MNRAS 000, 1–19 (2025) Preprint 12 June 2025 Compiled using MNRAS LATEX style file v3.0

Modelling the impact of quasar redshift errors on the full-shape analysis of correlations in the Lyman-α forest.

Calum Gordon[1]*, Andrei Cuceu[2,3]†, Andreu Font-Ribera[1], Hiram K. Herrera-Alcantar[3,4], Jessica Nicole Aguilar,[2] Steven Ahlen[5], Davide Bianchi[6,7], David Brooks,[8] Todd Claybaugh,[2] Shaun Cole[9], Axel de la Macorra[10], Biprateep Dey[11,12], Peter Doel,[8] Jaime E. Forero-Romero[13,14], Enrique Gaztañaga,[15,16,17] Satya Gontcho A Gontcho[2], Gaston Gutierrez,[18] Julien Guy[2], Klaus Honscheid[19,20,21], Mustapha Ishak[22], Robert Kehoe,[23] David Kirkby[24], Theodore Kisner[2], Anthony Kremin[2], Martin Landriau[2], Laurent Le Guillou[25], Michael Levi[2], Marc Manera[26,1], Paul Martini[19,27,21], Ramon Miquel,[28,1] John Moustakas[29], Seshadri Nadathur[16], Gustavo Niz[30,31], Nathalie Palanque-Delabrouille[4,2], Will Percival[32,33,34], Francisco Prada[35], Ignasi Pérez-Ràfols[36], Graziano Rossi,[37] Eusebio Sanchez[38], David Schlegel,[2] Michael Schubnell,[39,40] Hee-Jong Seo[41], Joseph Harry Silber[2], David Sprayberry,[42] Gregory Tarlé[40], Benjamin Alan Weaver,[42] Rongpu Zhou[2], Hu Zou[43]

Affiliations are listed at the end of the paper.

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

In preparation for the first cosmological measurements from the full-shape of the Lyman-α (Lyα) forest from DESI, we must carefully model all relevant systematics that might bias our analysis. It was shown in Youles et al. (2022) that random quasar redshift errors produce a smoothing effect on the mean quasar continuum in the Lyα forest region. This in turn gives rise to spurious features in the Lyα auto-correlation, and its cross-correlation with quasars. Using synthetic data sets based on the DESI survey, we confirm that the impact on BAO measurements is small, but that a bias is introduced to parameters which depend on the full-shape of our correlations. We combine a model of this contamination in the cross-correlation (Youles et al. 2022) with a new model we introduce here for the auto-correlation. These are parametrised by 3 parameters, which when included in a joint fit to both correlation functions, successfully eliminate any impact of redshift errors on our full-shape constraints. We also present a strategy for removing this contamination from real data, by removing ~0.3% of correlating pairs.

Key words: (cosmology:) large-scale structure of Universe – (cosmology:) distance scale – (cosmology:) dark energy

1 INTRODUCTION

One of the major objectives in modern cosmology is to understand the nature of dark energy. The discovery of the accelerating expansion of the universe from measurements of distant supernovae (Riess et al. 1998; Perlmutter et al. 1999) introduced the need for some form of dark energy, in the simplest case a cosmological constant. Studies of anisotropies in the Cosmic Microwave Background (CMB) with experiments like Planck (Planck Collaboration et al. 2020) have at the same time provided percent-level constraints on other key cosmological parameters. Another more recent probe of cosmic expansion history is the measurement of Baryon Acoustic Oscillations (BAO), first in galaxy clustering (Cole et al. 2005; Eisenstein et al. 2005) and later with the Lyman-α (Lyα) forest (Busca et al. 2013; Bautista et al. 2017; du Mas des Bourboux et al. 2020). In each case the observable is a feature imprinted in the clustering measurements at comoving separation ~100 h^{-1}Mpc. This feature was created at the epoch of recombination, where sound waves that were propagating through the dense photon-baryon plasma ceased to travel because of proton-electron recombination. The overdense regions at wave peaks gave rise to the preferential clustering scale we observe as BAO today. This scale can be measured in the CMB (Planck Collaboration et al. 2020), and as such can be used as a standard ruler to measure the expansion of the universe (Adame et al. 2025b).

The Dark Energy Spectroscopic Instrument (DESI) (Levi et al. 2013; DESI Collaboration et al. 2016a,b, 2022; Miller et al. 2024; Poppett et al. 2024; Guy et al. 2023) survey will measure the spectra and redshifts of more than 40 million galaxies (Hahn et al. 2023; Raichoor et al. 2023; Zhou et al. 2023; Schlafly et al. 2023) and 3 million quasars (Chaussidon et al. 2023) over the course of a 5-year survey, covering 14 000 deg^2 of the sky. Precise measurements of

*E-mail: cgordon@ifae.es
†NASA Einstein Fellow

large-scale structure have already been made in galaxy and quasar clustering for the first year of observations of DESI (Adame et al. 2025c; DESI Collaboration et al. 2024), comprising spectra of targets between redshifts $0.1 < z < 2.1$, and the Lyα forest at $z > 2.1$ (Adame et al. 2025a). The Lyα forest is a series of absorption lines present in the spectra measured from very distant quasars. They are created when light from these quasars intercepts neutral hydrogen in the intergalactic medium (IGM), and as such, are an effective tracer of large-scale structure. A BAO measurement from the Lyα forest, made using a combination of the autocorrelation of the Lyα flux-transmission field and its cross-correlation with quasars (Font-Ribera et al. 2014, 2013), has a higher effective redshift than DESI galaxies (2.33 (Adame et al. 2025a)). This allows us to probe an earlier stage of the Universe than the DESI galaxy sample, which additionally helps to break degeneracies in cosmological constraints (Adame et al. 2025b).

It is possible to drastically improve the cosmological constraining power of the Lyα forest (and galaxy) clustering by including information from the Alcock-Paczynski (AP) effect (Cuceu et al. 2021, 2023a). It was projected in Cuceu et al. (2021) that including AP information from both the BAO peak and the smooth component of the Lyα clustering measurement will give 2.5% and 1% constraints on $\Omega_{\rm m}$ and $H_0 r_{\rm d}$, from the full DESI survey. This represents a factor of $\sim$2 improvement over BAO-only analyses. Additionally, a sub-10% level measurement of $f\sigma_8$ (at $z \sim 2.33$) should be possible when combining AP + BAO measurements from the full set of Lyα correlation functions with redshift-space distortion (RSD) constraints derived from the quasar autocorrelation function.

It was demonstrated in Adame et al. (2025a) that the analysis method used to measure BAO is quite robust to various systematic effects, in large part due to the decomposition of the peak and smooth components of the correlation function (Kirkby et al. 2013). However, since we are interested in constraining both the peak and smooth components of the AP effect, we expect to be more sensitive to contaminants than the BAO-only measurement. The main contaminants usually considered are continuum fitting distortion (Slosar et al. 2011), high column density (HCD) systems (Font-Ribera & Miralda-Escudé 2012; Rogers et al. 2018; Pérez-Ràfols et al. 2023), broad absorption line (BAL) features (Ennesser et al. 2022), metal absorption (Font-Ribera & Miralda-Escudé 2012; Pieri et al. 2014) and redshift errors in quasars arising from both astrophysical peculiar velocities and measurements Font-Ribera et al. (2013). Both HCDs and BALs are identified using neural-network–based finder algorithms (Guo & Martini 2019; Wang et al. 2022) and masked to reduce their impact on the correlation functions. The former is also modelled to account for bias introduced by lower-density, undetected systems. Metals are modelled successfully by treating each absorber as a biased tracer of structure, where the biases are marginalised during the BAO fits.

Random velocities of galaxies or quasars distort clustering measurements along the line of sight, producing anisotropies relative to the transverse direction — an effect known as Fingers of God (FoGs). This effect will be present in the cross-correlation of Lyα forests with quasars (Font-Ribera et al. 2013), which we will marginalise over using a version of the model from Percival & White (2009) (see section 2). In the Lyα autocorrelation, since the redshifts of Lyα absorbers are set by intergalactic hydrogen and not by quasar motion, we do not need to model this effect.

Errors in redshift measurement produce both the FoG effect and a second, distinct effect on our correlation functions. When computing the Lyα transmission field δ, we use a weighted mean continuum of all forests (equations 1, 2). As first described in Youles et al. (2022), this mean continuum is smoothed by quasar redshift errors, causing spurious features in our correlation functions. They also presented a model for this contamination in the cross-correlation function, and showed that it had a small impact on the accuracy of their BAO constraints. This was also confirmed as a source of systematic bias in DESI DR1 (Cuceu et al. 2025b) and in DESI DR2 (Casas et al. 2025). In this paper we extend this work by modelling the same contamination in the Lyα autocorrelation, and introducing free parameters that we marginalise over in our full-shape analysis.

To measure redshifts of DESI quasar spectra, a package called `Redrock` is used (Bailey et al. 2025). It is a template-fitting algorithm that uses Principal Component Analysis (PCA)-based quasar templates to determine the best-fit redshift in a chi-squared (χ^2) minimisation. We expect a degree of systematic error in these estimations due to spectral variation and limited performance of templates, which was found to be ~ 340 km/s (Alexander et al. 2023) for DESI survey validation. These variations are caused by Doppler shift of quasar emission lines due to the physical processes happening in the line-emitting region (Shen et al. 2016). The typical level of shift away from the systemic quasar redshift depends on the line, but highly ionised, broad emission lines tend to exhibit the largest shifts. Recent work on new templates have significantly reduced this error for DESI Data Release 1 (DR1; DESI Collaboration et al. (2025b)), showing it to be $\sim$50% smaller than survey validation (Bault et al. 2025; Brodzeller et al. 2023). This was partly done by incorporating precise measurements of the Mg II line, which has relatively low bias with respect to the systemic redshift of the quasar. Broad emission lines within the Lyα forest, listed in table 1 (Harris et al. 2016), can also be Doppler-shifted, potentially contributing to errors in the mean continuum estimation (see section 2.2).

This paper is organised as follows. In section 2, we describe the mock datasets used to develop and test our model, the method for generating the flux transmission field from quasar spectra, and the procedures for computing and modelling our correlation functions. In section 3, we show measurements of the Lyα auto- and Lyα –quasar cross-correlations from different mock datasets, both with and without redshift errors. Section 4 presents the redshift error model for both the auto- and cross-correlation functions, and section 5 shows results from AP and isotropic BAO fits on our mock datasets. In section 6, we discuss the application of our model to real data and how the contamination can be mitigated using data cuts. Finally, section 7 summarises our findings.

2 METHOD

In this section we present the mock datasets we use (section 2.1), and the analysis techniques employed to derive measurements of BAO and AP parameters from synthetic quasar spectra. The latter part begins with the flux transmission field measurement in section 2.2, followed by a description of how our correlation functions are constructed in section 2.3. Finally, we describe the model of the correlation functions and how we constrain BAO and AP parameters in sections 2.4 and 2.5.

2.1 Data sets

To test our model, we use synthetic data sets developed to validate the Lyα BAO analysis in Adame et al. (2025a), described in detail in Cuceu et al. (2025b); Herrera-Alcantar et al. (2025). These are based

Line	Wavelength [Å]	Equivalent Width [Å]	σ [Å]
SiIV	1064/1074	2.9/0.7	7.7/3.5
NII	1083	1.32	5.3
PV	1118/1128	0.76/0.46	5.3/4.1
CIII*	1175	2.49	7.7

Table 1. The emission lines within the Lyα forest included in our synthetic spectra (Harris et al. 2016). We show from left to right, their rest-frame wavelength (in Angstroms), the mean equivalent widths and mean FWHM.

on the `CoLoRe`[1] (Ramírez-Pérez et al. 2022) suite, which generates log-normal density fields with quasars distributed according to a specific biasing model. Lines-of-sight (also called skewers) are drawn from quasar positions in this field, with modified small-scale power, and converted to transmitted flux fraction using `LyaCoLore`[2] (Farr et al. 2020a). The clustering of these noiseless skewers is designed to be realistic at scales relevant for BAO systematics studies (~100 h^{-1}Mpc). `LyaCoLore` also stores skewers of metal absorption and can include high column-density (HCD) systems, both of which are major contaminants in standard Lyα BAO analyses.

To generate realistic quasar spectra from the flux transmission skewers, we use a final package called `desisim`[3] (Herrera-Alcantar et al. 2025). This takes the transmission skewers of `LyaCoLore` as an input, multiplies them by a quasar continuum template, and adds noise to mimic the observing conditions and the instrumental model of DESI. Continuum templates are generated using `SIMQSO`[4] (McGreer et al. 2021), which combines a broken power law with a series of Gaussian emission lines. The slopes of the power law are sampled from a Gaussian, with mean and dispersion tuned to better reflect the continuum shape and variability of quasars in the eBOSS DR16 dataset (du Mas des Bourboux et al. 2020). The Lyα forest emission lines are simulated from the composite model of BOSS spectra (Harris et al. 2016), with line diversity is drawn from the distribution of equivalent widths (EWs). In table 1, we show the properties of these emission lines.

To emulate the effect of redshift errors in our mock datasets, we add Gaussian-distributed velocities dv with zero mean and a dispersion of 400 km/s to each quasar redshift $(1 + z_q = (1 + z_q^0)(1 + dv/c))$. The reason for doing this rather than using the error from a redshift-fitting algorithm is that the latter has been shown to perform better on simulated spectra than real data (Farr et al. 2020b). In Youles et al. (2022), they differentiate between the Fingers of God (FoG) effect and redshift errors that affect the mean continuum. In this paper we focus on the latter.

We use two synthetic datasets: one has redshift errors added only to the tracer quasar catalogue, emulating the effect of FoGs. In the other we add redshift errors immediately after quasar spectra are generated, which will additionally give rise to mean continuum errors (see section 4). Note that in our measured correlation functions, the latter set will exhibit the combined effect of FoGs and mean continuum errors. This means that later in section 3 when we compare the differences between the two sets, we will observe only the contamination caused by continuum errors.

Each set consists of a stack of 100 mocks designed to emulate the survey conditions of DESI DR1 (DESI Collaboration et al. 2025b), including target density, footprint and signal-to-noise ratio. We will refer to these as DR1 mocks going forward. They also include all of the major contaminants affecting the analysis on data: HCDs, metal absorption and broad absorption lines (BALs). Using this synthetic dataset, we can validate our model and test correlations between various parameters in a controlled environment.

[1] https://github.com/damonge/CoLoRe
[2] https://github.com/igmhub/LyaCoLoRe
[3] https://github.com/desihub/desisim
[4] https://github.com/imcgreer/simqso

2.2 Continuum fitting

The flux transmission field used to perform our clustering analysis is defined as:

$$1 + \delta(\lambda) = \frac{f(\lambda)}{\overline{F}(\lambda) C_q(\lambda)}, \tag{1}$$

where $f(\lambda)$ is observed flux, $\overline{F}(z)$ is the mean transmission and $C_q(\lambda)$ is the unabsorbed quasar continuum. Our analysis is discretised into "pixels" at the dispersion of the DESI spectrograph, 0.8Å. At large enough scales, the delta field is a linear tracer of the underlying dark matter, but at small-scales the relationship becomes more complex due to non-linear structure formation and gas dynamics. Since we can't measure the unabsorbed continuum $C_q(\lambda)$ directly, we re-characterise its product with the mean transmission $\overline{F}(z)$:

$$\overline{F}(\lambda) C_q(\lambda) = \overline{C}(\lambda^{\mathrm{rf}}) \left(a_q + b_q \frac{\log\lambda - \log\lambda_{\min}}{\log\lambda_{\max} - \log\lambda_{\min}} \right). \tag{2}$$

The term $\overline{C}(\lambda^{\mathrm{rf}})$, referred to as the "mean continuum", is the weighted mean of all forests in our analysis. This is multiplied by a first degree polynomial in $\log\lambda$ for each quasar q in our sample, to account for continuum variability. The free parameters a_q and b_q are estimated by "continuum fitting", a process that involves maximising the log-likelihood:

$$2\ln L = -\sum_q \frac{(f(\lambda) - \overline{F}(z) C_q(\lambda)(\lambda, a_q, b_q)}{\sigma_q^2(\lambda)} - \sum_q \sigma_q^2(\lambda), \tag{3}$$

where $\sigma_{\mathrm{q}}^2(\lambda)$ is the flux variance, a combination of pipeline noise and large-scale structure. The mean continuum $\overline{C}(\lambda^{\mathrm{rf}})$ is also computed during this maximisation. For further details on exactly how this is done, see Ramírez-Pérez et al. (2024).

A notable side effect of fitting $(a_{\mathrm{q}}, b_{\mathrm{q}})$ using data from entire forests is that the measured $\delta_q(\lambda)$ will be distorted. In other words, our measured $\delta(\lambda_i)$ in pixel i will be a linear combination of all true underlying $\delta(\lambda_j)$, such that $\delta'_i = \eta_{ij}\delta_j$. The solution to this, described in detail in du Mas des Bourboux et al. (2017); Bautista et al. (2017); Adame et al. (2025a), is to apply a linear projection operator to our measured δ' that explicitly removes the mean and slope of each quasar forest from the field:

$$\tilde{\delta}(\lambda) = \delta'(\lambda) - \sum_j \alpha_{qij}\delta_{q,j}, \tag{4}$$

where

$$\alpha_{qij} = \frac{w_{q,j}}{\sum_k w_{q,k}} + \frac{w_{q,j}\Lambda_{\mathrm{q,i}}\Lambda_{q,j}}{\sum_k w_{q,k}\Lambda_{q,k}^2}. \tag{5}$$

Here, $\Lambda_i = \log(\lambda_i) - \overline{\log\lambda_i}$, where i,j and k represent forest pixels. The reason for doing this is that the projected field $\tilde{\delta}$ has the same distortion as our projected true field. We measure the correlations of our projected field $\tilde{\delta}$ and, as we will show in section 2.5, project the model correlation function in the same way.

2.3 Correlation functions

In this section we will show how we measure the Lyα auto-correlation and the Lyα-quasar cross-correlations from the projected transmission field outlined in the previous section. Firstly, we measure our correlations in configuration space as a function of transverse and line-of-sight separation $(r_\perp, r_\parallel)$. To convert from redshift and angular separation we do the following:

$$r_\perp = (d_{\rm m}(z_i) + d_{\rm m}(z_j))\ \sin(\theta_{ij}/2), \tag{6}$$
$$r_\parallel = (d_{\rm c}(z_i) - d_{\rm c}(z_j))\ \cos(\theta_{ij}/2), \tag{7}$$

where $d_{\rm m}$ is the transverse comoving distance and $d_{\rm c}$ is the comoving distance. A fiducial cosmology is required to compute both distances in terms of redshift, so we use the same Planck 2018 (hereby `P18`; Planck Collaboration et al. (2020)) cosmology as Adame et al. (2025a); Gordon et al. (2023). The correlation within each bin $(r_\perp, r_\parallel) \in A$ is then determined by a simple weighted average:

$$\xi_A = \frac{\sum_{i,j} w_i w_j \delta_i \delta_j}{\sum_{i,j} w_i w_j}, \tag{8}$$

where the weights are defined as:

$$w_i = \frac{1}{\sigma_i^2}\,(1+z_i)^{\kappa-1}\,. \tag{9}$$

Here σ_i is the variance from equation 3, but with an extra term to increase the large-scale structure contribution. This term was implemented in Ramírez-Pérez et al. (2024), to maximise the signal-to-noise of our correlation functions. The second term in equation 9 takes into account the correlation function amplitude scaling with redshift. For Lyα pixels, $\kappa = 2.9$ and for quasars $\kappa = 1.44$. Note that for the cross-correlation, index j in equation 8 represents a quasar, which we treat as points that effectively have $\delta_j = 1$. This is equivalent to computing the weighted mean of δ_α at set distances from a quasar. In both cases we choose a bin size of 4 h^{-1}Mpc following Adame et al. (2025a); Gordon et al. (2023). In the autocorrelation we measure between 0 and 200 h^{-1}Mpc in both directions, giving a total of 2500 bins. The cross-correlation however is not symmetric along the line-of-sight under permutation of pixels and quasars. Therefore, we define $r_\parallel$ between -200 and 200 h^{-1}Mpc, where negative separations correspond to a quasar being behind the Lyα pixel with respect to the observer and vice-versa for positive separations. This gives us 5000 bins in total for the cross-correlation.

The covariance of our measurement, explained in detail in du Mas des Bourboux et al. (2020), is estimated by splitting our dataset into sub-samples and computing:

$$C_{AB} = \frac{\sum_s W_A^s W_B^s (\xi_A^s \xi_B^s - \xi_A \xi_B)}{W_A W_B}, \tag{10}$$

where A and B are bins of our correlation function, and s is a sub-sample. In our case, each Lyα forest or quasar is uniquely set in one `HEALpix` sample s. The covariance is then smoothed by averaging all off-diagonal elements of the correlation matrix $(C_{AB}/\sqrt{(C_{AA}C_{BB})})$ with the same $\Delta r_\perp = r_\perp^A - r_\perp^B$ and $\Delta r_\parallel = r_\parallel^A - r_\parallel^B$. Furthermore, as introduced in Adame et al. (2025a), we now include the cross-covariance between the two correlation functions when doing fits to the model.

In Adame et al. (2025a) and Cuceu et al. (2025a) they use two additional correlation functions that measure the autocorrelation of Lyα absorption in a bluer region of the quasar ("region B", between 920-1020 ÅDESI Collaboration et al. (2025a)) and its cross-correlation with quasars. For simplicity, and because the relative contribution to cosmological constraining power is small ($\sim$ 10%), we choose not to include them in our analysis.

2.4 BAO and full-shape information

In this section we outline the model used to extract cosmological information and capture behaviour of contaminants in our analysis. The full-shape analysis that we outline here was first described for the Lyα forest in Cuceu et al. (2021), and performed on eBOSS data in Cuceu et al. (2023a). We follow the method of Kirkby et al. (2013) and employ a template power spectrum, based on `P18` cosmology, that we allow to vary from the measured cosmology by re-scaling our co-ordinate grid:

$$r'_\parallel = q_\parallel r_\parallel, \qquad r'_\perp = q_\perp r_\perp. \tag{11}$$

We wish to distinguish between isotropic re-scaling α of the correlation function and anisotropic re-scaling ϕ (AP effect), therefore we define:

$$\alpha(z) = \sqrt{q_\parallel q_\perp} \quad \text{and} \quad \phi(\mathrm{z}) = \frac{\mathrm{q}_\perp}{\mathrm{q}_\parallel} \tag{12}$$

The first step in standard Lyα BAO analyses is to take the linear isotropic input power spectrum from `P18`, $P_{\rm L}(k)$, and decompose it into peak ($P_{\rm L}^{\rm peak}(k)$) and smooth ($P_{\rm L}^{\rm sm}(k)$) components following Kirkby et al. (2013). This is possible because the BAO is a distinct feature in the full correlation function, which also makes the standard peak-only analysis more robust to systematics affecting the smooth component. In the full-shape analysis we keep this peak-smooth decomposition, but introduce scaling parameters α and ϕ for each component:

$$\xi(r_\perp, r_\parallel) = \xi_{\rm p}(r_\perp, r_\parallel, \alpha_{\rm p}, \phi_{\rm p}) + \xi_{\rm s}(r_\perp, r_\parallel, \alpha_{\rm s}, \phi_{\rm s}), \tag{13}$$

where "s" and "p" refer to the smooth and peak components respectively. The peak component is defined in terms of distances and the sound horizon scale $r_{\rm d}$ as:

$$\alpha_{\rm p} = \sqrt{\frac{d_{\rm m}(z) d_h(z)/r_{\rm d}^2}{[d_{\rm m}(z) d_h(z)/r_{\rm d}^2]_{\rm fid}}} \tag{14}$$

where d_h is the Hubble distance and "fid" is the fiducial `P18` cosmology. From this one can constrain the combination of $H_0 r_{\rm d}$ and $\Omega_{\rm m}$, where H_0 is the Hubble constant. Both AP components are equivalent to $\phi = d_{\rm m}(z)H(z)/[d_{\rm m}(z)H(z)]_{\rm fid}$. From this we can directly measure $\Omega_{\rm m}$, and therefore break the degeneracy between $\Omega_{\rm m}$ and $H_0 r_{\rm d}$. A more detailed description of this is given in Cuceu et al. (2021).

In Cuceu et al. (2023b), it was shown that including the smooth component AP parameter $\phi_{\rm s}$ doubled the precision of constraints on $\Omega_{\rm m}$. Note that in Cuceu et al. (2023b) and in this work, we treat $\alpha_{\rm s}$ as a nuisance parameter since it is not clear how to extract cosmological information from it, and it's more strongly correlated with the Lyα bias. Finally, while the information gain of including the smooth-component is substantial, it is also more susceptible to contaminants, thus motivating us to create the model presented in this paper.

2.5 Correlation model

The Lyα and Lyα-quasar cross power spectra are given by linear perturbation theory, with a Kaiser (Kaiser 1987) model of RSD:

$$P_\alpha(k,\mu_k,z) = b_\alpha^2(1+\beta_\alpha\mu_k^2)^2 F_{sm} P_L(k,z), \quad (15)$$
$$P_X(k,\mu_k,z) = b_\alpha b_q(1+\beta_\alpha\mu_k^2)(1+\beta_q\mu_k^2) F_{sm} F_{NL,q} P_L(k,z), \quad (16)$$

where b_i are the cosmological biases and β_i are the RSD parameters for tracer i (quasar or Lyα pixel), and μ_k is the cosine of the k vector with respect to the line-of-sight direction. The cross-correlation term $F_{NL,q}$ accounts for quasar peculiar velocities, and is given in Percival & White (2009):

$$F_{NL,q}(k_\parallel) = \sqrt{\frac{1}{1+(k_\parallel\sigma_v)^2}}. \quad (17)$$

This produces a smoothing effect at large $k_\parallel$, where σ_v is a free parameter in our fit. This is a distinct effect from the redshift error continuum smoothing we model in this paper. We also account for any systematic quasar redshift errors by introducing a free parameter to our cross-correlation model: $\Delta r_\parallel = r_{\parallel,\rm true} - r_{\parallel,\rm measured}$.

F_{sm} is a Gaussian smoothing term that accounts for our log-normal simulations having limited grid size (~2.4 h^{-1}Mpc). This has two free parameters $(\sigma_\perp, \sigma_\parallel)$ which we marginalise over in our analysis.

In the contaminated mock dataset we introduce metal absorption and HCD contamination. For HCDs, we mask systems with column density $> 2\times10^{20}$cm^{-2}, following the analysis of Adame et al. (2025a). The remaining HCDs, which in real data are too small to detect for masking, are biased tracers of large-scale structure. Following Font-Ribera & Miralda-Escudé (2012), we treat re-write the Lyα bias and RSD parameters as effective combinations of signal from Lyα forest and HCDs:

$$b'_\alpha = b_\alpha + b_{HCD}\exp\left(-L_{HCD}k_\parallel\right), \quad (18)$$
$$b'_\alpha\beta_\alpha = b_\alpha\beta_\alpha + b_{HCD}\beta_{HCD}\exp\left(-L_{HCD}k_\parallel\right), \quad (19)$$

where L_{HCD} is the typical length-scale of the unmasked HCDs. In our fits we marginalise over this and the bias and RSD of the HCD systems.

In our analysis, we assume all of the absorption within the Lyα forest is due to the Lyα transition. However, in some cases we are mistaking Lyα absorption with absorption from different metals. This leads to a contamination which we correct for using the same process outlined in Adame et al. (2025a). In our fits, we marginalise over a set of bias parameters for 4 metal lines: SiIII(1207), SiII(1190), SiII(1193), and SiII(1260) (section 5).

As mentioned in section 2.2, distortions introduced to δ during continuum fitting lead us to use a "projected" field $\tilde{\delta}$ (equation 4). This is designed such that the projected distorted field has the same distortion as the projected true field. We write the distorted model (A) in terms of the undistorted model (B) using a "distortion matrix": $\xi_A = \sum_{AB} D_{AB}\xi_B$. D can be written as:

$$D_{AB} = \frac{1}{W_A}\sum_{ij\in A} w_i w_j \left(\sum_{i'j'\in B}\eta_{ii'}\eta_{jj'}\right), \quad (20)$$

where $\eta_{mn} = \delta^K_{mn} - \alpha_{mn}$. Here, δ^K is the Kronecker delta, and α is the projection matrix of equation 5. Note for the cross-correlation $\eta_{jj'} = 1$, since there is only one correlating forest. To compute the auto- and cross-correlation models in our analysis we use the package `Vega`[5]. This is also used to fit cosmological parameters, shown in section 5.

3 IMPACT OF REDSHIFT ERRORS

In this section we will show the impact of continuum redshift errors on the measured correlation functions of our mock datasets. Then, in section 4, we will discuss the origin of this contamination and propose a model for it.

We use the two synthetic datasets described in section 2.1, both of which are stacks of 100 DESI DR1 mocks. The "uncontaminated" set has redshift errors added to the tracer quasar catalogue, emulating FoGs. The "contaminated" set has errors added to the spectra, emulating pipeline redshift errors. Using this method we isolate the effect of redshift errors which enter during the continuum fitting process (section 2.2). We refer to this contamination as continuum redshift errors, to differentiate it from the effect of FoGs.

We then compute the Lyα autocorrelation and its cross-correlation with quasars using the method outlined in the previous section. For each of the two sets, we compute the weighted mean and covariance of the 100 correlation functions. We use these two stacks to model and fit the effect of continuum redshift errors in our correlation functions in the following sections. The auto- and cross-correlation functions from the uncontaminated and contaminated sets are shown in figure 1. These are plotted as a function of the line-of-sight separation, averaged over the first 4 bins in transverse separation ([0,16] h^{-1}Mpc) where the effect is strongest. We can also clearly see a peak at ~ 100 h^{-1}Mpc, a combination of BAO and SiII(1260) contamination, and other peaks caused by metal contamination (section 2.5).

The contamination of continuum redshift errors (the difference between the two datasets) in the auto- and cross-correlations is shown in figure 2. As discussed in Youles et al. (2022), there is a strong dependence of redshift error distortion in the cross-correlation on the small-scale quasar autocorrelation, which is significantly weaker for positive $r_\parallel$ where host quasars of Lyα pixels are much further from the correlating quasars. Therefore, as we see in figure 2, the distortion from redshift errors is only visible for negative $r_\parallel$. This dependence on the quasar autocorrelation is also the reason the contamination becomes negligible at transverse separations $r_\perp \gtrsim 20$ h^{-1}Mpc. The behaviour is similar in the Lyα autocorrelation, except the dependence is now on the small-scale Lyα-quasar cross-correlation. In section 4, this dependence can be seen explicitly in our model. The distortions in figure 2 are also localised around specific values of $r_\parallel$. In section 4.1 we show how this pattern arises due to smoothing of the mean continuum.

We can see that the continuum redshift errors will contaminate our full-shape measurement, but also that part of the contamination overlaps with the fiducial BAO position (~100 h^{-1}Mpc). In section 5, we show the effect of this contamination on the set of cosmological parameters measured from the BAO peak (α_p, ϕ_p) and the broadband component (ϕ_s, $f\sigma_8$) of our correlation functions.

4 MODELLING THE CONTAMINATION

In this section we will outline our model for continuum redshift error contamination. We start describing the effect of redshift errors on the

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

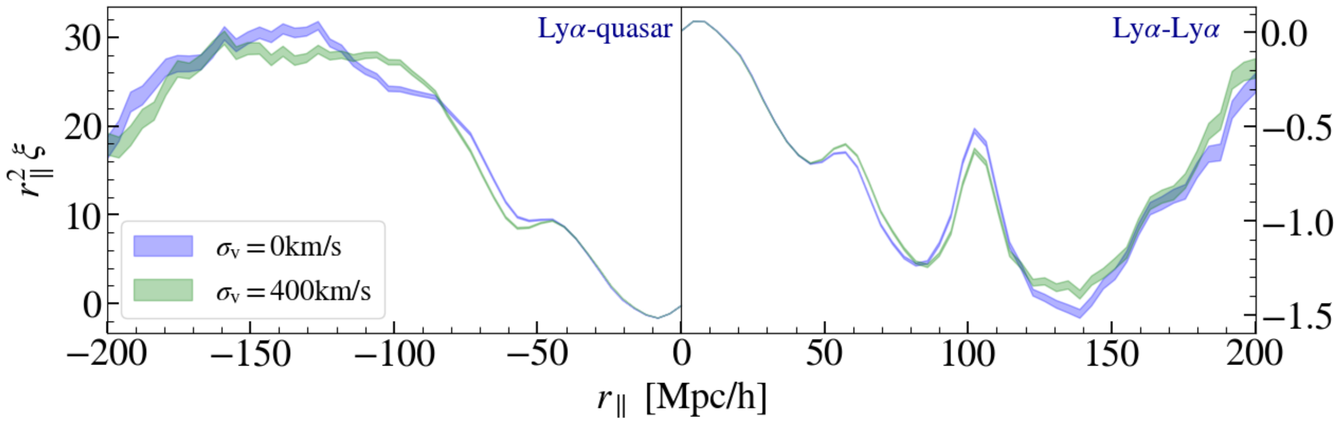

Figure 1. The effect of continuum redshift errors on the Lyα-quasar cross-correlation (left) and the Lyα-Lyα auto-correlation (right) computed from our mock datasets. Each dataset is the stack of 100 DESI DR1 mocks. We take the weighted average over the first 4 bins in transverse separation $r_\perp$ ([0,16] h^{-1}Mpc) where the effect of redshift errors is strongest, and plot as function of line-of-sight separation $r_\parallel$. We only include negative $r_\parallel$ for the cross-correlation, since as shown in figure 2 there is no impact at $r_\parallel > 0\ h^{-1}$Mpc. Note that at $\sim \pm 60\ h^{-1}$Mpc we see bumps in the correlations due to SiII(1190/1193) lines.

mean continuum and flux-transmission field (following Youles et al. (2022)), and carry this through to a full model of the contamination in the autocorrelation. We also highlight the cross-correlation redshift error model of Youles et al. (2022).

In Youles et al. (2022), redshift errors were identified to smooth the mean continuum (equation 2) that is fit during the transmission field estimation. They subsequently defined a term $\gamma(\lambda^{\rm rf})$:

$$\gamma(\lambda^{\rm rf}) = \frac{\hat{\overline{C}}(\lambda^{\rm rf})}{\overline{C}(\lambda^{\rm rf})} - 1, \tag{21}$$

where $\hat{\overline{C}}$ is the mean continuum estimate with redshift errors and $\overline{C}$ is the true mean continuum. The Lyα forest region of our mock continua consists of a power law smooth component, superimposed with many Gaussian emission lines. If quasar redshift errors are randomly distributed, the mean continuum will be smoothed around emission lines as shown in figure 3. In appendix B, we write an expanded version of $\gamma(\lambda^{\rm rf})$, accounting for the fact that the wavelength grid $\lambda^{\rm rf}$ for each forest is shifted from the true grid, due to errors in $z_{\rm q}$.

In the top panel of figure 3, we show a measurement of $\gamma(\lambda^{\rm rf})$ formed using mean continua from our mock datasets with redshift errors of 400 kms^{-1} and 0 kms^{-1}. Combining equations 21 and 1 we get:

$$\hat{\delta}(\lambda) = \frac{1+\delta(\lambda)}{1+\gamma(\lambda^{\rm rf})} - 1. \tag{22}$$

Assuming we can ignore second-order terms:

$$\hat{\delta}(\lambda) \approx \delta(\lambda) - \gamma(\lambda^{\rm rf}). \tag{23}$$

Thus the measured autocorrelation function can be written as:

$$\langle\hat{\delta}(\lambda)\hat{\delta}(\lambda)\rangle = \langle\delta(\lambda)\delta(\lambda)\rangle - 2\langle\delta(\lambda)\gamma(\lambda^{\rm rf})\rangle + \langle\gamma(\lambda^{\rm rf})\gamma(\lambda^{\rm rf})\rangle. \tag{24}$$

To determine the most dominant term in this equation, we directly compute $\langle\delta\gamma\rangle$ and $\langle\gamma\gamma\rangle$ using δ from mock datasets with no added redshift errors, and the measured gamma in figure 3. The result is shown in figure 4, where both terms are plotted over the measured contamination in the autocorrelation. Clearly, $\langle\gamma\gamma\rangle$ is negligible. Therefore, we only model the contribution of the $\langle\delta\gamma\rangle$ term. Noticeably, the trough in the correlation difference at $\sim 25\,h^{-1}$Mpc is not captured by either term. In the following section we discuss the behaviour of this feature.

4.1 Variation with wavelength

Spurious features in $\langle\delta\gamma\rangle$ caused by redshift errors in the Lyα continuum appear at specific separations. These separations (in comoving distance) are determined by the difference in rest-frame wavelength between the Lyα line and strong features in $\gamma(\lambda^{\rm rf})$:

$$r^{\rm feature} \approx \frac{c(1+\overline{z})}{H(\overline{z})}\frac{(\lambda^{\rm feature} - \lambda^{\alpha})}{\lambda^{\alpha}}, \tag{25}$$

where $\overline{z}$ is the effective redshift of the dataset and H is the Hubble parameter. In the top panel of figure 3, we highlight key features in Roman numerals: i at 1174 Å, ii at 1187 Å, iii at 1205.7 Å and iv at 1211.6 Å, the latter two of which fall outside of the fiducial rest-frame limits (1205 Å) of our clustering analysis. Using these values in equation 25 with an effective redshift of 2.3 predicts spurious correlation features at separations of 99 h^{-1}Mpc, 70 h^{-1}Mpc, 26 h^{-1}Mpc and 9 h^{-1}Mpc respectively. In figure 5, we show the contamination in the Lyα autocorrelation as a function of the upper rest-frame wavelength limit of the forest. We also include Roman numerals here, indicating the aforementioned comoving separations, and see clearly that they align with the strongest spurious correlation features.

Looking specifically at the rest-frame limit of our analysis, 1205 Å(red line in figure 5), we see a trough at $\sim 25\,h^{-1}$Mpc. This roughly overlaps with the feature iii (at 1205.7 Å) in figure 3, but is a trough where we expect a peak (peaks in γ produce peaks in $\hat{\xi}_{\alpha\alpha}$). Therefore we posit that this trough is instead caused by feature iv in figure 3, at a rest-frame wavelength of 1211.6 Å. This is possible because quasar redshift errors not only smooth the mean continuum, but also shift the rest-frame wavelength grid with respect to the truth. For any forest in our dataset, we apply the upper rest-frame wavelength limit by converting observed wavelength λ into rest-frame, where $\lambda^{\rm rf} = \lambda/(1+\hat{z}_{\rm q})$ for a measured quasar redshift $\hat{z}_{\rm q}$. If $\hat{z}_{\rm q} > z_{\rm q}$, pixels which have $\lambda^{\rm rf} > 1205$Å will fall into the accepted Lyα region. Consequently, values of $\gamma(\lambda^{\rm rf} > 1205$ Å) will contaminate our analysis.

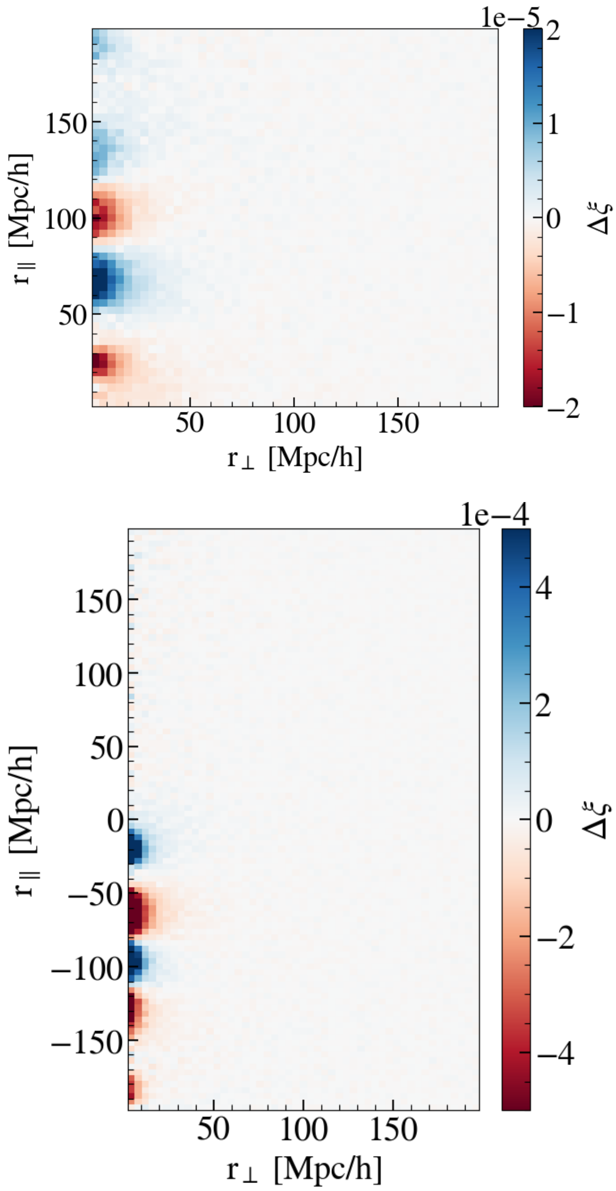

Figure 2. (Top) difference in the Lyα auto-correlation measured from contaminated (with continuum redshift errors) and uncontaminated datasets, as a function of $(r_\perp, r_\parallel)$. (Bottom) difference in the Lyα-quasar cross-correlation from contaminated and uncontaminated datasets.

In figure 6, we show (blue dashed line) the result of removing pixels (for each forest in our correlation measurement) which are in reality above the Lyα rest-frame wavelength limit, but are accepted into the sample because of redshift errors. In this case, the trough feature is no longer visible, and the contamination at small-scales is now roughly consistent with 0. We have also verified that including pixels that are in reality below 1205 Å, but are excluded due to redshift errors, has no effect on the shape of the contamination.

In appendix B, we make the shifting of the rest-frame wavelength grid where we evaluate each δ explicit, by expanding $\hat{\hat{C}}$ about small redshift deviations. This produces extra terms that contaminate the correlation functions, but to 1st order they do not capture the trough at 25 h^{-1}Mpc in our fiducial analysis. We also try to empirically model the trough feature (equation B8), which works well in a direct fit (see figure 8), but biases our full-shape analysis.

4.2 Auto-correlation model

To model the $\langle\delta(\lambda)\gamma(\lambda^{\rm rf})\rangle$ term in equation 24, we follow a similar approach to that used in Youles et al. (2022) to model the $\langle\gamma(\lambda^{\rm rf})\rangle$ term in the Lyα-quasar cross-correlation model.

We begin by considering a general expression for the expectation value of $\langle\delta\gamma\rangle$:

$$\langle\delta\gamma\rangle_A = \int_{-1}^{\infty} d\delta \int_{-1}^{\infty} d\gamma P(\delta, \gamma|\vec{r}^A)\delta\gamma \tag{26}$$

where $P(\delta,\gamma)$ is the probability of having a particular δ and γ, for separation $\vec{r}^A = \{\vec{r_\parallel}^A, \vec{r_\perp}^A\}$ between two Lyα pixels. This can be equivalently expressed as:

$$\langle\delta\gamma\rangle_A = \int_{-1}^{\infty} d\delta \int_{\lambda^{\rm rf,min}}^{\lambda^{\rm rf,max}} d\lambda^{\rm rf} P(\delta, \lambda^{\rm rf})\delta\gamma, \tag{27}$$

since γ is uniquely defined by $\lambda^{\rm rf}$, and the integration limits are given by the fact that we only consider pixels within the defined Lyα forest rest-frame region. We can write the quasar rest-frame wavelength of any pixel as:

$$\lambda^{\rm rf} = \lambda^{\alpha}(1+z^{\alpha})/(1+z^{\rm Q}), \tag{28}$$

where λ^{α} is the rest-frame wavelength of the Lyα transition, and z^{α} is the redshift of a Lyα pixel in the Lyα forest region of a quasar with redshift $z^{\rm Q}$. Therefore, for a given $\lambda_i^{\rm rf}$ for forest i, we have a set of possible $\{z_i^{\rm Q}, z_i^{\alpha}\}$. Furthermore, for separations $r_\parallel \in A$ and pixel redshift z_i^{α}, we have a unique z_j^{α} in forest j. Thus, we re-write equation 27 as:

$$\langle\delta\gamma\rangle_A = \int_{-1}^{\infty} d\delta_j \int_{z^{\rm Q,min}}^{z^{\rm Q,max}} \int_{z^{\alpha,\rm min}}^{z^{\alpha,\rm max}} dz_i^{\rm Q} dz_j^{\alpha} P(\delta_j, z_j^{\alpha}, z_i^{\rm Q}|r_\parallel^A)\delta_j\gamma_i. \tag{29}$$

We can expand this further as:

$$\langle\delta\gamma\rangle_A = \int_{-1}^{\infty} d\delta_j \int_{z^{\rm Q,min}}^{z^{\rm Q,max}} \int_{z^{\alpha,\rm min}}^{z^{\alpha,\rm max}} dz_i^{\rm Q} dz_j^{\alpha} P(\delta_j|z_j^{\alpha}, z_i^{\rm Q}) \times P(z_j^{\alpha}|z_i^{\rm Q})P(z_i^{\rm Q})\delta_j\gamma_i, \tag{30}$$

where we have made the dependence on $r_\parallel^A$ explicit. Expanding the first term in this equation, we get:

$$P(\delta_j|z_j^{\alpha}, z_i^{\rm Q}) = \frac{P(z_i^{\rm Q}|\delta_j, z_j^{\alpha})P(\delta_j|z_j^{\alpha})}{P(z_i^{\rm Q}|z_j^{\alpha})} \tag{31}$$

$$= \frac{P(z_i^{\rm Q}|\delta_j, z_j^{\alpha})P(\delta_j|z_j^{\alpha})P(z_j^{\alpha})}{P(z_j^{\alpha}|z_i^{\rm Q})P(z_i^{\rm Q})}. \tag{32}$$

Now we will make some approximations to evaluate the first two terms in the numerator of this equation - note that the first term in the denominator cancels with the second term in equation 30. We begin by writing $P(\delta_j|z_j^{\alpha}) \approx P(\delta_j)$. This is equivalent to ignoring the redshift evolution of δ, normally proportional to $(1+z_j^{\alpha})^{\kappa}$, which we find does not change the shape of $\langle\delta\gamma\rangle_A$.

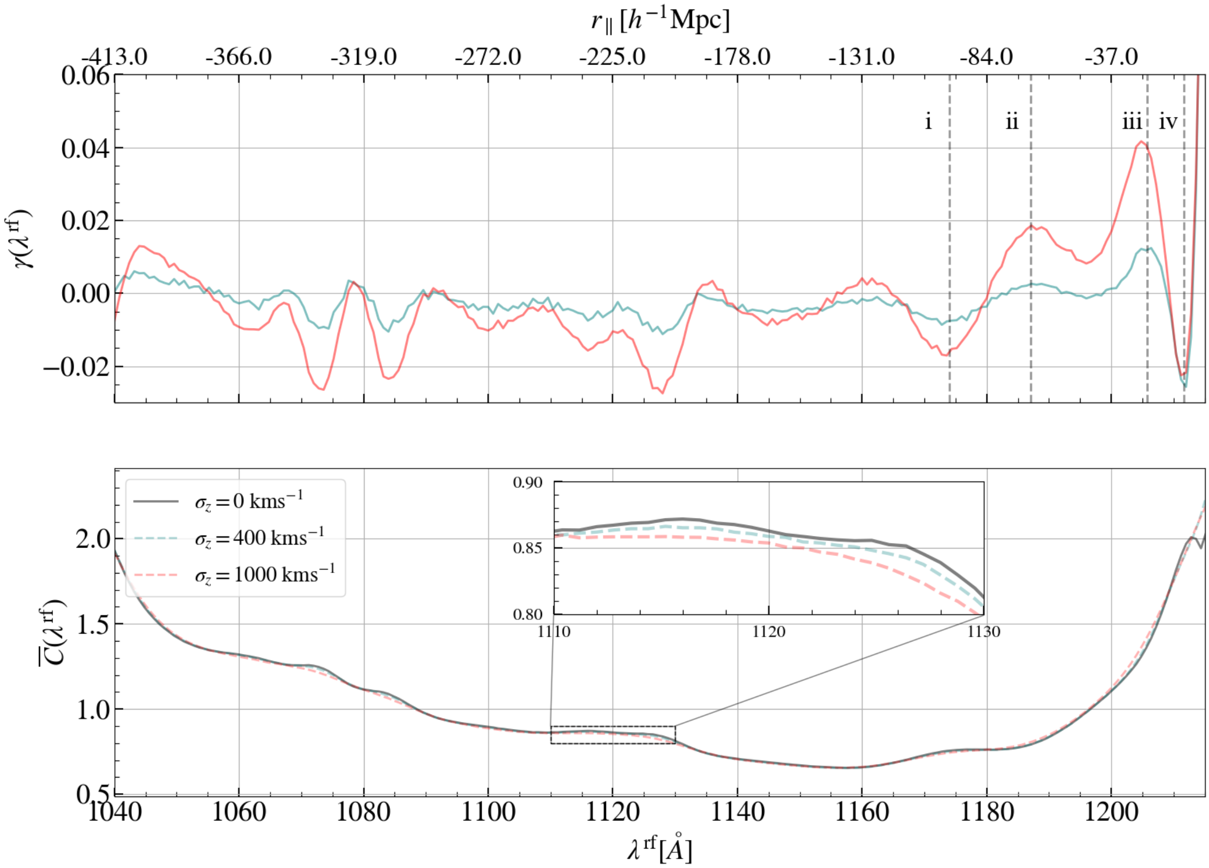

Figure 3. (Top) the mean continuum distortion function $\gamma = \hat{\overline{C}}/\overline{C} - 1$ as a function of rest-frame wavelength. $\hat{\overline{C}}(\lambda_{\rm rf}, \sigma_{\rm v})$ is the mean continuum with redshift errors $\sigma_{\rm v}$, as shown in the plot above. Roman numerals mark the location of prevalent features in γ, which contaminate our correlation functions. We include the approximate comoving distance between a quasar at $z = 2.3$ and pixels in its forest, along the top axis. (Bottom) the Lyα forest mean continuum (equation 2) of our mock datasets with 0 kms^{-1}, 400 kms^{-1} and 1000 kms^{-1} of redshift errors added. The later is for visualisation while 400 kms^{-1} is used in the actual analysis.

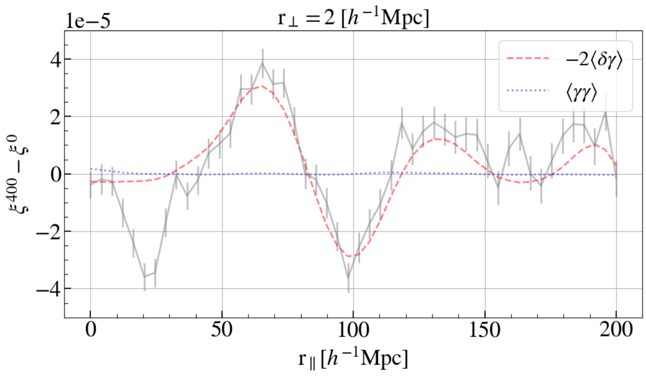

Figure 4. The difference in the Lyα autocorrelation function between contaminated (with continuum redshift errors) and uncontaminated datasets (grey), overlaid with direct measurements of $\langle\delta\gamma\rangle$ (red dashed) and $\langle\gamma\gamma\rangle$ (blue dotted). We plot only the first $r_\perp$ bin where the contamination is strongest.

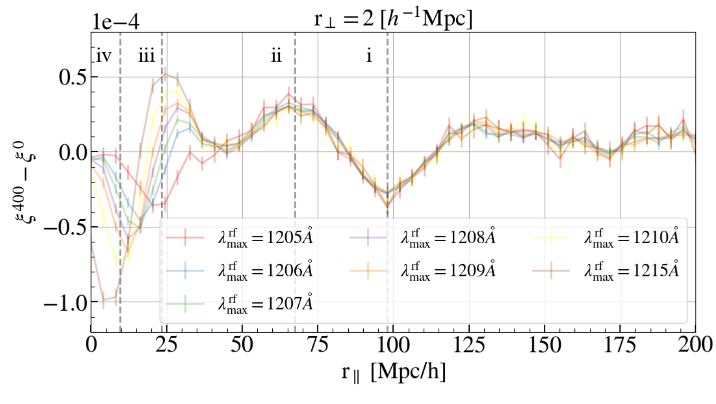

Figure 5. The evolution of continuum redshift error contamination with the maximum rest-frame wavelength of the Lyα forest. The 1205Ålimit (red) is the limit used in the analysis of this paper. We include Roman numerals to indicate the features in γ (figure 3) which correspond to spurious correlations in this figure.

Next, we express the first term in equation 31 as $P(z_i^{\rm Q}|\delta_j, z_j^\alpha) = P(z_i^{\rm Q})(1 + \delta_j \xi^{\rm X}(\vec{r}^{\rm X}))$, where $\vec{r}^{\rm X}$ is the separation between z_j^α and $z_i^{\rm Q}$, and $\xi^{\rm X}$ is the Lyα-quasar cross-correlation. Note that the dependence of $P(z_i^{\rm Q}|\delta_j, z_j^\alpha)$ on z_j^α is accounted for by the redshift evolution of the cross-correlation. This is not an analytically derived expression, but rather an ansatz that produces the expected behaviour of $P(z_i^{\rm Q}|\delta_j, z_j^\alpha)$. The definition of $\xi^{\rm X}$ here is somewhat analogous to the definition of the quasar autocorrelation as a measure of excess

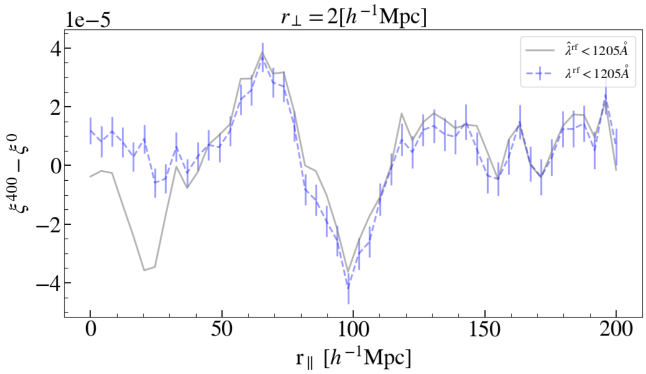

Figure 6. Contamination caused by redshift errors (grey), and the same contamination after removing pixels with true rest-frame wavelength greater than 1205 Å (blue dashed). After removing these pixels the trough feature at $\sim 25\,h^{-1}$Mpc disappears, while the rest of the contamination is unaffected.

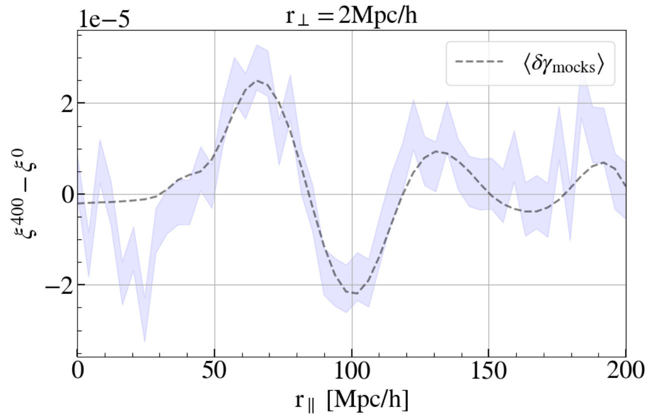

Figure 7. The contamination introduced by continuum redshift errors (blue shaded) as a function of $r_\parallel$, for the first (most contaminated) $r_\perp$ bin. The black dashed line is the contamination model (equation 34), estimated using a measurement of $\gamma(\lambda^{\rm rf})$.

probability, $\xi^{\rm Q} = P(z_j^{\rm Q}|z_i^{\rm Q})/P(z_j^{\rm Q}) - 1$, but with a modulating δ term. Making the subsequent substitutions into equation 30 gives:

$$\langle \delta\gamma \rangle_A = \int_{-1}^{\infty} d\delta_j \int_{z^{\rm Q,min}}^{z^{\rm Q,max}} \int_{z^{\alpha,\rm min}}^{z^{\alpha,\rm max}} dz_i^{\rm Q} dz_j^{\alpha} P(\delta_j) P(z_j^{\alpha}) P(z_i^{\rm X}) \times (1 + \delta_j \xi^{\rm X}(\vec{r}^{\rm X}))\, \delta_j \gamma_i. \quad (33)$$

Multiplying the terms in brackets results in two parts (with and without the cross-correlation). The integration over δ is independent of the other two variables, and since it has mean 0, the part with one δ_j term vanishes. In the second term, $\int_{-1}^{\infty} P(\delta_j)\delta_j^2 d\delta_j$ evaluates to ~ 0.6. We introduce a free amplitude for this model $A^{\rm auto}_{\rm cont}$, which will absorb this value, but here we call this factor a_δ, and write our final expression as:

$$\langle \delta\gamma \rangle_A = a_\delta A^{\rm auto}_{\rm cont} \int_{z^{\rm Q,min}}^{z^{\rm Q,max}} \int_{z^{\alpha,\rm min}}^{z^{\alpha,\rm max}} dz_i^{\rm Q} dz_j^{\alpha}\, P(z_j^{\alpha}) P(z_i^{\rm Q}) \xi^{\rm X}(\vec{r}^{\rm X}) \gamma_i. \quad (34)$$

Qualitatively, we evaluate this model by iterating over correlation function bins $(r_\perp, r_\parallel)$. For a bin A (width $4\,h^{-1}$Mpc in our analysis), we have N bins of $z_i^{\rm Q}$ covering the redshift range of the quasars in the dataset. This results in an array of N×M pixel redshifts z_i^{α} (equation 28), using M bins of rest-frame wavelength in the Lyα forest region ($\lambda^{\rm rf} \in [1040, 1205]$Å). With these values we can compute $P(z_i^{\rm Q})$ from our normalised quasar redshift distribution, and the γ_i term for each rest-frame wavelength in M. Then, given z_i^{α} and $r_\parallel \in A$, we can derive z_j^{α} and the vector $\vec{r}^{\rm X}$ between z_j^{α} and $z_i^{\rm Q}$. Finally, we compute the probability of z_j^{α} using a normalised pixel redshift distribution, and evaluate the cross-correlation $\xi(\vec{r}^{\rm X})$.

In figure 7, we present the model compared to the difference of correlation functions with and without redshift errors. To compute this, we directly measure $\gamma(\lambda)$ from the mock datasets (via equation 21), and input it to equation 34. In theory, one could also use a measurement of γ from mock datasets to fit our model on data. However, this relies on knowing the amount of redshift error present in the data, so we opt instead for building a simple model described in the next section.

4.3 Modelling γ

In real data we cannot directly measure $\gamma(\lambda^{\rm rf})$, since we don't know the true continuum (without redshift errors). Therefore, we need to construct a model of $\gamma(\lambda^{\rm rf})$.

Recalling equation 21, we see γ contains the ratio of the mean continuum in the presence of redshift errors ($\hat{\overline{C}}$), to the true mean continuum ($\overline{C}$). Therefore, to construct a model of γ, we need a template quasar continuum and a parameter that controls the amount of redshift error to input. The continuum template is a convolution of a set of emission lines and a smooth broadband component. For the former we use the emission line properties of the composite model of BOSS spectra DR9 (Harris et al. 2016). We then approximate the smooth component as flat, motivated by the fact that the majority of the contamination we observe comes from the smoothing of emission lines. To produce $\hat{\overline{C}}$ from our template $\overline{C}$, we need to emulate the same smoothing. We do this by broadening the emission lines in our template $\overline{C}$ (listed in table 1), via a parameter $\sigma_{\rm cont}$ representing the amount of the velocity dispersion in kms^{-1} [6]. The width of each emission line is then broadened according to:

$$\hat{\sigma}_{\rm line} = \sqrt{\sigma_{\rm line}^2 + \left(\frac{\lambda_{\rm line}\sigma_{\rm cont}}{c}\right)^2}, \quad (35)$$

where $\sigma_{\rm line}$ is the intrinsic line width, $\lambda_{\rm line}$ is the rest-frame wavelength of the line, and c is the speed of light in kms^{-1}. Note that since the emission lines are normalised Gaussians, the line amplitudes are $\propto 1/\hat{\sigma}_{\rm line}$. With the broadened emission lines we can construct $\hat{\overline{C}}(\lambda, \sigma_{\rm cont})$, and obtain $\gamma(\lambda)$. As mentioned earlier, this is robust to the choice of smooth component but highly dependent on the emission line model.

The result of fitting our model directly to the redshift error contamination is shown in figure 8 (using a model for γ). We show fits with the small-scale correction of appendix B (green solid line) and without (red dashed line), which we see are identical with the exception of the small-scale trough. In the full analysis (section 5) we choose not to use the small-scale correction model, to avoid biasing our results. We are still able to fit our model and recover the true

[6] The velocity dispersion here is related to redshift errors.

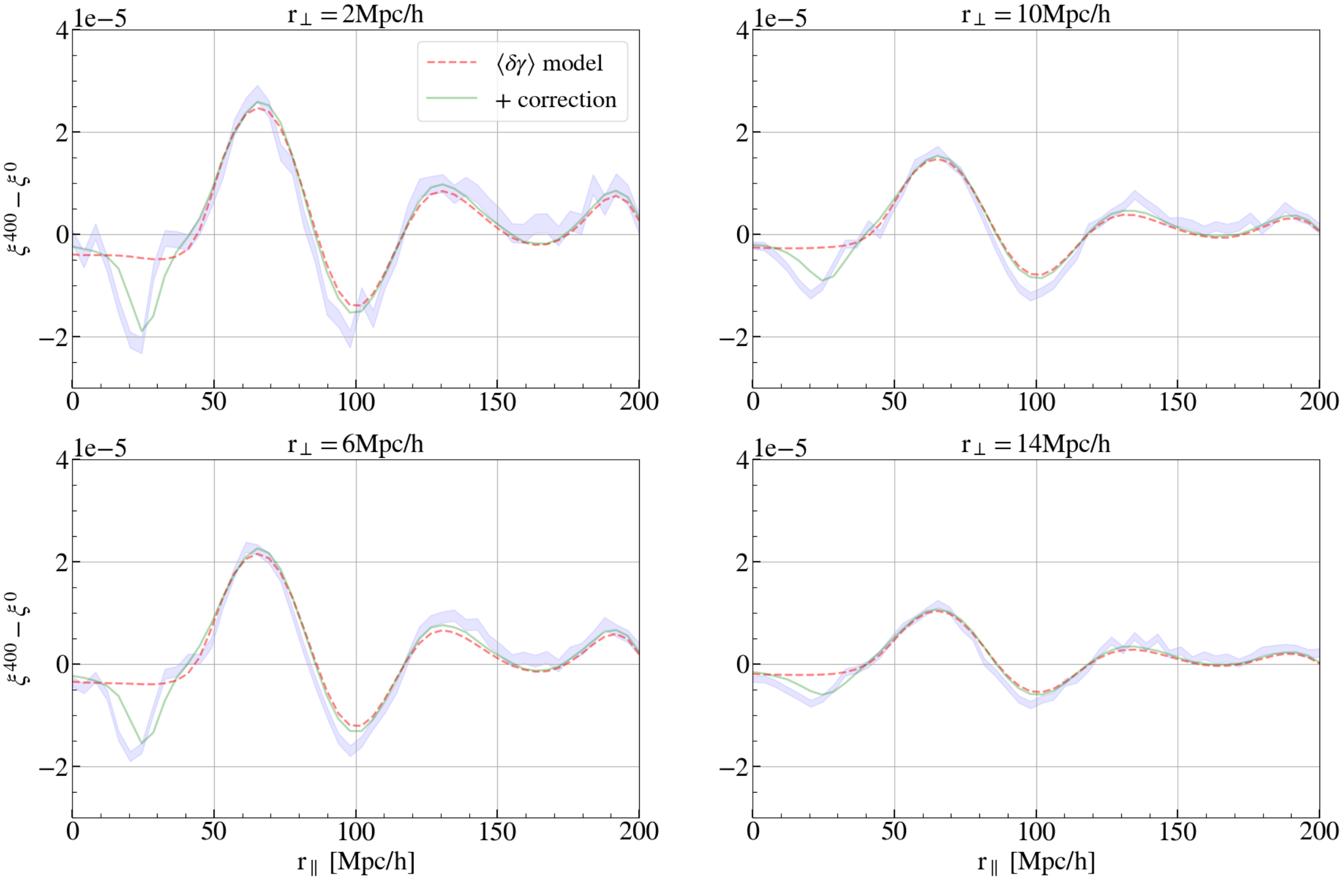

Figure 8. Direct fits to the contamination introduced by redshift errors, or the difference in the contaminated and uncontaminated Lyα autocorrelation, for the first 4 bins in $r_\perp$. We plot our model with (+correction) and without ($\langle \delta\gamma \rangle$ model) the addition of small-scale correction described in appendix B. Note we use the $\gamma(\lambda^{\rm rf})$ model described in section 4.3 to compute $\langle \delta\gamma \rangle$, rather than a measurement from mocks. In this plot we show the stack of 100 DESI DR1 mocks, used throughout this analysis.

input $\sigma_{\rm cont}$ without it, but may fail to capture the full impact of the trough at $\sim 25\,h^{-1}$Mpc. We see that the contamination is well fit at all scales, and that it drops off quickly with transverse separation.

4.4 Cross-correlation model

In Youles et al. (2022) they introduced a model of continuum redshift errors for the Lyα-quasar cross-correlation. The contamination in this case is characterised as: $\langle\hat{\delta}\rangle = \langle\delta\rangle - \langle\gamma(\lambda^{\rm rf})\rangle$, where $\langle\hat{\delta}\rangle = \hat{\xi}^{\rm X}$. The model of $\gamma(\lambda^{\rm rf})\rangle$ is then constructed using a similar method to the one in section 4.2, resulting in:

$$\langle\gamma\rangle_A = A^{\rm cross}_{\rm cont} \int\int dz_i^{\rm Q} dz_j^{\rm Q} P(z_i^{\rm Q}) P(z_j^{\rm Q}) \xi^{\rm Q}(\vec{r}^{\rm Q})\gamma_i, \tag{36}$$

where there is now dependence on the quasar autocorrelation $\xi^{\rm Q}$, rather than the Lyα-quasar cross-correlation. To include the Youles et al. (2022) model in our fits, we use the same γ model as in the autocorrelation (section 4.3). For this we have two free parameters, $\sigma_{\rm cont}$ and $A^{\rm cross}_{\rm cont}$. Note that in the cross-correlation the model of Youles et al. (2022) also fails to capture the behaviour of redshift errors at $\sim 25\,h^{-1}$Mpc, where instead of a trough we now see a peak (visible in figure 10). However, because we limit our cross-correlation to scales above $40\,h^{-1}$Mpc (see section 5), the impact of this peak is much smaller.

We ultimately combine both correlation functions in a joint fit, shown in section 5. This allows us to break degeneracies between certain nuisance parameters and key parameters like $\phi_{\rm s}$, $\alpha_{\rm s}$. As explained in Cuceu et al. (2021), we can also measure the combination $f\sigma_8$ from the joint fit. We expect γ to be the same for the auto- and cross-correlations, so we use only one parameter $\sigma_{\rm cont}$ for both. We have verified that allowing each correlation function to have a different set of parameters give perfectly consistent results.

5 RESULTS

In this section we will explain our model fitting procedure, and show the impact of redshift errors on the cosmological parameters of our full-shape analysis. In particular, we discuss how large this contamination is with respect to the precision of DESI DR1 and DR2. We will then present the full-shape fits including the redshift errors model for the auto- (section 4.2) and cross-correlations (section 4.4). We show results of fits to these correlations individually, and together in a joint fit which gives constraints on the combination $f\sigma_8$.

5.1 Parameter degeneracies

The RSD parameter of the Lyα forest is given by:

$$\beta_\alpha = \frac{b_{\eta,\alpha(z)} f(z)}{b_\alpha(z)}, \tag{37}$$

where $b_{\eta,\alpha}$ is the velocity divergence bias of the Lyα forest and f is the logarithmic growth-rate. In practice we must assume a linear matter power spectrum template (Planck Collaboration et al. 2020), which has a fixed normalisation proportional to $\sigma_8(z)$ - the amplitude of perturbations in $8h^{-1}$Mpc spheres. In linear theory f and σ_8 are fully degenerate (Percival & White 2009), so we are sensitive to the combinations $b_{\eta,\alpha} f\sigma_8$ and $b_\alpha\sigma_8$. Since $b_{\eta,\alpha}$ is unknown, we treat β_α as a nuisance parameter to marginalise over.

For the cross-correlation, we are sensitive to both the quasar and Lyα RSD and bias parameters. Since for quasars $\beta_q = f/b_q$, we are sensitive to $\{b_\alpha\sigma_8, b_{\eta,\alpha} f\sigma_8, b_q\sigma_8, f\sigma_8\}$. $f\sigma_8$ is difficult to constrain from the cross-correlation alone, since it's dependent on several biases ($b_{\eta,\alpha}, b_\alpha, b_q$). But we break these degeneracies by combining both correlation functions in a joint fit, as the autocorrelation provides precise measurements of $b_{\eta,\alpha} f\sigma_8$ and $b_\alpha\sigma_8$.

5.2 Fits

Following the full-shape analysis of DESI year-1 Lyα data (Cuceu et al. 2025a), we restrict the autocorrelation to $r \in [25, 180]\ h^{-1}$Mpc, and cross-correlation to $r \in [40, 180]\ h^{-1}$Mpc. This is a conservative minimum scale cut designed to limit the impact of increasingly non-linear scales and systematics (including redshift errors). We highlight at this point that there are some key differences between the analysis we perform here, and that of Cuceu et al. (2025a). First of all, as mentioned in section 2, Cuceu et al. (2025a) use two additional correlation functions measured from Lyα absorption in Lyβ region (between 920-1020 Å). These are the autocorrelation of Lyα absorption in the Lyβ region, and its cross-correlation with quasars. We omit these correlations from our analysis for simplicity, since we would need to extend our continuum redshift errors model to incorporate this new region. We also do not expect redshift errors in the Lyβ region to have a large impact overall, since these correlation functions have low statistical power relative to the two we analyse in this paper. Furthermore, in this paper we measure the AP effect separately on the broadband (ϕ_s) and peak component (ϕ_p) of the correlation function, as opposed to Cuceu et al. (2025a) who measure the AP effect across the full correlation function with one parameter (ϕ_f). We opt for the former to characterise the impact of redshift errors more precisely, but it may result in a larger overall shift. Finally, we use a Gaussian prior on $L_{\rm HCD}$, compared to Cuceu et al. (2025a) who use a wide uniform prior. We do this to make fitting our new redshift error model parameters easier, but it may also lead to slight differences in our results.

Since we are only interested in our set of cosmological parameters $\{\alpha_p, \phi_p, \phi_s, f\sigma_8\}$ and the parameters of the model introduced in this paper $\{\sigma_{\rm cont}, A^{\rm auto}_{\rm cont}, A^{\rm cross}_{\rm cont}\}$, we treat all other parameters as nuisance. Note that we have not included α_s in our parameters of interest, since it is not simple to extract cosmological from it.

In the autocorrelation fits, we have 15 free parameters when using the redshift errors model, and 13 without it. We fix $\beta_{\rm HCD}$ and $L_{\rm HCD}$ to values constrained in the joint fit, and set $f\sigma_8$ to value of the fiducial cosmology. For our model (equation 34), we also need the redshift distributions of Lyα pixels and quasars in our dataset, our gamma model $\gamma(\lambda^{\rm rf})$ and an input cross-correlation function $\xi^{\rm X}$. Ideally, one would use a model (i.e. one that we fit in this analysis), but for simplicity and to reduce computing time we will use the measured cross-correlation from our mock dataset (figure 1). This is a reasonable approximation since the stack of mocks has very high signal-to-noise. For the DESI year-1 dataset, as we discuss in section 6, the cross-correlation is noisy and therefore is a not a good substitute for a model.

For the cross-correlation we have 16 free parameters with the redshift errors model and 14 without. We additionally fit for systematic redshift errors ($\Delta r_\parallel$), and smoothing at large $k_\parallel$ due to redshift errors and peculiar velocities ($\sigma_{\rm v}$). However, we now fix $\sigma_\parallel$ [7] to the joint fit value since it is degenerate with $\sigma_{\rm v}$. For this we also provide a quasar redshift distribution, the continuum smoothing function $\gamma(\lambda^{\rm rf})$ (see section 4.3) and the quasar autocorrelation $\xi^{\rm Q}$ (section 4.4). Again we make a simplification by using a measurement of the quasar autocorrelation in-place of a model.

The joint fits have have 22 free parameters, combining the auto- and cross-correlations. Our full list of priors for all free parameters in the joint fit is given in table A1 in appendix A. In most cases we use a wide uniform prior, except for the HCD length-scale ($L_{\rm HCD}$) and RSD parameter ($\beta_{\rm HCD}$), which where assign Gaussian priors based on previous studies (Pérez-Ràfols et al. 2018; Adame et al. 2025a). Our joint fits in both cases have an effective redshift z ∼ 2.3. We compute our full-shape model using the `Vega`[8] library, which also contains a set of analysis tools. For the auto- and cross-correlation fits, we use the `Vega` interface for `Iminuit`, which performs a fast χ^2 minimisation. For the joint correlations we use the nested sampler `Polychord`[9](Handley et al. 2015) to sample the posterior space of our parameters. The latter is slower, but is able to capture any non-Gaussianity in our posterior space, and gives better error estimates than a χ^2 minimisation.

In figure 9 we present the posteriors of $\phi, \alpha, \phi_s, f\sigma_8$, from nested sampler runs on the joint correlation function for both the uncontaminated and contaminated datasets. For the latter we include two sets of contours, for runs including (red dashed) and not including (blue solid) the redshift errors model presented in this paper. We first note that the errors on these differences are very small, because we are analysing the stack of 100 DESI year-1 mocks. Later, when looking at the full set of results in table 2, we will compare our results to the precision of DESI DR1 (Cuceu et al. 2025a). We also note that the uncontaminated contours are slightly shifted from our fiducial cosmology (Planck Collaboration et al. 2020), but we will not discuss this here since we are specifically concerned about the impact of redshift errors (shift from green to blue). In appendix A, we show the full triangle plot of our joint correlation run, and discuss what effect our model has on other nuisance parameters. We see that, in general, the shifts introduced by continuum redshift errors are small compared to the precision of DESI DR1. In this paper, we consider any shift greater than $\sigma_{\rm DR1}/3$ to be significant, in-line with the requirements of Adame et al. (2025a). It should be noted however, that the shift we refer to here is only due to the impact of continuum redshift errors. In other words, this shift is not equivalent to a bias with respect to the fiducial cosmology (dashed crosshairs in figure 9).

From figure 9, we see ϕ_p is more affected than α_p, but the shift is still minor. The smooth AP parameter (ϕ_s) is most biased by continuum redshift errors because it depends on the broadband component of our correlation functions, which is more affected by contaminants in general. The observed shift may also be smaller because of the conservative minimum separations we choose for the auto- (25 h^{-1}Mpc)

[7] The line-of-sight smoothing due to limited simulation grid size.
[8] https://github.com/andreicuceu/vega
[9] https://github.com/PolyChord/PolyChordLite

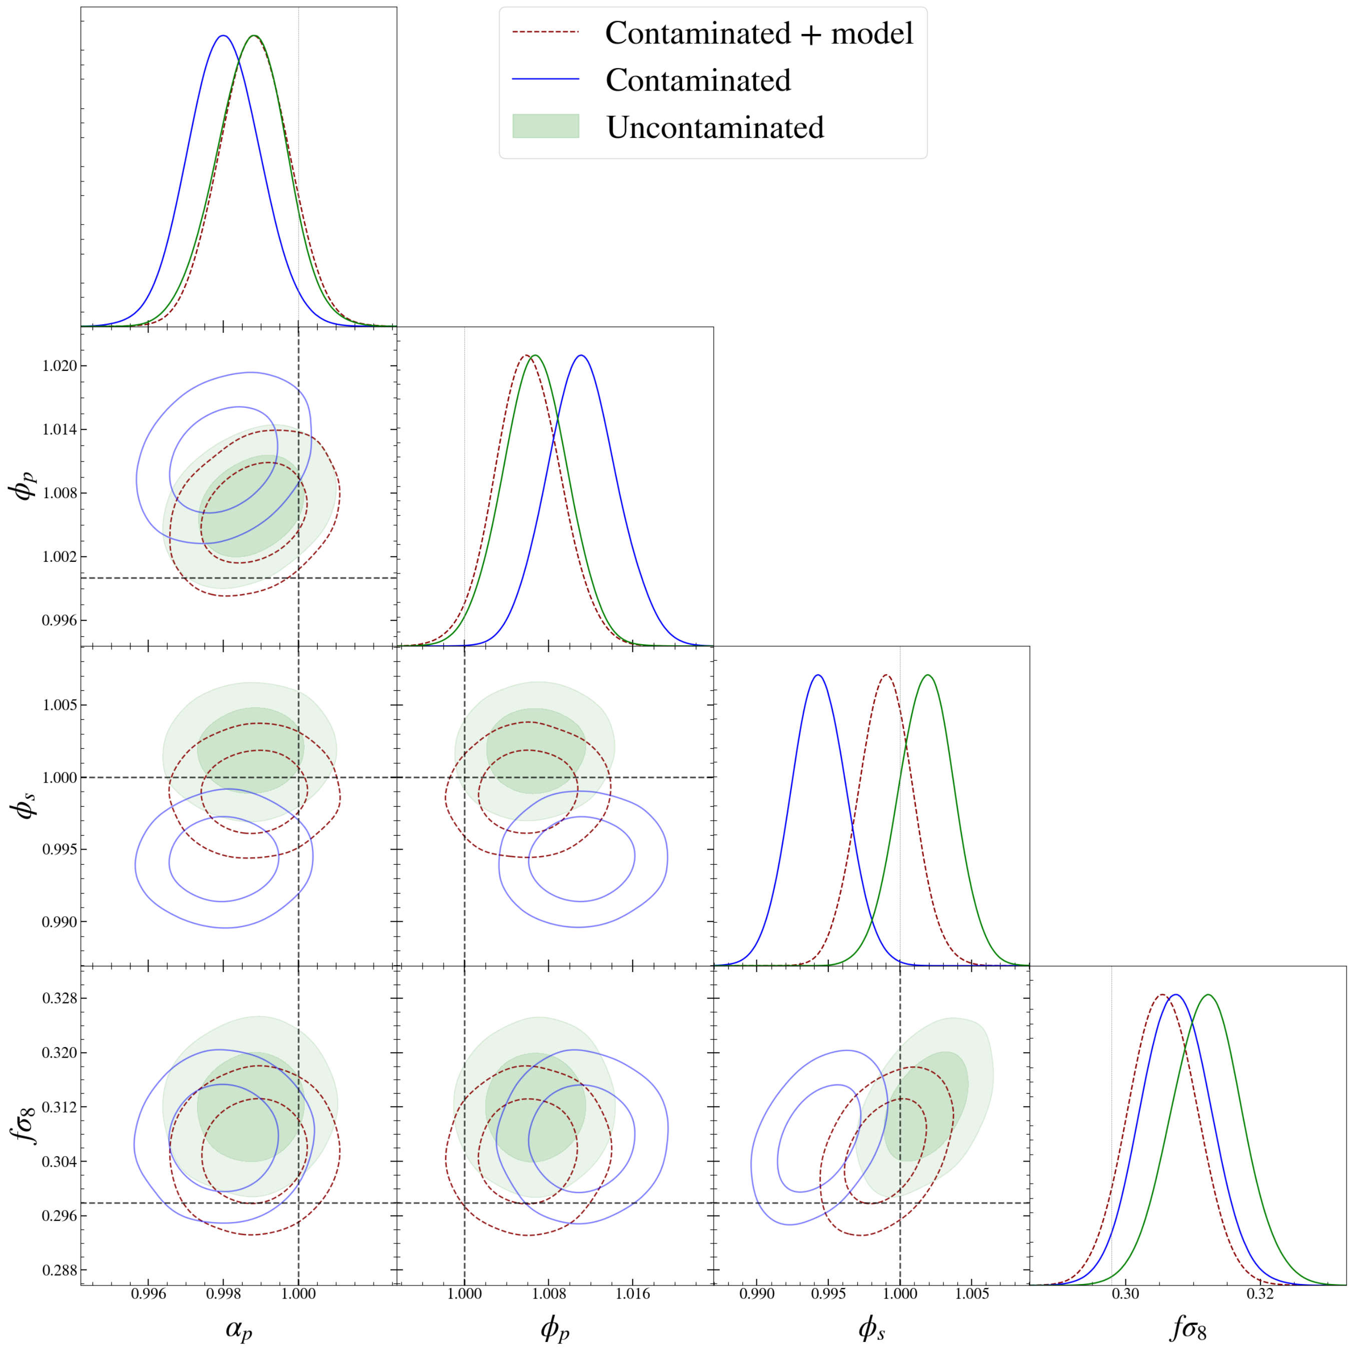

Figure 9. Full-shape posteriors from fits to a stack of 100 DESI DR1 mocks. We compare fit results from an uncontaminated (without redshift errors) dataset (green shaded), to fits on a contaminated dataset with (red dashed) and without our model (blue solid). Our fiducial cosmology is also indicated with black dashed crosshairs.

and cross-correlations ($40\,h^{-1}$Mpc). A small but insignificant shift occurs in the $f\sigma_8$ posterior, which is also sensitive to the broadband of the correlation functions.

After introducing our model (red dashed contours), we see an impact on each of our posteriors. Firstly, we successfully recover the BAO peak parameters (α_p and ϕ_p) from the uncontaminated case. This result shows that our model can be used to correct for continuum redshift errors on BAO measurements, without considering a small-scale correction (e.g. equation B8).

The shift introduced by continuum redshift errors on ϕ_s is reduced significantly upon introduction of our model, but still differs slightly from the uncontaminated case. As discussed in section 4, our model does not capture the redshift error contamination present at $\sim 25\,h^{-1}$Mpc, which is responsible for the remainder of the shift. The $f\sigma_8$ constraint is shifted very marginally further from the uncontaminated case when using our model, towards the fiducial value (black dashed crosshairs). Indeed, it seems that all of our posteriors are shifted marginally towards the fiducial value, which suggests our new model is capturing some imperfections in the existing model. However, as we will show in appendix A, our model parameters are not correlated with the cosmological parameters of interest, thus we are fitting for some other "nuisance" effect. The shift in $f\sigma_8$ is still small and consistent with the uncontaminated case to within 1-sigma.

In table 2 we show the differences in the cosmological parameters of our full-shape model, between analyses on uncontaminated and contaminated mocks. The columns under "+Model" give the same differences after adding our models of the contamination (section 4) to the full-shape fits. We ignore nuisance parameters constraints, except those of the model introduced in this paper, which we show in table 3. As noted earlier, we are analysing the stack of 100 DESI DR1 mocks with high signal-to-noise. For that reason, we include the projected precision of the DESI year-1 full-shape analyses on data in the final column of the table (see Cuceu et al. (2025a)).

The shift on α_p (isotropic BAO parameter, equation 14) is never more than $\sim$ 10% of the DR1 68% confidence level ($0.1\sigma_\mathrm{DR1}$). For ϕ_p (AP effect on the BAO peak, section 2.4) and the growth-rate, we also see shifts of $\sim 0.1\sigma_\mathrm{DR1}$ in the joint fit. For each of the latter 3 parameters, the shifts introduced by redshift errors are notably consistent with zero to $\sim 1\sigma$. However, the shift on ϕ_s is $0.44(\pm 0.13)\sigma_\mathrm{DR1}$, which is larger than the $\sigma_\mathrm{DR1}/3$ criterion we set earlier. Due to large computing time, errors on these shifts are estimated from the stack of mocks, rather than the distribution of shifts across the 100 individual mocks. This approximation is justified because the shift is systematic in nature, and we are using the same mock realisations on both sides of the comparison, so the impact of noise is minimised.

While we have only included DESI DR1 projections, it is important to note that the relative sizes of the shifts shown here will increase significantly for the DESI DR2 and future datasets. The DR2 constraints for example should be approximately ~40% tighter than DR1, making the ϕ_s shift $\sim 0.7\sigma_\mathrm{DR2}$. It is also important to stress that this shift is not a bias with respect to the fiducial cosmology, rather it characterises the impact of a single contaminant. It may be the case that parametrising the AP effect over the full correlation function instead of the peak and broadband components separately, as is done in Cuceu et al. (2025a), reduces the impact of redshift errors.

The next 3 columns (+ model) show the same differences as the first 3, but including the continuum redshift errors model. For the peak component parameters ($\alpha_\mathrm{p}, \phi_\mathrm{p}$), our model is very successful in removing the bias for all 3 fit cases. For ϕ_s, our model successfully reduces the bias in the joint case to $0.19(\pm 0.13)\sigma_\mathrm{DR1}$, now consistent with zero to 1.5σ. Our model has almost no impact on $f\sigma_8$, but the shift is small $\sim 0.1\sigma_\mathrm{DR1}$ and consistent with zero to $\sim 1\sigma$. The parameter values of our model constrained in the analysis are shown in the section below.

As previously mentioned, we have a single set of parameters controlling the shape of γ: σ_cont, $A^\mathrm{auto}_\mathrm{cont}$ (autocorrelation amplitude) and $A^\mathrm{cross}_\mathrm{cont}$ (cross-correlation amplitude), which have values given in table 3. In each case we are able to recover the input redshift error ($400\,\mathrm{km\,s^{-1}}$) to within a 1σ confidence interval, and have strong detections of the other 3 parameters. The cross-correlation amplitude parameter is consistent between the cross and joint cases; however, this is not true for the auto-correlation amplitude. The tension between the auto and joint cases may reflect the incompleteness of the model at small-scales. For each of these parameters we apply wide uniform priors, shown in table A1 with the rest of the parameters in the joint fit.

The χ^2 probabilities of the analyses both with and without our model are very low. This is expected, since we are fitting an extremely high signal-to-noise dataset with a linear theory model (section 2.5) that is suboptimal. Furthermore, it has been shown in Cuceu et al. (2025a) that our model is suitable for at least the level of precision of DESI DR1. We see a large improvement in χ^2 value when using our model in the autocorrelation ($\Delta\chi^2$ = 655, 1564 bins, 15 parameters), the cross-correlation ($\Delta\chi^2$ = 726, 3030 bins, 16 parameters) and the joint ($\Delta\chi^2$ = 970, 4594 bins, 22 parameters) correlation runs.

6 DISCUSSION

In the previous section, we showed and discussed the results of fits to synthetic data sets with the model introduced in this paper. In section 6.1 we will discuss the application of this model to real data, and highlight potential differences with the synthetic datasets we use. We will also show in section 6.2, that the effect of continuum redshift errors can be mitigated by removing pairs where the quasar auto-correlation (for contamination in the Lyα-quasar cross-correlation) or the Lyα-quasar cross-correlation (for contamination in the Lyα autocorrelation) is large.

6.1 Contamination on real data

In Youles et al. (2022), they found a strong dependence of $\langle\gamma\rangle$ on small-scale quasar clustering. Likewise, we verify that $\langle\delta\gamma\rangle$ has strong dependence on the small-scale cross-correlation function. This is discussed further in section 6.2. We also find that the quasar clustering in our lognormal mocks deviates more from linear theory at small-scales than more realistic n-body simulations (see figure 6 of Youles et al. (2022)), which is likely to increase the amount of contamination relative to real data.

As mentioned in section 5.2, we use a measurement of the cross-correlation in our $\langle\delta\gamma\rangle$ model, and a measurement of the quasar autocorrelation in $\langle\gamma\rangle$. For our mock datasets, where we have high signal-to-noise these measurements are smooth, and are a better alternative to using a model of the respective correlation functions. However, in e.g. DESI DR1, our correlations are relatively noisy and, in the case of $\hat{\xi}^X$, and are themselves affected by redshift errors. It is also important to note that because of our strong dependence on small-scale clustering, using a linear model of ξ^X in equation 34 (or ξ^Q in equation 36) would likely be a bad approximation. For these reasons, we choose to leave a full analysis on data for future work where we have a higher signal-to-noise data set, or a cross-correlation (and quasar autocorrelation) model that is realistic at small-scales. At the level of precision of DESI DR1, our model is also not yet necessary to perform an un-biased analysis (i.e. Cuceu et al. (2025a)).

Another contributing factor in our analysis realism is our synthetic spectra, or specifically the quasar continua that we use. As we discussed in section 2, we generate quasar spectral templates using convolutions of power law smooth components, and emission lines (table 1) from the composite spectrum of BOSS quasars (McGreer et al. 2021; Harris et al. 2016). We sample from a Gaussian distribution of power law slopes, tuned to have mean and scatter that better reflected the eBOSS DR16 dataset (du Mas des Bourboux et al. 2020). Likewise, for the Lyα forest emission lines we sample from a Gaussian distribution of EWs, tuned on BOSS DR9 data. The distribution of emission line EWs is particularly important for our study, since the dominant contribution to the contamination is around the positions of these. If, for example, the mean EW in our spectra was higher than in e.g. DESI DR1, we would observe a higher level of contamination. Also, since the scales of the spurious correlation features are directly related to the rest-frame wavelength of the emis-

Parameter	Contaminated fit			+ Model			σ_{Y1}
	auto	cross	joint	auto	cross	joint	
$10^3\Delta\alpha_{\mathrm{p}}$	0.1 ± 1.4	1 ± 1.2	0.8 ± 1.0	-0.2 ± 1.4	0.03 ± 1.2	-0.1 ± 0.9	11
$10^3\Delta\phi_{\mathrm{p}}$	-8 ± 5	-4 ± 4	-4 ± 3	-5 ± 5	1 ± 4	-0.7 ± 3	38
$10^3\Delta\phi_{\mathrm{s}}$	4 ± 2.4	10 ± 4	-7 ± 2	3 ± 2.4	10 ± 3	-3 ± 2	16
$10^2\Delta f\sigma_8$	-	-	0.5 ± 0.5	-	-	0.6 ± 0.5	6.2

Table 2. Differences in parameter constraints on fits to mocks with and without redshift errors. We fit a stack of 100 DESI DR1 contaminated mocks, without (Contaminated fit) and with (+ Model) our continuum redshift errors model, and measure the shift with respect to the uncontaminated set. The last column shows the projected errors (68% confidence interval) on the same parameters for the joint constraints from DESI DR1 dataset. Note that the values presented here are shifts introduced solely by continuum redshift errors, not the bias with respect to the fiducial cosmology.

Parameter	auto	cross	joint
$\sigma_{\mathrm{cont}}\,[\mathrm{kms}^{-1}]$	401 ± 19	406 ± 11	411 ± 18
A_γ^{auto}	-6.6 ± 0.4	-	-8.9 ± 0.4
$A_\gamma^{\mathrm{cross}}$	-	-6.2 ± 0.3	-5.8 ± 0.2

Table 3. Constraints on free parameters of the model we introduce to capture continuum redshift errors (section 4), from a fit to a stack of 100 DESI DR1 contaminated mocks. We show the constraints (section 4) for the auto-, cross- and joint correlations separately. The input amount of redshift error was $\sigma = 400\,\mathrm{kms}^{-1}$, which we recover to within 1σ in the each case.

sion lines in our Lyα forests (see section 4.1), we are sensitive to the relative strength of individual lines.

The way we introduce redshift errors into our spectra should also reflect the magnitude and nature of redshift errors in real data. However, this is difficult to emulate because there are several ways in which redshift errors are produced. We add errors drawn from a Gaussian with dispersion $\sigma = 400\,\mathrm{kms}^{-1}$ to each quasar in our sample, representing errors which might occur in our redshift estimation pipeline `Redrock`. However, as mentioned in section 1, it is also typical for emission lines to shift away from the systemic redshift of the quasar. The level of shift depends on the emission line, but is larger in general for broad, high-ionisation lines like CIV, SiIV and CIII. Therefore, the level of error depends on the set of lines being used to estimate the quasar redshift. There has been progress made in this area recently, with work on improving spectral templates in DESI (Bault et al. 2025; Brodzeller et al. 2023).

There is further complication due to the fact that the contaminating lines within the Lyα forest are high-ionisation. If these are strongly correlated with the high-ionisation lines used to estimate quasar redshift, then there would be little impact on the mean continuum. The latter point requires further studies like the one of Shen et al. (2016); Bault et al. (2025); Brodzeller et al. (2023), but is difficult in practice because Lyα forest lines are less prominent, and visible at a higher redshift ($z \gtrsim 2.1$). One could improve the way we add redshift errors to mocks using these studies to input typical relative line shifts into our spectra, and use a redshift estimator (i.e. `Redrock`) to estimate our redshifts.

6.2 Mitigating contamination in real data

Youles et al. (2022) identified that the level of redshift error contamination in the cross-correlation between Lyα and quasars, depends on the correlation between the host quasar of the forest and the correlating quasar. This correlation is larger for neighbouring pairs of quasars, and avoiding these configurations significantly reduces the contamination. Here we extend this further by explicitly removing Lyα-quasar pairs where the forest host-quasar and correlating quasar are separated by less than a given separation. The situation is similar for the Lyα autocorrelation, except that, as we showed in section 4.2, the contamination is now dependent on the Lyα-quasar cross-correlation.

The result of removing contributions to our model from pairs where the host quasar of a Lyα pixel and the correlating quasar are separated by less than r_{q}, is shown in figure 10 (bottom). We choose two cuts at $5\,h^{-1}$Mpc and $10\,h^{-1}$Mpc, which removes only $\sim 0.1-0.3$% of our total number of pairs, but $\sim$ 10% in the first 5 transverse bins (up to $r_\perp \sim 20\,h^{-1}$Mpc). From figure 10, we see that by $10\,h^{-1}$Mpc, the contamination is effectively negligible.

In the top panel of figure 10 we show an analogous test with the contamination in the Lyα autocorrelation. Here we remove pairs where the host quasar of one forest is separated from the correlating pixel in the other forest by less than $r_{\mathrm{X}} = 5, 10\,h^{-1}$Mpc. In this case, we see a similar same level of success in removing the contamination. The fraction of correlating pairs that is removed in this test is also sub-percent level overall, but $\sim$ 10% in the most line-of-sight bins.

These tests were undertaken in Casas et al. (2025); DESI Collaboration et al. (2025a) and shown to have very little impact on the precision of BAO constraints, where the most line-of-sight bins are at much smaller separation than the BAO feature. However, there was a significant impact on the Lyα bias, which may suggest this approach is not suitable for full-shape analyses. We leave this to be studied in future work, as a possible alternative to the model we present in this paper.

7 SUMMARY

Random errors in our quasar redshift measurements smooth the mean continuum used to estimate the Lyα flux transmission field. This smoothing gives rise to spurious features across the length of the Lyα forest, which in-turn distort the full-shape of the Lyα autocorrelation and its cross-correlation with quasars.

In this paper we presented a model of these spurious correlations for the Lyα autocorrelation function, building upon work in Youles et al. (2022), where they presented an equivalent model for the Lyα-quasar cross-correlation. We then created a model for the distortion as a function of rest-frame wavelength ($\gamma(\lambda^{\mathrm{rf}}, \sigma_{\mathrm{cont}})$), and introduced

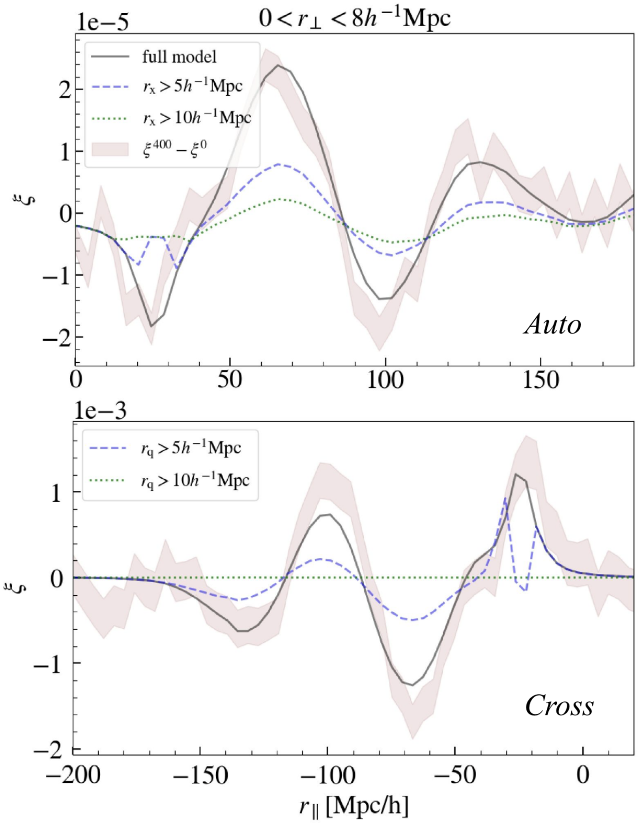

Figure 10. (Top) the contamination in the Lyα autocorrelation introduced by redshift continuum errors and our model of it (black solid). Also plotted is the same model evaluated only for separations r_X greater than $5\,h^{-1}$Mpc (blue dashed) and $10\,h^{-1}$Mpc (green dotted), where r_X is the separation between one pixel and the host quasar of the other pixel. (Bottom) continuum redshift errors contamination in the Lyα-quasar cross-correlation, over-plotted with the model of the same contamination. The model is evaluated with the same cuts as the autocorrelation, but with r_q now referring to the separation between a correlating quasar and the host quasar of the correlating Lyα forest pixel. In both cases we have used the small-scale correction of appendix B.

three new parameters to our standard full-shape model. These control the amount of smoothing introduced by redshift errors ($\sigma_{\rm cont}$), and the amplitude of the effect in the auto- ($A_\gamma^{\rm auto}$) and cross-correlation functions ($A_\gamma^{\rm cross}$).

We show using a high signal-to-noise synthetic dataset that continuum redshift errors shift the measurements of the growth rate ($f\sigma_8$), isotropic BAO parameter ($\alpha_{\rm p}$), anisotropic BAO parameter ($\phi_{\rm p}$) and anisotropic broadband parameter ($\phi_{\rm s}$) to varying degrees, with respect to a dataset which does not contain redshift errors. At the level of precision of DESI DR1, this shift is relatively minor in the first 3 parameters (∼10%). However, we find that $\phi_{\rm s}$ is shifted from the uncontaminated dataset (without redshift errors) value by $0.44(\pm0.13)\sigma_{\rm DR1}$, where $\sigma_{\rm DR1}$ is projected the 68% confidence region of the DESI DR1 constraint. Note that this relative shift will increase as analyses on future datasets (e.g. DESI DR2) improve the constraining power on these parameters. We demonstrated that our model reduces the shift introduced to $\phi_{\rm s}$ by ∼ 60%, and completely removes any shift on $\alpha_{\rm p}$ and $\phi_{\rm p}$. We also recover to within 1σ the input dispersion of the Gaussian distribution of redshift errors we add ($400\,{\rm km\,s^{-1}}$), and strongly detect each of the amplitude parameters of the model.

On real data, one should use models of the Lyα-quasar cross-correlation and the quasar autocorrelation as inputs for the redshift errors model in the autocorrelation ($\langle\delta\gamma\rangle$) and the cross-correlation ($\langle\gamma\rangle$) respectively. This is because measurements of these functions are noisy compared to the high signal-to-noise stack of mocks we used throughout this work. We also show that continuum redshift error contamination could be mitigated in real data by removing pairs in the cross-correlation where the quasar autocorrelation is strongest, and removing pairs in the autocorrelation where the cross-correlation is strongest. In section 6.2, we showed that removing pairs where $r_{\rm q}, r_{\rm X}$[10] $< 10\,h^{-1}$Mpc is enough to make $\langle\gamma\rangle$ and $\langle\delta\gamma\rangle$ negligible. This corresponds to ∼ 0.3% of total correlating pairs, but ∼ 10% of the most line-of-sight (up to $r_\perp \sim 20\,h^{-1}$Mpc) pairs, meaning this cut may be suitable for removing the effect of redshift errors on BAO constraints, but not for full-shape analyses.

ACKNOWLEDGEMENTS

CG is partially supported by the Spanish Ministry of Science and Innovation (MICINN) under grants PGC-2018-094773-B-C31 and SEV-2016-0588. AC acknowledges support provided by NASA through the NASA Hubble Fellowship grant HST-HF2-51526.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. AFR from the European Union's Horizon Europe research and innovation programme (COSMO-LYA, grant agreement 101044612). AFR acknowledges financial support from the Spanish Ministry of Science and Innovation under the Ramon y Cajal program (RYC-2018-025210) and the PGC2021-123012NB-C41 project. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya.

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE–AC02–05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science and Technology of Mexico (CONAHCYT); the Ministry of Science, Innovation and Universities of Spain (MICIU/AEI/10.13039/501100011033), and by the DESI Member Institutions: `https://www.desi.lbl.gov/collaborating-institutions`. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U. S. National Science Foundation, the U. S. Department of Energy, or any of the listed funding agencies.

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

[10] The separation between quasars in the quasar auto-correlation, and Lyα pixels and quasars in the cross-correlation respectively.

DATA AVAILABILITY

The data points from the figures in this paper will be available in a Zenodo repository when it is accepted for publication.

REFERENCES

Adame A. G., et al., 2025a, J. Cosmology Astropart. Phys., 2025, 124
Adame A. G., et al., 2025b, J. Cosmology Astropart. Phys., 2025, 021
Adame A. G., et al., 2025c, J. Cosmology Astropart. Phys., 2025, 012
Alexander D. M., et al., 2023, AJ, 165, 124
Bailey S. J., et al., 2025
Bault A., et al., 2025, J. Cosmology Astropart. Phys., 2025, 130
Bautista J. E., et al., 2017, A&A, 603, A12
Brodzeller A., et al., 2023, AJ, 166, 66
Busca N. G., et al., 2013, A&A, 552, A96
Casas L., et al., 2025, arXiv e-prints, p. arXiv:2503.14741
Chaussidon E., et al., 2023, ApJ, 944, 107
Cole S., et al., 2005, MNRAS, 362, 505
Cuceu A., Font-Ribera A., Joachimi B., Nadathur S., 2021, MNRAS, 506, 5439
Cuceu A., Font-Ribera A., Nadathur S., Joachimi B., Martini P., 2023a, Phys. Rev. Lett., 130, 191003
Cuceu A., et al., 2023b, MNRAS, 523, 3773
Cuceu A., et al., 2025a, in prep.
Cuceu A., et al., 2025b, J. Cosmology Astropart. Phys., 2025, 148
DESI Collaboration et al., 2016a, arXiv e-prints, p. arXiv:1611.00036
DESI Collaboration et al., 2016b, arXiv e-prints, p. arXiv:1611.00037
DESI Collaboration et al., 2022, AJ, 164, 207
DESI Collaboration et al., 2024, arXiv e-prints, p. arXiv:2411.12022
DESI Collaboration et al., 2025a, arXiv e-prints, p. arXiv:2503.14739
DESI Collaboration et al., 2025b, arXiv e-prints, p. arXiv:2503.14745
Eisenstein D. J., et al., 2005, ApJ, 633, 560
Ennesser L., Martini P., Font-Ribera A., Pérez-Ràfols I., 2022, MNRAS, 511, 3514
Farr J., et al., 2020a, J. Cosmology Astropart. Phys., 2020, 068
Farr J., Font-Ribera A., Pontzen A., 2020b, J. Cosmology Astropart. Phys., 2020, 015
Font-Ribera A., Miralda-Escudé J., 2012, J. Cosmology Astropart. Phys., 2012, 028
Font-Ribera A., et al., 2013, J. Cosmology Astropart. Phys., 2013, 018
Font-Ribera A., et al., 2014, J. Cosmology Astropart. Phys., 2014, 027
Gordon C., et al., 2023, J. Cosmology Astropart. Phys., 2023, 045
Guo Z., Martini P., 2019, ApJ, 879, 72
Guy J., et al., 2023, AJ, 165, 144
Hahn C., et al., 2023, AJ, 165, 253
Handley W. J., Hobson M. P., Lasenby A. N., 2015, MNRAS, 453, 4384
Harris D. W., et al., 2016, AJ, 151, 155
Herrera-Alcantar H. K., et al., 2025, J. Cosmology Astropart. Phys., 2025, 141
Kaiser N., 1987, MNRAS, 227, 1
Kirkby D., et al., 2013, J. Cosmology Astropart. Phys., 2013, 024
Levi M., et al., 2013, arXiv e-prints, p. arXiv:1308.0847
McGreer I., Moustakas J., Schindler J., 2021, simqso: Simulated quasar spectra generator, Astrophysics Source Code Library, record ascl:2106.008
Miller T. N., et al., 2024, AJ, 168, 95
Percival W. J., White M., 2009, Monthly Notices of the Royal Astronomical Society, 393, 297
Pérez-Ràfols I., et al., 2018, MNRAS, 473, 3019
Pérez-Ràfols I., Pieri M. M., Blomqvist M., Morrison S., Som D., Cuceu A., 2023, MNRAS, 524, 1464
Perlmutter S., et al., 1999, ApJ, 517, 565
Pieri M. M., et al., 2014, MNRAS, 441, 1718
Planck Collaboration et al., 2020, A&A, 641, A6
Poppett C., et al., 2024, AJ, 168, 245
Raichoor A., et al., 2023, AJ, 165, 126
Ramírez-Pérez C., Sanchez J., Alonso D., Font-Ribera A., 2022, J. Cosmology Astropart. Phys., 2022, 002
Ramírez-Pérez C., et al., 2024, MNRAS, 528, 6666
Riess A. G., et al., 1998, AJ, 116, 1009
Rogers K. K., Bird S., Peiris H. V., Pontzen A., Font-Ribera A., Leistedt B., 2018, MNRAS, 476, 3716
Schlafly E. F., et al., 2023, AJ, 166, 259
Shen Y., et al., 2016, ApJ, 831, 7
Slosar A., et al., 2011, J. Cosmology Astropart. Phys., 2011, 001
Wang B., et al., 2022, ApJS, 259, 28
Youles S., et al., 2022, MNRAS, 516, 421
Zhou R., et al., 2023, AJ, 165, 58
du Mas des Bourboux H., et al., 2017, A&A, 608, A130
du Mas des Bourboux H., et al., 2020, ApJ, 901, 153

APPENDIX A: FULL MODEL RESULTS

In section 5 we showed the posteriors of our set of cosmological parameters from nested sampler runs, for three different cases: uncontaminated (no redshift continuum errors, baseline model), contaminated (baseline model) and contaminated (baseline + redshift errors model). We showed that the model introduced in this paper (see section 4) successfully removed the effect of continuum redshift errors on the BAO peak parameters (ϕ_p,α_p). It also removes the majority of the shift on ϕ_s, but has little impact on $f\sigma_8$.

In table A1, we show the priors on parameters in the joint fit. In the first section, we show the cosmological parameters of interest, and in the second we show the nuisance parameters of the baseline model. In the final section, we show the parameters introduced in this paper to capture the effect of continuum redshift errors. For most of the parameters, we choose conservative uniform priors, except in the case of β_HCD and L_HCD, where we follow Adame et al. (2025a) and place more informative priors based on HCD studies to improve the performance of the sampler. The priors placed on parameters in our model are physically motivated, i.e. we can't have negative velocity dispersions, or positive values for the amplitude of the effect (see section 4). The value of σ_cont directly reflects the size of the redshift errors in our dataset, meaning the upper bound of 2000 $\mathrm{km s^{-1}}$ would be very extreme.

In figure A1, we show the triangle plot of posteriors for each parameter in our sample runs. Looking first at the parameters introduced in this paper ($A^\mathrm{cross}_\mathrm{cont}$, $A^\mathrm{auto}_\mathrm{cont}$, σ_cont), we see that they are mostly uncorrelated with all of the other parameters. There is a correlation between σ_cont and the amplitude parameters ($A^\mathrm{cross}_\mathrm{cont}$, $A^\mathrm{auto}_\mathrm{cont}$), due to the fact that broadening the (Gaussian) emission line profiles in the Lyα forest also reduces their amplitude. We can also see that while our model successfully reduces the bias on the key cosmological parameters, it often fails to recover the value of other nuisance parameters in the uncontaminated case, and in one case ($b_{\mathrm{SiII}(1206)}$), is significantly more discrepant than the fit to the contaminated dataset without our model. One of the main reasons for both of these issues is the large parameter space with, in some cases, high levels of degeneracy. These differences in nuisance parameters result in the very marginal over-fitting we see in figure 9, but really only serve to indicate that there are some systematic signals that we do not account for with our previous baseline.

APPENDIX B: MEAN CONTINUUM EXPANSION

For a forest of a given quasar q, the delta field is evaluated at λ in the observed frame. The mean continuum $\overline{C}$ is a function of rest-frame, which we can write as a function of λ and z_q as:

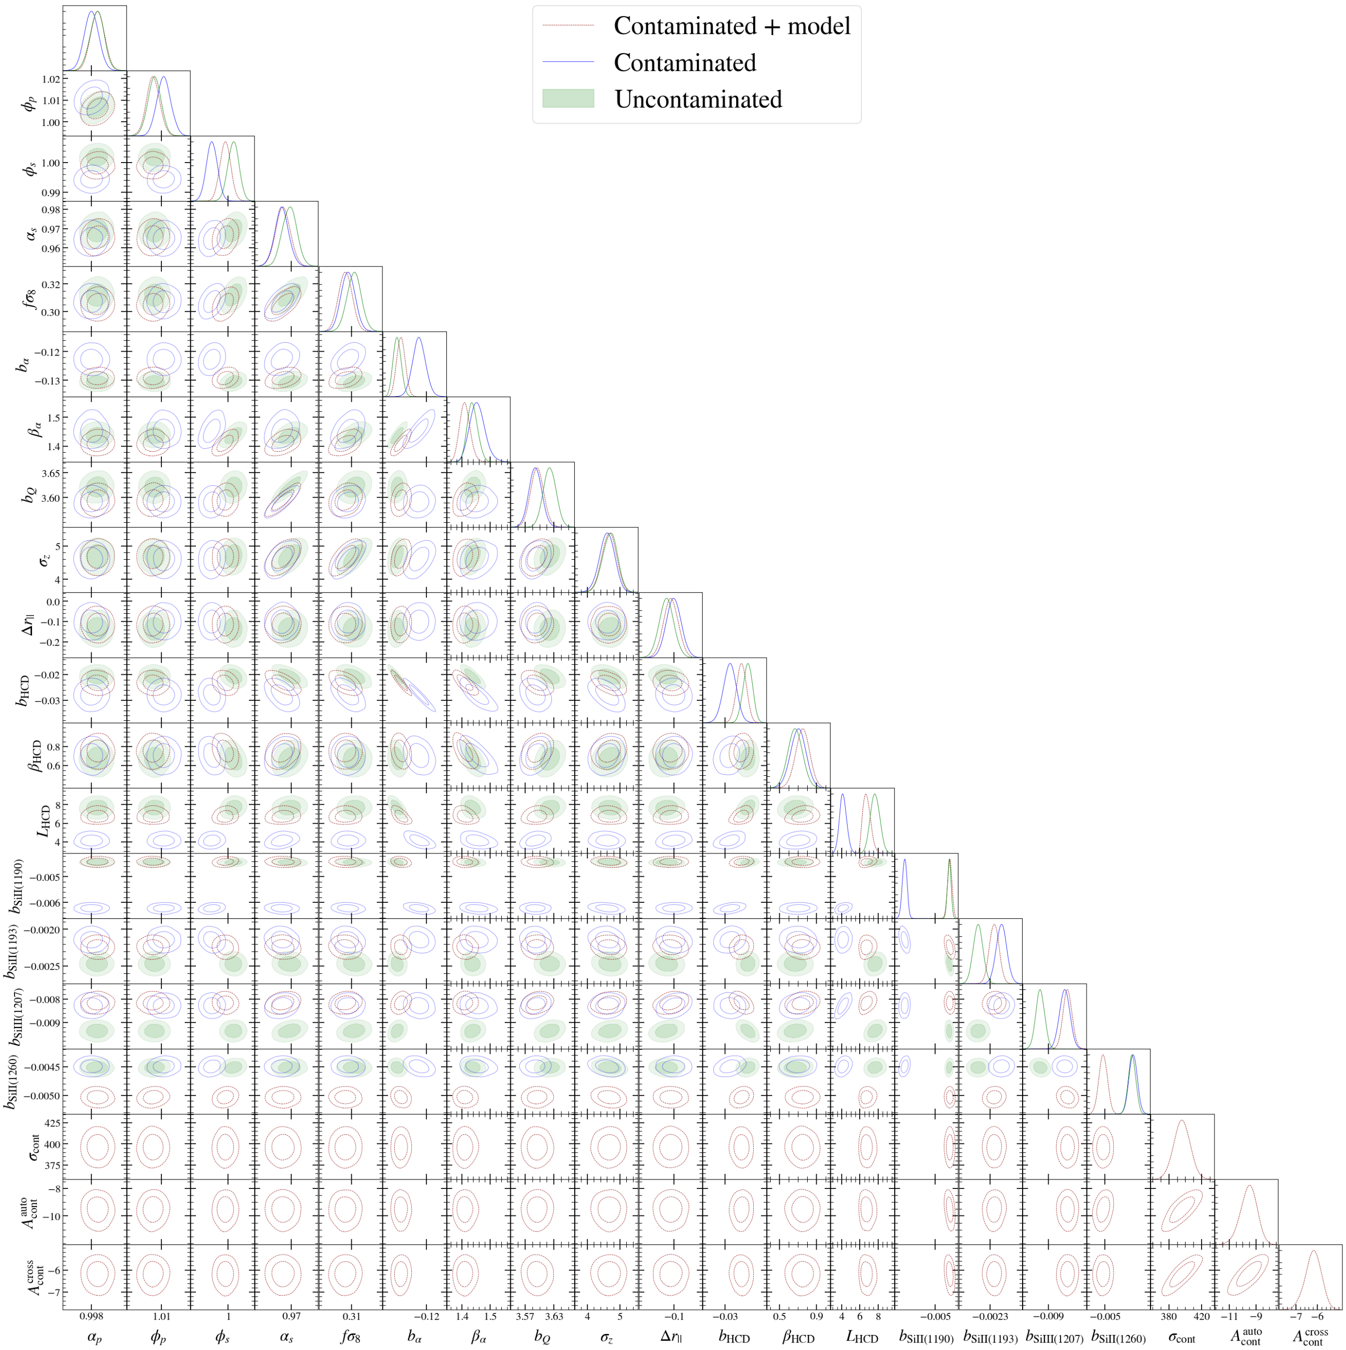

Figure A1. Posteriors of all free parameters in our sampler runs on the joint auto+cross correlation function, measured from a stack of 100 DESI DR1 mocks. The uncontaminated (without continuum redshift errors) run is shown in green, and the contaminated runs with and without the redshift errors model are shown in dark red and blue respectively. We exclude the smoothing parameters $\sigma_{\parallel}$ and $\sigma_{\perp}$, since they marginalise the effect of grid size in our simulated data sets, and are not relevant to real data.

$$1 + \delta(\lambda) = \frac{f(\lambda)}{\hat{\overline{C}}(\frac{\lambda}{1+\hat{z}_{\rm q}})(a_{\rm q} + b_{\rm q}\Lambda)}, \tag{B1}$$

where Λ is a function of $\log\lambda$ given in equation 2, and $\hat{\overline{C}}$ is the measured mean continuum in the presence of redshift errors. We can see from this equation that δ will be contaminated by the smoothed $\hat{\overline{C}}$, but also that the conversion from $\lambda^{\rm rf}$ to λ will be incorrect, given we measure a quasar redshift that has some error, $\hat{z}_{\rm q} = z_{\rm q} + \epsilon$. We can therefore expand the mean continuum about the true redshift as:

$$\hat{\overline{C}}(\hat{\lambda}^{\rm rf}) = \hat{\overline{C}}(\lambda^{\rm rf}) + \epsilon\frac{d\hat{\overline{C}}}{dz_{\rm q}} + \frac{\epsilon^2}{2}\frac{d^2\hat{\overline{C}}}{dz_{\rm q}^2} + \ \ldots \tag{B2}$$

Now, we re-write the differential with redshift as:

Parameter	Prior
α, ϕ*	$\mathcal{U}$[0.01,2]
$f\sigma_8$	$\mathcal{U}$[0,2]
b_α	$\mathcal{U}$[-2, 0]
β_α	$\mathcal{U}$[0, 5]
$b_{\rm q}$	$\mathcal{U}$[0, 10]
$\sigma_{\rm v}$	$\mathcal{U}$[0, 15]
$\Delta r_\parallel$	$\mathcal{U}$[-3, 3]
$b_{\rm HCD}$	$\mathcal{U}$[-0.2, 0]
$\beta_{\rm HCD}$	$\mathcal{N}$[0.5, 0.09]
$L_{\rm HCD}$	$\mathcal{N}$[5, 1]
$b_{\rm SiII(1190)}$	$\mathcal{U}$[-0.02, 0.02]
$b_{\rm SiII(1193)}$	$\mathcal{U}$[-0.02, 0.02]
$b_{\rm SiIII(1207)}$	$\mathcal{U}$[-0.02, 0.02]
$b_{\rm SiII(1260)}$	$\mathcal{U}$[-0.02, 0.02]
$\sigma_\parallel$	$\mathcal{U}$[0, 10]
$\sigma_\perp$	$\mathcal{U}$[0, 10]
$\sigma_{\rm cont}$	$\mathcal{U}$[0, 2e3]
$A^{\rm auto}_{\rm cont}$	$\mathcal{U}$[-50,0]
$A^{\rm cross}_{\rm cont}$	$\mathcal{U}$[-50,0]

Table A1. Priors on free parameters used in the joint fits to our stack of 100 DESI DR1 contaminated mocks. *Recall that there are two α and two ϕ parameters, corresponding to the BAO peak and broadband of the correlation function.

$$\frac{d\hat{\overline{C}}}{dz_{\rm q}} = \frac{d\lambda^{\rm rf}}{dz_{\rm q}}\frac{d\hat{\overline{C}}}{d\lambda^{\rm rf}} = -\frac{\lambda^{\rm rf}}{1+z_{\rm q}}\frac{d\hat{\overline{C}}}{d\lambda^{\rm rf}}. \tag{B3}$$

Substituting equations B2 and B3 into equation 21, we can derive a new expression for our measured transmission field $\hat{\delta}_{\rm q}$:

$$\hat{\delta}_{\rm q} = \delta_{\rm q} - \gamma + \gamma_{\rm z}, \tag{B4}$$

where as usual we ignore 2nd order terms, and

$$\gamma_{\rm z} = \frac{\lambda^{\rm rf}\epsilon}{\overline{C}(1+z_{\rm q})}\frac{d\hat{\overline{C}}}{d\lambda^{\rm rf}}. \tag{B5}$$

Now, as we did in section 4, we propagate our expression for our measured delta through to the cross- and autocorrelation functions:

$$\langle\hat{\delta}\rangle = \langle\delta\rangle - \langle\gamma\rangle + \langle\gamma_{\rm z}\rangle \tag{B6}$$

$$\langle\hat{\delta}\hat{\delta}\rangle = \langle\delta\delta\rangle - 2\langle\delta\gamma\rangle + 2\langle\delta\gamma_{\rm z}\rangle - 2\langle\gamma\gamma_{\rm z}\rangle + \langle\gamma\gamma\rangle + \langle\gamma_{\rm z}\gamma_{\rm z}\rangle. \tag{B7}$$

Modelling the additional terms analytically may prove to be difficult, but as we did for figure 4, we can use mock datasets to measure each term in equation B6. In figure B1, From this measurement, we see that only $\langle\delta\gamma_{\rm z}\rangle$ is non-negligible, and that this is roughly constant across $r_\parallel$ for the usual range of $r_\perp (\lesssim 20\,h^{-1}{\rm Mpc})$. An explanation for the trough feature $\sim 25\,h^{-1}$Mpc is therefore still required, where perhaps higher-order terms cannot be ignored.

B1 Small-scale contamination model

Here, we try to model the $25\,h^{-1}$Mpc contaminating trough feature of our fiducial analysis (see e.g. figure 4) empirically. We do this by adding an extra term to the original γ function (equation 21):

$$\gamma(\lambda^{\rm rf}) = \overline{\gamma_0}(\lambda^{\rm rf}, \sigma_{\rm cont}) + \frac{a_1}{(\lambda^{\rm rf} - \lambda_1^{\rm rf})^2}, \tag{B8}$$

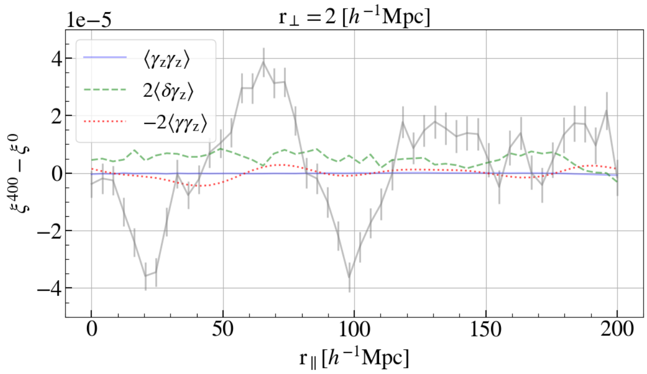

Figure B1. Contaminating terms in the Lyα autocorrelation, due to rest-frame wavelength grids shifts in the presence of quasar redshift errors. $\gamma_{\rm z}$ is defined in equation B5, and is the result of expanding the measured mean continuum $\hat{\overline{C}}$ about small redshift errors ϵ. The contamination in the first transverse bin is shown in grey, as a function of line-of-sight separation ($r_\parallel$).

where the function is offset by $\lambda_1^{\rm rf} = 1207.6$ Å, and $\overline{\gamma_0}$ is the function we model in the previous section. This value was constrained in a direct fit to the difference in the Lyα autocorrelation with and without continuum redshift errors, together with a_1, the redshift error parameter ($\sigma_{\rm cont}$, section 4.2) and a free amplitude. Note this function diverges for $\lambda^{\rm rf} = \lambda_1^{\rm rf} = 1207.6$ Å, but this is not an issue since in our analysis we always have $\lambda^{\rm rf} < \lambda_1^{\rm rf}$.

We show in figure 8 that in a direct fit to the redshift error contamination, our additional correction captures the trough feature very well. However, when performing a full-shape fit to our auto- and cross-correlations we find that a_1 takes much larger than expected values [11]. This consequently biased our measurements of $\phi_{\rm s}$, so it was decided to proceed for now without the additional correction.

AFFILIATIONS

[1]Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Edifici Cn, Campus UAB, 08193, Bellaterra (Barcelona), Spain
[2]Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA
[3]Institut d'Astrophysique de Paris. 98 bis boulevard Arago. 75014 Paris, France
[4]IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France
[5]Department of Physics, Boston University, 590 Commonwealth Avenue, Boston, MA 02215 USA
[6]Dipartimento di Fisica "Aldo Pontremoli", Università degli Studi di Milano, Via Celoria 16, I-20133 Milano, Italy
[7]INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy
[8]Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, UK
[9]Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK
[10]Instituto de Física, Universidad Nacional Autónoma de México, Circuito de la Investigación Científica, Ciudad Universitaria, Cd. de

[11] Our expected value of a_1 is computed in a direct fit to the contamination, where the only free parameters are $\sigma_{\rm cont}$, $A_{\rm cont}$ and a_1.

México C. P. 04510, México
[11]Department of Astronomy & Astrophysics, University of Toronto, Toronto, ON M5S 3H4, Canada
[12]Department of Physics & Astronomy and Pittsburgh Particle Physics, Astrophysics, and Cosmology Center (PITT PACC), University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, PA 15260, USA
[13]Departamento de Física, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia
[14]Observatorio Astronómico, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia
[15]Institut d'Estudis Espacials de Catalunya (IEEC), c/ Esteve Terradas 1, Edifici RDIT, Campus PMT-UPC, 08860 Castelldefels, Spain
[16]Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth, PO1 3FX, UK
[17]Institute of Space Sciences, ICE-CSIC, Campus UAB, Carrer de Can Magrans s/n, 08913 Bellaterra, Barcelona, Spain
[18]Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, USA
[19]Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA
[20]Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA
[21]The Ohio State University, Columbus, 43210 OH, USA
[22]Department of Physics, The University of Texas at Dallas, 800 W. Campbell Rd., Richardson, TX 75080, USA
[23]Department of Physics, Southern Methodist University, 3215 Daniel Avenue, Dallas, TX 75275, USA
[24]Department of Physics and Astronomy, University of California, Irvine, 92697, USA
[25]Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), FR-75005 Paris, France
[26]Departament de Física, Serra Húnter, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain
[27]Department of Astronomy, The Ohio State University, 4055 McPherson Laboratory, 140 W 18th Avenue, Columbus, OH 43210, USA
[28]Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain
[29]Department of Physics and Astronomy, Siena College, 515 Loudon Road, Loudonville, NY 12211, USA
[30]Departamento de Física, DCI-Campus León, Universidad de Guanajuato, Loma del Bosque 103, León, Guanajuato C. P. 37150, México
[31]Instituto Avanzado de Cosmología A. C., San Marcos 11 - Atenas 202. Magdalena Contreras. Ciudad de México C. P. 10720, México
[32]Department of Physics and Astronomy, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada
[33]Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada
[34]Waterloo Centre for Astrophysics, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada
[35]Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain
[36]Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 10, 08930 Barcelona, Spain
[37]Department of Physics and Astronomy, Sejong University, 209 Neungdong-ro, Gwangjin-gu, Seoul 05006, Republic of Korea
[38]CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain
[39]Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, USA
[40]University of Michigan, 500 S. State Street, Ann Arbor, MI 48109, USA
[41]Department of Physics & Astronomy, Ohio University, 139 University Terrace, Athens, OH 45701, USA
[42]NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA
[43]National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Road, Chaoyang District, Beijing, 100101, P. R. China

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