arXiv:1702.03314](#)] [[INSPIRE](#)].
- [11] V.K. Narayanan, D.N. Spergel, R. Dave and C.-P. Ma, *Constraints on the mass of warm dark matter particles and the shape of the linear power spectrum from the Ly α forest*, *Astrophys. J. Lett.* **543** (2000) L103 [[astro-ph/0005095](#)] [[INSPIRE](#)].
- [12] U. Seljak, A. Makarov, P. McDonald and H. Trac, *Can sterile neutrinos be the dark matter?*, *Phys. Rev. Lett.* **97** (2006) 191303 [[astro-ph/0602430](#)] [[INSPIRE](#)].
- [13] M.-Y. Wang et al., *Lyman- α forest constraints on decaying dark matter*, *Phys. Rev. D* **88** (2013) 123515 [[arXiv:1309.7354](#)] [[INSPIRE](#)].
- [14] V. Iršič et al., *First constraints on fuzzy dark matter from Lyman- α forest data and hydrodynamical simulations*, *Phys. Rev. Lett.* **119** (2017) 031302 [[arXiv:1703.04683](#)] [[INSPIRE](#)].
- [15] A. Boyarsky, J. Lesgourgues, O. Ruchayskiy and M. Viel, *Lyman-alpha constraints on warm and on warm-plus-cold dark matter models*, *JCAP* **05** (2009) 012 [[arXiv:0812.0010](#)] [[INSPIRE](#)].

- [16] M. Viel, G.D. Becker, J.S. Bolton and M.G. Haehnelt, *Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data*, *Phys. Rev. D* **88** (2013) 043502 [[arXiv:1306.2314](#)] [[INSPIRE](#)].
- [17] J. Baur et al., *Lyman-alpha Forests cool Warm Dark Matter*, *JCAP* **08** (2016) 012 [[arXiv:1512.01981](#)] [[INSPIRE](#)].
- [18] V. Iršič et al., *New Constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman- α forest data*, *Phys. Rev. D* **96** (2017) 023522 [[arXiv:1702.01764](#)] [[INSPIRE](#)].
- [19] B. Villasenor, B. Robertson, P. Madau and E. Schneider, *New constraints on warm dark matter from the Lyman- α forest power spectrum*, *Phys. Rev. D* **108** (2023) 023502 [[arXiv:2209.14220](#)] [[INSPIRE](#)].
- [20] DESI collaboration, *The DESI Experiment, a whitepaper for Snowmass 2013*, [arXiv:1308.0847](#) [[INSPIRE](#)].
- [21] DESI collaboration, *The DESI Experiment Part I: Science, Targeting, and Survey Design*, [arXiv:1611.00036](#) [[INSPIRE](#)].
- [22] DESI collaboration, *Data Release 1 of the Dark Energy Spectroscopic Instrument*, [arXiv:2503.14745](#) [[INSPIRE](#)].
- [23] EBOSS collaboration, *The SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data*, *Astron. J.* **151** (2016) 44 [[arXiv:1508.04473](#)] [[INSPIRE](#)].
- [24] C. Ravoux et al., *DESI DR1 Ly α 1D power spectrum: the Fast Fourier Transform estimator measurement*, *JCAP* **11** (2025) 079 [[arXiv:2505.09493](#)] [[INSPIRE](#)].
- [25] N.G. Karaçaylı et al., *DESI DR1 Ly α 1D power spectrum: the optimal estimator measurement*, *JCAP* **10** (2025) 004 [[arXiv:2505.07974](#)] [[INSPIRE](#)].
- [26] SDSS collaboration, *The Lyman-alpha forest power spectrum from the Sloan Digital Sky Survey*, *Astrophys. J. Suppl.* **163** (2006) 80 [[astro-ph/0405013](#)] [[INSPIRE](#)].
- [27] N.G. Karaçaylı et al., *Optimal 1D Ly α Forest Power Spectrum Estimation – III. DESI early data*, *Mon. Not. Roy. Astron. Soc.* **528** (2024) 3941 [[arXiv:2306.06316](#)] [[INSPIRE](#)].
- [28] N.G. Karaçaylı, A. Font-Ribera and N. Padmanabhan, *Optimal 1D Ly α forest power spectrum estimation – I. DESI-lite spectra*, *Mon. Not. Roy. Astron. Soc.* **497** (2020) 4742 [[arXiv:2008.06421](#)] [[INSPIRE](#)].
- [29] A.S. Bolton and D.J. Schlegel, *Spectro-Perfectionism: An Algorithmic Framework for Photon Noise-Limited Extraction of Optical Fiber Spectroscopy*, *Publ. Astron. Soc. Pac.* **122** (2010) 248 [[arXiv:0911.2689](#)] [[INSPIRE](#)].
- [30] DESI collaboration, *The Spectroscopic Data Processing Pipeline for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **165** (2023) 144 [[arXiv:2209.14482](#)] [[INSPIRE](#)].
- [31] DESI collaboration, *The Early Data Release of the Dark Energy Spectroscopic Instrument*, *Astron. J.* **168** (2024) 58 [[arXiv:2306.06308](#)] [[INSPIRE](#)].
- [32] H.G. Bi, G. Boerner and Y. Chu, *An alternative model for the Ly-alpha absorption forest*, *Astron. Astrophys.* **266** (1992) 1.
- [33] I. McGreer, J. Moustakas and J. Schindler, *simqso: Simulated quasar spectra generator*, *Astrophysics Source Code Library*, record ascl:2106.008, (2021).
- [34] D. Kirkby, T. Robitaille, Julienguy, E. Tollerud, B. Sipőcz, M. Droettboom et al., *desihub/specsim: August 2020 Release*, (2021), <https://zenodo.org/records/4566008>.

- [35] H.K. Herrera-Alcantar et al., *Synthetic spectra for Lyman- α forest analysis in the Dark Energy Spectroscopic Instrument*, *JCAP* **01** (2025) 141 [[arXiv:2401.00303](#)] [[INSPIRE](#)].
- [36] C. Ravoux et al., *The Dark Energy Spectroscopic Instrument: one-dimensional power spectrum from first Ly α forest samples with Fast Fourier Transform*, *Mon. Not. Roy. Astron. Soc.* **526** (2023) 5118 [[arXiv:2306.06311](#)] [[INSPIRE](#)].
- [37] J.X. Prochaska, J.F. Hennawi and S. Herbert-Fort, *The SDSS-DR5 Survey for Proximate Damped Ly-alpha Systems*, *Astrophys. J.* **675** (2008) 1002 [[astro-ph/0703594](#)] [[INSPIRE](#)].
- [38] J.X. Prochaska, P. Madau, J.M. O'Meara and M. Fumagalli, *Towards a unified description of the intergalactic medium at redshift $z \approx 2.5$* , *Mon. Not. Roy. Astron. Soc.* **438** (2014) 476 [[arXiv:1310.0052](#)] [[INSPIRE](#)].
- [39] P. McDonald et al., *Physical effects on the Ly-alpha forest flux power spectrum: Damping wings, ionizing radiation fluctuations, and galactic winds*, *Mon. Not. Roy. Astron. Soc.* **360** (2005) 1471 [[astro-ph/0407378](#)] [[INSPIRE](#)].
- [40] K.K. Rogers et al., *Simulating the effect of high column density absorbers on the one-dimensional Lyman α forest flux power spectrum*, *Mon. Not. Roy. Astron. Soc.* **474** (2018) 3032 [[arXiv:1706.08532](#)] [[INSPIRE](#)].
- [41] BOSS collaboration, *The one-dimensional Ly-alpha forest power spectrum from BOSS*, *Astron. Astrophys.* **559** (2013) A85 [[arXiv:1306.5896](#)] [[INSPIRE](#)].
- [42] EBOSS collaboration, *The one-dimensional power spectrum from the SDSS DR14 Ly α forests*, *JCAP* **07** (2019) 017 [[arXiv:1812.03554](#)] [[INSPIRE](#)].
- [43] S. Filbert et al., *Broad absorption line quasars in the Dark Energy Spectroscopic Instrument Early Data Release*, *Mon. Not. Roy. Astron. Soc.* **532** (2024) 3669 [[arXiv:2309.03434](#)] [[INSPIRE](#)].
- [44] P. Martini et al., *Validation of the DESI 2024 Lyman alpha forest BAL masking strategy*, *JCAP* **01** (2025) 137 [[arXiv:2405.09737](#)] [[INSPIRE](#)].
- [45] DESI collaboration, *Validation of the DESI 2024 Ly α forest BAO analysis using synthetic datasets*, *JCAP* **01** (2025) 148 [[arXiv:2404.03004](#)] [[INSPIRE](#)].
- [46] EBOSS collaboration, *The Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations with Ly α Forests*, *Astrophys. J.* **901** (2020) 153 [[arXiv:2007.08995](#)] [[INSPIRE](#)].
- [47] DESI collaboration, *The Lyman- α forest catalog from the Dark Energy Spectroscopic Instrument Early Data Release*, *Mon. Not. Roy. Astron. Soc.* **528** (2024) 6666 [[arXiv:2306.06312](#)] [[INSPIRE](#)].
- [48] N.G. Karaçaylı, *QSONic: fast quasar continuum fitting*, *J. Open Source Softw.* **9** (2024) 6373.
- [49] A.J.S. Hamilton, *Towards optimal measurement of power spectra I: minimum variance pair weighting and the fisher matrix*, *Mon. Not. Roy. Astron. Soc.* **289** (1997) 285 [[astro-ph/9701008](#)] [[INSPIRE](#)].
- [50] M. Tegmark, A. Taylor and A. Heavens, *Karhunen-Loeve eigenvalue problems in cosmology: How should we tackle large data sets?*, *Astrophys. J.* **480** (1997) 22 [[astro-ph/9603021](#)] [[INSPIRE](#)].
- [51] M. Tegmark et al., *Measuring the galaxy power spectrum with future redshift surveys*, *Astrophys. J.* **499** (1998) 555 [[astro-ph/9708020](#)] [[INSPIRE](#)].
- [52] U. Seljak, *Weak lensing reconstruction and power spectrum estimation: minimum variance methods*, *Astrophys. J.* **506** (1998) 64 [[astro-ph/9711124](#)] [[INSPIRE](#)].

- [53] N.G. Karaçaylı et al., *Optimal 1D Ly α forest power spectrum estimation – II. KODIAQ, SQUAD, and XQ-100*, *Mon. Not. Roy. Astron. Soc.* **509** (2022) 2842 [[arXiv:2108.10870](#)] [[INSPIRE](#)].
- [54] J.A. Hanley and B. MacGibbon, *Creating non-parametric bootstrap samples using Poisson frequencies*, *Computer Methods and Programs in Biomedicine* **83** (2006) 57.
- [55] N. Chamandy, O. Muralidharan, A. Najmi and S. Naidu, *Estimating uncertainty for massive data streams*, Tech. Rep. Google, 1600 Amphitheatre Pkwy, Building 40, Mountain View, CA 94043 (2012).
- [56] P. DIERCKX, *An Algorithm for Surface-Fitting with Spline Functions*, *IMA J. Numer. Anal.* **1** (1981) 267.
- [57] P. Dierckx, *Curve and Surface Fitting with Splines*, Monographs on Numerical Analysis, Clarendon Press, (1995).
- [58] P. McDonald, *Toward a measurement of the cosmological geometry at Z 2: predicting lyman-alpha forest correlation in three dimensions, and the potential of future data sets*, *Astrophys. J.* **585** (2003) 34 [[astro-ph/0108064](#)] [[INSPIRE](#)].
- [59] B. Wang et al., *Deep Learning of Dark Energy Spectroscopic Instrument Mock Spectra to Find Damped Ly α Systems*, *Astrophys. J. Supp.* **259** (2022) 28 [[arXiv:2201.00827](#)] [[INSPIRE](#)].
- [60] M.-F. Ho, S. Bird and R. Garnett, *Damped Lyman- α absorbers from Sloan digital sky survey DR16Q with Gaussian processes*, *Mon. Not. Roy. Astron. Soc.* **507** (2021) 704 [[arXiv:2103.10964](#)].
- [61] DESI collaboration, *Construction of the damped Ly α absorber catalog for DESI DR2 Ly α BAO*, *Phys. Rev. D* **112** (2025) 083510 [[arXiv:2503.14740](#)] [[INSPIRE](#)].
- [62] M. Lokken, A. Font-Ribera and P. McDonald, *Weighted FFT estimators for 1D and 3D correlations of the Lyman- α forest*, *JCAP* **11** (2025) 070 [[arXiv:2505.02904](#)] [[INSPIRE](#)].
- [63] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken et al. *GNU scientific library: Reference manual (release 2.7)*, <https://www.gnu.org/software/gsl/>, (2021).
- [64] M. Frigo and S.G. Johnson, *The Design and Implementation of FFTW3*, *IEEE Proc.* **93** (2005) 216 [[INSPIRE](#)].
- [65] K.M. Górski et al., *HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[INSPIRE](#)].
- [66] ASTROPY collaboration, *Astropy: A Community Python Package for Astronomy*, *Astron. Astrophys.* **558** (2013) A33 [[arXiv:1307.6212](#)] [[INSPIRE](#)].
- [67] ASTROPY collaboration, *The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package*, *Astron. J.* **156** (2018) 123 [[arXiv:1801.02634](#)] [[INSPIRE](#)].
- [68] ASTROPY collaboration, *The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package*, *Astrophys. J.* **935** (2022) 167 [[arXiv:2206.14220](#)] [[INSPIRE](#)].
- [69] C.R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357 [[arXiv:2006.10256](#)] [[INSPIRE](#)].
- [70] A. Zonca et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python*, *J. Open Source Softw.* **4** (2019) 1298 [[INSPIRE](#)].
- [71] J.D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9** (2007) 90 [[INSPIRE](#)].

Author List

Naim Göksel Karaçaylı^{1,2,3}, Corentin Ravoux⁴, Paul Martini^{1,2}, Jean-Marc Le Goff⁵,
 Eric Armengaud⁵, M. Abdul-Karim⁵, J. Aguilar⁶, S. Ahlen⁷, A. Anand⁶, S. BenZvi⁸,
 D. Bianchi^{9,10}, D. Brooks¹¹, T. Claybaugh⁶, A. Cuceu^{6,12}, A. de la Macorra¹³,
 Biprateep Dey^{14,15}, P. Doel¹¹, S. Ferraro^{6,16}, A. Font-Ribera¹⁷, J. E. Forero-Romero^{18,19},
 E. Gaztañaga^{20,21,22}, S. Gontcho A Gontcho^{6,23}, G. Gutierrez²⁴, H. K. Herrera-Alcantar^{25,5},
 K. Honscheid^{1,3}, M. Ishak²⁶, J. Jimenez¹⁷, R. Joyce²⁷, D. Kirkby²⁸, T. Kisner⁶,
 A. Kremin⁶, O. Lahav¹¹, M. Landriau⁶, L. Le Guillou²⁹, M. Manera^{30,17}, A. Meisner²⁷,
 R. Miquel^{31,17}, S. Nadathur²¹, G. Niz^{32,33}, N. Palanque-Delabrouille^{5,6},
 W. J. Percival^{34,35,36}, C. Poppett^{6,37,16}, F. Prada³⁸, I. Pérez-Ràfols³⁹, G. Rossi⁴⁰,
 E. Sanchez⁴¹, D. Schlegel⁶, M. Schubnell^{42,43}, H. Seo⁴⁴, J. Silber⁶, D. Sprayberry²⁷, T. Tan⁵,
 G. Tarlé⁴³, M. Walther^{45,46}, B. A. Weaver²⁷, H. Zou⁴⁷

¹ Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, U.S.A.

² Department of Astronomy, The Ohio State University, 4055 McPherson Laboratory, 140 W 18th Avenue, Columbus, OH 43210, U.S.A.

³ Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, U.S.A.

⁴ Université Clermont-Auvergne, CNRS, LPCA, 63000 Clermont-Ferrand, France

⁵ IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France

⁶ Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, U.S.A.

⁷ Department of Physics, Boston University, 590 Commonwealth Avenue, Boston, MA 02215, U.S.A.

⁸ Department of Physics & Astronomy, University of Rochester, 206 Bausch and Lomb Hall, P.O. Box 270171, Rochester, NY 14627-0171, U.S.A.

⁹ Dipartimento di Fisica “Aldo Pontremoli”, Università degli Studi di Milano, Via Celoria 16, I-20133 Milano, Italy

¹⁰ INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy

¹¹ Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, U.K.

¹² NASA Einstein Fellow

¹³ Instituto de Física, Universidad Nacional Autónoma de México, Circuito de la Investigación Científica, Ciudad Universitaria, Cd. de México C. P. 04510, México

¹⁴ Department of Astronomy & Astrophysics, University of Toronto, Toronto, ON M5S 3H4, Canada

¹⁵ Department of Physics & Astronomy and Pittsburgh Particle Physics, Astrophysics, and Cosmology Center (PITT PACC), University of Pittsburgh, 3941 O’Hara Street, Pittsburgh, PA 15260, U.S.A.

¹⁶ University of California, Berkeley, 110 Sproul Hall #5800 Berkeley, CA 94720, U.S.A.

¹⁷ Institut de Física d’Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Edifici Cn, Campus UAB, 08193, Bellaterra (Barcelona), Spain

¹⁸ Departamento de Física, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia

¹⁹ Observatorio Astronómico, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia

²⁰ Institut d’Estudis Espacials de Catalunya (IEEC), c/ Esteve Terradas 1, Edifici RDIT, Campus PMT-UPC, 08860 Castelldefels, Spain

²¹ Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth, PO1 3FX, U.K.

²² Institute of Space Sciences, ICE-CSIC, Campus UAB, Carrer de Can Magrans s/n, 08913 Bellaterra, Barcelona, Spain

²³ University of Virginia, Department of Astronomy, Charlottesville, VA 22904, U.S.A.

- ²⁴ *Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, U.S.A.*
- ²⁵ *Institut d'Astrophysique de Paris. 98 bis boulevard Arago. 75014 Paris, France*
- ²⁶ *Department of Physics, The University of Texas at Dallas, 800 W. Campbell Rd., Richardson, TX 75080, U.S.A.*
- ²⁷ *NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, U.S.A.*
- ²⁸ *Department of Physics and Astronomy, University of California, Irvine, 92697, U.S.A.*
- ²⁹ *Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), FR-75005 Paris, France*
- ³⁰ *Departament de Física, Serra Hünter, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain*
- ³¹ *Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain*
- ³² *Departamento de Física, DCI-Campus León, Universidad de Guanajuato, Loma del Bosque 103, León, Guanajuato C. P. 37150, México*
- ³³ *Instituto Avanzado de Cosmología A. C., San Marcos 11 - Atenas 202. Magdalena Contreras. Ciudad de México C. P. 10720, México*
- ³⁴ *Department of Physics and Astronomy, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*
- ³⁵ *Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada*
- ³⁶ *Waterloo Centre for Astrophysics, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*
- ³⁷ *Space Sciences Laboratory, University of California, Berkeley, 7 Gauss Way, Berkeley, CA 94720, U.S.A.*
- ³⁸ *Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain*
- ³⁹ *Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 10, 08930 Barcelona, Spain*
- ⁴⁰ *Department of Physics and Astronomy, Sejong University, 209 Neungdong-ro, Gwangjin-gu, Seoul 05006, Republic of Korea*
- ⁴¹ *CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain*
- ⁴² *Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, U.S.A.*
- ⁴³ *University of Michigan, 500 S. State Street, Ann Arbor, MI 48109, U.S.A.*
- ⁴⁴ *Department of Physics & Astronomy, Ohio University, 139 University Terrace, Athens, OH 45701, U.S.A.*
- ⁴⁵ *Excellence Cluster ORIGINS, Boltzmannstrasse 2, D-85748 Garching, Germany*
- ⁴⁶ *University Observatory, Faculty of Physics, Ludwig-Maximilians-Universität, Scheinerstr. 1, 81677 München, Germany*
- ⁴⁷ *National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Road, Chaoyang District, Beijing, 100101, China*