

Measuring and unbiased the BAO shift in the Ly α forest with ABACUSSUMMIT

Boryana Hadzhiyska ^{1,2}★†, Roger de Belsunce ^{1,3}★, A. Cuceu ^{1,4}, J. Guy,¹ M. M. Ivanov,^{5,6}
H. Coquiot¹ and A. Font-Ribera ⁷

¹Lawrence Berkeley National Laboratory, Physics Department, 1 Cyclotron Road, Berkeley, CA 94720, USA

²University of California, Physics Department, 110 Sproul Hall #5800, Berkeley, CA 94720, USA

³Berkeley Center for Cosmological Physics, Department of Physics, UC Berkeley, CA 94720, USA

⁴Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

⁵Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁶The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, Cambridge, MA 02139, USA

⁷Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, E-08193 Bellaterra Barcelona, Spain

Accepted 2025 May 16. Received 2025 May 15; in original form 2025 March 18

ABSTRACT

The Dark Energy Spectroscopic Instrument (DESI) places sub- per cent constraints on measurements of the Baryon Acoustic Oscillation (BAO) scaling parameters from the Ly α forest. However, no systematic error budget stemming from non-linearities in the three-dimensional clustering of the Ly α forest is included in the DESI-Ly α analysis. In this work, we measure the size of the shift of the BAO peak using large Ly α forest mocks produced on the N -body simulation suite ABACUSSUMMIT, which adopt the Fluctuating–Gunn–Peterson Approximation (FGPA). Specifically, we measure the Ly α autocorrelation and the Ly α -quasar cross-correlation functions. To mitigate the noise, we adopt a linear control variates technique, reducing the error bars by a factor of up to $\sim \sqrt{50}$ on large scales. From the autocorrelation, we detect a small positive shift in radial direction of $\Delta\alpha_{\parallel} = 0.35$ per cent at the 3σ level and virtually no shift in the transverse direction, α_{\perp} . From the cross-correlation, we see a similar shift to $\Delta\alpha_{\parallel}$, albeit with larger error bars, and a small negative shift, $\Delta\alpha_{\perp} = \sim -0.25$ per cent, at the 2σ level. We also make a connection with the Ly α forest effective field theory (EFT) framework and find that the one-loop EFT power spectrum yields unbiased measurements of the BAO shift parameters in radial and transverse direction for Ly α auto- and the Ly α -quasar cross-correlation measurements. When using the one-loop EFT framework, we find that we can recover the BAO parameters without a shift, which has important implications for future Ly α forest analyses based on EFT. This work paves the way for novel full-shape analyses of the currently observing DESI and future surveys such as the PFS, WEAVE-QSO, and 4MOST.

Key words: (galaxies:) quasars: absorption lines – (cosmology:) dark energy – (cosmology:) large-scale structure of Universe – cosmology: theory.

1 INTRODUCTION

By serving as a ‘standard ruler’, the Baryon Acoustic Oscillations (BAOs) provide a powerful tool to measure cosmic distances and thus the expansion history of the Universe. These cosmological features, seen as an excess clustering at a comoving scale of approximately $100 h^{-1}$ Mpc, appear in both the two-point correlation function, as a single peak, and the power spectrum, as a series of peaks and troughs, for all tracers of the matter field. BAOs originate from early Universe acoustic density waves and photon–plasma interactions, and by observing them at various redshifts, we can measure cosmological parameters more precisely, especially when combined with

additional probes such as cosmic microwave background (CMB) data. Thus, the BAO scale contributes to refining the standard Λ CDM model and aids in constraining potential extensions of this framework (Weinberg et al. 2013).

The first independent detections of the BAO peak by the *Sloan Digital Sky Survey* (SDSS) and the *2dF Galaxy Redshift Survey* (2dFGRS) collaborations using galaxy clustering (Cole et al. 2005; Eisenstein et al. 2005) validated predictions from standard cosmological theory. Since then, multiple detections of the acoustic peak in large-scale structure using various tracers have reinforced this framework (Beutler et al. 2011; Anderson et al. 2012; Busca et al. 2013; Slosar et al. 2013; Font-Ribera et al. 2014; Delubac et al. 2015; Kitaura et al. 2016). Recently, the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration 2016), a 5-yr survey aiming to measure BAO over a wide range of redshifts by obtaining accurate redshifts for over 40 million galaxies and quasars, has published

* E-mail: boryanah@alumni.princeton.edu (BH); rbelsunce@berkeley.edu (RDB)

† Both authors contributed equally to this work.

seven BAO measurements at different redshifts as part of the Data Release 1 (DR1) analysis, including one from the Ly α forest (DESI Collaboration 2024b).

While the BAO peak position can be accurately predicted in linear theory by means of a Boltzmann solver (e.g. Lewis, Challinor & Lasenby 2000; Blas, Lesgourgues & Tram 2011a), the gravitational growth of structure and redshift space distortions (RSDs) induce non-linear effects on the BAO peak position, which are more challenging to compute from first principle. In particular, the non-linear evolution of the density field broadens the BAO, thus making the measurement less precise, and introduces a systematic sub-percent shift in the acoustic peak position, the magnitude and sign of which depend on the cosmological tracer under study. To tackle these problems, a plethora of reconstruction techniques, based on shifting cosmic structures to their non-evolved positions, have been developed and shown to successfully reduce the uncertainty on the BAO peak position without biasing it (e.g. Eisenstein et al. 2007b; Schmittfull et al. 2015). As for the shift to the BAO peak due to non-linear growth, various works have studied it both from a theoretical (Eisenstein, Seo & White 2007a; Crocce & Scoccimarro 2008; Padmanabhan & White 2009; Sherwin & Zaldarriaga 2012; McQuinn & White 2016; Blas et al. 2016a; Chen, Vlah & White 2022; Chen et al. 2024; de Belsunce et al. 2025) and from an empirical (Achitouv & Blake 2015; Kitaura et al. 2016; Neyrinck et al. 2018; Hernández-Aguayo et al. 2020) perspective.

The Ly α forest, resulting from absorption features in the spectra of quasars due to the presence of neutral hydrogen along the line of sight, offers access to the high-redshift Universe: on scales of tens of kiloparsecs, gas pressure dominates the distribution of neutral hydrogen (McQuinn 2016), while on larger scales the Ly α forest traces the density fluctuations in the matter distribution (Slosar et al. 2011). Thus, through measurements of the BAO scale, it provides an invaluable lever arm of dark energy at cosmic noon. The first BAO measurement from the autocorrelation of the Ly α forest was presented in Busca et al. (2013) and Slosar et al. (2013), whereas the cross-correlation of quasars and the Ly α forest was presented in Font-Ribera et al. (2014). Subsequent SDSS data releases updated these measurements (Delubac et al. 2015; Bautista et al. 2017; du Mas des Bourboux et al. 2017; Blomqvist et al. 2019; de Sainte Agathe et al. 2019; du Mas des Bourboux et al. 2020). The DESI Ly α forest BAO measurement, which was presented in DESI Collaboration (2024b), doubled the number of lines of sight used in the last SDSS measurement and further reduced the uncertainties, bringing to the forefront the question of whether a systematic shift is present in the Ly α BAO peak. Recently, Sinigaglia et al. (2024b) presented simulation-based measurements of the BAO feature using a large suite of 1000 Ly α forest simulations with $1 h^{-1}$ Gpc in volume centred at redshift $z = 2.0$ and found a large negative shift in the BAO peak at the ~ 1 per cent (~ -0.3 per cent)-level in redshift (real) space, respectively. In addition, Belsunce et al. (2025) analytically derive using perturbation theory how non-linearities in the three-dimensional (3D) clustering of the Ly α forest introduce biases in measurements of the BAO scaling parameters. Based on measurements of the bias parameters they *predict* a theoretical error budget of -0.2 per cent (-0.3 per cent) for the radial (transverse) parameters at the same redshift. This work aims to bridge the gap between these two measurements.

In this work, we study the shift in the BAO peak inferred from the clustering of the Ly α forest, adopting the high-resolution Ly α forest mocks built on the full N -body simulation suite ABACUSSUMMIT. Namely, we perform fits to measurements of the two-point auto- and cross-correlation function of the Ly α forest and quasars employing

the VEGA model, used in the analysis of the DESI Ly α forest data (DESI Collaboration 2024b). To decrease the noise in our measurements of the correlation function, we also apply a linear control variates (LCV) technique similar to the Zel'dovich control variates presented in Kokron et al. (2022), DeRose et al. (2023), and Hadzhiyska et al. (2023a). Finally, we make a connection with the effective field theory (EFT) of large-scale structure (Ivanov 2024). We fit the Ly α forest auto and Ly α -quasar cross-spectrum using the one-loop EFT prescription and demonstrate that the EFT can jointly fit for the bias and BAO scaling parameters in a self-consistent manner.

This paper is organized as follows. In Section 2, we briefly describe the ABACUSSUMMIT Ly α forest and quasar mocks, the VEGA model and fitting procedure, and the LCV technique. In Section 3, we summarize our main results from performing fits on the auto- and cross-correlation function for a number of scale cut choices. In Section 4, we draw a connection with EFT and test the ability of the model to obtain unbiased constraints on the BAO shift parameters. Finally, in Section 5, we offer a discussion of our results in a broader context and conclude with an outlook to the upcoming fits to Ly α forest data from DESI.

2 METHODS

2.1 ABACUSSUMMIT

ABACUSSUMMIT is a suite of high-performance cosmological N -body simulations, which was designed to meet and exceed the Cosmological Simulation Requirements of the DESI survey (Maksimova et al. 2021). The simulations were run with ABACUS (Garrison, Eisenstein & Pinto 2019; Garrison et al. 2021), a high-accuracy cosmological N -body simulation code, optimized for GPU architectures and for large-volume simulations, on the Summit supercomputer at the Oak Ridge Leadership Computing Facility.

The majority of the ABACUSSUMMIT simulations are made up of the base resolution boxes, which house 6912^3 particles in a $2 h^{-1}$ Gpc box, each with a mass of $M_{\text{part}} = 2.1 \times 10^9 h^{-1} M_{\odot}$. While the ABACUSSUMMIT suite spans a wide range of cosmologies, here we focus on the fiducial outputs with *Planck* 2018 cosmology: $\Omega_b h^2 = 0.02237$, $\Omega_c h^2 = 0.12$, $h = 0.6736$, $A_s = 2.0830 \times 10^{-9}$, $n_s = 0.9649$, $w_0 = -1$, $w_a = 0$. For full details on the simulation products, see Maksimova et al. (2021).

In this work, we utilize the Ly α forest mocks generated on the ABACUSSUMMIT suite of simulations (Hadzhiyska et al. 2023b). These mocks are generated on the six base boxes `AbacusSummit_base_c000_ph{000-005}`, for which there is full particle outputs at $z = 2.5$, which falls within the redshift of interest of the DESI Ly α forest analysis ($z \sim 2$ to 3). In Section 2.3, we offer a brief summary of the methods for generating these mocks.

2.2 Quasar catalogue

One of the major observables for Ly α forest analyses is the cross-correlation function of the Ly α forest with quasars. Therefore, studying the shift to the BAO peak for this observable is of great importance. In this work, we utilize the quasar catalogues generated as part of the Ly α forest mock release on ABACUSSUMMIT (Hadzhiyska et al. 2023b). These are obtained via ABACUSHOD, an ornamented halo occupation distribution (HOD) model, which incorporates various extensions affecting both the one- and two-halo terms. The full model is described in detail in Yuan et al. (2022).

Table 1. Specifications of the four models used in the creation of the Ly α forest synthetic catalogues (Hadzhiyska et al. 2023b). In particular, we indicate the model parameters, γ , n , k_1 , τ_0 , and σ_ϵ , as well as the clustering bias parameters, $b_{\text{Ly}\alpha}$ and $\beta_{\text{Ly}\alpha}$.

Model #	Method	Fit	$\langle F \rangle$	$\sqrt{\text{Var}[F]}$	τ_0	σ_ϵ	γ	n	k_1	$\chi_{1\text{D}}^2$	$\chi_{3\text{D}}^2$	$b_{\text{Ly}\alpha}$	$\beta_{\text{Ly}\alpha}$
1	Method I	P3D	0.801	0.168	0.387	0.000	1.650	–	–	242.276	61.138	–0.146	0.920
2	Method I	P1D + P3D	0.811	0.187	0.391	0.772	1.450	1.000	0.063	27.536	104.887	–0.129	0.949
3	Method II	P3D	0.825	0.212	0.385	1.696	1.500	1.500	1.000	770.975	25.186	–0.130	2.022
4	Method II	P1D + P3D	0.810	0.187	0.654	2.116	1.550	0.000	–	131.403	100.800	–0.126	2.330

The quasi-stellar object (QSO) mocks used here adopt a simple HOD model without any decorations:

$$\bar{n}_{\text{cent}}^{\text{QSO}}(M) = \frac{\text{ic}}{2} \text{erfc} \left[\frac{\log_{10}(M_{\text{cut}}/M)}{\sqrt{2}\sigma} \right], \quad (1)$$

$$\bar{n}_{\text{sat}}^{\text{QSO}}(M) = \left[\frac{M - \kappa M_{\text{cut}}}{M_1} \right]^\alpha \bar{n}_{\text{cent}}^{\text{QSO}}(M), \quad (2)$$

where M_{cut} characterizes the minimum halo mass to host a central galaxy, M_1 is the typical halo mass that hosts one satellite galaxy, σ is the steepness of the transition from 0 to 1 in the number of central galaxies, α is the power-law index on the number of satellite galaxies, ic is the incompleteness parameter, and κM_{cut} gives the minimum halo mass to host a satellite galaxy. The parameters we choose for our QSO catalogues are in units of $h^{-1}M_\odot$

$$\begin{aligned} \log_{10}(M_{\text{cut}}) &= 13.2, \quad \kappa = 1.11, \quad \sigma = 0.65, \\ \log_{10}(M_1) &= 13.8, \quad \alpha = 0.8, \quad \text{ic} = 1.0, \end{aligned} \quad (3)$$

which have been selected so as to yield a linear bias of about $b_{\text{QSO}} \approx 3.3$, roughly matching the quasar bias in du Mas des Bourboux et al. (2020). The number density of the catalogue is about $1.75 \times 10^{-4} [h^{-1}\text{Mpc}]^{-3}$.

2.3 Ly α forest mocks

Here, we summarize the methods used to generate the ABACUS-SUMMIT Ly α forest mocks, the goal of which is to support the full-shape analysis of the Ly α forest power spectrum, planned to be conducted as part of the DESI Y3 Ly α science programme.

These mocks take the approach of painting gas on the fully evolved ABACUS-SUMMIT N -body simulations via the approximate Fluctuating Gunn–Peterson Approximation (FGPA; Croft et al. 1998), calibrated to Ly α forest skewers extracted from the IllustrisTNG hydro simulation (Qezlou et al. 2022). The resolution of the mocks is 6912^3 cells per box, corresponding to an average of one particle per cell and a mean interparticle distance of $0.29 h^{-1}\text{Mpc}$. The resolution is chosen to be comparable to the Jeans length at that redshift ($\sim 100 \text{ kpc}/h$). Due to the simulation resolution, the power on scales below $\sim 0.3 h^{-1}\text{Mpc}$ is suppressed. For this reason, we boost the power spectrum by adding small-scale fluctuations to the density field.

The optical depth is obtained by converting the dark matter density field into a neutral gas density field. Specifically, we adopt two different, but closely related approaches. The first one of them, **Method I**, follows the standard FGPA prescription:

In the FGPA approach (Gunn & Peterson 1965), the relationship between density and temperature is well approximated by (Hui & Gnedin 1997)

$$T(\mathbf{x}) \propto \rho(\mathbf{x})^{\gamma-1}, \quad (4)$$

where γ is the slope of the temperature–density relation. The temperature of the gas and the number of neutral hydrogen atoms are

connected by

$$n_{\text{HI}}(\mathbf{x}) \propto \rho(\mathbf{x})^2 T(\mathbf{x})^{-0.7}. \quad (5)$$

Here, ρ is the baryonic matter density (Hui, Gnedin & Zhang 1997). Combining these two equations and noting that the optical depth, τ , is proportional to the neutral hydrogen column density, n_{HI} , we arrive at the final form of the FGPA method (Bi & Davidsen 1997; Croft et al. 1998)

$$\tau(\mathbf{x}) = \tau_0 \rho(\mathbf{x})^\alpha, \quad (6)$$

where τ_0 is the overall normalization resulting from the temperature–density and photoionization rate assumptions and $\alpha \equiv 2 - 0.7(\gamma - 1)$. In our analysis, τ_0 and γ are free parameters chosen to fit the one-dimensional (1D) and 3D power spectra of TNG300-1.

The second method, Method II, introduces a small modification to the FGPA approach [see Hadzhiyska et al. (2023b) for details]. While the second method also assumes that the optical depth is related to the density field as $\tau(\mathbf{x}) \propto \rho(\mathbf{x})^\alpha$, here we go directly from the particle positions and their velocities to the final ‘observed’ optical depth in redshift space. In particular, we compute a weight for each particle given by $\rho_{\text{dm}}(\mathbf{x})^{\alpha-1} \times [1 + \delta_{\text{in}}(\mathbf{x})]^\alpha$, where $\rho_{\text{dm}}(\mathbf{x})$ is the dark matter density field in real space and $\delta_{\text{in}}(\mathbf{x})$ is the lognormal noise field in real space. We then displace the line-of-sight coordinate component of each particle according to its peculiar velocity as follows:

$$s = x + v_r(1+z)/H(z), \quad (7)$$

where $H(z)$ is the Hubble parameter at redshift z . Adopting TSC interpolation, we deposit the weighted and displaced particles on the 3D grid to obtain the observed optical depth $\tau(s)$. Thus, this method yields the optical depth directly in redshift-space and as such is less computationally intensive.

The main difference between the two methods is that Method I treats the individual grid cells as Ly- α absorption clouds with a velocity determined by the mean in the cell, whereas in Method II, each particle is approximated as an individual absorber of Ly α photons.

In both cases, we add small-scale noise to compensate for the lack of small-scale power. We finally convert them to transmission flux spectra. Having constructed these two methods, we now fit the free parameters to match the power spectrum measured from skewers extracted from TNG300-1. In particular, we fit to the 3D power spectrum only (Models 1 and 3), or to both the 1D and 3D power spectra (Models 2 and 4). The parameter values of the four models can be found in Table 1. We stress that conclusions about the shift to the BAO peak shift will depend on the consistency between the true underlying bias parameters of the Ly α forest field in the Universe and those of the Ly α forest mocks, be it obtained through the physics model of the TNG300-1 simulation or that of the FGPA approach even if it were the case that the FGPA-based ABACUS-SUMMIT mocks were to faithfully reproduce the bias parameters of the TNG300-1 Ly α forest field. None the less, the assumption made in this work is that while quantitatively we might not be able to match the true

values of the bias parameters, this analysis should give us a rough estimate of how significant we expect the shift to the BAO peak to be.

While the methods of Hadzhiyska et al. (2023b) are simplistic, they offer a fast and straightforward way of connecting the matter density to that of the neutral hydrogen. Examples of more complex techniques include the Ly α Mass Association Scheme [LyMAS; Peirani et al. (2014, 2022)], the Iteratively Matched Statistics (IMS; Sorini et al. (2016) method, Hydro-BAM (Sinigaglia et al. 2022), and cosmic-web-dependent FGPA (Sinigaglia et al. 2024a). These use a variety of approaches tuned using smaller hydro simulations compared with the typical N -body boxes¹ that range from matching the Ly α forest probability distribution function and/or power spectrum to using a supervised machine-learning method. We leave the application of more complex recipes for future work (see e.g. Iršič & McQuinn 2018).

2.4 Correlation function measurements

To compute the correlation function efficiently and accurately, we make use of the saved complex fields for each of the four models and six simulation boxes. In particular, those are low-pass filtered Fourier transforms of the full-resolution Ly α forest and QSO fields up to $k_{\parallel} = 4 \text{ h/Mpc}$ and $k_{\perp} = 2 \text{ h/Mpc}$. As such, they contain information down to the Mpc scale. We then compute the 3D auto- and cross-power spectrum of the Ly α forest with the quasars, zeroing out the $\mathbf{k} = 0$ mode.² Looping over each mode in the line-of-sight direction, we pad the 3D power spectrum with zeros in the remaining two directions and inverse Fourier transform the field, saving the thus-obtained half-transformed 3D correlation function (also known as the cross-spectrum) only up to $r_{\perp} = 200 \text{ Mpc/h}$. We repeat the same operation for the line-of-sight direction and finally, bin the 3D correlation function into r_{\perp} , r_{\parallel} bins from 0 to 4 Mpc/h with bin size of $\Delta r = 4 \text{ Mpc/h}$. The end product is, thus, a (50, 50) matrix for each of the 4 models, 6 simulations, and 2 lines-of-sight directions, with the measured $\xi(r_{\perp}, r_{\parallel})$ correlation functions, which matches the format of the Ly α measurements of both eBOSS and DESI.

Apart from the complex fields, full 3D maps (6912³) of the redshift-space flux are also available. However, we opt not to use them, as that would make our measurements either less efficient or less accurate. In order to brute-force compute the correlation function using the configuration space fields, we would need to downsample the field (due to the large number of cells). However, we found that this induces sharp features into the measured correlation function. Since the fields are on a regular grid, other clever techniques could be employed such as identifying all possible pairs through geometric arguments. However, if the focus is on scales beyond 1 Mpc/h, such approaches would still be more computationally expensive than the Fourier transform approach. For measurements that require one to probe the sub-Mpc regime, those might be more suitable, as these small-scale Fourier modes are not saved for the ABACUSUMMIT mocks.

¹We note that Sinigaglia et al. (2024a) uses two hydro simulations, each with a side of 500 Mpc/h.

²We found this step to be very important, as the $\delta_F = F/\bar{F} - 1$ step performed when obtaining the complex δ_F field has some small vestigial noise due to the piecewise manner in which we obtain \bar{F} , such that $\langle \delta_F \rangle$ is small but non-zero.

2.5 VEGA model

We use the VEGA³ package to model the Ly α forest correlation functions and measure the BAO position. This is the same package used by DESI, and our model is a simplified version of the one used in DESI Collaboration (2024b), that only includes Ly α forest and quasar clustering, and does not include the effect of contaminants. The model is based on a template approach that starts with a linear isotropic matter power spectrum, $P_{\text{fid}}(k)$ computed using a fiducial cosmology based on the Planck CMB results (Planck Collaboration 2020). This power spectrum is decomposed into a peak (or wiggles) component, $P_{\text{fid}}^p(k)$, and a smooth (or no-wiggles) component, $P_{\text{fid}}^s(k)$, following Kirkby et al. (2013). These two components undergo the rest of the modelling process independently and are combined at the end.

To model the effect of broadening of the BAO peak due to non-linear evolution, we follow Kirkby et al. (2013) and multiply the peak component, $P_{\text{fid}}^p(k)$, by a Gaussian smoothing term $\exp[-k_{\parallel}^2 \Sigma_{\parallel}^2/2 - k_{\perp}^2 \Sigma_{\perp}^2/2]$. We treat Σ_{\parallel} and Σ_{\perp} as nuisance parameters and allow them to vary in the fits.

The anisotropic model power spectrum is then given by

$$P_{A \times B}(k, \mu_k, z) = b_A b_B (1 + \beta_A \mu_k^2) (1 + \beta_B \mu_k^2) P_{QL}(k, z) F_{NL}(k, \mu_k, z), \quad (8)$$

where (A, B) are either (Ly α , Ly α) for the autocorrelation, or (Ly α , QSO) for the cross-correlation. The linear bias and RSD parameters, b and β , are treated as nuisance parameters and marginalized over. For the peak component, P_{QL} is given by the product of $P_{\text{fid}}^p(k)$ with the Gaussian smoothing term that models the BAO broadening, while for the smooth component $P_{QL} = P_{\text{fid}}^s$. The F_{NL} term only applies to the Ly α autocorrelation functions, and represents the model for small-scale non-linearities. We follow (DESI Collaboration 2024b) and use the empirical model developed by Arinyo-i Prats et al. (2015).

The resulting anisotropic power spectrum model is decomposed into even multipoles up to $\ell = 6$, which are then Fourier transformed into the equivalent correlation function multipoles. Finally, these multipoles are combined to compute the correlation function on the two-dimensional (2D) grid in r_{\parallel} and r_{\perp} , along with the appropriate window function that reflects the rectangular bins of the 2D grid.

The BAO position is fitted when combining the peak and smooth components into the full correlation function model:

$$\xi(r_{\parallel}, r_{\perp}) = \xi_s(r_{\parallel}, r_{\perp}) + \xi_p(\alpha_{\parallel} r_{\parallel}, \alpha_{\perp} r_{\perp}), \quad (9)$$

where ξ_s and ξ_p are the model correlation functions for the smooth and peak components, respectively. The BAO parameters, $(\alpha_{\parallel}, \alpha_{\perp})$, rescale the coordinates of the peak component and are the target parameters we wish to measure.

Besides the physical model described above, we also add broadband polynomials to ensure the robustness of our BAO measurements. We follow DESI Collaboration (2024b), and use Legendre polynomials $L_j(\mu)$ of order $j = 0, 2, 4$ and 6, divided by powers of r^i with $i = 0, 1, 2$ (corresponding to a parabola in $r^2 \xi(r)$). These polynomials are added to the final correlation function model, and have an extra 12 free parameters for each correlation function.

2.6 Fitting procedure

Our fitting procedure follows that of DESI Collaboration (2024b), as implemented in the VEGA package. We use a Gaussian likelihood and the `iminuit` minimizer to obtain the best-fitting parameter values. Our parameter vector includes the BAO parameters, $(\alpha_{\parallel}, \alpha_{\perp})$, and 16

³Publicly available at <https://github.com/andreiceuceu/vega>.

nuisance parameters: 12 associated with the broad-band polynomials, and 4 associated with the Ly α model, $(b_{\text{Ly}\alpha}, \beta_{\text{Ly}\alpha}, \Sigma_{\parallel}, \Sigma_{\perp})$. We use wide uninformative priors for all parameters.

As we do not have estimates of the covariance matrices of our correlation function measurements, we use covariances from real data for the fitting process, and then compute uncertainties using the population of BAO fits we obtain. The covariance matrices we use were measured by du Mas des Bourboux et al. (2020) from the SDSS DR16 data set. The reason we can use the SDSS covariance instead of one constructed using the ABACUSSUMMIT simulations is that it captures well the covariances between the different radial bins in the transverse and perpendicular directions. We have also tested that using a different covariance matrix (e.g. just the diagonal entries) makes no difference to the final results. We note that for convenience we normalize the covariance matrix such that $\chi^2/\text{d.o.f.} \sim 1$, but note that this normalization has no bearings on the constraints we get for the BAO shift parameters, as those are obtained by performing jackknives of the measured α parameters.

2.7 Linear control variates

To mitigate the noise of the measured correlation function, we employ an analytical implementation of the linear control variate technique (LCV) following Hadzhiyska et al. (2023a), which is a simplified version of the Zel'dovich CV (Kokron et al. 2022; DeRose et al. 2023; Hadzhiyska et al. 2023a). The LCV technique allows us to subtract the noise from an expensive observable of interest as long as we can find an alternative (surrogate) quantity that is highly correlated with the quantity of interest. In our case, the surrogate is the linear density field at the initial conditions of the simulation, whereas the quantity of interest is the auto- and cross-correlation of the Ly α forest with the quasars. In this work, we pioneer an application of analytical CV to the cross-correlation between Ly α forest and quasars.

In the case of the autocorrelation of the Ly α forest, we start with the LCV equation in Fourier space (as that is easier to manipulate):

$$\hat{P}_{\ell}^{*,tt}(k) = \hat{P}_{\ell}^{tt}(k) - \beta_{\ell}(k) (\hat{P}_{\ell}^{ll}(k) - P_{\ell}^{ll}(k)), \quad (10)$$

where $\hat{P}_{\ell}^{tt}(k)$ is the measured Ly α forest power spectrum, whereas $P_{\ell}^{ll}(k)$ and $\hat{P}_{\ell}^{ll}(k)$ are the analytical and measured power spectra predicted by linear theory.

As done in Kokron et al. (2022), DeRose et al. (2023), and Hadzhiyska et al. (2023a), we adopt the disconnected approximation for $\beta_{\ell}(k)$:

$$\beta_{\ell}(k) = \frac{[\hat{P}_{\ell}^{tl}(k)]^2}{[\hat{P}_{\ell}^{ll}(k)]}, \quad (11)$$

where \hat{P}^{tl} is the measured cross-power spectrum between the true and the modelled reconstructed fields. We apply damping on small scales and smoothing with a Savitsky-Golay (S-G) filter. For estimates of the bias parameters, we use the VEGA fits to each raw measurement; $\beta_{\ell}(k)$ is assumed to be uncorrelated with the measured Ly α forest and linear theory power spectra due to the smoothing with the S-G filter, which removes any dependence on the noise realization of the simulation as demonstrated in DeRose et al. (2023). This is also evident in Fig. A1, which shows that indeed $\beta_{\ell}(k)$ is smooth.

To make the linear theory prediction of the Ly α forest power spectrum, we use the CLASS-generated (Lesgourgues 2011) linear power spectrum for the ABACUSSUMMIT initial conditions, $P_{\text{lin}}(k)$, and compute the theory-predicted multipoles, truncating at the hexadecapole, $\ell = 4$,

$$P_{\ell}^{ll}(k) = b_{\text{Ly}\alpha}^2 C_{\ell}(\beta_{\text{Ly}\alpha}) P_{\text{lin}}(k) \quad (12)$$

where

$$C_{\ell}(\beta_{\text{Ly}\alpha}) \equiv \frac{2\ell + 1}{2} \int_{-1}^1 d\mu (1 + \beta_{\text{Ly}\alpha} \mu^2)^2 \mathcal{L}_{\ell}(\mu) \\ = \begin{cases} 1 + \frac{2}{3}\beta_{\text{Ly}\alpha} + \frac{1}{5}\beta_{\text{Ly}\alpha}^2 & \ell = 0 \\ \frac{4}{3}\beta_{\text{Ly}\alpha} + \frac{4}{7}\beta_{\text{Ly}\alpha}^2 & \ell = 2 \\ \frac{8}{35}\beta_{\text{Ly}\alpha}^2 & \ell = 4 \end{cases}. \quad (13)$$

Finally, we need to convert to configuration space in order to obtain the correlation function, $\xi(r_{\parallel}, r_{\perp})$. As in the case of Hadzhiyska et al. (2023a), we work with 3D quantities, and express the LCV-reduced power spectrum measurement as follows:

$$\hat{P}^{*,tt}(\mathbf{k}) = \hat{P}^{tt}(\mathbf{k}) - \beta(\mathbf{k}) (\hat{P}^{ll}(\mathbf{k}) - P^{ll}(\mathbf{k})), \quad (14)$$

where $\beta(\mathbf{k})$ and $P^{ll}(\mathbf{k})$ are expanded into the 3D \mathbf{k} -grid from their multipole counterparts. We then apply an inverse Fourier transform to the LCV-reduced 3D field

$$\hat{\xi}^{*,tt}(\mathbf{r}) = \text{IFT}[\hat{P}^{*,tt}(\mathbf{k})] \quad (15)$$

and bin into $\xi^{tt} = (r_{\parallel}, r_{\perp})$ with $\Delta r = 4 \text{ Mpc}/h$ from 0 to 200 Mpc/h .

The procedure for obtaining the LCV-mitigated measurements of the cross-correlation function is analogous, with two main differences. The first one is that the linear theory prediction of the cross-power spectrum with QSOs is given by the following equation:

$$P_{\times,\ell}(k) = b_{\text{Ly}\alpha} b_{\text{QSO}} C_{\times,\ell}(\beta_{\text{Ly}\alpha}, \beta_{\text{QSO}}) P_{\text{lin}}(k) \quad (16)$$

with

$$C_{\times,\ell}(\beta, \beta_q) \equiv \begin{cases} 1 + \frac{1}{3}\beta + \frac{1}{3}\beta_q + \frac{1}{5}\beta\beta_q & \ell = 0 \\ \frac{2}{3}\beta + \frac{2}{3}\beta_q + \frac{4}{7}\beta\beta_q & \ell = 2 \\ \frac{8}{35}\beta\beta_q & \ell = 4 \end{cases} \quad (17)$$

where $\beta \equiv \beta_{\text{Ly}\alpha}$ and $\beta_q \equiv \beta_{\text{QSO}}$.

The second difference is that $\beta_{\times,\ell}(k)$, which encodes the amount of correlation between the observable of interest, $P^{ab}(k)$ (with ‘a’ being Ly α and ‘b’ QSO), and its surrogate, $P^{cd}(k)$ (with ‘c’ being the linear theory prediction for Ly α and ‘d’ the same thing for QSO), takes the following form in the disconnected approximation using the Knox formula (Knox 1995):

$$\beta_{\times,\ell}(k) = \frac{\hat{P}_{\ell}^{ac}(k)\hat{P}_{\ell}^{bd}(k) + \hat{P}_{\ell}^{ad}(k)\hat{P}_{\ell}^{bc}(k)}{\hat{P}_{\ell}^{cc}(k)\hat{P}_{\ell}^{dd}(k) + \hat{P}_{\ell}^{cd}(k)^2}. \quad (18)$$

3 RESULTS

In this section, we describe our main findings from performing fits to the Ly α forest auto- and cross-correlation function with quasars. In particular, we report the values of the α_{\perp} and α_{\parallel} parameters as measured in the Ly α forest mocks on AbacusSummit. We note that the values of α_{\perp} and α_{\parallel} reported here will match those in the real Universe as long as the bias parameters b_1 and b_2 , which as discussed below (also, see Section 4) virtually determine the BAO shift, match the true b_1 and b_2 of the Ly α forest.

3.1 Ly α forest autocorrelation scatter

First, we study the fits to the raw measurements of the Ly α forest autocorrelation function on the ABACUSSUMMIT mock catalogues.

In the left panel of Fig. 1, we show a scatter plot of the BAO parameters, α_{\parallel} and α_{\perp} , for Model 1, measured from the 6 ABACUSSUMMIT simulations with lines of sight along the z and the y axes, for a total of 12 measurements. Thus, the effective volume




Figure 1. Scatter plot of the BAO parameters, α_{\parallel} and α_{\perp} , for Model 1, measured from the six ABACUSSUMMIT simulations with lines of sight along the z - and the y -axes. Fits are performed down to minimum isotropic separation of $r_{\min} = 30 \text{ Mpc}/h$. *Left panel:* We see that the scatter between them is large, but the average value of the two parameters is consistent with no shift. *Right panel:* Same set of fits but using the LCV technique to mitigate the noise on the measurements. The scatter in both parameters is significantly reduced, *i.e.* the error bars on the measured α 's are reduced by a factor of ~ 3 .




Figure 2. Similar to Fig. 1 except here, we show results from the cross-correlation function of the Ly α forest with quasars. Fit to six simulations of model 1 with lines of sight y and z . *Left panel:* The BAO parameter values are consistent with there being no shift, with α_{\perp} preferring slightly lower values, but still within the error bars, and α_{\parallel} preferring no shift. *Right panel:* Fits of the cross-correlation function of the Ly α forest with the quasars are done using the LCV technique to mitigate the noise on the measurements. The scatter is significantly reduced, following baseline expectation.

of these mocks is in fact larger than the combined volume of the six simulations, $\sim 50 \text{ (Gpc}/h)^3$. We see that the scatter between them is large, but the average value of the two parameters are consistent with 1, *i.e.* no shift. The mean value across all simulations with its uncertainty is shown in black. We note that the size of the error bars, around 0.25 per cent, which is computed via the jackknife method, is larger than this shift when no noise reduction is applied. In this figure, we perform fits of the autocorrelation function of the Ly α forest using the VEGA model down to minimum isotropic separation of $r_{\min} = 30 \text{ Mpc}/h$. For details on the VEGA model, see Section 2.5.

3.2 Ly α forest \times QSO cross-correlation scatter

Next, we move to the cross-correlation between Ly α and quasars. We first study a scatter plot of the BAO parameters for all 12

measurements in the left panel of Fig. 2. We see that the values of both parameters are consistent with there being no shift, with α_{\perp} preferring slightly lower values, and α_{\parallel} preferring no shift. We note that the shifts of the two BAO parameters in the ‘cross’ case do not need to match those in the ‘auto’ case, as in the case of ‘cross’ they depend on the behaviour of both tracers: the Ly α forest and the quasars, whereas in the case of ‘auto’ they depend solely on the Ly α forest (through predominantly the so-called b_2 bias parameter⁴) (Chen et al. 2024; Ivanov 2024; Belsunce et al. 2025). None the

⁴The non-linear shift of the BAO parameters in the Ly α –quasar cross-correlation is driven by the ratio of the linear to the non-linear bias parameter

$$\Delta\alpha \sim \left(\frac{b_2^q}{b_1^q} + \frac{b_2}{b_1} \right) \sigma_d^2. \quad (19)$$




Figure 3. Constraints on the BAO parameters α_{\parallel} and α_{\perp} , for all four models, as a function of the minimum separation, r_{\min} , used in the VEGA fits of the autocorrelation function of the Ly α forest, with noise mitigated by the LCV technique. 12 measurements are made for each of the 6 ABACUSSUMMIT simulations with lines-of-sight along the z - and the y -axes. There is a $\sim 3\sigma$ shift of about 0.35 per cent to the α_{\parallel} , while α_{\perp} differs from 1 within 1–1.8 σ . The most likely reason for the difference in α_{\perp} between the models is a difference in the values of the bias parameters, b_1 and b_2 , resulting from the different methods adopted in generating the mocks. The mild dependence that we see in α_{\perp} between $r_{\min} = 10$ Mpc/ h and the larger r_{\min} values can be attributed to issues with the small-scale modelling, which is expected to fail at $r \lesssim 10$ Mpc/ h . In Appendix A, we test in more detail the LCV outputs and demonstrate that the error bars shrink by almost a factor of 10 on large scales.

less, the fact that the values of the α 's are broadly similar reveals that neither of the two tracers is significantly biased at this level of precision.

3.3 Linear control variates fits

In this section, we apply the LCV method to the correlation function measurements, as described in Section 2.7, to obtain tighter constraints on α_{\parallel} and α_{\perp} . Comparing the left to the right panel of Fig. 1, we see that the scatter is much smaller, demonstrating that the LCV technique has successfully reduced the noise in the inferred α parameters.

We next study, in Fig. 3, how the constraints on the BAO parameters α_{\parallel} and α_{\perp} , for all four models, change as a function of minimum separation, r_{\min} , used in the VEGA fits of the autocorrelation function of the Ly α forest. Interestingly, the behaviour of all four models is largely consistent with each other despite the different values of $b_{\text{Ly}\alpha}$ and $\beta_{\text{Ly}\alpha}$ that each of them exhibits (see Table 1). We also note that Models 3 and 4 are constructed using a modified

The mode-coupling shift term is given in parentheses and divided by the size of the linear BAO signal with the non-linear bias parameters b_2 (auto) and b_2^{c} (cross).

version of FGPA, but none the less, they yield very similar values to those extracted for Models 1 and 2.

Looking at the α values, we see a $\sim 3\sigma$ shift to the α_{\parallel} , while α_{\perp} is consistent with one within $\sim 1\sigma$. The shift to α_{\parallel} is in a consistent direction with what we expect from the left panel of Fig. 1. On average, we see that it is consistent for all four models and it is about 0.35 per cent with an error of around 0.11 per cent (see Table 2).

Additionally, we note that the values of the BAO parameters show very weak dependence on r_{\min} . This is expected, as the BAO peak occurs on much larger scales, $r \sim 100$ Mpc/ h and therefore should not be sensitive to the choice of r_{\min} . In App. B, we test that if we switch off the broad-band parameters, the dependence on r_{\min} is a bit more pronounced. This shows that the broad-band model manages to absorb some of the issues on small scales: one of the most prominent one being that we keep the ‘Arinyo model’ parameters fixed. In the future, we plan to study this in more detail through comparisons with more physically motivated frameworks such as EFT (Ivanov 2024).

We see that α_{\perp} appears to be mildly more sensitive to the model choice (in particular, between Models 1 and 2 and Models 3 and 4) and this sensitivity is underscored when adopting the LCV method, but still within 1σ of each other. The mild dependence that we see in α_{\perp} between $r_{\min} = 10$ Mpc/ h and the larger r_{\min} values can be attributed to issues with the small-scale modelling, which is expected to fail at $r \lesssim 10$ Mpc/ h . For this reason, when quoting the final errors

Table 2. Shifts from the linear theory values of the BAO parameters, α_{\perp} and α_{\parallel} , defined as $\Delta\alpha \equiv \alpha - 1$. A positive value of $\Delta\alpha$, therefore, corresponds to a preference for $\alpha > 1$ and vice versa. Values of the best-fitting α 's are obtained by using the VEGA model down to a minimum isotropic separation of r_{\min} , with noise mitigation obtained via the LCV technique. We adopt four different models, each with distinct values of the $b_{Ly\alpha}$ and $\beta_{Ly\alpha}$ parameters (see Table 1). We see 3σ significant shifts to the value of α_{\parallel} for all values of r_{\min} and all models of about 0.35 ± 0.11 per cent. The shift of α_{\perp} , on the other hand, is consistent with zero. We note that we disregard the $r_{\min} = 10$ Mpc/h case, as those values are affected by shortcomings of the model on small scales.

$\Delta\alpha_{\parallel}, \Delta\alpha_{\perp}$	Model 1	Model 2	Model 3	Model 4
$r_{\min} = 10$ Mpc/h	$0.43 \pm 0.10 \%$, $0.20 \pm 0.08 \%$	$0.42 \pm 0.10 \%$, $0.18 \pm 0.08 \%$	$0.38 \pm 0.12 \%$, $0.38 \pm 0.11 \%$	$0.36 \pm 0.12 \%$, $0.41 \pm 0.12 \%$
$r_{\min} = 20$ Mpc/h	$0.37 \pm 0.10 \%$, $-0.16 \pm 0.13 \%$	$0.37 \pm 0.10 \%$, $-0.17 \pm 0.08 \%$	$0.34 \pm 0.12 \%$, $0.05 \pm 0.11 \%$	$0.33 \pm 0.12 \%$, $0.09 \pm 0.12 \%$
$r_{\min} = 30$ Mpc/h	$0.37 \pm 0.10 \%$, $-0.13 \pm 0.07 \%$	$0.36 \pm 0.10 \%$, $-0.14 \pm 0.08 \%$	$0.34 \pm 0.12 \%$, $0.06 \pm 0.11 \%$	$0.33 \pm 0.12 \%$, $0.08 \pm 0.12 \%$
$r_{\min} = 40$ Mpc/h	$0.37 \pm 0.12 \%$, $-0.13 \pm 0.13 \%$	$0.37 \pm 0.10 \%$, $-0.15 \pm 0.08 \%$	$0.37 \pm 0.12 \%$, $0.02 \pm 0.11 \%$	$0.36 \pm 0.12 \%$, $0.02 \pm 0.12 \%$

on the shifts, we quote them at $r = 30$ Mpc/h, which does not appear to be affected by this. In Appendix A, we test in more detail the LCV outputs. We note that the effective volume increase is nearly a factor of 50 on large scales (note that the error bars go as $\sqrt{V_{\text{eff}}}$), with this number going down on small scales.

Next, we move on to applications of the LCV approach to the cross-correlation function with the quasars. Reassuringly, comparing the left to the right panel in Fig. 2, we see that the scatter is much smaller, demonstrating that the LCV technique has successfully reduced the noise in the inferred α parameters. Looking at the plot of the BAO parameters as a function of minimum separation, r_{\min} , the size of the error bars has shrunk significantly in Fig. 4 and Table 3. Similarly to the autocorrelation case, we see that there is a small shift towards values larger than one in α_{\parallel} of about 0.3–0.4 per cent (though with slightly larger error bars), whereas α_{\perp} is slightly shifted (at the 0.2–0.3 per cent level) towards values smaller than one. Neither of these is at a high significance likely due to the larger noise in the cross-correlation (LCV works less well and the quasars yield a noisier correlation function measurement due to their sparsity and larger non-linearities).

4 CONNECTION TO EFFECTIVE FIELD THEORY

The accessible amount of cosmological information from the Ly α forest can be increased by (i) cross-correlating the Ly α forest with quasars (Font-Ribera et al. 2013; du Mas des Bourboux et al. 2020) and (ii) including information beyond the BAO feature either through compressed statistics such as redshift-space distortions and the Alcock-Paczyński effect (AP) (Cuceu et al. 2021) or by using the broad-band shape (i.e. the full-shape) of the correlations (Gerardi et al. 2022). However, only recent advances in theoretical modelling using the EFT of large-scale structure (EFT; Baumann et al. 2012; Carrasco, Hertzberg & Senatore 2012) extended to the Ly α forest (Ivanov 2024; Chudaykin & Ivanov 2025) allow for a consistent model spanning large-to-intermediate scales to jointly extract the BAO information as well as the shape of the power spectrum (or correlation function).

In the first part of the present work, we quantify the systematic error budget pertaining to the BAO peak shift stemming from using linear-theory templates to fit the BAO peak on ABACUSSUMMIT simulations. In the second part (the present section), we introduce the EFT model for the one-loop Ly α forest autocorrelation and for the cross-correlation with quasars (or halos) in Sections 4.1–4.4. For the first time, we use EFT to fit the large and small scales of the Ly α

forest directly whilst including information from the BAO feature.⁵ Therefore, we jointly fit the BAO scaling and EFT parameters for the auto and subsequently the auto and cross-correlation in Sections 4.5–4.6. We present two tests to connect both approaches in this work in Section 4.7: First, we use the quasi-linear theory fitting procedure introduced in Section 2 to fit the best-fitting EFT spectra obtained from fits to the ABACUSSUMMIT simulations. Secondly, we use EFT to *theoretically predict* the physical shift of the BAO feature in the autocorrelation based on measurements of the bias parameters from ABACUSSUMMIT simulations (Chen et al. 2024; Belsunce et al. 2025).

4.1 One-loop Ly α forest EFT model

An alternative approach to using phenomenological fitting functions from hydrodynamical simulations to model the Ly α forest, as described in Section 2.5, is to model the Ly α forest using the EFT of large-scale structure (EFT) (Baumann et al. 2012; Carrasco et al. 2012; Ivanov 2023). The high redshift ($2 \leq z \leq 4$) of the Ly α forest yields access to more (quasi)-linear modes than, e.g. galaxy surveys, rendering it particularly fruitful for perturbative frameworks such as EFT (see e.g. McDonald & Roy 2009; Baumann et al. 2012; Carrasco et al. 2012). EFT for the Ly α forest fluctuations describes large-scale dynamics using only the equivalence principle and rotations around the line of sight \hat{z} , denoted by the $SO(2)$ group (Givans & Hirata (2020); Chen, Vlah & White (2021); Ivanov (2024); de Belsunce & Sullivan (2025)). In the following, we briefly summarize the one-loop Ly α EFT model. For a fuller presentation, we refer the reader to Desjacques, Jeong & Schmidt (2018a); Ivanov (2024); Ivanov, Toomey & Göksel Karaçaylı (2024a); Belsunce et al. (2025).

Schematically, the one-loop EFT model of the Ly α forest for the autocorrelation consists of four components

$$P^{\text{th}}(k, \mu) = P^{\text{tree}}(k, \mu) + P^{1\text{-loop}}(k, \mu) + P^{\text{ct}}(k, \mu) + P^{\text{st}}(k, \mu), \quad (20)$$

where k is the Fourier wavenumber, μ the angle of $k = \{k_{\parallel}, k_{\perp}\}$ to the line of sight, $\mu \equiv k_{\parallel}/k$, $P^{\text{tree}}(k)$ is the infrared resummed linear theory power spectrum⁶⁷ (details of IR resummation will be

⁵Previous EFT analyses used hydrodynamical simulations with access to fewer quasi-linear modes than ABACUSSUMMIT (Givans et al. 2022; Ivanov 2024; Belsunce et al. 2025; Chudaykin & Ivanov 2025).

⁶The linear power spectrum as well as the decomposition into an oscillatory and a smooth part, which are required for IR resummation, are obtained using the Boltzmann solver CLASS-PT (Blas, Lesgourgues & Tram 2011b; Chudaykin et al. 2020). We perform an anisotropic one-loop IR resummation of the power spectrum.

⁷We remind the reader that our convention of b_{η} is related to the literature by a negative sign.




Figure 4. Similar to Fig. 3. The difference is that here, we show the BAO parameters as a function of minimum separation, r_{\min} , measured from fits of the cross-correlation function of the Ly α forest with noise mitigated by the LCV technique. The size of the error bars has shrunk significantly. There is a small shift towards values larger than one in α_{\parallel} of about 0.3–0.4 per cent (though with slightly larger error bars), whereas α_{\perp} is slightly shifted (at the 0.2–0.3 per cent level) towards values smaller than one.

Table 3. Same as Table 2, except for the reported values are for the case of fits to the cross-correlation function with noise mitigation obtained via the LCV technique. We see that the α_{\parallel} shifts are a bit larger compared with the autocorrelation case, but the error bars are also about twice as big. The shift of α_{\perp} is also a bit more pronounced, about -0.3 per cent at around 3σ . We note that we disregard the $r_{\min} = 10$ Mpc/h case, as those values are affected by shortcomings of the model on small scales.

$\Delta\alpha_{\parallel}, \Delta\alpha_{\perp}$	Model 1	Model 2	Model 3	Model 4
$r_{\min} = 10$ Mpc/h	$0.42 \pm 0.17 \%$, $-0.09 \pm 0.08 \%$	$0.41 \pm 0.18 \%$, $-0.09 \pm 0.09 \%$	$0.39 \pm 0.20 \%$, $-0.12 \pm 0.10 \%$	$0.38 \pm 0.20 \%$, $-0.12 \pm 0.11 \%$
$r_{\min} = 20$ Mpc/h	$0.48 \pm 0.17 \%$, $-0.27 \pm 0.08 \%$	$0.48 \pm 0.18 \%$, $-0.28 \pm 0.09 \%$	$0.44 \pm 0.20 \%$, $-0.24 \pm 0.10 \%$	$0.43 \pm 0.23 \%$, $-0.24 \pm 0.12 \%$
$r_{\min} = 30$ Mpc/h	$0.55 \pm 0.18 \%$, $-0.34 \pm 0.08 \%$	$0.55 \pm 0.18 \%$, $-0.35 \pm 0.09 \%$	$0.50 \pm 0.20 \%$, $-0.29 \pm 0.10 \%$	$0.48 \pm 0.20 \%$, $-0.29 \pm 0.11 \%$
$r_{\min} = 40$ Mpc/h	$0.56 \pm 0.18 \%$, $-0.32 \pm 0.17 \%$	$0.63 \pm 0.26 \%$, $-0.36 \pm 0.10 \%$	$0.54 \pm 0.24 \%$, $-0.29 \pm 0.12 \%$	$0.51 \pm 0.21 \%$, $-0.31 \pm 0.13 \%$

presented in Section 4.4), which connects the redshift-space flux power to the linear matter power spectrum through (Kaiser 1987; McDonald et al. 2000; McDonald 2003)

$$P^{\text{tree}}(k, \mu) = K_1^2(\mathbf{k})P_{\text{lin}}(k), \quad K_1(\mathbf{k}) \equiv (b_1 - b_n f \mu^2), \quad (21)$$

where f is the (linear) growth rate, and $P^{1\text{-loop}}(k)$ is the one-loop Ly α power spectrum

$$P^{1\text{-loop}}(k, \mu) = 2 \int_{\mathbf{q}} K_2^2(\mathbf{q}, \mathbf{k} - \mathbf{q})P_{\text{lin}}(|\mathbf{k} - \mathbf{q}|)P_{\text{lin}}(q) + 6K_1(\mathbf{k})P_{\text{lin}}(k) \int_{\mathbf{q}} K_3(\mathbf{k}, -\mathbf{q}, \mathbf{q})P_{\text{lin}}(q). \quad (22)$$

with higher order redshift-space kernels, $K_{2,3}$ [see equation (3.19) in Ivanov 2024] and we use the notation $\int_{\mathbf{q}} \equiv \int \frac{d^3q}{(2\pi)^3}$ to denote the 3D

integral over \mathbf{q} . The counter terms

$$P^{\text{ct}}(k, \mu) = -2(c_0 + c_2\mu^2 + c_4\mu^4)K_1(\mathbf{k})k^2P_{\text{lin}}(k), \quad (23)$$

scale as $k^2P_{\text{lin}}(k)$ and the stochastic contributions

$$P^{\text{st}}(k, \mu) = P_{\text{shot}} + a_0 \frac{k^2}{k_{\text{NL}}^2} + a_2 \frac{k^2\mu^2}{k_{\text{NL}}^2}, \quad (24)$$

scale as a constant shot noise piece and a scale- and angle-dependent term capturing small-scale clustering. The non-linear scale k_{NL} is defined as the point where the dimensionless power spectrum reaches unity, i.e. $\Delta^2 = k^3P_{\text{lin}}(k, z)/(2\pi^2) \sim 1$ which corresponds to $k_{\text{NL}} = 4.1 h \text{ Mpc}^{-1}$ at the redshift of the ABACUSSUMMIT simulation. Beyond these scales the one-loop EFT model should be treated as phenomenological. The parameters $c_{0,2,4}$ control the angular dependence of the counter terms and $a_{0,2}$ are the so-called Wilson coefficients cancelling the UV sensitivity (Ivanov 2024). The final

EFT model is evaluated at the redshift of the simulation, here $z = 2.5$. For conciseness, we suppress the explicit time dependence.

The one-loop Ly α power spectrum is described by the following set of nuisance parameters

$$\{b_1, b_\eta, b_2, b_{\mathcal{G}_2}, b_{(KK)_\parallel}, b_{\Pi_\parallel^{[2]}}, b_{\delta\eta}, b_{\eta^2}\}, \quad (25)$$

in addition to the cubic EFT terms

$$\{b_{\Pi_\parallel^{[3]}}, b_{(K\Pi^{[2]})_\parallel}, b_{\delta\Pi_\parallel^{[2]}}, b_{\eta\Pi_\parallel^{[2]}}, b_{\Gamma_3}\}, \quad (26)$$

and the counter and stochastic terms

$$\{c_{0,2,4}, a_{0,2}, P_{\text{shot}}\}, \quad (27)$$

yielding a total of 19 parameters. We adopt the prior choices of table 1 in Belsunce et al. (2025) and analytically marginalize over the parameters in equations (26)–(27).

4.2 EFT model for Ly α – quasar cross-correlation

We extend the one-loop EFT model to include cross-correlations of the Ly α forest with quasars (or halos) (Chudaykin & Ivanov 2025).⁸ In linear theory, this model has been computed as the geometric mean of the tree-level for both tracers, *i.e.*

$$P_{\times}^{\text{tree}}(k, \mu) = (b_1 - b_\eta f \mu^2)(b_1^q + f \mu^2) P_{\text{lin}}(k), \quad (28)$$

where the subscript \times denotes the cross-correlation, and superscript q represents the quasar (or halo) tracer, b_1^q is the linear bias parameter for quasars. To include higher-order contributions, we follow Chudaykin & Ivanov (2025) and use the same structure as the auto-correlation in equation (20) for the one-loop Ly α – quasar cross-correlation

$$\begin{aligned} P_{\times}^{\text{th}}(k, \mu) &= K_1(\mathbf{k}) K_1^q(\mathbf{k}) P_{\text{lin}}(k) \\ &+ 2 \int_{\mathbf{q}} K_2(\mathbf{q}, \mathbf{k} - \mathbf{q}) K_2^q(\mathbf{q}, \mathbf{k} - \mathbf{q}) P_{\text{lin}}(|\mathbf{k} - \mathbf{q}|) P_{\text{lin}}(q) \\ &+ 3 P_{\text{lin}}(k) \int_{\mathbf{q}} [K_1(\mathbf{k}) K_3^q(\mathbf{k}, -\mathbf{q}, \mathbf{q}) \\ &+ K_1^q(\mathbf{k}) K_3(\mathbf{k}, -\mathbf{q}, \mathbf{q})] P_{\text{lin}}(q) \\ &- (c_0 + c_2 \mu^2 + c_4 \mu^4) K_1^q(\mathbf{k}) k^2 P_{\text{lin}}(k) \\ &- (c_0^q + c_2^q \mu^2 + c_4^q \mu^4) K_1(\mathbf{k}) k^2 P_{\text{lin}}(k) \\ &- c_x^q (f \mu k)^4 (K_1^q(\mathbf{k}))^2 P_{\text{lin}}(k), \end{aligned} \quad (29)$$

where $K_1^q(\mathbf{k}) \equiv (b_1^q + f \mu^2)$ and $K_{2,3}^q$ are the standard redshift space kernels for galaxies (see e.g. Bernardeau et al. 2002).⁹ The last term c_x^q marginalizes over the estimated uncertainty due to the fingers-of-God modelling (Ivanov, Simonović & Zaldarriaga 2020; Chudaykin, Ivanov & Simonović 2021). The cross-correlation analysis is described by including three additionally sampled quasar bias parameters

$$\{b_1^q, b_2^q, b_{\mathcal{G}_2}^q\}. \quad (30)$$

in addition to EFT terms that are analytically marginalized over, *i.e.*

$$\{b_{\Gamma_3}^q, c_{0,2,4}^q, P_{\text{shot}}^q, c_x^q\}, \quad (31)$$

⁸For Ly α – halo cross-correlation analyses on hydrodynamic simulations using heuristic fitting functions (see Givans et al. 2022).

⁹To recover the expression for galaxies, one can remove the so-called line-of-sight operators from the Ly α one-loop model (Desjacques et al. 2018a, b), yielding the description for quasars (or galaxies) by setting $b_\eta = -1$, $b_{\eta^2} = 1$, $b_{\delta\eta} = -b_1$, $b_{(KK)_\parallel} = b_{\Pi_\parallel^{[2]}} = 0$, in addition to removing the cubic EFT terms in equation (26).

yielding a total of nine parameters. For the cross-bias terms, we follow the prior choices of Chudaykin & Ivanov (2023), *i.e.* we use $b_1^q \in \mathcal{U}(0, 4)$ and $\{b_{\mathcal{G}_2}^q, b_{\Gamma_3}^q\} \in \mathcal{N}(0, 2^2)$. We additionally, impose a simulation-based prior for $b_2^q \in \mathcal{U}(1.5, 1^2)$ (Ivanov et al. 2024c). For the counter terms we use $c_{0,2,4}^q/[h^{-1}\text{Mpc}]^2 \in \mathcal{N}(0, 10^2)$, and for the stochastic terms $P_{\text{shot}}^q, a_{0,2}^q \in \mathcal{N}(0, 1^2) \times \bar{n}^{-1}$ (note that we subtract the Poisson piece from the power spectrum, *i.e.* P_{shot} captures corrections to the Poisson sampling).

4.3 BAO scaling parameters

The key quantity of interest in the present work is to measure the position of the BAO peak (see e.g. DESI Collaboration 2024a, b). By fitting (linear) BAO templates to two-point clustering measurements and re-scaling these templates in radial and transverse directions (Padmanabhan & White 2008), we extract the BAO scaling parameters, parametrized by $\alpha \equiv \{\alpha_\parallel, \alpha_\perp\}$,¹⁰

$$\alpha_\parallel \equiv \frac{H^{\text{fid}}(z) r_s^{\text{fid}}(z_d)}{H(z) r_s(z_d)}, \quad \alpha_\perp \equiv \frac{D_A(z) r_s^{\text{fid}}(z_d)}{D_A^{\text{fid}}(z) r_s(z_d)}, \quad (32)$$

where the superscript ‘fid’ denotes the fiducial value at a reference cosmology. The BAO scaling parameters originate from a mismatch between the fiducial and true cosmologies giving rise to additional anisotropies when converting observed to projected coordinates, *i.e.* the so-called Alcock–Paczynski effect (Alcock & Paczynski 1979). Both parameters measure distortions in radial and transverse directions encoding the Hubble parameter $H(z)$, the sound horizon at the baryon-drag epoch $r_s(z_d)$ and the angular diameter distance $D_A(z)$ at the effective redshift z of the sample. This amounts to performing a re-mapping for the observed wavenumber

$$k \rightarrow k' \equiv k \left[\left(\frac{H_{\text{true}}}{H_{\text{fid}}} \right)^2 \mu^2 + \left(\frac{D_{A,\text{fid}}}{D_{A,\text{true}}} \right)^2 (1 - \mu^2) \right]^{1/2}, \quad (33)$$

and for the angle to the line of sight

$$\mu \rightarrow \mu' \equiv \mu \left(\frac{H_{\text{true}}}{H_{\text{fid}}} \right) \left[\left(\frac{H_{\text{true}}}{H_{\text{fid}}} \right)^2 \mu^2 + \left(\frac{D_{A,\text{fid}}}{D_{A,\text{true}}} \right)^2 (1 - \mu^2) \right]^{-1/2}, \quad (34)$$

where unprimed quantities are those measured observationally. In the present analysis, this corresponds to fitting the power spectrum predictions as a function of k' and μ' , *i.e.*,

$$P(k, \mu) = \frac{1}{\alpha_\perp^2 \alpha_\parallel} P(k', \mu') \quad (35)$$

which is evaluated at the fiducial redshift of the simulations ($z = 2.5$). The power spectrum decomposition into a smooth (nw) and oscillatory (w) component is done using a peak-average split in time-sliced perturbation theory which we introduce in the following (Blas et al. 2016a; Chudaykin et al. 2020).

4.4 Non-linear evolution of the BAO and IR-resummation

Eulerian perturbation theory fails to capture the non-linear suppression of the BAO wiggles in the statistics of clustering observables. This happens because the Eulerian perturbative expansion contains the linear displacement, which is enhanced by the soft (infrared, IR) modes, making the convergence of the perturbative series slow (Crocco & Scoccimarro 2008). Individual diagrams in Eulerian

¹⁰These can also be recast into an isotropic and anisotropic component.

standard perturbation theory contain IR enhanced terms, which cancel when added together (Blas, Garny & Konstantin 2013). The cancellation happens due to the equivalence principle, which dictates that physical observables are free from IR divergences (Blas et al. 2016b). The cancellation is, however, not complete because of the BAO feature (Baldauf et al. 2015; Blas et al. 2016a). The procedure to perform resummation of the large IR terms in the BAO is called ‘IR resummation’ (Senatore & Zaldarriaga 2015). IR resummation depends on the line of sight because the mapping from the rest frame of the tracer to the observer’s frame (commonly known as the redshift space mapping) depends on the displacement along the line-of-sight, which is IR enhanced. In the context of the Ly α forest, the line-of-sight operators are present in the bias expansion to begin with, and therefore it may not be immediately obvious that IR resummation will only depend on operators generated by the redshift space mapping. The bias expansion for the Ly α forest, however, depends only on the operators allowed by the equivalence principle, and these operators are IR safe. The line-of-sight displacements relevant for IR resummation are IR unsafe, and therefore they can only stem from the redshift space mapping of the rest-frame Ly α forest field. Thus, IR resummation for the Ly α forest depends on the same operators as a usual bias tracer without selection effects. This means that the non-linear damping of the BAO in the Ly α forest will have the same structure as that of galaxies. Specifically, using *time-sliced perturbation theory* at one-loop order we obtain (Ivanov & Sibiryakov 2018)

$$P_{FF} = P_{\text{tree}}^{FF} [e^{-\mathcal{S}(k,\mu)}(1 + \mathcal{S}(k,\mu))P_w + P_{nw}] + P_{1\text{-loop}}^{FF} [e^{-\mathcal{S}(k,\mu)}P_w + P_{nw}], \quad (36)$$

where the damping functions are given by

$$\begin{aligned} \mathcal{S}(k,\mu) &= k^2 (\Sigma^2(1 + f\mu^2(2 + f)) + \delta\Sigma^2 f^2 \mu^2 (\mu^2 - 1)), \\ \Sigma^2 &\equiv \frac{1}{6\pi^2} \int_0^{\Lambda_{\text{IR}}} dp P_{\text{lin}}(p, z) [1 - j_0(r_s p) + 2j_2(r_s p)], \\ \delta\Sigma^2 &= \int_0^{\Lambda_{\text{IR}}} \frac{dp}{2\pi^2} P_{\text{lin}}(p, z) j_2(pr_s). \end{aligned} \quad (37)$$

where r_s is comoving the BAO radius, $j_{0,2}(x)$ are spherical Bessel functions of the 0th and second order, while Λ_{IR} denotes the formal scale separating the domain of IR modes. We have tested that our final results do not depend on the particular choice of Λ_{IR} . In practice, we use $\Lambda_{\text{IR}} = 0.2 \text{ hMpc}^{-1}$ as in Blas et al. (2016a). As for the one-loop power spectrum, we approximate the last term in equation (36) as

$$\begin{aligned} P_{1\text{-loop}}^{FF} [e^{-\mathcal{S}(k,\mu)}P_w + P_{nw}] &\approx P_{1\text{-loop}}^{FF} [P_{nw}] + P_{1\text{-loop}}^{FF} \Big|_w, \\ P_{1\text{-loop}}^{FF} \Big|_w &\equiv e^{-\mathcal{S}(k,\mu)} (P_{1\text{-loop}}^{FF} [P_{\text{lin}}] - P_{1\text{-loop}}^{FF} [P_{nw}]), \end{aligned} \quad (38)$$

which allows us to significantly speed up the calculation of loop integrals. The approximation in equation (38) is within the regime of validity of the assumptions made to derive equation (37) so formally it does not introduce any additional errors at the given order in the IR resummation power counting. Let us note that in principle, one could derive analogues of equations (38) without explicitly separating the linear power spectrum into wiggly and non-wiggly components, see the discussion in Ivanov et al. (2024b).

Our expression for the IR resummed cross-spectrum is completely analogous to equation (36), with the obvious substitution $P^{FF} \rightarrow P^{FQ}$.

To derive the BAO parameters from the EFT power spectrum (in Section 4.3), we adopt a fitting methodology analogous to DESI Collaboration (2024a): we AP rescale only the arguments of the

wiggly part of the total power spectrum $P_w^{\text{tree}} + P^{1\text{-loop}} \Big|_w$, which ensures that the information we extract comes specifically from the BAO wiggles, and not from the broad-band power. Any shift in α_{\parallel} , α_{\perp} obtained this way should be interpreted as a shift of the BAO scale.

4.5 Fitting methodology

We calibrate the EFT bias parameters for the one-loop Ly α power spectrum and BAO scaling parameters by fitting our model, P^{model} , introduced in Sections 4.1–4.2 to the measured power spectra from ABACUSSUMMIT simulations, P_i^{data} , by sampling the χ^2 function¹¹

$$\chi^2 = \sum_i \frac{[P_i^{\text{data}} - P^{\text{model}}(k_i, \mu_i)]^2}{2(P_i^{\text{data}})^2 / N_i}, \quad (39)$$

where N_i are the Fourier modes per bin, k_i the Fourier wavenumbers with the cosine of the angle to the line-of-sight, $\mu = k_{\parallel}/k$. The data spectra are measured with $\Delta k = 0.005 \text{ h Mpc}^{-1}$ and in ten angular bins with $\Delta\mu = 0.1$. The data vector consists of the flux–flux spectrum, P_{FF} , and, when including the cross-correlation, the flux–quasar spectrum, P_{FQ} . For the fits we assume a Gaussian diagonal covariance: $\mathbf{C} = \text{diag}(\mathbf{C}^{\text{FF}}, \mathbf{C}^{\text{FQ}})$. The auto covariance is given by $C_{ii}^{\text{FF}} = 2(P_i^{\text{FF}})^2 / N_i$ and for the cross-spectrum covariance we use the form $C_{ii}^{\text{FQ}} = N_i^{-1} [(P_i^{\text{FQ}})^2 + P_i^{\text{FF}} P_i^{\text{QQ}}]$ where the quasar–quasar power spectrum is denoted by P^{QQ} .¹² For our baseline results, we use $k_{\text{max}} = 0.6 \text{ h Mpc}^{-1}$ for the autocorrelation and $k_{\text{max}}^{\times} = 0.3 \text{ h Mpc}^{-1}$ for the cross-correlation.¹³

For the fits to the autocorrelation, we sample over ten parameters, eight of which are EFT parameters and the two BAO scaling parameters, and analytically marginalize over the remaining 11 parameters. For jointly fitting the auto- and cross-correlation, we sample over 12 parameters (two of which are the BAO scaling parameters) and analytically marginalize over 18 parameters. The priors are given in Section 4.1 for the auto and Section 4.2 for the analysis including the cross-correlation. For the BAO scaling parameters α_{\parallel} and α_{\perp} we use $\alpha_{\parallel}, \alpha_{\perp} \in \mathcal{N}(1.0, 0.5^2)$ as priors.

In the present analysis, we fix the cosmological parameters of our EFT model to the input values of the ABACUSSUMMIT simulation, described in Section 2.1 for which the linear input power spectrum is generated using the Boltzmann solver CLASS-PT (Blas et al. 2011b; Chudaykin et al. 2020). We explore the parameter space using the Markov Chain Monte Carlo (MCMC) sampler Cobaya (Torrado & Lewis 2021). Each fit to an ABACUSSUMMIT simulation takes ~ 0.20 CPU-hours (using one AMD Milan CPU on the Perlmutter computer at NERSC) and is considered converged once the Gelman–Rubin diagnostic reaches $R - 1 < 0.01$ for all parameters (Gelman & Rubin 1992).

¹¹Note that we do not include a noise floor as in e.g. Givans et al. (2022) to use the full sensitivity of the simulations to constrain the bias parameters.

¹²We leave including the cross-covariance $C_{ii}^{\times} = 2(P_i^{\text{FF}} P_i^{\text{FQ}})^2 / N_i$ to future work (see e.g. Chudaykin & Ivanov 2025).

¹³We verified that varying the scale cuts by $\Delta k_{\text{max}} = \pm 0.2 \text{ h Mpc}^{-1}$ yielded consistent constraints on our final result on $\Delta\alpha_{\{\perp, \parallel\}}$, which following the scaling Universe argument (Ivanov 2024) and the non-linear scale of the simulations (Belsunce et al. 2025), defined as $\Delta^2(k) = k^3 P_{\text{lin}}(k, z) / (2\pi^2) \approx 1$, is located at $k_{\text{NL}} \approx 2.2 \text{ h Mpc}^{-1}$, is a conservative scale cut. Beyond which the EFT should be treated as a purely phenomenological model.

Table 4. Mean best-fitting values for the one-loop EFT parameters obtained from the ABACUSSUMMIT simulation ‘one’ and line-of-sight parallel to the y -axis using models one to four. The default fit is performed with $k_{\text{max}} = 0.6 h \text{ Mpc}^{-1}$ for the autocorrelation. The last two rows of the upper section of the table quotes the BAO scaling parameters which are consistent with unity, *i.e.* the EFT model debiases the BAO measurement. We analytically marginalize over the parameters shown in the bottom part of the table and recover their posteriors from the chains a posteriori. Note that the counter terms ($c_{0,2,4}$) are divided by $(h \text{ Mpc}^{-1})^2$. The resulting χ^2 and reduced χ^2_{ν} for the best-fitting linear parameters are quoted in the last two rows for 1175 data points and 18 degrees of freedom. In Appendix C, we show the triangle plot for the sampled parameters for the first column. In Table C1, we show the corresponding EFT bias parameters for the fits to the mean data vector.

$b_{\mathcal{O}}$	Model 1	Model 2	Model 3	Model 4
b_1	$-0.1463^{+0.0016}_{-0.0014}$	$-0.1298^{+0.0015}_{-0.0013}$	-0.1311 ± 0.0012	-0.1273 ± 0.0011
b_{η}	0.1448 ± 0.0027	0.1298 ± 0.0026	0.2734 ± 0.0030	0.3073 ± 0.0031
b_2	$-0.0292^{+0.1963}_{-0.1749}$	$-0.1211^{+0.1861}_{-0.1823}$	$0.0304^{+0.1145}_{-0.1050}$	-0.0061 ± 0.1039
$b_{\mathcal{G}_2}$	-0.0364 ± 0.1421	$0.0241^{+0.1346}_{-0.1782}$	0.0326 ± 0.0784	0.0175 ± 0.0735
b_{η^2}	-0.1199 ± 0.0556	$-0.3537^{+0.0778}_{-0.0546}$	-0.1253 ± 0.0539	-0.0800 ± 0.0543
$b_{\delta\eta}$	$-0.3543^{+0.1384}_{-0.1096}$	$-0.6961^{+0.1612}_{-0.0513}$	$-0.1398^{+0.0965}_{-0.0786}$	$-0.1195^{+0.0978}_{-0.0786}$
$b_{(KK)_{\parallel}}$	$-0.0403^{+0.3311}_{-0.2811}$	$-0.3424^{+0.4923}_{-0.2546}$	-0.0785 ± 0.1947	$0.0007^{+0.1799}_{-0.1986}$
$b_{\Pi_{\parallel}^{[2]}}$	$-0.3146^{+0.0583}_{-0.0469}$	-0.2479 ± 0.0531	$-0.2950^{+0.0590}_{-0.0527}$	-0.3425 ± 0.0588
α_{\parallel}	1.0040 ± 0.0094	1.0020 ± 0.0102	0.9994 ± 0.0084	0.9994 ± 0.0081
α_{\perp}	$1.0011^{+0.0039}_{-0.0034}$	$1.0004^{+0.0048}_{-0.0035}$	$1.0014^{+0.0038}_{-0.0033}$	1.0005 ± 0.0032
$b_{\Pi_{\parallel}^{[3]}}$	0.0420 ± 0.0486	-0.0392 ± 0.0443	0.7338 ± 0.0447	0.8240 ± 0.0450
$b_{\delta\Pi_{\parallel}^{[2]}}$	-1.6886 ± 0.1442	-1.4214 ± 0.1335	-0.9226 ± 0.1722	-1.1403 ± 0.1795
$b_{(K\Pi^{[2]})_{\parallel}}$	-1.1586 ± 0.1654	-0.8222 ± 0.1479	-1.5634 ± 0.1531	-1.9723 ± 0.1530
$b_{\eta\Pi_{\parallel}^{[2]}}$	-0.9338 ± 0.4428	-0.4698 ± 0.4156	-1.9866 ± 0.5684	-2.2649 ± 0.5922
b_{Γ_3}	0.4228 ± 0.0836	0.4802 ± 0.0763	-0.3260 ± 0.0947	-0.3680 ± 0.0971
P_{shot}	-0.3066 ± 0.0970	0.1080 ± 0.0757	0.3826 ± 0.0379	0.2364 ± 0.0328
a_0	-0.0109 ± 0.8768	-0.2250 ± 0.7127	-0.5768 ± 0.7067	-0.0894 ± 0.6710
a_2	-0.6411 ± 2.6056	-2.6171 ± 2.2136	-4.3836 ± 2.9013	-5.6704 ± 2.9223
c_0	0.1425 ± 0.0351	-0.0449 ± 0.0301	-0.1689 ± 0.0171	-0.1158 ± 0.0163
c_2	-0.0882 ± 0.0602	0.1304 ± 0.0558	0.3216 ± 0.0578	0.3072 ± 0.0573
c_4	-0.1559 ± 0.0278	-0.2293 ± 0.0259	-0.4255 ± 0.0334	-0.4418 ± 0.0346
χ^2	1202	1255	1276	1281

4.6 BAO and EFT parameters from ABACUSSUMMIT

In this section, we present results on fits to the *Ly* α forest autocorrelation and *Ly* α –quasar cross-correlation measured on ABACUSSUMMIT simulations, introduced in Section 2.3. The large size of the simulation boxes at redshift $z = 2.5$ gives access to many quasi-linear modes which, in turn, yield tight constraints on the (linear) bias parameters. We present the results for all four models for simulation ‘one’ and line-of-sight y on the autocorrelation in Table 4 and for the cross-correlation in Table 5. Our measured linear bias parameters are consistent with the linear-theory results presented in Hadzhiyska et al. (2023b) at the $\sim 1\sigma$ level. For both series of fits, we find BAO scaling parameters that are consistent with unity showing that the EFT model debiases the BAO fit. In agreement with the autocorrelation analysis on ACCEL² hydrodynamical simulations of box size $L = 160 h^{-1} \text{ Mpc}$ with a physical resolution down to $25 h^{-1} \text{ kpc}$ in Belsunce et al. (2025), we detect the shot noise for all four models whilst we do not find evidence for the Wilson coefficients a_0 and a_2 which cancel the UV sensitivity of the loop integrals. However, we detect the $c_{0,2,4}$ counter terms which account for the short range non-locality (McDonald & Roy 2009). We note that since the quasar field does not change between the four models, we would

expect the counter terms $c_{0,2,4}^q$ to be in agreement with each other. Whilst the counter terms between models 1 and 2 as well as 3 and 4 are consistent with each other, we find a tension between those two sets of models (up to $\sim 3\sigma$). The presence of this tension suggests that the counter terms are affected by degeneracies which are not (fully) removed by including the *Ly* α – quasar cross-correlation. We leave including the quasar-autocorrelation to future work.

For the autocorrelation the quality of the fits ($\chi^2 = 1202, 1255, 1276, 1281$) is acceptable for 1175 data points with an increasing χ^2 for the models one to four. The quality of the fits for the cross-correlation is somewhat worse ($\chi^2 = 1791, 1839, 1894, 1901$) for 1755 data points.^{14,15}

For completeness, we give the numerical values of the EFT fits to the autocorrelation mean data vector in Table C1 and for the joint

¹⁴We have verified that removing the counter terms from the cross-correlation fits did not substantially affect the quality of the fits.

¹⁵We caution the reader that since we are analytically marginalizing over the parameters that linearly enter the EFT model (see equation 27) the counting of degrees of freedom requires some care (Bridle et al. 2002; Taylor & Kitching 2010).

Table 5. Same as Table 4 but for the joint fits of the auto- and cross-correlation. The top section shows the sampled parameters and the middle (bottom) section the ones that we analytically marginalize over for the auto- (cross-) correlation. We use 1755 data points, sample over 12 parameters and analytically marginalize over 18.

$b_{\mathcal{O}}$	Model 1	Model 2	Model 3	Model 4
b_1^q	3.3105 ± 0.0427	3.3461 ± 0.0441	3.3906 ± 0.0365	3.3666 ± 0.0356
b_2^q	1.6487 ± 0.9992	1.3933 ± 1.0026	$2.3774^{+1.0202}_{-1.0408}$	$1.2312^{+0.9905}_{-1.0822}$
b_1	-0.1433 ± 0.0011	-0.1264 ± 0.0011	-0.1277 ± 0.0010	-0.1234 ± 0.0009
b_η	0.1492 ± 0.0018	0.1352 ± 0.0017	0.2818 ± 0.0023	0.3180 ± 0.0024
b_2	$-0.2162^{+0.1696}_{-0.1281}$	$-0.5015^{+0.1268}_{-0.0849}$	$-0.1706^{+0.1168}_{-0.1064}$	-0.1298 ± 0.1058
$b_{\mathcal{G}_2}$	$-0.1581^{+0.1141}_{-0.0908}$	$-0.3255^{+0.1287}_{-0.0814}$	-0.0912 ± 0.0710	-0.0718 ± 0.0668
b_{η^2}	-0.1786 ± 0.0529	$-0.3289^{+0.0524}_{-0.0455}$	-0.2074 ± 0.0489	-0.1712 ± 0.0503
$b_{\delta\eta}$	$-0.4206^{+0.1422}_{-0.1148}$	$-0.5092^{+0.1589}_{-0.0873}$	$-0.1835^{+0.0811}_{-0.0715}$	$-0.1084^{+0.0737}_{-0.0657}$
$b_{(KK)_\parallel}$	$0.3246^{+0.1906}_{-0.2469}$	$0.5375^{+0.2356}_{-0.2848}$	$0.4644^{+0.1303}_{-0.1508}$	$0.4604^{+0.1242}_{-0.1478}$
$b_{\Pi_\parallel^{[2]}}$	$-0.2596^{+0.0573}_{-0.0479}$	$-0.2056^{+0.0536}_{-0.0437}$	-0.1921 ± 0.0536	-0.2034 ± 0.0551
α_\parallel	0.9997 ± 0.0068	1.0021 ± 0.0072	0.9991 ± 0.0059	0.9992 ± 0.0057
α_\perp	1.0013 ± 0.0034	$1.0024^{+0.0041}_{-0.0038}$	1.0010 ± 0.0029	0.9995 ± 0.0026
$b_{\Pi_\parallel^{[3]}}$	-0.1537 ± 0.0316	-0.1089 ± 0.0299	0.2844 ± 0.0340	0.2747 ± 0.0350
$b_{\delta\Pi_\parallel^{[2]}}$	-1.2028 ± 0.1269	-0.7561 ± 0.1163	-1.1539 ± 0.1269	-1.4458 ± 0.1297
$b_{(K\Pi^{[2]})_\parallel}$	-1.5418 ± 0.1351	-2.0036 ± 0.1225	-1.9400 ± 0.1102	-2.2706 ± 0.1089
$b_{\eta\Pi_\parallel^{[2]}}$	0.2632 ± 0.3354	1.9228 ± 0.3112	-2.5832 ± 0.3933	-3.0661 ± 0.4080
b_{Γ_3}	0.7170 ± 0.0629	1.4109 ± 0.0581	-0.1916 ± 0.0687	-0.2057 ± 0.0703
P_{shot}	-0.4064 ± 0.0735	-0.2247 ± 0.0567	0.4116 ± 0.0212	0.2750 ± 0.0175
a_0	1.1312 ± 3.1873	0.4227 ± 2.7668	-4.1192 ± 2.9160	-2.6957 ± 2.8277
a_2	0.5094 ± 4.8389	-1.3978 ± 4.6306	-1.5596 ± 4.7806	-1.3415 ± 4.8413
c_0	0.1863 ± 0.0316	0.1179 ± 0.0272	-0.1787 ± 0.0141	-0.1207 ± 0.0134
c_2	-0.2061 ± 0.0365	-0.1210 ± 0.0341	0.1979 ± 0.0313	0.1442 ± 0.0316
c_4	-0.0610 ± 0.0221	-0.0862 ± 0.0203	-0.3384 ± 0.0235	-0.3535 ± 0.0243
$b_{\mathcal{G}_2}^q$	0.1427 ± 0.6497	$-0.3223^{+0.6085}_{-0.7342}$	0.1116 ± 0.5026	0.0458 ± 0.4734
P_{shot}^q	-0.8379 ± 2.8662	2.1759 ± 2.6402	-5.1053 ± 2.2466	-5.9015 ± 2.0804
c_0^q	0.3062 ± 0.4629	-0.1038 ± 0.4762	1.7626 ± 0.3687	1.8137 ± 0.3436
c_2^q	1.2548 ± 0.5366	0.7582 ± 0.5564	1.8479 ± 0.6097	1.6705 ± 0.6144
c_4^q	0.3197 ± 0.3279	1.2225 ± 0.3336	0.0871 ± 0.3936	0.0509 ± 0.4096
$b_{\Gamma_3}^q$	0.0361 ± 0.0102	0.0514 ± 0.0091	-0.1879 ± 0.0098	-0.2424 ± 0.0098
c_x	0.2753 ± 0.9774	0.2030 ± 0.9631	0.6029 ± 0.9657	0.5385 ± 0.9779
χ^2	1791	1839	1894	1901

fits of the auto- and cross-correlation in Table C2. We note that we find consistent, yet tighter, results than fits to individual simulations given in Table 4. The fits to the mean data vector can not only differentiate between the linear biases but also between the non-linear bias parameters. In particular, we find clear differences in b_2 between the four models which, in turn, is a driver for the non-linear shift of the BAO peak.¹⁶

The best-fitting auto spectra are shown in the left panel of Fig. 5 for four angular bins. The corresponding best-fitting cross-spectrum obtained from the joint fit of the auto- and cross-spectrum is shown in the right panel. (For conciseness, we omit the auto spectrum from the joint fit as the quality is identical by eye.) The bottom two panels show the residuals between the model and the data: (i) the middle panel

shown the residuals divided by the error bars, illustrating that, even with the small error bars from ABACUSSUMMIT, we find agreement at the $\sim 2\sigma$ -level at *all* scales; (ii) the bottom panel shows that the deviations at small scales are for the auto alone sub-2 per cent whilst for the joint auto- and cross-correlation fit around 3 – 5 per cent. Note that the quality of the fit degrades along the line-of-sight due to the velocity field.

To assess at which scales the EFT corrections become important, we show the loop corrections, the ratio of the EFT model to the tree-level power spectrum, in the left for the auto and in the right panel the cross-spectrum obtained from the joint auto- and cross-correlation fit of the of Fig. 6. Following baseline expectation, the large amount of available (quasi-) linear modes results in the EFT corrections to the tree-level power spectrum becoming relevant at $k \gtrsim 0.2 h \text{ Mpc}^{-1}$, especially for transverse modes (blue). As expected the perturbative reach of the EFT models is smaller for quasars (or halos) than for

¹⁶Since the error bars are divided by an additional factor of $\sqrt{6}$, the resulting χ^2 values should be taken with a grain of salt.




Figure 5. EFT fits to the Ly α autocorrelation (left panel) and Ly α -quasar cross-correlation (right panel) function in four (out of ten) angular bins, μ , with $k_{\max} = 0.6 h \text{ Mpc}^{-1}$ for the first simulation of model one along the y-axis. The best-fitting parameters are shown in the first column of Table 4. Best-fitting EFT model (solid line) compared to P^{tree} (dashed line) and data points. The error bars are obtained assuming a diagonal Gaussian covariance based on the number of expected Fourier modes per bin $P(k, \mu)\sqrt{2/N(k, \mu)}$. The residuals between the model and the data are shown in the bottom two panels. The upper one shows the difference between model and data divided by the error bars with a grey band indicating the 1σ error region to guide the eye. The bottom panel shows the ratio of the two, with a 2 per cent error band to guide the eye.

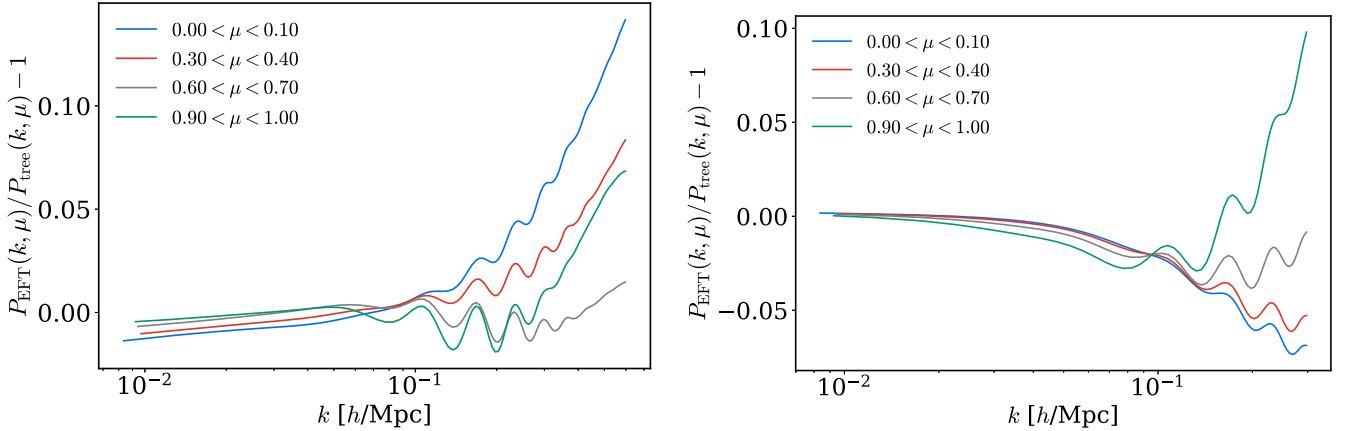


Figure 6. Size of the one-loop corrections for the Ly α autocorrelation (left panel) and Ly α -quasar cross-correlation (right panel) model obtained from Fig. 5 by comparing the IR-resummed linear theory power spectrum, P_{tree} , to the best-fitting one-loop power spectrum, P_{EFT} .

the Ly α forest (see, e.g. Maus et al. (2025) for a discussion of EFT full-shape analyses in the context of DESI for low-redshift galaxies and quasars). It is interesting to note that for the cross-correlation the loop corrections are more pronounced for the line-of-sight modes than for the transverse ones. Note that the unity-crossing of the loop corrections is an indicator as to when EFT formally breaks down which we do not reach here.¹⁷

The measured values for the BAO scaling parameters for the Ly α autocorrelation (joint auto- and cross-correlation) are shown in the left (right) panel of Fig. 7 with the corresponding mean values of the BAO parameters tabulated in Table 6. We find that the EFT model debiases the α parameters and our fits are consistent with a value of

$\alpha_{\parallel, \perp} = 1$ across all four models. We find a mean cross-correlation coefficient over the 48 realizations for the α 's of $\rho = -0.39 \pm 0.04$. Similar to the case for the Ly α autocorrelation, we find for the cross-correlation that the EFT debiases the BAO fit with a mean cross-correlation coefficient of $\rho^{\times} = -0.43 \pm 0.03$.

4.6.1 EFT fits to the mean data vector

To stress-test the fitting of the EFT model using a low(er) noise power spectrum, we average the measured spectra over simulation realizations and lines-of-sight, resulting in reduced error bars by a factor of $\sqrt{6}$. Again, we jointly fit the BAO scaling and EFT bias parameters and find no evidence for a shift in the BAO peak. For the autocorrelation, we show for model 1 the evolution of the BAO parameter as a function of scale-cut k_{\max} in Fig. 8. It is

¹⁷See, however, the discussion in Belsunce et al. (2025).

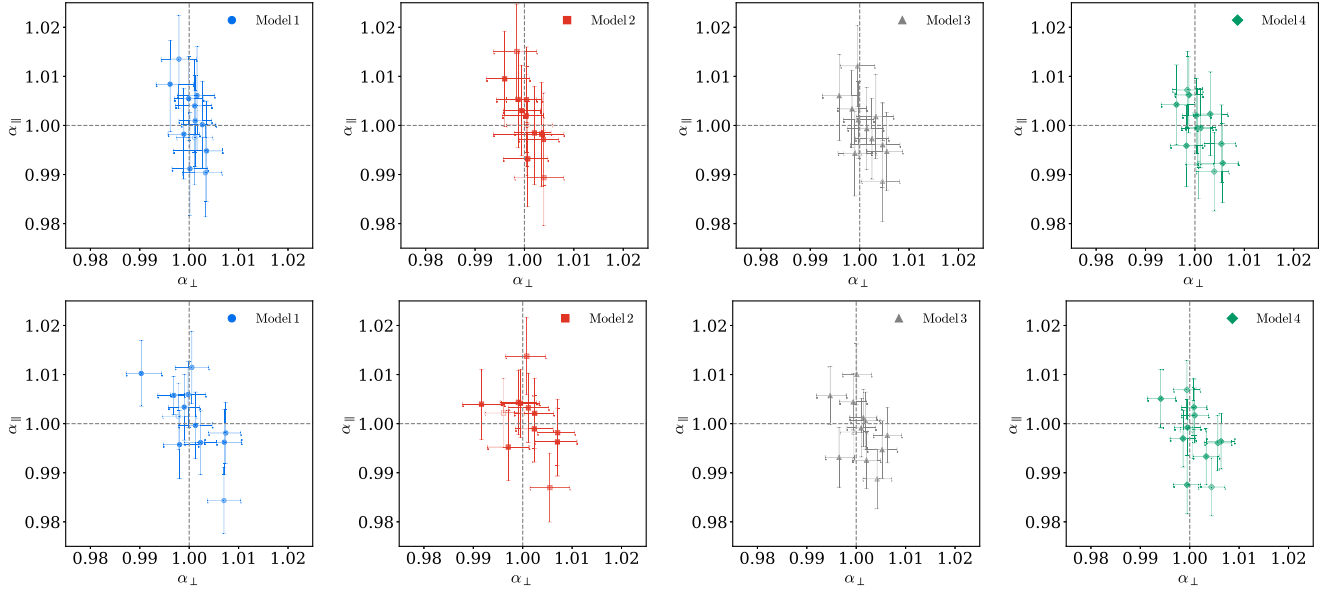


Figure 7. BAO scaling parameters obtained from a joint fit of the α 's and EFT bias parameters to each simulation, model and line-of-sight which are represented by a different colour-shape combination: 1 (blue, circle), 2 (red, square), 3 (grey, triangle), and 4 (green, diamond). Filled (empty) markers represent the line-of-sight y -axis (z -axis). We show the Ly α autocorrelation in the top row and the joint fit to the Ly α auto and Ly α -quasar cross-correlation in the bottom row. The average cross-correlation coefficients are $\bar{\rho} = -0.39 \pm 0.04$ and $\bar{\rho}^\times = -0.42 \pm 0.03$, for auto alone and auto and cross jointly, respectively. The EFT model debiases the BAO fit and we find a value for BAO scaling parameters consistent with unity with the corresponding mean values for the α 's in Table 6.

Table 6. Radial (α_{\parallel}) and transverse (α_{\perp}) BAO scaling parameters obtained from jointly fitting the BAO and EFT bias parameters. *Top panel:* mean of fits to $N_{\text{sim}} = 12$ simulations of each model, illustrated in Fig. 7. The radial (transverse) mean error bars are divided by $\sqrt{N_{\text{sim}}}$ ($\sqrt{N_{\text{sim}}/2}$). *Bottom panel:* Fits to the mean data vector averaged over 12 realizations for each model. The left (right) columns are the values for the auto (auto + cross) correlation. We apply scale cuts of $k_{\text{max}} = 0.6 h \text{ Mpc}^{-1}$ to P^{FF} and $k_{\text{max}} = 0.3 h \text{ Mpc}^{-1}$ to P^{FQ} .

Model	Auto		Auto + Cross	
	α_{\parallel}	α_{\perp}	α_{\parallel}	α_{\perp}
Mean of fits				
1	$1.0009^{+0.0027}_{-0.0028}$	$1.0006^{+0.0014}_{-0.0016}$	$1.0007^{+0.0019}_{-0.0019}$	$1.0006^{+0.0010}_{-0.0010}$
2	$1.0014^{+0.0029}_{-0.0029}$	$1.0006^{+0.0016}_{-0.0020}$	$1.0008^{+0.0020}_{-0.0020}$	$1.0008^{+0.0011}_{-0.0012}$
3	$0.9992^{+0.0024}_{-0.0024}$	$1.0012^{+0.0014}_{-0.0016}$	$0.9988^{+0.0017}_{-0.0017}$	$1.0010^{+0.0008}_{-0.0008}$
4	$0.9989^{+0.0023}_{-0.0023}$	$1.0010^{+0.0013}_{-0.0013}$	$0.9979^{+0.0017}_{-0.0017}$	$1.0011^{+0.0008}_{-0.0008}$
Fits to stacked spectra				
1	$1.0014^{+0.0039}_{-0.0041}$	$1.0002^{+0.0015}_{-0.0017}$	$1.0015^{+0.003}_{-0.003}$	$1.0003^{+0.0013}_{-0.0013}$
2	$1.0025^{+0.0038}_{-0.0039}$	$1.0016^{+0.0024}_{-0.0025}$	$1.0006^{+0.003}_{-0.0031}$	$1.0015^{+0.0019}_{-0.0018}$
3	$0.9994^{+0.0034}_{-0.0034}$	$1.0014^{+0.0015}_{-0.0015}$	$1.0000^{+0.0029}_{-0.003}$	$1.0008^{+0.0014}_{-0.0013}$
4	$0.9988^{+0.0033}_{-0.0034}$	$1.0011^{+0.0015}_{-0.0014}$	$0.9994^{+0.0028}_{-0.0028}$	$1.0007^{+0.0013}_{-0.0014}$

interesting (and reassuring) to note that the error bars only slightly reduce above $k \sim 0.3 h \text{ Mpc}^{-1}$ as this indicates that we can isolate the information stemming from the BAO feature and not pick up full-shape information. This figure is an additional confirmation of our conservative choice of scale cut-off $k_{\text{max}} = 0.6 h \text{ Mpc}^{-1}$. The corresponding numerical values for all four models are given in Table 6 for our baseline scale cut. The results for the joint fit to the auto- and cross-correlation are given in the bottom panel of the same table. Both series of fits are consistent with BAO scaling parameters of $\alpha_{\parallel} = \alpha_{\perp} = 1$.




Figure 8. Mean values of α_{\parallel} (black dashed line) and α_{\perp} (red dotted line) obtained from fits to the mean of six simulations along the y -axis for model 1 with the 1σ error bars. The fits are performed for increasing values of $k_{\text{max}} = 0.3 - 1.0 h \text{ Mpc}^{-1}$. We give the numerical values for our baseline value of $k_{\text{max}} = 0.6 h \text{ Mpc}^{-1}$ in the first row of the bottom panel in Table 6. This plot shows that the EFT fits can isolate the BAO information and we use $k_{\text{max}} = 0.6 h \text{ Mpc}^{-1}$ as a conservative scale cut for our analysis.

4.7 Consistency tests

4.7.1 Quasi-linear theory fits to the EFT best-fitting data vector

As a consistency check and to reconcile the analyses in the first part of this work to the present section, we replace the ABACUSSUMMIT data vector with the EFT best-fitting spectra using the VEGA fitting pipeline introduced in Section 2.5.¹⁸ First, we transform the EFT best-fitting theory predictions for $P(k, \mu)$ to anisotropic spectra using the

¹⁸Note that we fix $\alpha_{\parallel} \equiv \alpha_{\perp} \equiv 1$.




Figure 9. Fitting the EFT best-fitting Ly α autocorrelation spectra (with BAO scaling parameters fixed to unity) instead of ABACUSSUMMIT simulation data. We show the BAO parameters as a function of minimum separation, r_{\min} , whilst the EFT templates are obtained from fits with $k_{\max} = 0.6 h \text{ Mpc}^{-1}$ corresponding to $r \approx 10.5 h^{-1} \text{ Mpc}$. We obtain consistent results when replacing the ABACUSSUMMIT data with the EFT templates, *i.e.* we detect a shift of 0.3–0.4 per cent for α_{\parallel} , and a smaller shift of α_{\perp} which is consistent with zero within $\sim 1-2\sigma$ for $r_{\min} \geq 20 h^{-1} \text{ Mpc}$. We attribute the larger measured shift at $r_{\min} = 10 h^{-1} \text{ Mpc}$ to issues with the small-scale modelling when using the quasi-linear VEGA model.

standard multipole decomposition

$$P_{\ell}(k) \equiv \frac{2\ell + 1}{2} \int_{-1}^1 d\mu \mathcal{L}_{\ell}(\mu) P(k, \mu), \quad (40)$$

where \mathcal{L}_{ℓ} are the Legendre moments in multipoles ℓ . Secondly, we perform the same analysis as in Section 3.1 and measure the BAO scaling parameters from our EFT templates using the empirical model of Prats et al. (2015). The results are shown in Fig. 9. Following the baseline expectations, we obtain consistent results, *i.e.* the VEGA pipeline measures the same biases in α when replacing the ABACUSSUMMIT data by the best-fitting EFT spectra – this confirms that the EFT accurately describes the data.

4.7.2 Fisher forecasts for the non-linear BAO shift

Another avenue to quantify the magnitude of the BAO shift stemming from non-linearities in the 3D clustering of the Ly α forest, is to theoretically predict it using perturbation theory. We follow the approach developed in the context of galaxies (Chen et al. 2024) and in the context of the Ly α forest (Belsunce et al. 2025) and predict the shift in the BAO scaling parameters when using a linear theory model. Whilst the ABACUSSUMMIT simulations resolve the BAO feature, they have much lower physical resolution than hydrodynamical simulations (Bolton et al. 2017; Chabanier et al. 2024). Thus, the accuracy of the ABACUSSUMMIT simulations should be confined to larger scales and the exact magnitude of the non-linear bias parameters should be treated as indicatively, *e.g.* Chabanier et al. (2024) showed that the physical resolution of the simulation affects the velocity bias. Therefore, this section serves to generate some intuition for the reader and we compare the Fisher prediction, briefly introduced in the following, to fits of the ABACUSSUMMIT spectra using a linear theory model in Fourier space (*i.e.* this amounts to setting all the higher order bias parameters to zero bar b_1 and b_{η} in Section 4.1).

Conceptually, in real space, the BAO feature contracts isotropically around overdensities and expands around underdensities. However,

Table 7. Comparison of the Fisher prediction (left column) for the BAO parameter in radial (α_{\parallel}) and transverse (α_{\perp}) direction to measurements of the BAO parameters using linear theory (right column) for the Ly α autocorrelation. The corresponding bias parameters are given in Table C1 and obtained from fitting the mean data vector (averaged over 12 simulations per model).

Model	Fisher prediction		Linear theory fit (auto)	
	$\Delta\alpha_{\parallel}$ [%]	$\Delta\alpha_{\perp}$ [%]	$\Delta\alpha_{\parallel}$ [%]	$\Delta\alpha_{\perp}$ [%]
1	-0.08 ± 0.07	-0.01 ± 0.01	-0.08 ± 0.19	0.16 ± 0.14
2	-0.12 ± 0.09	0.04 ± 0.02	-0.23 ± 0.19	0.22 ± 0.15
3	-0.33 ± 0.11	0.03 ± 0.01	-0.63 ± 0.19	-0.07 ± 0.14
4	-0.43 ± 0.13	0.07 ± 0.04	-0.59 ± 0.18	-0.22 ± 0.14

the Ly α forest is suppressed around overdensities resulting in an expanded BAO radius. In general, the shift is sourced by long-wavelength modes (wavelengths longer than the distance a sound wave travels prior to the baryon drag epoch, post-recombination, denoted by r_d), and is captured in perturbation theory by isolating the $P^{(22)}$ term (see *e.g.* Chen et al. 2024, and references therein). Schematically, the BAO shift in the autocorrelation is driven by (i) the relative amplitude of the linear (b_1) and quadratic contributions (b_2) to the Ly α clustering; and (ii) the amplitude of density fluctuations on BAO scales which is reduced by the growth factor $D(z)^2$.¹⁹

We use a Fisher formalism to predict the BAO shift using measurements of the bias parameters on ABACUSSUMMIT

$$\Delta\{\alpha_{\parallel}, \alpha_{\perp}\} = F_{\{\alpha_{\parallel}, \alpha_{\perp}\}ab}^{-1} t_i^b C_{ij}^{-1} \epsilon_j, \quad (41)$$

where C is the covariance matrix, t_i^b is the linear template $dP_i/d\theta_b$ corresponding to fitting parameter θ_b , and F_{ab} is the Fisher matrix. This set-up assumes that, in addition to the BAO scales, we also marginalize over a set of other parameters including the linear density and velocity bias, and BAO damping parameters. For the covariance matrix C , we use the Gaussian covariance matrix from Section 4.5, and ϵ , are the ‘wedges’ defined in total separation $k = (k_{\parallel}^2 + \mathbf{k}_{\perp}^2)^{1/2}$ and in μ bins for the analytic description of the BAO shift, multiplied by the mean-square linear density fluctuation in a sphere of radius r_d and the response of the power spectrum, $\frac{\partial P}{\partial \ln k}$. The fit is performed with $k_{\min} = 0.02 h \text{ Mpc}^{-1}$ and $k_{\max} = 0.3 h \text{ Mpc}^{-1}$ with a linear $\Delta k = 0.005 h \text{ Mpc}^{-1}$ spacing.

In Table 7, we compare the results on the biases of the BAO scaling parameters obtained when performing linear theory fits in Fourier space (first column) to the Fisher forecasted BAO shifts (second column). As expected, both approaches agree at the 1–2 σ level. The linear theory fits show a negatively biased scaling parameter in radial direction ($-0.2 - 0.5$ per cent) and a shift that is consistent with zero in the transverse direction at the $\sim 1.5\sigma$ level, depending on the chosen model. The predictions and linear theory fits are obtained using $k_{\max} = 0.3 h \text{ Mpc}^{-1}$. Varying k_{\max} does not change our conclusions. We emphasize that these BAO shifts cannot directly be extrapolated to observational data, since we do not expect the (higher-order) biases to agree with the data. Therefore, we refer the reader to Belsunce et al. (2025) where they use hydrodynamical simulations which are expected to better represent the underlying physics.

To understand the differences between the results obtained using the VEGA fitting pipeline in Section 3.3 and the linear theory fits

¹⁹The picture for the cross-correlation with quasars is analogous. The relative amplitude of the linear and quadratic contribution is added the one of the autocorrelation (Belsunce et al. 2025; Chudaykin & Ivanov 2025).

in Fourier space we perform a series of tests: (i) To assess that both models agree, we perform a similar fit as in Fig. 9. We fit with VEGA the linear theory best-fitting power spectrum (with BAO scaling parameters set to unity) instead of the ABACUSSUMMIT data. We find matching constraints as the ones shown in Fig. 3, albeit with smaller error bars since the EFT best-fitting data vector is noiseless, substantially reducing the scatter in the measurement. This test confirms that the models are consistent with each other. (ii) To assess the validity of the covariance matrix, we compute the Jacobian matrix that converts the Gaussian covariance from Fourier to configuration space in $\{r_p, r_t\}$ space and fit the LCV ABACUSSUMMIT data vector. We obtain consistent results with the eBOSS covariance matrix results, following baseline expectation. (iii) Further, to be closer to the implementation in VEGA we remove the next-to-leading order correction of the one-loop IR resummation in time-sliced perturbation theory, discussed in Section 4.4. We find that the radial BAO scaling parameter shifts by $+(0.1 - 0.2)\%$. This series of tests demonstrates how sensitive the results are to exact analysis choices and we leave exploring the optimal modelling for linear theory analyses of the BAO scaling parameters to future work.

5 DISCUSSION AND CONCLUSIONS

5.1 Related literature

Adopting the optimal combination for the isotropic parameter (DESI Collaboration 2024b):

$$\alpha_{\text{iso}} = \alpha_{\perp}^{9/20} \alpha_{\parallel}^{11/20} \quad (42)$$

and the Alcock–Paczynski parameter:

$$\alpha_{\text{AP}} = \alpha_{\perp} / \alpha_{\parallel}, \quad (43)$$

we can convert our main result, the LCV-reduced α_{\parallel} and α_{\perp} constraints from the autocorrelation function into constraints on α_{iso} and α_{AP} , to obtain shifts of 0.18 ± 0.05 per cent and -0.8 ± 0.2 per cent, respectively. These are consistent with the measurements we make from the cross-correlation with quasars, namely: 0.25 ± 0.12 per cent and -0.85 ± 0.25 per cent, respectively.

We remark that these values of α_{iso} are markedly different from those reported in another recent work studying the shift of the BAO peak in the Ly α forest from the isotropic part (monopole) of the signal (Sinigaglia et al. 2024b). The authors of that work find a 1 per cent negative shift to the α_{iso} parameter (albeit with a slightly different parametrization) confirmed at the $\gtrsim 3\sigma$ level. While our work is broadly inconsistent with this result, we note that there are several important differences that could play a role in causing this discrepancy. Sinigaglia et al. (2024b) adopt a modification of the FGPA approach used to create the ABACUSSUMMIT Ly α forest mocks, which assigns a different connection between gas and dark matter depending on the cosmic web classification of the region. This novel method is motivated by the fact that the different phases of the neutral gas exhibit a different relation with the underlying matter (Furlanetto & Oh 2009) and could therefore yield more accurate behaviour on small scales where these astrophysical effects are highly important. At the same time, care should be taken as to how these effects are implemented, as they could cause a distortion at the BAO peak scale, which need to be robustly tested. We remind the reader that the shifts we are looking for in this work are at the 0.1 per cent level, so any small systematic effects could potentially lead to large differences at that level. Related to this is the fact that both works use different fitting techniques to obtain the α distortions. We leave

for future work the exploration of a more sophisticated small-scale model.

Our work is similar in spirit also to Belsunce et al. (2025), in which the authors utilize a state-of-the-art hydrodynamical simulation to measure the EFT bias parameters and translate them into α parameters via the EFT of Ly α forest (Ivanov 2024) (a direct fit of the BAO parameters is not possible due to the limited volume of the hydrodynamical simulation). They find a -0.2 per cent (-0.3 per cent) deviation in α_{\parallel} (α_{\perp}). We find good agreement with Belsunce et al. (2025) for similar parameter settings, as can be seen from Table 7. Specifically, we note a great consistency between the EFT predictions and the actual measurements of the BAO shift w.r.t. linear theory prediction within the measurement errors. We stress that actual values of the BAO shift depend on the particular ‘microscopic’ Ly α model, and hence the differences between the exact values of the shifts w.r.t. Belsunce et al. (2025) based on hydrodynamical simulations are to be expected, since each Ly α realization has different EFT parameters. Crucially, the observed shift of the BAO can be accurately predicted for a given Ly α model, which is manifest in Table 7. A more direct connection with this line of work and how to self-consistently model the BAO shift is presented in Section 4. In the future, we plan to address the limitation of employing a single hydro simulation and the FGPA approach when making this comparison.

5.2 Summary of findings

In this work, we aim to constrain the size of the shift of the BAO peak due to non-linear structure growth as seen in the Ly α forest. In the era of DESI, where the error bars on the BAO parameters from the Ly α forest measurements will be at the sub per cent level, it is of utmost importance to place constraints on the error budget due to various systematic effects. One of these is the intrinsic shift to the α parameters, which if unaccounted for, could lead to a biased inference on the cosmology. In this work, we utilize the large Ly α forest mocks produced on the N -body simulation suite ABACUSSUMMIT to measure the BAO shift, which features four different models with different prescriptions. Specifically, we make measurements of the auto- and cross- (with quasars) correlation function, $\xi(r_p, \pi)$, for the six available simulations, each with two lines of sight, along the z and the y directions, for a total of 12 measurements assumed to be independent. We employ the software developed specifically for fitting the DESI Ly α forest correlation functions as well as the publicly available covariance matrix from the eBOSS DR16 analysis (du Mas des Bourboux et al. 2020). In that way, our work validates both the analysis pipeline of DESI Ly α forest data and also tests for possible shifts to the BAO parameters. Our main findings can be summarized as follows:

(i) To mitigate the noise on these measurements, we also adopt an LCV technique meant to reduce the noise on large scales for which linear theory yields an accurate prediction of the two-point correlation function. Applying the LCV method to our problem yields tighter measurements of the α parameters both in the auto- and in the cross-correlation case. From the autocorrelation, we detect a small 0.35 per cent shift at the 3σ level to α_{\parallel} and virtually no shift to α_{\perp} . From the cross-correlation, we see a similar shift to α_{\parallel} , albeit with larger error bars, and a small negative shift, ~ 0.25 per cent, at the 2σ level. Similarly to the raw measurements, we see the same structure of r_{min} dependence that is attributable to the accuracy of the model on small scales. For the explicit values of the shifts, we direct the reader to Tables 2 and 3.

(ii) We find that the one-loop Ly α forest EFT model accurately fits the ABACUSUMMIT data. When jointly fitting the BAO and EFT parameters, we find unbiased fits for α_{\parallel} and α_{\perp} that are consistent with one, i.e. no shift in the BAO parameter. Crucially, a full one-loop EFT treatment mitigates and *automatically includes* this shift into the final parameter constraints, rendering unnecessary to account for this shift in the systematic error budget.

Overall, in this work we confirm the assumption of a number of Ly α forest analyses (e.g. DESI Collaboration 2024b) that the BAO peak is shifted by less than 0.5 per cent. None the less, we recommend that a 0.35 per cent shift be accounted for in the systematic errors budgeted by the Ly α forest analysis. Doing more detailed tests with a wide range of Ly α forest mocks [both applied to N -body as well as hydro simulations (Belsunce et al. 2025)] and improved models that provide an accurate description of the clustering down to small scales (e.g. EFT) will bring great insight and prove invaluable for the analysis of large-scale structure surveys in the future.

The validation of the EFT fitting procedure on large simulations paired with a quantification that the EFT debiases the BAO fitting paves the way to a number of future analyses such as full-shape analyses of the Ly α forest in Fourier space using the 3D power spectrum (Font-Ribera, McDonald & Slosar 2018; Belsunce et al. 2024; Horowitz, de Belsunce & Lukic 2024) with field-level simulation-based priors (Ivanov et al. 2024b, d), beyond BAO analyses using the configuration space correlation function (Cuceu et al. 2021), or (near optimal compression techniques of) the 3D bispectrum (Belsunce & Sullivan 2025).

ACKNOWLEDGEMENTS

The authors thank Stephen Chen, Pat McDonald, Martin White, Mark Maus, and Nathalie Palanque-Delabrouille for fruitful discussions.

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

DATA AVAILABILITY

We make all our synthetic maps and catalogues publicly available on Globus through NERSC SHARE at this link: https://app.globus.org/file-manager?origin_id=9ce29982-eed1-11ed-9bb4-c9bb788c490e&path=%2F under the name ‘Abacus-Summit Lyman Alpha Forest’.

REFERENCES

Achitouv I., Blake C., 2015, *Phys. Rev. D*, 92, 083523
 Alcock C., Paczynski B., 1979, *Nature*, 281, 358
 Anderson L. et al., 2012, *MNRAS*, 427, 3435
 Baldauf T., Mirbabayi M., Simonović M., Zaldarriaga M., 2015, *Phys. Rev. D*, 92, 043514
 Baumann D., Nicolis A., Senatore L., Zaldarriaga M., 2012, *JCAP*, 2012, 051
 Bautista J. E. et al., 2017, *A&A*, 603, A12
 de Belsunce R., Sullivan J. M., 2025, The Lyman- α Forest Skew Spectrum, in press
 de Belsunce R., Philcox O. H. E., Irsic V., McDonald P., Guy J., Palanque-Delabrouille N., 2024, *MNRAS*, 533, 3756
 de Belsunce R., Chen S.-F., Ivanov M. M., Ravoux C., Chabanier S., Sexton J., Lukic Z., 2025, *Phys. Rev. D*, 111, 063524
 Bernardeau F., Colombi S., Gaztañaga E., Scoccimarro R., 2002, *Phys. Rep.*, 367, 1

Beutler F. et al., 2011, *MNRAS*, 416, 3017
 Bi H., Davidsen A. F., 1997, *ApJ*, 479, 523
 Blas D., Lesgourgues J., Tram T., 2011b, *J. Cosmology Astropart. Phys.*, 2011, 034
 Blas D., Lesgourgues J., Tram T., 2011a, *J. Cosmology Astropart. Phys.*, 2011, 034
 Blas D., Garny M., Konstandin T., 2013, *JCAP*, 2013, 024
 Blas D., Garny M., Ivanov M. M., Sibiryakov S., 2016a, *JCAP*, 2016, 028
 Blas D., Garny M., Ivanov M. M., Sibiryakov S., 2016b, *JCAP*, 2016, 052
 Blomqvist M. et al., 2019, *A&A*, 629, A86
 Bolton J. S., Puchwein E., Sijacki D., Haehnelt M. G., Kim T.-S., Meiksin A., Regan J. A., Viel M., 2017, *MNRAS*, 464, 897
 Bridle S. L., Crittenden R., Melchiorri A., Hobson M. P., Kneissl R., Lasenby A. N., 2002, *Mon. Not. Roy. Astron. Soc.*, 335, 1193
 Busca N. G. et al., 2013, *A&A*, 552, A96
 Carrasco J. J. M., Hertzberg M. P., Senatore L., 2012, *JHEP*, 2012, 082
 Chabanier S., Ravoux C., Latrille L., Sexton J., Armengaud É., Bautista J., Dumerchat T., Lukic Z., 2024, *MNRAS*, 534, 2674
 Chen S.-F., Vlah Z., White M., 2021, *JCAP*, 2021, 053
 Chen S.-F., Vlah Z., White M., 2022, *J. Cosmology Astropart. Phys.*, 2022, 008
 Chen S. F. et al., 2024, *MNRAS*, 534, 544
 Chudaykin A., Ivanov M. M., 2023, *Phys. Rev. D*, 107, 043518
 Chudaykin A., Ivanov M. M., 2025, *Phys. Rev. D*, 111, 083515
 Chudaykin A., Ivanov M. M., Philcox O. H. E., Simonović M., 2020, *Phys. Rev. D*, 102, 063533
 Chudaykin A., Ivanov M. M., Simonović M., 2021, *Phys. Rev. D*, 103, 043525
 Cole S. et al., 2005, *MNRAS*, 362, 505
 Crocce M., Scoccimarro R., 2008, *Phys. Rev.*, D77, 023533
 Croft R. A. C., Weinberg D. H., Katz N., Hernquist L., 1998, *ApJ*, 495, 44
 Cuceu A., Font-Ribera A., Joachimi B., Nadathur S., 2021, *MNRAS*, 506, 5439
 DESI Collaboration, 2016, preprint (arXiv:1611.00036)
 DESI Collaboration, 2024a, *JCAP*, 2025, 012
 DESI Collaboration, 2024b, *JCAP*, 2025, 124
 DeRose J., Chen S.-F., Kokron N., White M., 2023, *J. Cosmology Astropart. Phys.*, 2023, 008
 Delubac T. et al., 2015, *A&A*, 574, A59
 Desjacques V., Jeong D., Schmidt F., 2018a, *JCAP*, 2018, 035
 Desjacques V., Jeong D., Schmidt F., 2018b, *Phys. Rept.*, 733, 1
 Eisenstein D. J. et al., 2005, *ApJ*, 633, 560
 Eisenstein D. J., Seo H.-j., White M. J., 2007a, *ApJ*, 664, 660
 Eisenstein D. J., Seo H.-j., Sirko E., Spergel D. N., 2007b, *ApJ*, 664, 675
 Font-Ribera A. et al., 2013, *JCAP*, 2013, 018
 Font-Ribera A. et al., 2014, *J. Cosmology Astropart. Phys.*, 2014, 027
 Font-Ribera A., McDonald P., Slosar A., 2018, *J. Cosmology Astropart. Phys.*, 2018, 003
 Furlanetto S. R., Oh S. P., 2009, *ApJ*, 701, 94
 Garrison L. H., Eisenstein D. J., Pinto P. A., 2019, *MNRAS*, 485, 3370
 Garrison L. H., Eisenstein D. J., Ferrer D., Maksimova N. A., Pinto P. A., 2021, *MNRAS*, 508, 575
 Gelman A., Rubin D. B., 1992, *Stat. Sci.*, 7, 457
 Gerardi F., Cuceu A., Font-Ribera A., Joachimi B., Lemos P., 2022, *MNRAS*, 518, 2567
 Givans J. J., Hirata C. M., 2020, *Phys. Rev. D*, 102, 023515
 Givans J. J. et al., 2022, *JCAP*, 2022, 070
 Gunn J. E., Peterson B. A., 1965, *ApJ*, 142, 1633
 Hadzhiyska B. et al., 2023a, *Open J. Astrophys.*, 6, 38
 Hadzhiyska B. et al., 2023b, *MNRAS*, 524, 1008
 Hernández-Aguayo C., Cautun M., Smith A., Baugh C. M., Li B., 2020, *MNRAS*, 494, 3120
 Horowitz B., de Belsunce R., Lukic Z., 2024, preprint(arXiv:)
 Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
 Hui L., Gnedin N. Y., Zhang Y., 1997, *ApJ*, 486, 599
 Iršič V., McQuinn M., 2018, *J. Cosmology Astropart. Phys.*, 2018, 026
 Ivanov M. M., 2023, *Effective Field Theory for Large-Scale Structure*, preprint (arXiv:2212.08488)
 Ivanov M. M., 2024, *Phys. Rev. D*, 109, 023570

- Ivanov M. M., Sibiryakov S., 2018, *JCAP*, 2018, 053
 Ivanov M. M., Simonović M., Zaldarriaga M., 2020, *J. Cosmology Astropart. Phys.*, 2020, 042
 Ivanov M. M., Toomey M. W., Göksel Karaçaylı N., 2024a, *Phys. Rev. Lett.*, 134, 091001
 Ivanov M. M., Obuljen A., Cuesta-Lazaro C., Toomey M. W., 2024b, *Phys. Rev. D*, 111, 063548
 Ivanov M. M. et al., 2024c, preprint (arXiv:2412.01888)
 Ivanov M. M., Cuesta-Lazaro C., Mishra-Sharma S., Obuljen A., Toomey M. W., 2024d, *Phys. Rev. D*, 110, 063538
 Kaiser N., 1987, *MNRAS*, 227, 1
 Kirkby D. et al., 2013, *J. Cosmology Astropart. Phys.*, 2013, 024
 Kitaura F.-S. et al., 2016, *Phys. Rev. Lett.*, 116, 171301
 Knox L., 1995, *Phys. Rev. D*, 52, 4307
 Kokron N., Chen S.-F., White M., DeRose J., Maus M., 2022, *J. Cosmology Astropart. Phys.*, 2022, 059
 Lesgourgues J., 2011, preprint (arXiv:1104.2932)
 Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
 Maksimova N. A., Garrison L. H., Eisenstein D. J., Hadzhiyska B., Bose S., Satterthwaite T. P., 2021, *MNRAS*, 508, 4017
 Maus M. et al., 2025, *JCAP*, 2025, 138
 McDonald P., 2003, *ApJ*, 585, 34
 McDonald P., Roy A., 2009, *JCAP*, 2009, 020
 McDonald P., Miralda-Escude J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., Ostriker J. P., 2000, *ApJ*, 543, 1
 McQuinn M., 2016, *ARA&A*, 54, 313
 McQuinn M., White M., 2016, *JCAP*, 2016, 043
 Neyrinck M. C., Szapudi I., McCullagh N., Szalay A. S., Falck B., Wang J., 2018, *MNRAS*, 478, 2495
 Padmanabhan N., White M. J., 2008, *Phys. Rev. D*, 77, 123540
 Padmanabhan N., White M., 2009, *Phys. Rev. D*, 80, 063508
 Peirani S., Weinberg D. H., Colombi S., Blaizot J., Dubois Y., Pichon C., 2014, *ApJ*, 784, 11
 Peirani S. et al., 2022, *MNRAS*, 514, 3222
 Planck Collaboration, 2020, *A&A*, 641, A6
 Arinyo-i Prats A., Miralda-Escudé J., Viel M., Cen R., 2015, *JCAP*, 2015, 017
 Qezlou M., Newman A. B., Rudie G. C., Bird S., 2022, *ApJ*, 930, 109
 de Sainte Agathe V. et al., 2019, *A&A*, 629, A85
 Schmittfull M., Feng Y., Beutler F., Sherwin B., Chu M. Y., 2015, *Phys. Rev. D*, 92, 123522
 Senatore L., Zaldarriaga M., 2015, *JCAP*, 2015, 013
 Sherwin B. D., Zaldarriaga M., 2012, *Phys. Rev. D*, 85, 103523
 Sinigaglia F., Kitaura F.-S., Balaguera-Antolínez A., Shimizu I., Nagamine K., Sánchez-Benavente M., Ata M., 2022, *ApJ*, 927, 230

- Sinigaglia F., Kitaura F. S., Nagamine K., Oku Y., Balaguera-Antolínez A., 2024a, *A&A*, 682, A21
 Sinigaglia F., Kitaura F.-S., Nagamine K., Oku Y., 2024b, *ApJ*, 971, L22
 Slosar A. et al., 2011, *J. Cosmology Astropart. Phys.*, 2011, 001
 Slosar A. et al., 2013, *J. Cosmology Astropart. Phys.*, 2013, 026
 Sorini D., Oñorbe J., Lukić Z., Hennawi J. F., 2016, *ApJ*, 827, 97
 Taylor A. N., Kitching T. D., 2010, *MNRAS*, 408, 865
 Torrado J., Lewis A., 2021, *JCAP*, 2021, 057
 Weinberg D. H., Mortonson M. J., Eisenstein D. J., Hirata C., Riess A. G., Rozo E., 2013, *Phys. Rep.*, 530, 87
 Yuan S., Garrison L. H., Hadzhiyska B., Bose S., Eisenstein D. J., 2022, *MNRAS*, 510, 3301
 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: CONTROL VARIATES VALIDATION

In this section, we validate the linear control variate (LCV) technique by studying three important metrics in the autocorrelation function: the ratio between theory and simulation, which needs to be unbiased for the technique to work; the coefficient $\beta_\ell(k)$, which gives us the scales over which LCV error mitigation is applied (for more details, see Section 2.7); the error reduction as a result of adopting LCV.

The left panel of Fig. A1 shows the ratio of the analytical model to the control variable (i.e. the linear theory approximation to the autopower spectrum of the Ly α forest). The fact that we see perfect sub-per cent level agreement down to very small scales ($k \sim 1$ h/Mpc) assures us that no bias to the two-point correlation function can be introduced through the LCV method.

The middle panel displays the β parameter (see equation 11) quantifying the amount of correlation between the control variable and the variable of interest (the autopower spectrum of the Ly α forest at $z = 2.5$). We see that the correlation is strongest on large scales up to $k \sim 0.1$ h/Mpc. Beyond that the linear theory approximation breaks down and we can no longer rely on this technique to reduce the variance for the measured signal. Additionally, we see that the theoretical model does a slightly better job of reproducing the monopole, which is expected as redshift-space effects due to the non-linear clustering can propagate to larger scales (as seen in the $\beta_{\ell=2}(k)$ curve). We emphasize that a low value of β_ℓ at some scale k simply implies that the LCV algorithm is not reducing the noise of




Figure A1. Validation of the linear control variate technique for the Ly α forest autopower spectrum. *Leftmost panel:* Ratio of the analytical model to the control variable (i.e. the linear theory approximation to the autopower spectrum of the Ly- α forest). Sub-per cent level agreement is seen down to small scales. *Middle panel:* The β parameter (see equation 11) quantifying the amount of correlation between the control variable and the variable of interest (the autopower spectrum of the Ly α forest at $z = 2.5$). The correlation is strongest on large scales up to $k \sim 0.1$ h/Mpc. *Rightmost panel:* Factor by which the error bars on the monopole and quadrupole of the Ly α forest autopower spectrum decrease as a function of scale. The largest reduction of error is on large scales, where the improvement factor is about 7 or 8. The error on the power spectrum is virtually unchanged beyond $k \gtrsim 2$ h/Mpc.

the power spectrum at those scales. However, that does not introduce biases to the LCV-reduced power spectrum.

Finally, in the rightmost panel, we show the factor by which the error bars on the monopole and quadrupole of the Ly α forest autopower spectrum decrease as a function of scale. Looking at the monopole, we see that the largest reduction of error is on large scales, where the improvement factor is about 7 or 8. The error on the power spectrum is virtually unchanged beyond $k \gtrsim 2$ h/Mpc. The quadrupole curve, on the other hand, is very noisy, but it follows the same trend. This is expected as linear theory describes the clustering most accurately on large scales.

APPENDIX B: EFFECT OF ADOPTING BROAD-BAND PARAMETERS

As specified in Section 2.5, for our default analysis we perform a fit to the clustering data with a model featuring broad-band polynomial parameters, aimed to make our model immune to limitations on small scales. In particular, the model employed in this study is the so-called Arinyo model, which is motivated by measurements of the clustering performed in a hydro simulation. In the main body of the paper, we fix the Arinyo parameters to their fiducial values and include a comprehensive set of broad-band parameters, which we vary. In this appendix, we are interested in exploring the case where we do not include broad-band polynomial templates and instead simply perform our fits with the default Arinyo parameters.

In Fig. B1, we see that the values of α_{\parallel} and α_{\perp} are broadly consistent, both in terms of the mean values as well as their error

bars, with the values in Fig. 3, which is reassuring. None the less, the fits do appear to be more noisy and also show a stronger dependence on r_{\min} than their counterparts using the broad-band parametrization. This shows that the broad-band model manages to absorb well many of the issues on small scales. Upon visual inspection of the fitted correlation functions, we furthermore find that when using only the Arinyo parameters without the broad-band templates, the fits are noticeably poorer. This is another reason we opt to keep the Arinyo parameters fixed and vary the broad-band parameters. We also tried varying the Arinyo model parameters, dropping the broad-band parameters, and found very similar results to our default case. However, we stress that this model is not available for modelling the cross-correlation function, and for this reason, we opt not to use it. In the future, we plan to study extensions and alternatives to the Arinyo model in more detail, exploring comparisons with more physically motivated frameworks such as EFT (Ivanov 2024).

APPENDIX C: TRIANGLE PLOTS FOR EFT FITS TO AUTO- AND CROSS-CORRELATION

In this appendix, we show the full marginalized posteriors for the sampled BAO and EFT bias parameters for model 1, simulation one and line-of-sight y . In Fig. C1, we show the results for the autocorrelation. The 1D marginalized posteriors are highly Gaussian and we only see clear degeneracies between $b_{G_2} - b_2$ and $b_{\Pi^{(2)}} - b_{\delta\eta}$. The picture is consistent when jointly fitting the auto- and cross-correlation, shown in Fig. C2.




Figure B1. Similar to Fig. 3. The difference is that here, we have switched off the broad-band parameters, and are instead adopting the ‘Arinyo’ model with parameters of the model fixed to their fiducial values. We see that the values of the α parameters are largely consistent with the fiducial analysis, which is encouraging. At the same time, the fits shown here exhibit a stronger dependence on r_{\min} and are more noisy.

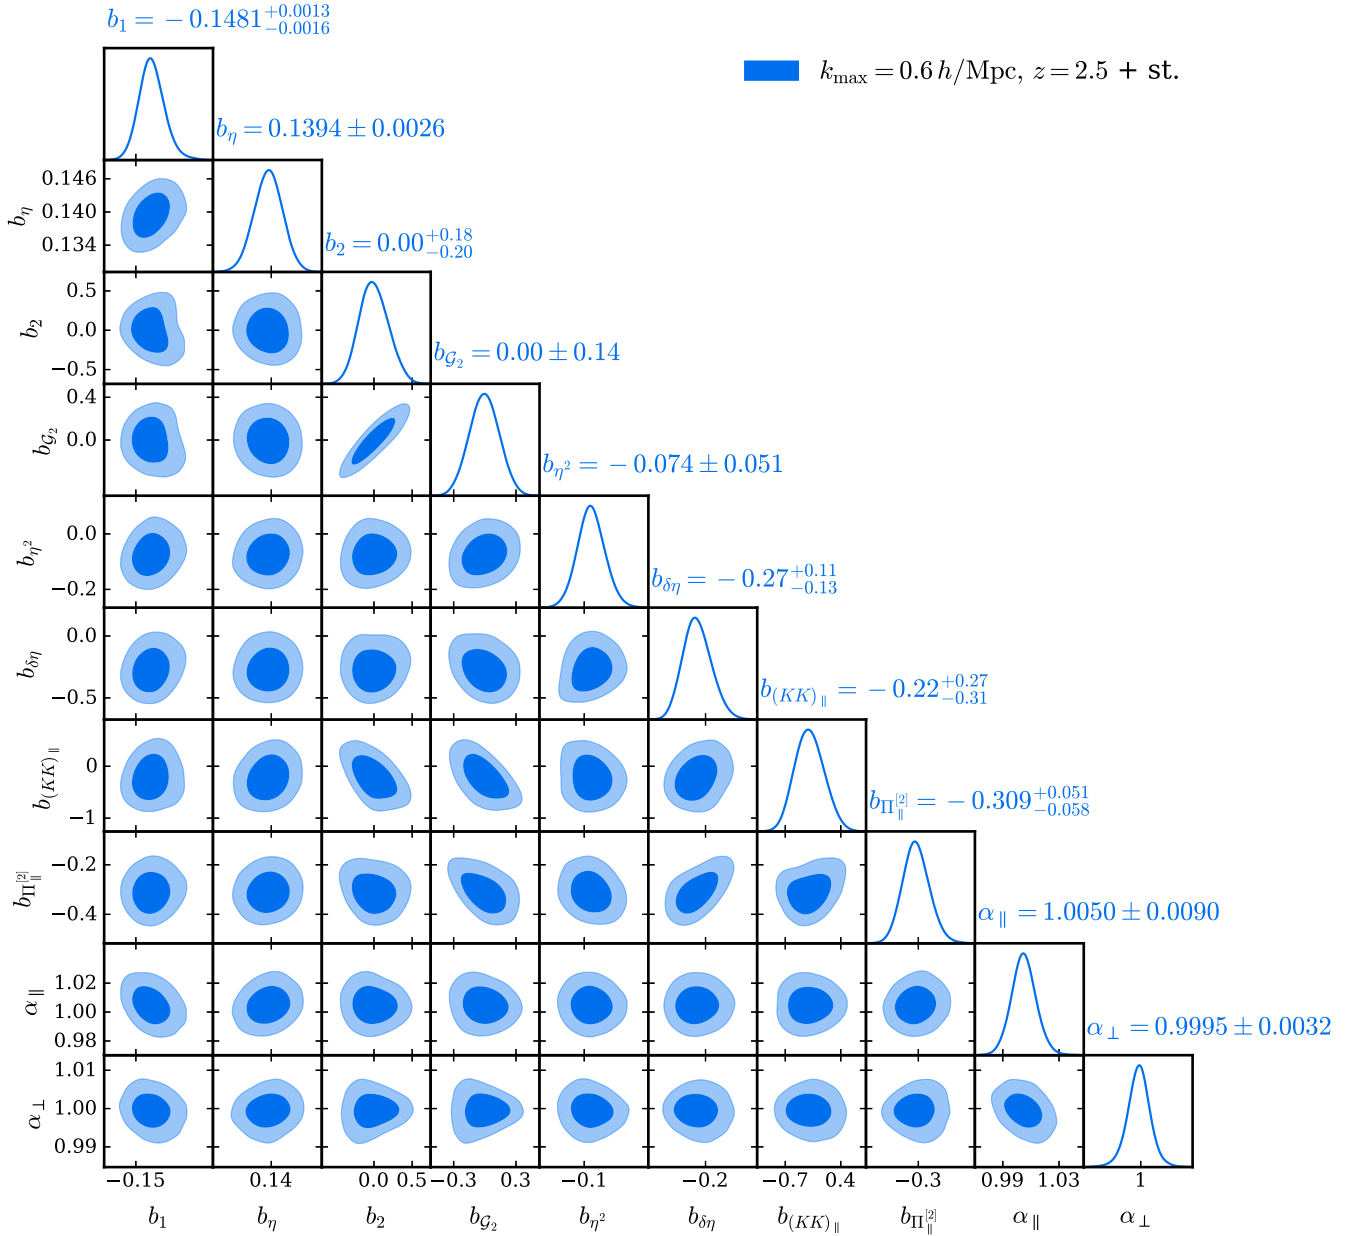


Figure C1. Triangle plot of the EFT bias and BAO scaling parameters obtained from fitting simulation ‘one’ and model 1. The corresponding numerical values are given in Table 4. We show the best-fitting 1D and 2D marginalized posteriors for the *sampled* parameters.




Table C1. EFT bias parameters for each model obtained from fits to the stacked data vector of 12 simulations for 1175 data points. Table 4 is the corresponding table for fits to the individual simulations. Note that the average error bars are divided by $\sqrt{6}$. Given the small error bars, we can differentiate between the linear bias parameters in addition to the b_2 parameter.

$b_{\mathcal{O}}$	Model 1	Model 2	Model 3	Model 4
b_1	-0.1485 ± 0.0006	$-0.1314^{+0.0007}_{-0.0006}$	-0.1337 ± 0.0005	-0.1295 ± 0.0005
b_η	0.1423 ± 0.0012	0.1295 ± 0.0011	0.2705 ± 0.0013	0.3041 ± 0.0014
b_2	-0.1231 ± 0.0942	$-0.3316^{+0.0897}_{-0.0770}$	-0.0353 ± 0.0505	-0.0400 ± 0.0483
$b_{\mathcal{G}_2}$	-0.0847 ± 0.0701	$-0.2765^{+0.0651}_{-0.0482}$	-0.0383 ± 0.0361	-0.0320 ± 0.0351
b_{η^2}	-0.1088 ± 0.0300	$-0.2631^{+0.0244}_{-0.0223}$	-0.0811 ± 0.0271	-0.0385 ± 0.0275
$b_{\delta\eta}$	$-0.3771^{+0.0557}_{-0.0568}$	$-0.3511^{+0.0610}_{-0.0743}$	-0.0633 ± 0.0429	-0.0167 ± 0.0422
b_{KK_\parallel}	$-0.0751^{+0.1499}_{-0.1329}$	$0.4841^{+0.1285}_{-0.1773}$	-0.0164 ± 0.0868	$0.0327^{+0.0883}_{-0.0864}$
$b_{\Pi_\parallel^{[2]}}$	-0.3046 ± 0.0339	-0.2943 ± 0.0473	-0.2429 ± 0.0353	-0.2578 ± 0.0379
$b_{\Pi_\parallel^{[3]}}$	0.1155 ± 0.0216	0.4038 ± 0.0192	0.7470 ± 0.0200	0.7984 ± 0.0200
$b_{\delta\Pi_\parallel^{[2]}}$	-1.1613 ± 0.0614	1.5483 ± 0.0553	-1.5111 ± 0.0800	-2.0542 ± 0.0846
$b_{(K\Pi^{[2]})_\parallel}$	-0.7105 ± 0.0694	-1.2533 ± 0.0622	-1.1988 ± 0.0675	-1.7604 ± 0.0676
$b_{\eta\Pi_\parallel^{[2]}}$	-0.8287 ± 0.1952	6.0565 ± 0.1771	-5.5202 ± 0.2661	-6.8086 ± 0.2816
b_{Γ_3}	0.1699 ± 0.0365	1.3522 ± 0.0329	-1.0584 ± 0.0429	-1.1835 ± 0.0445
P_{shot}	-0.6533 ± 0.0439	-0.1932 ± 0.0337	0.2082 ± 0.0181	0.0885 ± 0.0160
a_0	3.3854 ± 0.3902	2.2565 ± 0.3119	2.2597 ± 0.3040	2.5358 ± 0.2903
a_2	-3.4240 ± 1.2475	-4.9955 ± 1.0051	-11.3607 ± 1.4209	-12.4223 ± 1.4650
c_0	0.1763 ± 0.0150	0.0155 ± 0.0129	-0.1859 ± 0.0073	-0.1412 ± 0.0069
c_2	-0.0342 ± 0.0277	0.1095 ± 0.0248	0.5097 ± 0.0275	0.5065 ± 0.0277
c_4	-0.1394 ± 0.0123	-0.1347 ± 0.0110	-0.4875 ± 0.0156	-0.5160 ± 0.0164
χ^2	726.4	933.1	1032.7	1043.3

Table C2. Same as Table C1 for the joint fits of the auto- and cross-correlation with 1,755 data points.

$b_{\mathcal{O}}$	Model 1	Model 2	Model 3	Model 4
b_1^q	3.3160 ± 0.0435	3.3146 ± 0.0441	3.3730 ± 0.0368	3.3744 ± 0.0359
b_2^q	1.5437 ± 0.9904	1.1257 ± 0.9936	1.8083 ± 1.0180	$1.9024^{+0.9954}_{-1.0921}$
$b_{\mathcal{G}_2}^q$	0.0605 ± 0.6385	$-0.6220^{+0.6124}_{-0.7404}$	0.1710 ± 0.4942	-0.0383 ± 0.4716
b_1	$-0.1465^{+0.0012}_{-0.0011}$	-0.1294 ± 0.0011	-0.1316 ± 0.0009	-0.1284 ± 0.0009
b_η	0.1448 ± 0.0018	0.1315 ± 0.0017	0.2745 ± 0.0022	0.3102 ± 0.0024
b_2	$-0.3613^{+0.1664}_{-0.1298}$	$-0.4283^{+0.1288}_{-0.0855}$	$-0.1961^{+0.1188}_{-0.1048}$	-0.0885 ± 0.1060
$b_{\mathcal{G}_2}$	$-0.2598^{+0.1140}_{-0.0897}$	$-0.2186^{+0.1304}_{-0.0833}$	-0.0862 ± 0.0705	-0.0455 ± 0.0671
b_{η^2}	-0.2032 ± 0.0525	$-0.3448^{+0.0531}_{-0.0452}$	-0.1961 ± 0.0483	-0.1837 ± 0.0498
$b_{\delta\eta}$	$-0.4041^{+0.1402}_{-0.1163}$	$-0.6267^{+0.1596}_{-0.0856}$	$-0.2108^{+0.0811}_{-0.0710}$	$-0.1393^{+0.0763}_{-0.0681}$
$b_{K K_\parallel}$	$0.4594^{+0.1919}_{-0.2422}$	$0.3074^{+0.2361}_{-0.2866}$	$0.4538^{+0.1286}_{-0.1548}$	$0.3918^{+0.1279}_{-0.1427}$
$b_{\Pi_\parallel^{[2]}}$	$-0.2123^{+0.0577}_{-0.0486}$	$-0.2376^{+0.0532}_{-0.0429}$	-0.2407 ± 0.0531	-0.2078 ± 0.0571
$b_{\Pi_\parallel^{[3]}}$	-0.2223 ± 0.0313	-0.2287 ± 0.0295	0.2143 ± 0.0341	0.2681 ± 0.0353
$b_{\delta\Pi_\parallel^{[2]}}$	-1.7482 ± 0.1283	-1.3608 ± 0.1184	-1.4153 ± 0.1267	-1.6799 ± 0.1298
$b_{(K\Pi^{[2]})_\parallel}$	-1.7220 ± 0.1328	-1.5296 ± 0.1244	-1.8891 ± 0.1107	-2.2383 ± 0.1098
$b_{\eta\Pi_\parallel^{[2]}}$	-1.2416 ± 0.3383	-0.2563 ± 0.3164	-3.2567 ± 0.3889	-3.5330 ± 0.4075
b_{Γ_3}	0.7481 ± 0.0628	0.9753 ± 0.0589	-0.2344 ± 0.0680	-0.2553 ± 0.0706
P_{shot}	-0.3132 ± 0.0732	-0.0226 ± 0.0569	0.4352 ± 0.0214	0.3174 ± 0.0177
a_0	-0.8737 ± 3.1579	-1.4900 ± 2.7583	-4.6427 ± 2.9631	-3.6087 ± 2.8547
a_2	0.4249 ± 4.7041	-0.8907 ± 4.6365	-0.9101 ± 4.7460	-1.2032 ± 4.8208
c_0	0.1571 ± 0.0314	0.0252 ± 0.0275	-0.1853 ± 0.0142	-0.1364 ± 0.0135
c_2	-0.1610 ± 0.0362	-0.0280 ± 0.0341	0.2055 ± 0.0310	0.1768 ± 0.0316
c_4	-0.1091 ± 0.0222	-0.1539 ± 0.0205	-0.3657 ± 0.0232	-0.3803 ± 0.0243
c_x	-0.9703 ± 2.8711	1.6564 ± 2.6310	-6.0331 ± 2.2181	-6.3074 ± 2.0911
c_0^q	0.0726 ± 0.0102	0.0961 ± 0.0092	-0.1810 ± 0.0099	-0.2006 ± 0.0097
c_2^q	0.2795 ± 0.9659	0.1858 ± 0.9687	0.5637 ± 0.9728	0.5738 ± 0.9829
χ^2	1790.6621	1842.2431	1892.3472	1898.7141

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.