# ForestFlow: predicting the Lyman- $\alpha$ forest clustering from linear to nonlinear scales

J. Chaves-Montero<sup>1,\*</sup>, L. Cabayol-Garcia<sup>1,2</sup>, M. Lokken<sup>1</sup>, A. Font-Ribera<sup>1</sup>, J. Aguilar<sup>3</sup>, S. Ahlen<sup>4</sup>,

D. Bianchi<sup>5</sup>, D. Brooks<sup>6</sup>, T. Claybaugh<sup>3</sup>, S. Cole<sup>7</sup>, A. de la Macorra<sup>8</sup>, S. Ferraro<sup>3,9</sup>, J. E. Forero-Romero<sup>10,11</sup>, E. Gaztañaga<sup>12,13</sup>, S. Gontcho A Gontcho<sup>3</sup>, G. Gutierrez<sup>14</sup>, K. Honscheid<sup>15,16,17</sup>, R. Kehoe<sup>18</sup>, D. Kirkby<sup>19</sup>, A. Kremin<sup>3</sup>, A. Lambert<sup>3</sup>, M. Landriau<sup>3</sup>, M. Manera<sup>20,1</sup>, P. Martini<sup>15,21,17</sup>, R. Miquel<sup>22,1</sup>,
A. Muñoz-Gutiérrez<sup>8</sup>, G. Niz<sup>23,24</sup>, I. Pérez-Ràfols<sup>25</sup>, G. Rossi<sup>26</sup>, E. Sanchez<sup>27</sup>, M. Schubnell<sup>28</sup>, D. Sprayberry<sup>29</sup>,

G. Tarlé<sup>28</sup>, and B. A. Weaver<sup>29</sup>

(Affiliations can be found after the references)

Received 28 August 2024 / Accepted 7 January 2025

### ABSTRACT

On large scales, the Lyman- $\alpha$  forest provides insights into the expansion history of the Universe, while on small scales, it imposes strict constraints on the growth history, the nature of dark matter, and the sum of neutrino masses. This work introduces ForestFlow, a novel framework that bridges the gap between large- and small-scale analyses, which have traditionally relied on distinct modeling approaches. Using conditional normalizing flows, ForestFlow predicts the two Lyman- $\alpha$  linear biases ( $b_{\delta}$  and  $b_{\eta}$ ) and six parameters describing small-scale deviations of the three-dimensional flux power spectrum  $(P_{3D})$  from linear theory as a function of cosmology and intergalactic medium physics. These are then combined with a Boltzmann solver to make consistent predictions, from arbitrarily large scales down to the nonlinear regime, for  $P_{3D}$  and any other statistics derived from it. Trained on a suite of 30 fixed-and-paired cosmological hydrodynamical simulations spanning redshifts from z = 2to 4.5, ForestFlow achieves 3 and 1.5% precision in describing  $P_{3D}$  and the one-dimensional flux power spectrum ( $P_{1D}$ ) from linear scales to  $k = 5 \,\mathrm{Mpc}^{-1}$  and  $k_{\parallel} = 4 \,\mathrm{Mpc}^{-1}$ , respectively. Thanks to its conditional parameterization, ForestFlow shows similar performance for ionization histories and two ACDM model extensions - massive neutrinos and curvature - even though none of these are included in the training set. This framework will enable full-scale cosmological analyses of Lyman- $\alpha$  forest measurements from the DESI survey.

**Key words.** cosmological parameters – cosmology: theory – large-scale structure of Universe

# 1. Introduction

The Lyman- $\alpha$  forest refers to absorption lines in the spectra of high-redshift quasars resulting from Lyman- $\alpha$  absorption by neutral hydrogen in the intergalactic medium (IGM; for a review, see McQuinn 2016). Even though quasars can be observed at very high redshifts with relatively short exposure times, the scarcity of these sources limits their direct use for precision cosmology. Conversely, Lyman- $\alpha$  forest measurements from a single quasar spectrum provide information about density fluctuations over hundreds of megaparsecs along the line of sight, making this observable an excellent tracer of large-scale structure at high redshifts.

Cosmological analyses of the Lyman- $\alpha$  forest rely on either three-dimensional correlations of the Lyman- $\alpha$  transmission field ( $\xi_{3D}$ ; e.g., Slosar et al. 2011) or correlations along the line-of-sight of each quasar (i.e., the one-dimensional flux power spectrum  $P_{1D}$ , Croft et al. 1998; McDonald et al. 2000). The first analyses set constraints on the expansion history of the Universe by measuring baryonic acoustic oscillations (BAO; e.g., Busca et al. 2013; Slosar et al. 2013; du Mas des Bourboux et al. 2020), for which linear or perturbation theory is accurate enough. On the other hand,  $P_{1D}$  analyses measure the small-scale amplitude and slope of the linear power spectrum (e.g., Croft et al. 1998; McDonald et al. 2000, 2005; Zaldarriaga et al. 2001; Viel et al. 2004), the nature of

dark matter (e.g., Seljak et al. 2006a; Viel et al. 2013; Iršič et al. 2017; Palanque-Delabrouille et al. 2020; Rogers & Peiris 2021a; Iršič et al. 2024), the thermal history of the IGM (e.g., Viel & Haehnelt 2006; Bolton et al. 2008; Lee et al. 2015; Walther et al. 2019; Boera et al. 2019; Gaikwad et al. 2020, 2021) and the reionization history of the Universe (see the reviews Meiksin 2009; McQuinn 2016). In combination with cosmic microwave background constraints, P<sub>1D</sub> analyses also set tight constraints on the sum of neutrino masses and the running of the spectral index (e.g., Spergel et al. 2003; Verde et al. 2003; Viel et al. 2004; Seljak et al. 2005, 2006b; Palanque-Delabrouille et al. 2015, 2020).

Unlike  $\xi_{3D}$  studies,  $P_{1D}$  analyses go deep into the nonlinear regime and require time-demanding hydrodynamical simulations (e.g., Cen et al. 1994; Miralda-Escudé et al. 1996; Meiksin et al. 2001; Lukić et al. 2015; Bolton et al. 2017; Walther et al. 2021; Chabanier et al. 2023; Puchwein et al. 2023; Bird et al. 2023). Naive analyses would demand running millions of hydrodynamical simulations, which is currently unfeasible. Rather, the preferred solution is constructing fast surrogate models that make precise predictions across the input parameter space using simulation measurements as the training set. The main advantage of these surrogate models, known as emulators, is reducing the number of simulations required for Bayesian inference from millions to dozens or hundreds. In the context of Lyman- $\alpha$  forest studies, the first  $P_{1D}$  emulators involved simple linear interpolation (McDonald et al. 2006)

<sup>\*</sup> Corresponding author; jchaves@ifae.es

Open Access article, published by EDP Sciences, under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. This article is published in open access under the Subscribe to Open model. Subscribe to A&A to support open access publication

and progressively moved toward using Gaussian processes (GPs; Sacks et al. 1989; MacKay et al. 1998) and neural networks (NNs; McCulloch & Pitts 1943); for instance, Bird et al. (2019, 2023), Rogers et al. (2019), Walther et al. (2019), Pedersen et al. (2021), Takhtaganov et al. (2021), Rogers & Peiris (2021b), Fernandez et al. (2022), Molaro et al. (2023), and Cabayol-Garcia et al. (2023).

The primary purpose of this work is to provide consistent predictions for Lyman- $\alpha$  forest clustering statistics from linear to nonlinear scales. There are three main approaches to achieve this. The first relies on perturbation theory (e.g., Givans & Hirata 2020; Chen et al. 2021; Ivanov 2024), which delivers precise predictions on perturbative scales at the cost of marginalizing over a large number of free parameters. The second involves emulating power spectrum modes measured from a suite of cosmological hydrodynamical simulations, which provides precise predictions from quasilinear to nonlinear scales. The main limitation of this approach is that accessing the largest scales used in BAO analyses,  $r \simeq 300 \,\mathrm{Mpc}$ , would require hydrodynamical simulations at least three times larger than this scale (e.g., Angulo et al. 2008), which is currently unfeasible due to the computational demands of these simulations. Nonetheless, there are recent strides in this direction using approximated methods (e.g., Jacobus et al. 2023).

Instead, we adopt a third approach: emulating the best-fitting parameters of a physically motivated Lyman- $\alpha$  clustering model to measurements from a suite of cosmological hydrodynamical simulations as a function of cosmology and IGM physics. In what follows, we refer to this strategy as FORESTFLOW<sup>1</sup>. In this work, we use conditional normalizing flows (cNF; Winkler et al. 2019; Papamakarios et al. 2019) to emulate the two Lyman- $\alpha$ linear biases ( $b_{\delta}$  and  $b_{\eta}$ ), which completely determine the largescale behavior of  $P_{3D}$  in conjunction with the linear power spectrum, along with six parameters capturing small-scale deviations of  $P_{3D}$  from linear theory as a function of six parameters capturing the cosmological and IGM dependence of Lyman- $\alpha$  clustering (Pedersen et al. 2021). We show that this strategy enables precise  $P_{3D}$  predictions from nonlinear scales to arbitrarily large (linear) scales even when training on a suite of simulations with moderate size (Pedersen et al. 2021; Cabayol-Garcia et al. 2023). FORESTFLOW also allows for the prediction of any statistic derived from  $P_{3D}$  without requiring interpolation or extrapolation. For instance, we can compute  $\xi_{3D}$  by taking the Fourier transform of  $P_{3D}$  or derive  $P_{1D}$  by integrating its perpendicular modes

$$P_{1\mathrm{D}}(k_{\parallel}) = (2\pi)^{-1} \int_0^\infty \mathrm{d}k_{\perp} \, k_{\perp} \, P_{3\mathrm{D}}(k_{\parallel}, \, k_{\perp}), \tag{1}$$

where  $k_{\parallel}$  and  $k_{\perp}$  indicate parallel and perpendicular modes, respectively.

The release of FORESTFLOW is quite timely for BAO and  $P_{1D}$ analyses of the ongoing Dark Energy Spectroscopic Instrument survey (DESI; DESI Collaboration 2016), which will quadruple the number of line-of-sights observed by first the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013) and its extension (eBOSS; Dawson et al. 2016). DESI has already proven the constraining power of Lyman- $\alpha$  studies by measuring the isotropic BAO scale with  $\approx 1\%$  precision from the Data Release 1 (DESI Collaboration 2024) and  $P_{1D}$  at nine redshift bins with a precision of a few percent from the Early Data Release (Ravoux et al. 2023; Karaçayh et al. 2024). In addition to being used for BAO and  $P_{1D}$  studies, FORESTFLOW has the potential to extend toward nonlinear scales the current full-shape analyses of  $\xi_{3D}$  (Cuceu et al. 2023; Gerardi et al. 2023) and  $P_{3D}$  (Font-Ribera et al. 2018; de Belsunce et al. 2024; Horowitz et al. 2025), and can be used to interpret alternative three-dimensional statistics (Hui et al. 1999; Font-Ribera et al. 2018; Abdul Karim et al. 2024).

The structure of this paper is as follows: we describe the strategy adopted by FORESTFLOW, the input data for the cNF, and its architecture in Sects. 2, 3, and 4, respectively. Next, we assess the performance of FORESTFLOW in Sect. 5 and highlight some novel analyses enabled by this framework in Sect. 6. Finally, we summarize the main findings and conclude in Sect. 7. Throughout this paper, all statistics and distances are in comoving units.

# 2. Strategy adopted by FORESTFLOW

In FORESTFLOW, we emulate the parameters of a physically motivated parametric model for Lyman- $\alpha$  forest clustering as a function of parameters that capture the influence of cosmology and IGM physics on this observable. This section begins with an overview of the physically motivated model, followed by an introduction to the parameters used to characterize the dependence of Lyman- $\alpha$  forest clustering on cosmology and IGM physics.

## 2.1. Parametric model for P<sub>3D</sub>

We can express fluctuations in the Lyman- $\alpha$  forest flux as  $\delta_F(\mathbf{s}) = \overline{F}^{-1}(\mathbf{s})F(\mathbf{s}) - 1$ , where  $F = \exp(-\tau)$  and  $\overline{F}$  are the transmitted flux fraction and its mean, respectively,  $\tau$  is the optical depth to Lyman- $\alpha$  absorption, and  $\mathbf{s}$  is the redshift-space coordinate. On linear scales, these fluctuations depend upon the matter field as follows (e.g., McDonald 2003)

$$\delta_F = b_\delta \,\delta + b_\eta \,\eta,\tag{2}$$

where  $\delta$  refers to matter density fluctuations,  $\eta = -(aH)^{-1}(\partial v_r/\partial r)$  stands for the dimensionless line-of-sight gradient of radial peculiar velocities, *a* is the cosmological expansion factor, *H* is the Hubble expansion factor,  $v_r$  is the radial velocity, and *r* stands for the radial comoving coordinate. The linear bias coefficients  $b_{\delta}$  and  $b_{\eta}$  capture the response of  $\delta_F$  to large-scale fluctuations in the  $\delta$  and  $\eta$  fields, respectively.

Following McDonald (2003), we decompose the threedimensional power spectrum of  $\delta_F$  into three terms

$$P_{\rm 3D}(k,\,\mu) = (b_{\delta} + b_{\eta}\,f\,\mu^2)^2 D_{\rm NL}(k,\,\mu) P_{\rm lin}(k),\tag{3}$$

where  $f = d \log G/d \log a$  is the logarithmic derivative of the growth factor G,  $\mu = k^{-1}k_{\parallel}$  is the cosine of the angle between the Fourier mode and the line of sight,  $(b_{\delta} + b_{\eta} f \mu^2)^2$  accounts for both linear biasing and large-scale redshift space distortions (Kaiser 1987; McDonald et al. 2000),  $P_{\rm lin}$  is the linear matter power spectrum<sup>2</sup>, and  $D_{\rm NL}$  is a physically motivated parametric correction accounting for the nonlinear growth of the density field, nonlinear peculiar velocities, thermal broadening, and pressure.

The large-scale behavior of  $P_{3D}$  is set by the bias coefficients  $b_{\delta}$  and  $b_{\eta}$  together with the linear power spectrum, and the latter can be computed using a Boltzmann solver (e.g., Lewis et al. 2000; Lesgourgues 2011). Therefore, the emulation of the two

<sup>&</sup>lt;sup>1</sup> Publicly available at https://github.com/igmhub/ForestFlow

 $<sup>^2</sup>$  This is the linear power spectrum of cold dark matter and baryons even for cosmologies with massive neutrinos.

Lyman- $\alpha$  linear biases enables predicting  $P_{3D}$  on arbitrarily large (linear) scales<sup>3</sup>. In contrast with direct emulation of power spectrum modes, this approach only requires simulations large enough for measuring the two Lyman- $\alpha$  linear biases precisely.

Predicting  $P_{3D}$  on small scales is more challenging than on large scales due to the variety of effects influencing this statistic in the nonlinear regime (e.g., McDonald 2003). In this work, we capture these small-scale effects using the physically motivated Arinyo-i-Prats et al. (2015) parameterization

$$D_{\rm NL} = \exp\left\{\left(q_1\Delta^2 + q_2\Delta^4\right)\left[1 - \left(\frac{k}{k_{\rm v}}\right)^{a_{\rm v}}\mu^{b_{\rm v}}\right] - \left(\frac{k}{k_{\rm p}}\right)^2\right\},\tag{4}$$

where  $\Delta^2(k) \equiv (2\pi^2)^{-1}k^3 P_{\text{lin}}(k)$  is the dimensionless linear matter power spectrum and the free parameters  $k_v$  and  $k_p$  are in Mpc<sup>-1</sup> units throughout this work. The terms involving  $\{q_1, q_2\}$ ,  $\{k_v, a_v, b_v\}$ , and  $\{k_p\}$  account for nonlinear growth, peculiar velocities and thermal broadening, and gas pressure, respectively. The expression above does not account for a shot noise term (e.g., Iršič & McQuinn 2018). While Givans et al. (2022) successfully described  $P_{3D}$  and  $P_{1D}$  measurements down to highly nonlinear scales using this formulation, it is possible that the shot noise contribution was implicitly absorbed into their fit through free parameters representing other effects. A more detailed investigation of shot noise is deferred to future work.

In the top panel of Fig. 1, dotted lines show the ratio of measurements from the CENTRAL simulation at z = 3 and the linear power spectrum, while solid lines do so for the best-fitting model to these measurements (Eqs. (3) and (4)) and the linear power spectrum. See Sect. 3 for details about this simulation and the fitting procedure. The dashed lines depict the results for the best-fitting model when setting  $D_{\rm NL} = 1$  after carrying out the fit; in other words, the behavior of the best-fitting model on linear scales. We can readily see that nonlinear growth isotropically increases the power with growing k, while peculiar velocities and thermal broadening suppress the power of parallel modes as k increases. On even smaller scales, pressure takes over and causes an isotropic suppression. Nonlinear growth modifies the perpendicular power relative to linear theory by 10% for scales as large as  $k = 0.5 \text{ Mpc}^{-1}$ , indicating that small-scale corrections are important for most of the scales sampled by our simulations. Nevertheless, in Appendix A, we show that we can measure the two Lyman- $\alpha$  linear biases with percent precision from these simulations. Deviations from linear theory are less pronounced down to smaller scales for modes with  $\mu \simeq 0.5$  because nonlinear growth and the combination of peculiar velocities and thermal broadening tend to cancel each other out. As we can see, the parametric model achieves an average accuracy of 2% for  $k > 0.5 \,\mathrm{Mpc}^{-1}$ , supporting the validity of Eq. (4) for capturing small-scale deviations from linear theory.

On the largest scales, we find strong variations between consecutive k-bins for the same  $\mu$ -wedge. Some of these oscillations are driven by differences in the average value of  $\mu$  between consecutive bins due to the limited number of modes entering each bin on large scales. To ensure an accurate comparison between simulation measurements and model predictions, we individually evaluate the  $P_{3D}$  model for all the modes within each  $k - \mu$ bin from our simulation boxes. We then calculate the mean of the resulting distribution and assign this mean value to the bin, thereby mirroring the approach used to compute  $P_{3D}$  measurements from the simulations. This process is also important on



**Fig. 1.** Accuracy of the  $P_{3D}$  model (see Eqs. (3) and (4)) in reproducing measurements from the CENTRAL simulation at z = 3. In the top panel, dotted and solid lines show the ratio of simulation measurements and model predictions relative to the linear power spectrum, respectively. Dashed lines do so for the linear part of the best-fitting model ( $D_{NL} = 1$ ). Line colors correspond to different  $\mu$  wedges, and vertical dashed lines mark the minimum scale used for computing the best-fitting model,  $k = 5 \text{ Mpc}^{-1}$ . The bottom panel displays the relative difference between the best-fitting model and simulation measurements. The overall accuracy of the model is 2% on scales in which simulation measurements are not strongly affected by cosmic variance ( $k > 0.5 \text{ Mpc}^{-1}$ ; see text).

small scales, where the number of modes increases rapidly with k. Throughout this work, we adopt this approach to make predictions from the  $P_{3D}$  model.

Using this approach to generate theoretical predictions, the best-fitting model successfully captures most of the aforementioned large-scale oscillations. However, a fluctuation at  $k \simeq$  $0.25 \,\mathrm{Mpc}^{-1}$  in the  $0 < \mu < 0.25$  wedge remains unaccounted for by the model. The negligible differences between model predictions and simulation measurements in the adjacent bins suggest that this fluctuation is likely due to cosmic variance. We assess the impact of this source of uncertainty on  $P_{3D}$  in Appendix A, finding that it can induce fluctuations of up to 10% on scales  $k < 0.5 \,\mathrm{Mpc}^{-1}$ . As a result, cosmic variance limits our ability to evaluate the model's performance on the largest scales shown. However, this does not necessarily reflect reduced model accuracy but rather highlights the limitations of using our simulations for validating the model. Proper validation on large scales would require either a larger simulation or multiple simulations with different initial conditions.

#### 2.2. Input and output parameters

In addition to the density and velocity fields, the Lyman- $\alpha$  forest depends upon the ionization and thermal state of the IGM (e.g., McDonald 2003). Following Pedersen et al. (2021), we use six parameters to describe the dependency of this observable with cosmology and IGM physics:

 $<sup>^3</sup>$  Aside from nonlinear effects affecting the position and damping of BAO.

- Amplitude and slope of the linear matter power spectrum on small scales. We define the amplitude  $(\Delta_p^2)$  and slope  $(n_p)$  as

$$\Delta_{\rm p}^2(z) = (2\pi^2)^{-1} k^3 P_{\rm lin}(k_{\rm p}, z), \tag{5}$$

$$n_{\rm p}(z) = \left( d\log P_{\rm lin}/d\log k \right)|_{k=k_{\rm p}},\tag{6}$$

where we use  $k_p = 0.7 \,\mathrm{Mpc}^{-1}$  as the pivot scale because it is at the center of the range of interest for DESI smallscale studies. These parameters capture multiple physical effects modifying the linear power spectrum on small scales (see Pedersen et al. 2021, for a detailed discussion), including cosmological parameters such as the amplitude  $(A_s)$ and slope  $(n_s)$  of the primordial power spectrum, the Hubble parameter, and the matter density ( $\Omega_M$ ), or  $\Lambda CDM$ extensions such as curvature and massive neutrinos. The advantage of using this parameterization rather than ACDM parameters is twofold. First, it reduces the dimensionality of the input to the cNF (see Sect. 4), which decreases the number of simulations required for precise training. Second, the resulting emulator has the potential for making precise predictions for variations in cosmological parameters and ACDM extensions not considered in the training set (Pedersen et al. 2021, 2023; Cabayol-Garcia et al. 2023). Note that we do not consider cosmological parameters capturing changes in the growth rate or expansion history because the Lyman- $\alpha$  forest probes cosmic times during which the universe is practically Einstein de-Sitter, and both vary very little with cosmology in this regime.

- Mean transmitted flux fraction. The mean transmitted flux fraction  $(\bar{F})$  depends on the intensity of the cosmic ionizing background and evolves strongly with redshift. One of the advantages of using this parameter is that it encodes the majority of the redshift dependence of the signal, serving as a proxy for cosmic time.
- Amplitude and slope of the temperature-density relation. The thermal state of the IGM can be approximated by a power law on the densities probed by the Lyman- $\alpha$  forest (Lukić et al. 2015):  $T_0 \Delta_b^{\gamma-1}$ , where  $\Delta_b$  is the baryon overdensity,  $T_0$  is the gas temperature at mean density, and  $\gamma - 1$ is the slope of the relation. These parameters influence the ionization of the IGM, which is captured by  $\bar{F}$ , and the thermal motion of gas particles, which causes Doppler broadening that suppresses the parallel power. Instead of using  $T_0$ as an emulator parameter, we follow Pedersen et al. (2021) and use the thermal broadening scale in comoving units. First, we express the thermal broadening in velocity units as  $\tilde{\sigma}_{\rm T} = 9.1(T_0[{\rm K}]/10^4)^{1/2}$ , and then we convert it to comoving units,  $\sigma_{\rm T} = \tilde{\sigma}_{\rm T}(1 + z)H^{-1}$ .
- Pressure smoothing scale. Gas pressure supports baryons on small scales, leading to a strong isotropic power suppression in this regime that depends upon the entire thermal history of the gas (Gnedin & Hui 1998). We parameterize this effect using the pressure smoothing scale in units of comoving Mpc<sup>-1</sup>,  $k_{\rm F}$  (see Pedersen et al. 2021, for more details).

In summary, our cNF predicts the eight free parameters of the physically motivated model for Lyman- $\alpha$  clustering introduced by Eqs. (3) and (4),  $\mathbf{y} = \{b_{\delta}, b_{\eta}, q_1, q_2, k_v, a_v, b_v, k_p\}$ , as a function of the aforementioned six parameters capturing the cosmological and IGM dependence of the Lyman- $\alpha$  forest,  $\mathbf{x} = \{\Delta_p^2, n_p, \bar{F}, \sigma_T, \gamma, k_F\}$ .

## 3. Training and testing set

In this section, we describe how we generated the training and testing data for the cNF described in Sect. 4. In Sect. 3.1, we present a suite of cosmological hydrodynamical simulations from which we generated mock Lyman- $\alpha$  forest measurements, and we detail our approach for extracting  $P_{3D}$  and  $P_{1D}$  measurements from these simulations in Sect. 3.2. In Sect. 3.3, we compute the best-fitting parameters of the model introduced by Eqs. (3) and (4) to measurements of these statistics, and we evaluate the performance of the fits in Sect. 3.4.

## 3.1. Simulations

We extracted Lyman- $\alpha$  forest simulated measurements from a suite of simulations run with MP-GADGET<sup>4</sup> (Feng et al. 2018; Bird et al. 2019), a massively scalable version of the cosmological structure formation code GADGET-3 (last described in Springel 2005). This suite was first presented and used in Pedersen et al. (2021); we briefly describe it next. Each simulation tracked the evolution of 768<sup>3</sup> dark matter and baryon particles from z = 99 to z = 2 inside a box of L = 67.5 Mpc on a side, producing as output 11 snapshots uniformly spaced in redshift between z = 4.5 and 2. This configuration ensures convergence for  $P_{1D}$  measurements down to  $k_{\parallel} = 4$  Mpc<sup>-1</sup> (the smallest scale used in this work) at z = 2 and less than 10% errors for this scale at z = 4 (see Bolton et al. 2017, for more details). On the other hand, this configuration may cause non-negligible biases for  $P_{3D}$  at high redshift (Lukić et al. 2015).

Two realizations were run for each combination of cosmological and astrophysical parameters using the "fixedand-paired" technique (Angulo & Pontzen 2016; Pontzen et al. 2016), which significantly reduces cosmic variance for multiple observables, including the Lyman- $\alpha$  forest (Villaescusa-Navarro et al. 2018; Anderson et al. 2019). The initial conditions were generated using the following configuration of MP-GENIC (Bird et al. 2020): initial displacements produced using the Zel'dovich approximation and baryons and dark matter initialized on an offset grid using species-specific transfer functions. Some studies have suggested that this configuration might lead to incorrect evolution of linear modes (Bird et al. 2020). However, in a recent study, Khan et al. (2024) showed that variations in the specific setting of MP-GENIC initial conditions have a minimal impact on  $P_{1D}$  measurements across the range of redshifts and scales used in this work.

To increase computational efficiency, the simulations utilized a simplified prescription for star formation that turns regions of baryon overdensity  $\Delta_b > 1000$  and temperature  $T < 10^5$  K into collisionless stars (e.g., Viel et al. 2004), implemented a spatially uniform ultraviolet background (Haardt & Madau 2012), and did not consider active galactic nuclei (AGN) feedback (e.g., Chabanier et al. 2020). These approximations are justified because we focus on emulating the Lyman- $\alpha$  forest in the absence of astrophysical contaminants like AGN feedback, damped Lyman-alpha absorbers (DLAs), or metal absorbers, and we will model these before comparing our predictions with observational measurements (e.g., McDonald et al. 2005; Palanque-Delabrouille et al. 2015, 2020).

In Sect. 4, we train a cNF using data from 30 fixed-andpaired simulations from the previous suite, covering combinations of cosmological and astrophysical parameters selected via a Latin hypercube sampling method (McKay et al. 1979).

<sup>4</sup> https://github.com/MP-Gadget/MP-Gadget/

Hereafter, we refer to these as TRAINING simulations. The Latin hypercube spans the parameters { $\Delta_p^2(z = 3)$ ,  $n_p(z =$ 3),  $z_{\rm H}$ ,  $H_{\rm A}$ ,  $H_{\rm S}$ }, where we use z = 3 because it is approximately at the center of the range of interest for DESI studies (Ravoux et al. 2023; Karaçaylı et al. 2024), z<sub>H</sub> is the midpoint of hydrogen reionization, and the last two parameters rescale the total photoheating rate  $\epsilon_0$  as  $\epsilon = H_A \Delta_b^{H_S} \epsilon_0$  (Oñorbe et al. 2017). Cosmological parameters were generated within the ranges  $\Delta_p^2(z = 3) \in [0.25, 0.45], n_p(z = 3) \in [-2.35, -2.25]$  by exploring values of the amplitude and slope of the primordial power spectrum within the intervals  $A_s \in [1.35, 2.71] \times 10^{-9}$  and  $n_s \in$ [0.92, 1.02]. Any other ACDM parameter was held fixed to values approximately following Planck Collaboration VI (2020): dimensionless Hubble parameter h = 0.67, physical cold dark matter density  $\Omega_c h^2 = 0.12$ , and physical baryon density  $\Omega_b h^2 =$ 0.022. As for the IGM parameters, these explored the ranges  $z_{\rm H} \in [5.5, 15], H_{\rm A} \in [0.5, 1.5], \text{ and } H_{\rm S} \in [0.5, 1.5].$  All simulation pairs share the same set of initial Fourier phases, making their  $P_{3D}$  and  $P_{1D}$  subject to the same large-scale noise pattern.

We evaluated different aspects of the emulation strategy using six fixed-and-paired simulations with cosmological and astrophysical parameters not considered in the TRAINING simulations:

- The CENTRAL simulation uses cosmological and astrophysical parameters at the center of the TRAINING parameter space:  $A_s = 2.01 \times 10^{-9}$ ,  $n_s = 0.97$ ,  $z_H = 10.5$ ,  $H_A = 1$ , and  $H_S = 1$ . We use this simulation for an out-of-sample test to evaluate the performance of FORESTFLOW under optimal conditions, recognizing that the accuracy of machinelearning models generally declines near the boundaries of the convex hull defined by the training set.
- The SEED simulation uses the same parameters as the CEN-TRAL simulation while considering a different distribution of initial Fourier phases. Given that all TRAINING simulations use the same initial Fourier phases, SEED is useful to evaluate the impact of cosmic variance in the training set on FOREST-FLOW predictions.
- The GROWTH, NEUTRINOS, and CURVED simulations adopt the same values of  $\Delta_p^2(z = 3)$ ,  $n_p(z = 3)$ , physical cold dark matter and baryonic densities, and astrophysical parameters as the CENTRAL simulation. However, the GROWTH simulation uses 10% larger Hubble parameter (h = 0.74) and 18% smaller matter density ( $\Omega_M = 0.259$ ) while using the same value of  $\Omega_M h^2$  as the TRAINING simulations, the NEUTRINOS simulation includes massive neutrinos ( $\sum m_v = 0.3 \text{ eV}$ ), and the CURVED simulation considers an open universe ( $\Omega_k =$ 0.03). The NEUTRINOS and CURVED simulations also modify the value of the cosmological constant while holding fixed hto compensate for the increase in the matter density and the addition of curvature, respectively. We use the testing simulations to evaluate the performance of the emulation strategy for cosmologies not included in the training set.
- The REIONISATION simulation uses the same cosmological parameters as the CENTRAL simulation while implementing a distinct helium ionization history relative to the CENTRAL and TRAINING simulations (Puchwein et al. 2019). The main difference between the ionization histories of these simulations is that the one implemented in the REIONISATION simulation peaks at a later time than the others, leading to a significantly different thermal history. The REIONISA-TION simulation therefore tests the performance of FOREST-FLOW for thermal histories not considered in the TRAINING simulations.

### 3.2. Simulating Lyman- $\alpha$ forest data

To extract Lyman- $\alpha$  forest measurements from each simulation, we first selected one of the simulation axes as the line of sight and displace the simulation particles from real to redshift space along this axis. Then, we computed the transmitted flux fraction along 768<sup>2</sup> uniformly distributed line of sights along this axis using FSFE<sup>5</sup> (Bird 2017); these lines of sight are commonly known as skewers. Following Springel (2005), we computed pressure forces using the density-entropy formulation of smoothed particle hydrodynamics (SPH) with a cubic spline kernel and 33 neighbors<sup>6</sup>. We set the resolution of the skewers to 0.05 Mpc, which is enough to resolve the thermal broadening and pressure scales, and spaced these by 0.09 Mpc in the transverse direction. We checked that  $P_{3D}$  and  $P_{1D}$  measurements within the range of interest (see Sect. 3.3) do not vary by increasing the line-of-sight resolution or the transverse sampling. We repeated all the previous steps for the three simulation axes to extract further cosmological information, as each simulation axis samples the velocity field in a different direction. Finally, we scaled the effective optical depth of the skewers to 0.90, 0.95, 1.05, and 1.10 times its original value, which is equivalent to running simulations with different UV background photoionization rates (see Lukić et al. 2015).

Using this data as input, we measured  $P_{3D}$  by first computing the three-dimensional Fourier transform of the skewers. Then, we took the average of the square norm of all modes within 20 logarithmically spaced bins in wavenumber k from the fundamental mode of the box,  $k_{\min} = 2\pi L^{-1} \simeq 0.09 \,\mathrm{Mpc}^{-1}$ , to  $k_{\text{max}} = 40 \,\text{Mpc}^{-1}$  and 16 linearly spaced bins in the cosine of the angle between Fourier modes and the line of sight from  $\mu = 0$  to 1. We measured  $P_{1D}$  by first computing the onedimensional Fourier transform of each skewer without applying any binning, and then by taking the average of the square norm of all these Fourier transforms. The impact of cosmic variance on fixed-and-paired simulations is not straightforward (Maion et al. 2022), and thus we would ideally use multiple fixed-and-paired simulations with different initial distributions of Fourier phases to estimate the precision of  $P_{3D}$  and  $P_{1D}$  measurements from our simulations. However, such simulations are not available, and we instead relied on the comparison between two simulations with same configuration but different initial conditions (see Appendix A). We found that the impact of cosmic variance on  $P_{3D}$  and  $P_{1D}$  measurements can be as large as 10 and 1% on intermediate scales, respectively.

We measured  $P_{3D}$  and  $P_{1D}$  from the 30 TRAINING and the six test simulations, ending up with 2 (opposite Fourier phases) × 3 (simulation axes) × 11 (snapshots) × 5 (mean flux rescalings) = 330 measurements of each of these statistics per simulation. To reduce cosmic variance, we computed the average of measurements from different axes and phases of fixed-and-paired simulations, which decreased the number of measurements per simulation to 55. The training and testing sets of FORESTFLOW are thus comprised of 1650 and 330 Lyman- $\alpha$  power spectra measurements<sup>7</sup>, respectively.

<sup>5</sup> https://github.com/sbird/fake\_spectra

<sup>&</sup>lt;sup>6</sup> This approach may lead to biases in flux statistics (Chabanier et al. 2023).

<sup>&</sup>lt;sup>7</sup> These are publicly available at https://github.com/ igmhub/LaCE

#### 3.3. Fitting the parametric model

To generate training and testing data for our emulator, we computed the best-fitting parameters of Eqs. (3) and (4) to measurements from the simulations described in Sect. 3.1. We fitted the model using  $P_{3D}$  measurements from k = 0.09 to  $5 \text{ Mpc}^{-1}$  and  $P_{1D}$  measurements from  $k_{\parallel} = 0.09$  to  $4 \text{ Mpc}^{-1}$ . The size of our simulation boxes determines the largest scales used, while the smallest scales measured by DESI set the maximum wavenumbers (Ravoux et al. 2023; Karaçaylı et al. 2024). It is important to note that the two Lyman- $\alpha$  linear biases determine the largescale behavior of  $P_{3D}$  (see Eq. (3)), and thus FORESTFLOW could make accurate predictions for  $P_{3D}$  on arbitrarily large (linear) scales as long as the Lyman- $\alpha$  linear biases are measured accurately.

We computed the best-fitting value of model parameters  $\mathbf{y} = \{b_{\delta}, b_{\eta}, q_1, q_2, k_{v}, a_{v}, b_{v}, k_{p}\}$  to simulation measurements by minimizing the pseudo- $\chi^2$ :

$$\chi^{2}(\mathbf{y}) = \sum_{i}^{M_{3D}} w_{3D} \left[ P_{3D}^{\text{data}}(k_{i},\mu_{i}) - P_{3D}^{\text{model}}(k_{i},\mu_{i},\mathbf{y}) \right]^{2} + \sum_{i}^{M_{1D}} w_{1D} \left[ P_{1D}^{\text{data}}(k_{\parallel,i}) - P_{1D}^{\text{model}}(k_{\parallel,i},\mathbf{y}) \right]^{2},$$
(7)

where  $M_{3D} = 164$  and  $M_{1D} = 42$  were the number of  $P_{3D}$  and  $P_{1D}$  bins employed in the fit, respectively, the superscripts data and model refer to simulation measurements and model predictions, and  $w_{3D}$  and  $w_{1D}$  weighed the fit. We used the Nelder-Mead algorithm implemented in the routine MINIMIZE of SCIPY (Virtanen et al. 2020) to carry out the minimization<sup>8</sup>. The results of the fits are publicly accessible<sup>9</sup>.

Ideally, we would have used the covariance of  $P_{3D}$  and  $P_{1D}$  measurements to weigh the previous expression. However, estimating this covariance requires multiple realizations of the same simulation with different initial distributions of Fourier phases, and we do not have these simulations available. Instead, we disregarded correlations between  $P_{3D}$  and  $P_{1D}$  and weighed these by  $w_{3D} = N_{3D}(k,\mu)/(1 + \mu^2)^2$  and  $w_{1D} = \alpha(1 + k_{\parallel}/k_0)^2$ , where  $N_{3D}$  is the number of modes in each  $k - \mu$  bin and  $k_0 = 2 \text{ Mpc}^{-1}$ . The terms involving  $N_{3D}$ ,  $\mu$ , and  $k_0$  attempt to ensure an unbiased fit of  $P_{3D}$  and  $P_{1D}$  across the full range of scales used. The parameter  $\alpha = 8000$  controls the relative weight of  $P_{3D}$  and  $P_{1D}$  in the fit, and we set this value motivated by the different impact of cosmic variance on these statistics (see Appendix A).

We expect significant correlations between the best-fitting value of the parameters to measurements from relatively small simulation boxes. As shown by Arinyo-i-Prats et al. (2015), these correlations are especially significant for the parameters accounting for nonlinear growth of structure,  $q_1$  and  $q_2$ . Givans et al. (2022) advocated for setting  $q_2 = 0$  since this parameter is not necessary for describing  $P_{3D}$  at z = 2.8. However, we found non-zero values of this parameter indispensable for describing  $P_{3D}$  at redshifts below z = 2.5. This is not surprising since the gravitational evolution of density perturbations becomes increasingly more nonlinear as cosmic time progresses.



**Fig. 2.** Accuracy of the parametric model (see Eqs. (3) and (4)) in reproducing  $P_{3D}$  and  $P_{1D}$  measurements from all the TRAINING simulations. Lines and shaded areas show the mean and standard deviation of the relative difference between simulation measurements from the 1650 snapshots of the TRAINING simulations and best-fitting models to these, respectively. The accuracy of the model in recovering  $P_{3D}$  and  $P_{1D}$  is 2.4 and 0.6%, respectively, on scales not strongly affected by cosmic variance.

## 3.4. Accuracy of the parametric model

In the previous section, we computed the best-fitting parameters of the  $P_{3D}$  model to measurements from the TRAINING simulations. Two main sources of uncertainty can affect these fits: model inaccuracies and cosmic variance. The first relates to using a parametric model without enough flexibility to describe Lyman- $\alpha$  clustering accurately, while the second has to do with the limited size of the TRAINING simulations. The influence of cosmic variance on the training set is amplified because all the TRAINING simulations use the same initial distribution of Fourier phases, meaning all simulations are subject to the same largescale noise. We study this source of uncertainty in Appendix A, where we compared the best-fitting models to the CENTRAL and SEED simulations, whose only difference is in their initial distribution of Fourier phases. We proceed to study model inaccuracies next.

In Fig. 2, we show the performance of the parametric model in reproducing  $P_{3D}$  and  $P_{1D}$  measurements from the 1650 snapshots of the TRAINING simulations. As discussed in Sect. 2.1, cosmic variance limits our ability to evaluate the accuracy of the model for  $P_{3D}$  on scales  $k < 0.5 \,\mathrm{Mpc}^{-1}$ ; therefore, we quote the model accuracy from  $k = 0.5 \,\mathrm{Mpc}^{-1}$  down to the smallest scale used in the fit,  $k = 5 \,\mathrm{Mpc}^{-1}$ . In contrast, since cosmic variance has a much smaller impact on  $P_{1D}$ , we evaluate the model performance for this statistic using all scales considered in the fit  $(0.09 < k_{\parallel} [Mpc^{-1}] < 4)$ . We adopt the same approach when evaluating the performance of FORESTFLOW in Sect. 5. Under these considerations, the overall accuracy of the parametric model is 2.4 and 0.6% for  $P_{3D}$  and  $P_{1D}$ , respectively. Given that we estimate the accuracy of the parametric model using the TRAINING simulations, the previous numbers account for both the limited flexibility of such model and cosmic variance. As discussed in Appendix A, the impact of cosmic variance on measurements of  $P_{3D}$  and  $P_{1D}$  from these simulations is 1.3 and 0.5%, respectively. If we assume that cosmic variance and errors coming from the limited flexibility of the parametric model are

<sup>&</sup>lt;sup>8</sup> To ensure that this routine did not get stuck in a local minimum, we checked that the likelihood is unimodal in all cases using the Affine Invariant Markov chain Monte Carlo Ensemble sampler EMCEE (Foreman-Mackey et al. 2013).

<sup>%</sup> https://github.com/igmhub/ForestFlow

uncorrelated and add in quadrature, the second are responsible for 2.0 and 0.3% errors on  $P_{3D}$  and  $P_{1D}$ , respectively.

# 4. Emulator

In this section, we use a cNF to predict the two Lyman- $\alpha$  linear biases and six parameters describing small-scale deviations of  $P_{3D}$  from linear theory as a function of cosmology and IGM physics. We detail the architecture and implementation of this emulator in Sects. 4.1 and 4.2, respectively.

#### 4.1. Conditional normalizing flows

Normalizing flows (NFs; Jimenez Rezende & Mohamed 2015) are a class of machine-learning generative models designed to predict complex distributions by applying a sequence of bijective mappings to simple base distributions. A natural extension to this framework is conditional NFs (cNFs; Winkler et al. 2019; Papamakarios et al. 2019), a type of NFs that condition the mapping between the base and target distributions on a series of input variables. Given an input  $\mathbf{x} \in X$  and target  $\mathbf{y} \in Y$ , cNFs predict the conditional distribution  $p_{Y|X}(\mathbf{y}|\mathbf{x})$  by applying a parametric, bijective mapping  $f_{\phi} : Y \times X \to Z$  to a base distribution  $p_Z(\mathbf{z})$  as follows

$$p_{Y|X}(\mathbf{y}|\mathbf{x}) = p_Z(f_\phi(\mathbf{y}, \mathbf{x})|\mathbf{x}) \left| \frac{\partial f_\phi(\mathbf{y}, \mathbf{x})}{\partial \mathbf{y}} \right|, \tag{8}$$

where  $\phi$  are the parameters of the mapping, while the last term of the previous equation is the Jacobian determinant of the mapping. In our cNF, the input is given by the parameters capturing the dependence of the Lyman- $\alpha$  forest on cosmology and IGM physics,  $\mathbf{x} = \{\Delta_p^2, n_p, \bar{F}, \sigma_T, \gamma, k_F\}$ , the target by the parameters of the  $P_{3D}$  model,  $\mathbf{y} = \{b_{\delta}, b_{\eta}, q_1, q_2, k_v, a_v, b_v, k_p\}$ , and the base distribution is an eight-dimensional Normal distribution  $N^8(0, 1)$ , where the dimension is determined by the number of  $P_{3D}$  model parameters.

Once trained, cNFs are a generative process from **x** to **y**. In our implementation, we start by randomly sampling from the base distribution, and then we pass this realization through a sequence of mappings conditioned on a particular combination of cosmology and IGM parameters,  $\tilde{\mathbf{y}} = f_{\phi}^{-1}(p_Z(\mathbf{z}), \mathbf{x})$ , ending up with a prediction for the value of the  $P_{3D}$  model parameters. Repeating this process multiple times, the emulator yields a distribution of  $P_{3D}$  parameters  $p_{\tilde{Y}|X}$  that, for a sufficiently large number of samples, approaches the target distribution  $p_{Y|X}$ . The breadth of this distribution captures uncertainties arising from the limited size of the training set. Finally, outside the cNF, we use each combination of  $P_{3D}$  parameters to evaluate Eqs. (3) and (1), obtaining predictions and uncertainties for  $P_{3D}$  and  $P_{1D}$ .

The main challenge when using cNFs is finding the mapping between the target and the base distribution, typically done using an *N*-layer neural network with bijective layers. This process runs in reverse relative to the generating process: we start by applying the mapping  $f_{\phi}$  to the target data y conditioned on the input x, yielding z. Then, we optimize the model parameters by minimizing the loss function

$$\mathcal{L} = \frac{1}{2} \sum \mathbf{z}^2 - \log \left| \frac{\delta f_{\phi}(\mathbf{y}, \mathbf{x})}{\delta \mathbf{y}} \right|.$$
(9)

We carried out this optimization process using stochastic gradient descent applied to minibatches, a methodology commonly employed for training neural networks.

#### 4.2. Implementation

Neural Autoregressive Flows (Huang et al. 2018) use a series of invertible univariate operations to build a bijective transformation between a conditional distribution and a base distribution. In FORESTFLOW, we created a bijective mapping between the best-fitting parameters of the  $P_{3D}$  model and an eight-dimensional Normal distribution by applying  $N_{ACB} = 12$  consecutive Affine-Coupling Block (ACB; Dinh et al. 2016) conditioned on cosmology and IGM physics. The transformation goes from the best-fitting parameters of the  $P_{3D}$  model to the base distribution when training the model, and in the opposite direction when evaluating it.

Each ACB conducts a series of operations  $g_{i,\tilde{\phi}_i}$  on its input data  $\mathbf{w}_i$ , with *i* going from 1 to  $N_{ACB}$  and  $\tilde{\phi}_i$  standing for the parameters of the transformation. First, it splits the input data into two subsamples with approximately the same number of elements,  $\mathbf{w}'_i$  and  $\mathbf{w}''_i$ . Then, it applies an affine transformation to the first subsample  $\mathbf{w}'_i$ 

$$T(\mathbf{w}'_i) = \alpha_i \, \mathbf{w}'_i + \beta_i,\tag{10}$$

where  $\alpha_i$  and  $\beta_i$  are neural networks with a single hidden layer of 128 neuron units. Third, the ACB merges the output from the affine transformation and the unchanged subsample, and then it applies a permutation layer to randomly rearrange these elements, obtaining  $\tilde{\mathbf{w}}_i$ . Fourth, the ACB applies an affine transformation to this sample,  $\tilde{T}(\tilde{\mathbf{w}}_i)$ . The first and second affine transformations involve a subset of the training set and the entire training set, respectively, enabling the model to capture local and global features.

In Fig. 3, we show the architecture of the cNF. The blue arrow indicates the training direction, while the green arrow depicts the emulation direction. In the training direction, the input to the first ACB,  $\mathbf{u}_1 = \mathbf{w}_1$ , is a 1650-dimensional array composed of 14-dimensional vectors, where 1650 is the number of simulation snapshots in the training set. Each vector includes the eight best-fitting  $P_{3D}$  model parameters to each snapshot and the six parameters describing the cosmology and IGM physics of this snapshot. The input to the *i* ACB,  $\mathbf{u}_i$ , is a 1650-array containing 14-dimensional vectors with the output of the *i* – 1 ACB and, once again, the six parameters describing the cosmology and IGM physics of each snapshot. Each ACB applies a transformation  $f_{i,\phi_i} = g_{i,\bar{\phi}_i}$ , and the consecutive application of all ACBs results in the mapping between the target and the base distributions  $\mathbf{z} = f_{\phi}(\mathbf{y}, \mathbf{x})$ , where  $f_{\phi} = \prod_{i=1}^{N_{ACB}} f_{i,\phi_i}$ . In the emulation direction, the input to the first ACB,  $\mathbf{v}_1 =$ 

In the emulation direction, the input to the first ACB,  $\mathbf{v}_1 = \mathbf{w}_1$ , is a 14-dimensional vector containing random draws from an eight-dimensional Normal distribution and the six parameters describing the cosmology and IGM physics for which we want to obtain predictions. As in the training direction, the input to each subsequent ACB relies on the output from the previous ACB, each conditioned on cosmology and IGM physics. The ACBs apply the transformations  $f_{i,\phi_i}^{-1} = g_{i,\tilde{\phi}_i}$ , which are the inverse of the corresponding transformations in the training direction,  $f_{i,\phi_i}$ . The cNF makes predictions for  $P_{3D}$  model parameters by applying the composition of the inverse of all ACBs to random samples from the base distribution,  $\tilde{\mathbf{y}} = f_{\phi}^{-1}(p_Z(\mathbf{z}), \mathbf{x})$ , where  $f^{-1} = \Pi^{N_{ACB}} f^{-1}$ 

$$f_{\phi}^{T} = \prod_{i=1}^{i \neq ACB} f_{i,\phi_i}^{T}.$$

We implemented the emulator within the FreIA framework (Ardizzone et al. 2018-2022), which uses PyTorch (Ansel et al. 2024) in the backend. We trained it by minimizing Eq. (9) using an Adam optimizer (Kingma & Ba 2014) for 300 epochs with an initial learning rate of  $10^{-3}$ . We used the Optuna framework



**Fig. 3.** Architecture of the Lyman- $\alpha$  forest clustering emulator. The blue arrow indicates the training direction, where the cNF optimizes a bijective mapping between the best-fitting parameters of the  $P_{3D}$  model to measurements from the TRAINING simulations and an eight-dimensional Normal distribution. The mapping is conditioned on cosmology and IGM physics, and performed using 12 consecutive affine coupling blocks. The green arrow denotes the emulation direction, where the cNF applies the inverse of the mapping to random samples from the base distribution to predict the value of the  $P_{3D}$  model parameters. Outside the cNF, FORESTFLOW introduces these parameters in Eq. (3) and (1) to obtain predictions for  $P_{3D}$  and  $P_{1D}$ , respectively.

(Akiba et al. 2019) to select the number of ACBs and epochs, as well as the value of the learning rate. First, Optuna trains our cNF for a particular combination of these hyperparameters. Then, it computes the average value of Eq. (7) for all simulations in the training set. After that, depending on the goodness of the fit to  $P_{3D}$  and  $P_{1D}$  measurements, Optuna selects a new value of the hyperparameters. We iterated with Optuna 50 times through a hyperparameter grid, selecting the hyperparameters that yield the highest accuracy. We checked that the performance of the cNF depended weakly on small variations in the value of the hyperparameters.

## 5. Performance of FORESTFLOW

In Sect. 5.1, we analyze the performance of FORESTFLOW across the parameter space of the training set. Then, in Sect. 5.2, we test its accuracy using simulations with cosmologies and IGM models that are not part of the training set. All performance evaluations are conducted using the simulations described in Sect. 3.1, which were run using the same code and resolution. Before employing FORESTFLOW for cosmological inference, it will be crucial to validate it against large, high-resolution simulations produced with alternative codes. We defer this task to future work.

## 5.1. Cosmologies and IGM histories in the training set

In this section, we evaluate the performance of FORESTFLOW in recovering the two Lyman- $\alpha$  linear biases, which determine the behavior of  $P_{3D}$  on linear scales, as well as  $P_{3D}$  and  $P_{1D}$  measurements from simulations on the intervals  $0.5 < k [Mpc^{-1}] < 5$  and  $0.09 < k_{\parallel} [Mpc^{-1}] < 4$ , respectively. These are the ranges of scales used when fitting the parametric model in Sect. 3 that are not strongly affected by cosmic variance (see Sect. 2.1). We begin by assessing the accuracy of FORESTFLOW at the center of the training set, where machine-learning methods typically perform best, and then extend our evaluation across the entire input parameter space.

In Fig. 4, we compare measurements of  $P_{3D}$  and  $P_{1D}$  from the CENTRAL simulation at z = 3 with FORESTFLOW predictions. Dotted lines show simulation measurements, while solid lines and shaded areas display the average and 68% credible interval of FORESTFLOW predictions, respectively. We characterize the accuracy of the credible intervals in Appendix B. As we can see, FORESTFLOW captures the amplitude and scale-dependence of  $P_{3D}$  and  $P_{1D}$  precisely. In Fig. 5, we present the relative difference between  $P_{3D}$  and  $P_{1D}$  measurements from the CEN-TRAL simulation and FORESTFLOW predictions as a function of redshift. The model's accuracy remains consistent for redshifts above z = 2 but shows a slight decline at this redshift. This is likely because z = 2 is the lowest redshift included in the training set and is therefore near the boundary of the convex hull defined by the training data.

To better characterize the performance of FORESTFLOW, we compute the average accuracy of FORESTFLOW in recovering measurements from CENTRAL across redshift. We find that it is 1.2 and 0.3% for  $b_{\delta}$  and  $b_{\eta}$ , respectively, which translates into 1.1 and 1.2% for perpendicular and parallel  $P_{3D}$  modes on linear scales, and 2.6 and 0.8% for  $P_{3D}$  and  $P_{1D}$ . Note that cosmic variance hinders our ability to test the performance of the model; however, this does not necessarily indicate a decrease in the model's accuracy for  $P_{3D}$  on the largest scales sampled by our simulation.

We expect the efficiency of FORESTFLOW to decrease away from the center of the input space. We could assess its performance across the parameter space using the training simulations; however, the cNF has been optimized for these points, which introduces the risk of overfitting. Overfitting could result in high precision for these specific points of the parameter space but not for others nearby. As a result, we would ideally evaluate the performance of FORESTFLOW using multiple test simulations covering the entire parameter space, but such simulations are unavailable. Instead, we conduct leave-one-out tests, which are widely used to assess the performance of an emulator when the number of training points is insufficient for out-of-sample tests (e.g., Hastie et al. 2001). In a leave-one-out test, we optimize a cNF after removing a subsample from the



**Fig. 4.** Accuracy of FORESTFLOW in recovering  $P_{3D}$  and  $P_{1D}$  measurements from the CENTRAL simulation at z = 3. Dotted lines show measurements from simulations, solid lines and shaded areas display the average and 68% credible interval of FORESTFLOW predictions, respectively, and vertical dashed lines indicate the minimum scales considered for computing the training data for the cNF. The overall performance of FORESTFLOW in recovering  $P_{3D}$  is 2.0% on scales not strongly affected by cosmic variance and 0.6% for  $P_{1D}$ .

 $10^{0}$ 

 $k_{\parallel} \, [{\rm Mpc}^{-1}]$ 

-0.05

 $10^{-1}$ 

training set; for example, all measurements from one of the TRAINING simulations. We then check the accuracy of FOREST-FLOW for the new cNF using the subsample held back. The rationale is that the new emulator should closely approximate the original emulator everywhere in the parameter space except near the excluded simulation, and more importantly, there is no risk of overfitting. By repeating this process for other subsamples, we can estimate the performance of FORESTFLOW across the parameter space. Since each cNF is trained without using the entire dataset, leave-one-out tests provide a lower bound on FORESTFLOW performance. Additionally, leave-one-out tests may require extrapolating the predictions from the cNF, and it is widely known that machine-learning methods do not extrapolate well.

In the top panels of Fig. 6, lines and shaded areas display the average and standard deviation of 30 leave-simulation-out tests. Each test requires optimizing a cNF with 29 distinct TRAINING simulations, and then using the remaining simulation as the validation. Each panel shows the results for a different redshift, and we check that the results are similar for redshifts not shown. As we can see, the large-scale noise is similar for all TRAINING



**Fig. 5.** Accuracy of FORESTFLOW in recovering  $P_{3D}$  and  $P_{1D}$  measurements from the CENTRAL simulation as a function of redshift. The upper four panels show the results for  $P_{3D}$  across different  $\mu$  bins, while the bottom panel displays the results for  $P_{1D}$ . Each color represents a different redshift. The model's accuracy remains consistent for redshifts above z = 2 but exhibits a slight decline at this redshift.

simulations; this is because they use the same initial distribution of Fourier phases. The overall performance of FORESTFLOW in recovering  $b_{\delta}$  and  $b_{\eta}$  is 1.0 and 3.1%, respectively, which translates into 2.0 and 2.9% for perpendicular and parallel  $P_{3D}$  modes on linear scales, and 3.4 and 1.8% for  $P_{3D}$  and  $P_{1D}$ .

In Table 1, we gather the accuracy of FORESTFLOW at the center and across the parameter space, as well as the expected level of uncertainties due to cosmic variance and the limited flexibility of the  $P_{3D}$  model. Due to the limited size of our simulations, the maximum levels of accuracy we can test for  $P_{3D}$ and  $P_{1D}$  are 1.3 and 0.5% (see Appendix A), respectively. These levels would decrease by evaluating the accuracy of FOREST-FLOW using bigger simulations with the same resolution. On the other hand, the combined impact of impact of cosmic variance on the training data and the limited flexibility of the  $P_{3D}$  model are 2.4 and 0.6% for  $P_{3D}$  and  $P_{1D}$ , respectively, which is 1.1 and 0.1% worse than the minimum accuracy we can test for these statistics. At the center of the parameter space, the accuracy of FORESTFLOW for  $P_{3D}$  and  $P_{1D}$  is only 0.2% worse than the previous levels, letting us conclude that the primary factors limiting the performance of FORESTFLOW at the center of the parameter space are the size of the training simulations and model inaccuracies.

The efficiency of FORESTFLOW across the parameter space is 1.2 and 1.0% worse than at the center for  $P_{3D}$  and  $P_{1D}$ , respectively. Consequently, its accuracy would likely improve by increasing the number of training simulations. However, leaveone-out tests significantly underestimate the performance of an



Leave-simulation-out

**Fig. 6.** Accuracy of FORESTFLOW across the input parameter space estimated via leave-simulation-out (top panels) and leave-redshift-out tests (bottom panels). Top panels. Each leave-simulation-out test involves training one independent emulator with measurements from 29 distinct simulations, and then using the measurements from the remaining simulation as the validation set. Lines and shaded areas show the average and standard deviation of 30 leave-simulation-out tests, and each panel shows the results for a different redshift. Bottom panels. Leave-redshift-out tests require optimizing one emulator with all measurements but the ones at a particular redshift, and then using measurements from this redshift as validation. Each panel shows the results of a different test.

emulator at the edges of the training set, especially for a small number of simulations, because it often requires extrapolating the emulator's predictions. We can thus conclude that the quality of the training data, the accuracy of the model, and the number of training simulations have a similar impact on the performance of FORESTFLOW. Given that leave-simulation-out tests tend to provide a lower performance bound, we conclude that the overall accuracy of FORESTFLOW in predicting  $P_{3D}$  from linear scales to  $k = 5 \text{ Mpc}^{-1}$  is approximately 3%, and  $\approx 1.5\%$  for  $P_{1D}$  down to  $k_{\parallel} = 4 \text{ Mpc}^{-1}$ .

As discussed in Sect. 2.2, FORESTFLOW does not use as input "traditional" cosmological parameters such as  $\Omega_m$ ,  $A_s$ , or  $H_0$ . Instead, it uses a set of parameters measured from the outputs of individual simulation snapshots. This strategy enables training FORESTFLOW without specifying the input redshift and making predictions for redshifts not present in the training set. To test this assumption, we carry out two leave-redshift-out tests. The first involves optimizing one emulator with all TRAINING measurements but the ones at z = 2.5, and then validating it with data from this redshift. For the second, we follow the same approach but using measurements at z = 3.5. We display the results of these tests in the bottom panels of Fig. 6. The performance of FORESTFLOW is similar for leave-redshift-out and leave-simulation-out tests, validating the approach mentioned above. We find similar results for leave-redshift-out tests at other redshifts.

Туре	P <sub>3D</sub>	$P_{1D}$	$b_{\delta}$	$b_\eta$	$P_{\rm 3D,\perp}^{\rm lin}$	$P_{\mathrm{3D},\parallel}^{\mathrm{lin}}$
C. var. fit $(a)$	0.8	0.1	0.6	1.8	1.2	1.8
C. var. data <sup>(b)</sup>	1.3	0.5	-	_	_	_
C. var. & fit <sup>(c)</sup>	2.4	0.6	-	_	_	_
Emu. center <sup>(d)</sup>	2.6	0.8	1.2	0.3	1.1	1.2
Emu. overall <sup>(e)</sup>	3.4	1.8	1.0	3.1	2.0	2.9

 Table 1. Percentage impact of different sources of uncertainty on

 FORESTFLOW predictions and overall accuracy.

**Notes.** The second and third columns show the results for  $P_{3D}$  and  $P_{1D}$  over the intervals  $0.5 < k [Mpc^{-1}] < 5$  and  $0.09 < k_{\parallel} [Mpc^{-1}] < 4$  intervals, respectively, while the last two columns do so for perpendicular and parallel  $P_{3D}$  linear modes. <sup>(a)</sup>Impact of cosmic variance on the best-fitting  $P_{3D}$  model to simulation measurements (see Appendix A). <sup>(b)</sup>Impact of cosmic variance on simulation measurements (see Appendix A). <sup>(c)</sup>Joint impact of cosmic variance and the limited flexibility of the  $P_{3D}$  model (see Sect. 3.4). <sup>(d)</sup>Performance of FORESTFLOW at the center of the parameter space, estimated using the CENTRAL simulation (see Sect. 5.1). <sup>(e)</sup>Accuracy of FORESTFLOW actors the parameter space, estimated via leave-simulation-out tests (see Sect. 5.1).

## 5.2. Other cosmologies and IGM histories

In Fig. 7, we examine the accuracy of FORESTFLOW in reproducing  $P_{3D}$  and  $P_{1D}$  measurements from simulations not included in the training set. Lines indicate the redshift average of the relative difference between model predictions and simulation measurements. The first two rows show the results for the CENTRAL and SEED simulations, whose only difference is their initial distribution of phases. Consequently, the predictions of FORESTFLOW are the same for the two simulations. As we can see, these simulations present a different large-scale pattern of fluctuations, signaling that are caused by cosmic variance. Once we ignore these, we find that the performance of FORESTFLOW is practically the same for the two simulations. We can thus conclude that FORESTFLOW predictions are largely insensitive to the impact of cosmic variance on the training set.

In the third, fourth, and fifth rows of Fig. 7, we use the GROWTH, NEUTRINOS, and CURVED simulations to evaluate the accuracy of FORESTFLOW for three different scenarios not contemplated in the training set: different growth history, massive neutrinos, and curvature. As we can see, the performance of FORESTFLOW for all these simulations is approximately the same as for the CENTRAL simulation. These results support that using the small-scale amplitude and slope of the linear power spectrum to capture cosmological information enables setting precise constraints on growth histories and ACDM extensions not included in the training set (see also Pedersen et al. 2021, 2023; Cabayol-Garcia et al. 2023).

In the last row of Fig. 7, we examine the accuracy of FOREST-FLOW for the REIONISATION simulation, which employs a He II reionization history significantly different from those used by the TRAINING simulations. The performance of FORESTFLOW for this and the CENTRAL simulation is similar, which is noteworthy given that the performance of  $P_{1D}$  emulators for the REION-ISATION simulation is significantly worse than for the CENTRAL simulation (Cabayol-Garcia et al. 2023). The outstanding performance of FORESTFLOW is likely because the relationship between IGM physics and the parameters of the  $P_{3D}$  model is more straightforward than with  $P_{1D}$  variations.

## 6. Discussion

Cosmological analyses of the Lyman- $\alpha$  forest come in two flavors: one-dimensional studies focused on small, nonlinear scales and three-dimensional analyses of large, linear scales. With FORESTFLOW, we can now consistently model Lyman- $\alpha$  correlations from nonlinear to linear scales, enabling a variety of promising analyses that we discuss next.

#### 6.1. Connecting large-scale biases with small-scale physics

Small-scale Lyman- $\alpha$  analyses use emulators to predict  $P_{1D}$  as a function of cosmology and IGM physics (e.g., Cabayol-Garcia et al. 2023), while large-scale analyses use linear or perturbation theory models to predict  $\xi_{3D}$  together with Lyman- $\alpha$  linear bias parameters that need to be marginalized over. FORESTFLOW provides a relationship between IGM physics and linear biases, enabling the use of  $P_{1D}$  studies to inform three-dimensional analyses and vice versa.

We could use FORESTFLOW to set constraints on  $b_{\delta}$  and  $b_{\eta}$  by fitting  $P_{1D}$  measurements, and then use these constraints as priors in three-dimensional studies. As a result, we would break degeneracies between Lyman- $\alpha$  linear bias parameters and cosmology, allowing us to measure the amplitude of linear density and velocity fluctuations,  $\sigma_8(z)$  and  $f\sigma_8(z)$ , rather than  $b_{\delta}\sigma_8$  and  $b_n f\sigma_8$  like in traditional Lyman- $\alpha$  forest analyses. To illustrate this application, we proceed to compare measurements of  $b_{\delta}$  and  $\beta \equiv b_{\delta}^{-1}b_{\eta}f$  from BAO analyses with FORESTFLOW predictions for these parameters based on smallscale  $P_{1D}$  analyses. The analysis of BAO in the Lyman- $\alpha$  forest from the first data release of DESI yields  $b_{\delta} = -0.108 \pm 0.005$ and  $\beta = 1.74 \pm 0.09$  at z = 2.33 (DESI Collaboration 2024). On the other hand, FORESTFLOW predicts  $b_{\delta} = -0.118$  and  $\beta = 1.57$  at z = 2.33 for a *Planck* cosmology when using as input the best-fitting constraints on IGM parameters from Table 4 of Walther et al. (2019), which were derived from high-resolution  $P_{1D}$  measurements. The constraints on IGM parameters were derived using a  $P_{1D}$  emulator trained on a suite of simulations with the same input cosmology and possibly slightly different definitions of IGM parameters relative to those used in this work. Nonetheless, FORESTFLOW predictions and DESI measurements agree at the two sigma level, encouraging this new type of study.

In the left panels of Fig. 8, we display FORESTFLOW predictions for the response of the Lyman- $\alpha$  linear biases and  $\beta$ to variations in cosmology and IGM physics. The response of  $b_{\delta}$  to these changes is strong and has a different redshift dependence for cosmology and IGM parameters; therefore, we could use FORESTFLOW to analyze  $P_{3D}$  measurements from different redshifts to further break degeneracies between  $b_{\delta}$  and  $\sigma_8$ . On the other hand, the response of  $b_{\eta}$  to these changes is weak, and it is thus challenging to use this approach to break degeneracies between  $b_{\eta}$  and  $f\sigma_8$ . Note that the response of the Lyman- $\alpha$  linear biases and  $\beta$  to  $A_s$  variations broadly agrees with measurements from simulations run while only varying  $\sigma_8$ (Arinyo-i-Prats et al. 2015).

Similarly, we could use measurements of linear bias parameters from three-dimensional analyses (du Mas des Bourboux et al. 2020; DESI Collaboration 2024) to make predictions for IGM parameters, which could be used in  $P_{1D}$  studies to break degeneracies between cosmology and IGM physics. In the right panels of Fig. 8, we display FORESTFLOW predictions for the



**Fig. 7.** Performance of FORESTFLOW in recovering  $P_{3D}$  and  $P_{1D}$  for test simulations not included in the training set. Lines and shaded areas display the average and standard deviation of the results for 11 snapshots between z = 2 and 4.5, respectively. From top to bottom, the rows show the results for the CENTRAL, SEED, GROWTH, NEUTRINOS, CURVED, and REIONISATION simulations, where the CENTRAL and SEED simulations are at the center of the input parameter space and employ the same and different initial distribution of Fourier phases as the training simulations, respectively, the GROWTH and REIONISATION simulations use a different growth and reionization history relative to those used by the TRAINING simulations, and the NEUTRINOS and CURVED simulations consider massive neutrinos and curvature. The efficiency of FORESTFLOW is approximately the same for all simulations.

response of  $P_{3D}$  and  $P_{1D}$  to variations in cosmology and IGM physics. As we can see, the response of  $P_{1D}$  to  $A_s$  and  $\bar{F}$  variations is largely scale-independent down to  $k_{\parallel} = 1 \text{ Mpc}^{-1}$  where many other effects are at play, and thus these two parameters are largely degenerated. On the other hand, this is not the case for  $P_{3D}$ ; consequently, we could use information from  $P_{3D}$  analyses to break degeneracies in  $P_{1D}$  studies. Note that the response of  $P_{3D}$  and  $P_{1D}$  to  $A_s$ ,  $\bar{F}$ , and  $\sigma_T$  variations broadly agrees with measurements from simulations run varying only one of these parameters at a time (McDonald 2003; McDonald et al. 2005).

We also observe that  $P_{3D}$  and  $P_{1D}$  respond significantly to variations in  $\Omega_{\rm M}h^2$ , suggesting that the Lyman- $\alpha$  clustering is highly sensitive to the expansion and growth history. However, we find that variations in  $A_{\rm s}$  and  $n_{\rm s}$  can absorb the changes in  $P_{3D}$  and  $P_{1D}$  to the 2% level, and completely do so for  $P_{3D}$  at the pivot scale of the cosmological parameters of FORESTFLOW,  $k_{\rm p} = 0.7 \,{\rm Mpc}^{-1}$ . Furthermore, *Planck* measured  $\Omega_{\rm M}h^2$  with 0.8% precision Planck Collaboration VI (2020), and  $A_{\rm s}$  and  $n_{\rm s}$  absorb 1% variations in  $\Omega_{\rm M}h^2$  to the  $\approx 0.4\%$  level. This result supports the approach of not considering any cosmological parameter related to variations in the expansion of growth history as input for FORESTFLOW (see also Sect. 5.2).

## 6.2. Alcock-Paczyński on mildly nonlinear scales

Thanks to the increasing precision of galaxy surveys, there is a growing interest in extracting cosmological information from increasingly smaller scales in three-dimensional analyses. An avenue to do so is to analyze anisotropies in the correlation function Alcock & Paczynski (AP test; 1979), first proposed in the context of the Lyman- $\alpha$  forest by McDonald & Miralda-Escudé (1999) and Hui et al. (1999). Recently, Cuceu et al. (2023) followed this approach to analyze Lyman- $\alpha$  forest measurements from the Sloan Digital Sky Survey (SDSS) data release 16 (DR16; Ahumada et al. 2020), yielding constraints on some cosmological parameters a factor of two tighter than those from BAO-only analyses.

This study modeled three-dimensional correlations using linear theory, which restricted the range of scales analyzed to those larger than  $25 h^{-1}$  Mpc. We could significantly extend the range of scales used in this type of analysis by modeling threedimensional correlations using FORESTFLOW. As a result, the constraining power of AP analyses would be much larger. Furthermore, we could use FORESTFLOW to extract information from  $P_{1D}$  analyses to reduce degeneracies between cosmology and the parameters describing  $\xi_{3D}$  (see Sect. 6.1).



**Fig. 8.** Response of Lyman- $\alpha$  clustering to variations in cosmology and IGM physics according to FORESTFLOW. The top, middle, and bottom panels of the left column show the results for  $\beta$ ,  $b_{\delta}$ , and  $b_{\eta}$ , respectively, while those of the right column do so for the perpendicular modes of  $P_{3D}$ , the parallel modes of  $P_{3D}$ , and  $P_{1D}$ . Blue, orange, and red lines show the response of the previous quantities to a 5% increase in  $A_s$ ,  $\Omega_M h^2$ , and  $\sigma_T$ , respectively, while green lines do so for a 1% increase in  $\bar{F}$ .

## 6.3. Extending 3D analyses to the smallest scales

The ultimate goal of FORESTFLOW is to perform a joint analysis of one- and three-dimensional measurements from small to large scales. An interesting approach to do so is to measure the Lyman- $\alpha$  forest cross-spectrum ( $P_{\times}$ ; e.g., Hui et al. 1999; Font-Ribera et al. 2018), which captures the correlation between one-dimensional Fourier modes from two neighboring quasars separated by a transverse separation ( $r_{\perp}$ ). We can model this statistic by taking the inverse Fourier transform of  $P_{3D}$  only along the perpendicular directions

$$P_{\times}(k_{\parallel}, r_{\perp}) \equiv \frac{1}{(2\pi)^2} \int d\mathbf{k}_{\perp} e^{i \, \mathbf{k}_{\perp} \cdot \mathbf{r}_{\perp}} P_{3\mathrm{D}}(k_{\parallel}, k_{\perp})$$
$$= \frac{1}{2\pi} \int_0^\infty dk_{\perp} \, k_{\perp} \, J_0(k_{\perp} r_{\perp}) \, P_{3\mathrm{D}}(k_{\parallel}, k_{\perp}). \tag{11}$$

Comparing this equation with Eq. (1), it becomes clear that  $P_{1D}$  is a special case of  $P_{\times}$ , corresponding to the limit where the transverse separation is zero.

In Sect. 3.3, we optimize the  $P_{3D}$  model to describe measurements of  $P_{3D}$  and  $P_{1D}$  from the TRAINING simulations. Then, in Sect. 4, we use the distribution of best-fitting parameters as the

training set for FORESTFLOW, which predicts the value of  $P_{3D}$  model parameters as a function of cosmology and IGM physics. Even though neither the best-fitting model nor FORESTFLOW use  $P_{\times}$  for their optimization, we can make predictions of  $P_{\times}$  for the two. To do so, we first estimate  $P_{3D}$  using the value of the model parameters using Eq. (3), and then we integrate it using Eq. (11). We carry out the integration using the fast Hankel transform algorithm FFTlog (Hamilton 2000) implemented in the hankl package (Karamanis & Beutler 2021).

We use  $P_{\times}$  measurements from the simulations described in Sect. 3.1 to evaluate the accuracy of FORESTFLOW for this statistic. We first define four bins in  $r_{\perp}$ , the transverse separation between skewers in configuration space, with edges 0.13, 0.32, 0.80, 2, and 6 Mpc. Then, we measure  $P_{\times}$  using all pairs of skewers with  $r_{\perp}$  separation within the previous bins

$$P_{\times}(r_{\perp}, k_{\parallel}) = \left\langle \Re \left[ \tilde{\delta}_{i}(k_{\parallel}) \tilde{\delta}_{j}^{*}(k_{\parallel}) \right] \right\rangle$$
(12)

where  $\tilde{\delta}_i$  and  $\tilde{\delta}_j^*$  stand for the Fourier transform of a skewer *i* and the complex conjugate of its partner *j*, respectively, the average  $\langle \rangle$  includes all possible pairs in the bin without repetition or permutation, and  $\Re$  indicates that we only use the real part



**Fig. 9.** Accuracy of the parametric model and FORESTFLOW in describing  $P_{\times}$  measurements from the CENTRAL simulation at z = 3. Dots show simulation measurements, dashed lines depict predictions from the best-fitting parametric model to  $P_{3D}$  and  $P_{1D}$  measurements, and solid lines and shaded areas display the average and 68% credible interval of FORESTFLOW predictions. The color of the lines indicates the results for different bins in transverse separation  $r_{\perp}$ . The middle panel shows the residual between simulation measurements and the best-fitting parametric model, while the bottom panel displays the residual between predictions from the parametric model and FORESTFLOW. The performance of FORESTFLOW in reproducing simulation measurements is similar to that of the best-fitting model.

of the expression between brackets because the average of the imaginary part is zero. The  $r_{\perp}$  on the left-hand side denotes the effective center of the bin, accounting for the skewed distribution of  $r_{\perp}$  within each bin: the number of skewers separated by a small distance  $dr_{\perp}$  is proportional to  $r_{\perp}$ , and therefore the effective center is not at the halfway point. To compute  $P_{\times}$  at the effective center, we perform the integration using ten sub-bins within each  $r_{\perp}$  bin and calculate the average of these weighed by  $r_{\perp}$ .

In Fig. 9, we study the performance of FORESTFLOW in reproducing  $P_{\times}$  measurements from the CENTRAL simulation at z = 3. Dots display simulation measurements, dashed lines the best-fitting model to  $P_{3D}$  and  $P_{1D}$  measurements from this simulation, and the solid lines FORESTFLOW predictions. As we can see,  $P_{\times}$  decreases as the  $r_{\perp}$  separation increases; this is because more distant sightlines are sampling increasingly uncorrelated regions. In the middle panel, we examine the accuracy of the best-fitting model in describing simulation measurements, finding that it is better than 10% throughout all the scales shown. The performance of the model improves for smaller  $r_{\perp}$  separations. This is likely because the fit's likelihood function (Eq. (7)) considers  $P_{1D}$ , which is equivalent to  $P_{\times}$  at  $r_{\perp} = 0$  separation, but not  $P_{\times}$ . The bottom panel illustrates the performance of FOREST-FLOW relative to the best-fitting model, providing an approximate assessment of its ability to reproduce the training data. FORESTFLOW achieves an accuracy better than 5.

## 7. Conclusions

We present FORESTFLOW, a novel framework for predicting Lyman- $\alpha$  clustering from linear and nonlinear scales as a function of cosmology and IGM physics. FORESTFLOW employs conditional normalizing flows to emulate the eight parameters of a physically motivated model for Lyman- $\alpha$  clustering: the two linear Lyman- $\alpha$  biases ( $b_{\delta}$  and  $b_{\eta}$ ) and six parameters that capture small-scale deviations of the three-dimensional flux power spectrum ( $P_{3D}$ ) from linear theory. By combining this model with a Boltzmann solver, FORESTFLOW predicts  $P_{3D}$  and any derived statistics, including the two-point correlation function ( $\xi_{3D}$ , the primary statistic for large-scale analyses), the one-dimensional Lyman- $\alpha$  flux power spectrum ( $P_{1D}$ , central to small-scale studies), and the cross-spectrum ( $P_{x}$ , a promising tool for full-scale analyses).

We trained the emulator using the best-fitting parameters of the physically motivated model to measurements from a suite of 30 fixed-and-paired cosmological hydrodynamical simulations spanning 11 redshifts equally spaced between z = 2 and 4.5 (Pedersen et al. 2021). Despite the moderate size of these simulations, FORESTFLOW achieves an accuracy of 3% for  $P_{3D}$ from linear scales to  $k = 5 \text{ Mpc}^{-1}$  and 1.5% for  $P_{1D}$  down to  $k_{\parallel} = 4 \text{ Mpc}^{-1}$ . FORESTFLOW uncertainties arise from three comparable sources: the size of the training simulations, their number, and the limited flexibility of the physically motivated model. However, we only evaluate the accuracy of FORESTFLOW using simulations from our suite, and its performance may vary for higher-resolution simulations or those generated with other codes.

FORESTFLOW demonstrates comparable performance for two extensions to the ACDM model – massive neutrinos and curvature – and ionization histories not included in the training set. This generalization is made possible by emulating the parameters of the physically motivated model as a function of the small-scale amplitude and slope of the linear power spectrum, the mean transmitted flux fraction, the amplitude and slope of the temperature-density relation, and the pressure smoothing scale (see Pedersen et al. 2023; Cabayol-Garcia et al. 2023).

The release of FORESTFLOW is particularly timely for Lyman- $\alpha$  forest analyses with the ongoing Dark Energy Spectroscopic Instrument (DESI) survey. As discussed in Sect. 6, FORESTFLOW enables a range of novel analyses with DESI data, including linking large- and small-scale studies and extending three-dimensional analyses to smaller scales. However, before applying FORESTFLOW to cosmological inference, it will be necessary to model the effects of various astrophysical processes on Lyman- $\alpha$  clustering, including metal contamination, damped Lyman- $\alpha$  systems, and AGN feedback.

# Data availability

FORESTFLOW and all the notebooks used to generate the plots of this paper can be found in Github<sup>10</sup>, as well as all data points

<sup>&</sup>lt;sup>10</sup> https://github.com/igmhub/ForestFlow

shown in the published graphs. The simulations utilized for training and testing the emulator are also publicly accessible<sup>11</sup>.

Acknowledgements. We thank Eric Armengaud, Roger de Belsunce, Andrei Cuceu, Vid Irsic, Ignasi Pérez-Ràfols, and Michael Walther for their useful comments and helpful discussions. We also thank the referee for his detailed comments and suggestions. JCM, LCG, ML, and AFR acknowledge support from the European Union (ERC Consolidator Grant, COSMO-LYA, grant agreement 101044612). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. AFR acknowledges financial support from the Spanish Ministry of Science and Innovation under the Ramon y Cajal program (RYC-2018-025210) and the PGC2021-123012NB-C41 project. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya. The analysis has been performed at Port d'Informació Científica (PIC); we acknowledge the support provided by PIC in granting us access to their computing resources. This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science and Technology of Mexico (CONAHCYT); the Ministry of Science and Innovation of Spain (MICINN), and by the DESI Member Institutions (https://www.desi.lbl.gov/collaborating-institutions). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. National Science Foundation, the U.S. Department of Energy, or any of the listed funding agencies. The authors are honored to be permitted to conduct scientific research on Iolkam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation.

#### References

- Abdul Karim, M. L., Armengaud, E., Mention, G., et al. 2024, JCAP, 2024, 088 Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, ApJS, 249, 3
- Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. 2019, arXiv e-prints
- [arXiv:1907.10902]
- Alcock, C., & Paczynski, B. 1979, Nature, 281, 358
- Anderson, L., Pontzen, A., Font-Ribera, A., et al. 2019, ApJ, 871, 144
- Angulo, R. E., & Pontzen, A. 2016, MNRAS, 462, L1
- Angulo, R. E., Baugh, C. M., Frenk, C. S., & Lacey, C. G. 2008, MNRAS, 383, 755
- Ansel, J., Yang, E., He, H., Gimelshein, N., & Jain, A. 2024, in ASPLOS 24 (ACM)
- Ardizzone, L., Bungert, T., Draxler, F., et al. 2018-2022, Framework for Easily Invertible Architectures (FrEIA), https://github.com/vislearn/FrEIA
- Arinyo-i-Prats, A., Miralda-Escudé, J., Viel, M., & Cen, R. 2015, JCAP, 2015, 017
- Bird, S. 2017, Astrophysics Source Code Library [record ascl:1710.012] Bird, S., Rogers, K. K., Peiris, H. V., et al. 2019, JCAP, 2019, 050
- Bird, S., Feng, Y., Pedersen, C., & Font-Ribera, A. 2020, JCAP, 2020, 002
- Bird, S., Fernandez, M., Ho, M.-F., et al. 2023, JCAP, 2023, 037
- Boera, E., Becker, G. D., Bolton, J. S., & Nasir, F. 2019, ApJ, 872, 101
- Bolton, J. S., Viel, M., Kim, T. S., Haehnelt, M. G., & Carswell, R. F. 2008, MNRAS 386 1131
- Bolton, J. S., Puchwein, E., Sijacki, D., et al. 2017, MNRAS, 464, 897
- Busca, N. G., Delubac, T., Rich, J., et al. 2013, A&A, 552, A96
- Cabayol-Garcia, L., Chaves-Montero, J., Font-Ribera, A., & Pedersen, C. 2023, MNRAS, 525, 3499
- Cen, R., Miralda-Escudé, J., Ostriker, J. P., & Rauch, M. 1994, ApJ, 437, L9
- Chabanier, S., Bournaud, F., Dubois, Y., et al. 2020, MNRAS, 495, 1825
- Chabanier, S., Emberson, J. D., Lukić, Z., et al. 2023, MNRAS, 518, 3754
- Chen, S.-F., Vlah, Z., & White, M. 2021, JCAP, 2021, 053
- Croft, R. A. C., Weinberg, D. H., Katz, N., & Hernquist, L. 1998, ApJ, 495, 44 Cuceu, A., Font-Ribera, A., Nadathur, S., Joachimi, B., & Martini, P. 2023, Phys. Rev. Lett., 130, 191003
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, AJ, 145, 10
- Dawson, K. S., Kneib, J.-P., Percival, W. J., et al. 2016, AJ, 151, 44
- 11 https://github.com/igmhub/LaCE

- de Belsunce, R., Philcox, O. H. E., Iršič, V., et al. 2024, MNRAS, 533, 3756
- DESI Collaboration (Aghamousa, A., et al.) 2016, arXiv e-prints [arXiv:1611.00036]
- DESI Collaboration (Adame, A. G., et al.) 2024, JCAP, accepted [arXiv:2404.03001]
- Dinh, L., Sohl-Dickstein, J., & Bengio, S. 2016, arXiv e-prints [arXiv:1605.08803]
- du Mas des Bourboux, H., Rich, J., Font-Ribera, A., et al. 2020, ApJ, 901, 153
- Feng, Y., Bird, S., Anderson, L., Font-Ribera, A., & Pedersen, C. 2018, https: //doi.org/10.5281/zenodo.1451799
- Fernandez, M. A., Ho, M.-F., & Bird, S. 2022, MNRAS, 517, 3200
- Font-Ribera, A., McDonald, P., & Slosar, A. 2018, JCAP, 2018, 003
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306
- Gaikwad, P., Rauch, M., Haehnelt, M. G., et al. 2020, MNRAS, 494, 5091
- Gaikwad, P., Srianand, R., Haehnelt, M. G., & Choudhury, T. R. 2021, MNRAS, 506, 4389
- Gerardi, F., Cuceu, A., Font-Ribera, A., Joachimi, B., & Lemos, P. 2023, MNRAS, 518, 2567
- Givans, J. J., & Hirata, C. M. 2020, Phys. Rev. D, 102, 023515
- Givans, J. J., Font-Ribera, A., Slosar, A., et al. 2022, JCAP, 2022, 070
- Gnedin, N. Y., & Hui, L. 1998, MNRAS, 296, 44
- Haardt, F., & Madau, P. 2012, ApJ, 746, 125
- Hamilton, A. J. S. 2000, MNRAS, 312, 257
- Hastie, T., Tibshirani, R., & Friedman, J. 2001, in The Elements of Statistical Learning, (New York, NY, USA: Springer New York Inc.), Springer Series in Statistics
- Horowitz, B., de Belsunce, R., & Lukić, Z. 2025, MNRAS, 536, 845
- Huang, C. W., Krueger, D., Lacoste, A., & Courville, A. 2018, arXiv e-prints [arXiv:1804.00779]
- Hui, L., Stebbins, A., & Burles, S. 1999, ApJ, 511, L5
- Iršič, V., & McQuinn, M. 2018, JCAP, 2018, 026
- Iršič, V., Viel, M., Haehnelt, M. G., Bolton, J. S., & Becker, G. D. 2017, Phys. Rev. Lett., 119, 031302
- Iršič, V., Viel, M., Haehnelt, M. G., et al. 2024, Phys. Rev. D, 109, 043511
- Ivanov, M. M. 2024, Phys. Rev. D, 109, 023507
- Jacobus, C., Harrington, P., & Lukić, Z. 2023, ApJ, 958, 21
- Jimenez Rezende, D., & Mohamed, S. 2015, arXiv e-prints [arXiv:1505.05770] Kaiser, N. 1987, MNRAS, 227, 1
- Karaçaylı, N. G., Martini, P., Guy, J., et al. 2024, MNRAS, 528, 3941
- Karamanis, M., & Beutler, F. 2021, arXiv e-prints [arXiv:2106.06331]
- Khan, N. K., Kulkarni, G., Bolton, J. S., et al. 2024, MNRAS, 530, 4920
- Kingma, D. P., & Ba, J. 2014, arXiv e-prints [arXiv:1412.6980]
- Kokron, N., Chen, S.-F., White, M., DeRose, J., & Maus, M. 2022, JCAP, 2022, 059
- Lee, K.-G., Hennawi, J. F., Spergel, D. N., et al. 2015, ApJ, 799, 196
- Lesgourgues, J. 2011, arXiv e-prints [arXiv:1104.2932]
- Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473
- Lukić, Z., Stark, C. W., Nugent, P., et al. 2015, MNRAS, 446, 3697
- MacKay, D. J. C., et al. 1998, in Neural Networks and Machine Learning, ed.
- McCulloch, W. S., & Pitts, W. 1943, Bull. Math. Biophys., 5, 115

- McDonald, P., Seljak, U., Burles, S., et al. 2006, ApJS, 163, 80
- McKay, M. D., Beckman, R. J., & Conover, W. J. 1979, Technometrics, 21, 239 McQuinn, M. 2016, ARA&A, 54, 313
- Meiksin, A. A. 2009, Rev. Mod. Phys., 81, 1405
- Meiksin, A., Bryan, G., & Machacek, M. 2001, MNRAS, 327, 296
- Miralda-Escudé, J., Cen, R., Ostriker, J. P., & Rauch, M. 1996, ApJ, 471, 582
- Molaro, M., Iršič, V., Bolton, J. S., et al. 2023, MNRAS, 521, 1489
- Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2017, ApJ, 837, 106
- Palanque-Delabrouille, N., Yèche, C., Lesgourgues, J., et al. 2015, JCAP, 2015, 045
- Palanque-Delabrouille, N., Yèche, C., Schöneberg, N., et al. 2020, JCAP, 2020, 038
- Papamakarios, G., Nalisnick, E., Jimenez Rezende, D., Mohamed, S., & Lakshminarayanan, B. 2019, arXiv e-prints [arXiv:1912.02762]
- Pedersen, C., Font-Ribera, A., Rogers, K. K., et al. 2021, JCAP, 2021, 033
- Pedersen, C., Font-Ribera, A., & Gnedin, N. Y. 2023, ApJ, 944, 223
- Planck Collaboration VI. 2020, A&A, 641, A6
- Pontzen, A., Slosar, A., Roth, N., & Peiris, H. V. 2016, Phys. Rev. D, 93, 103519
- Puchwein, E., Haardt, F., Haehnelt, M. G., & Madau, P. 2019, MNRAS, 485, 47
- Puchwein, E., Bolton, J. S., Keating, L. C., et al. 2023, MNRAS, 519, 6162

- C. M. Bishop, NATO ASI Series (Berlin: Springer), 168, 133 Maion, F., Angulo, R. E., & Zennaro, M. 2022, JCAP, 2022, 036
- McDonald, P. 2003, ApJ, 585, 34
  - McDonald, P., & Miralda-Escudé, J. 1999, ApJ, 518, 24
  - McDonald, P., Miralda-Escudé, J., Rauch, M., et al. 2000, ApJ, 543, 1
  - McDonald, P., Seljak, U., Cen, R., et al. 2005, ApJ, 635, 761

- Ramachandra, N., Chaves-Montero, J., Alarcon, A., et al. 2022, MNRAS, 515, 1927
- Ravoux, C., Abdul Karim, M. L., Armengaud, E., et al. 2023, MNRAS, 526, 5118
- Rogers, K. K., & Peiris, H. V. 2021a, Phys. Rev. Lett., 126, 071302
- Rogers, K. K., & Peiris, H. V. 2021b, Phys. Rev. D, 103, 043526
- Rogers, K. K., Peiris, H. V., Pontzen, A., et al. 2019, JCAP, 2019, 031
- Sacks, J., Welch, W. J., Mitchell, T. J., & Wynn, H. P. 1989, Stat. Sci., 4, 409
- Seljak, U., Makarov, A., McDonald, P., et al. 2005, Phys. Rev. D, 71, 103515
- Seljak, U., Makarov, A., McDonald, P., & Trac, H. 2006a, Phys. Rev. Lett., 97, 191303
- Seljak, U., Slosar, A., & McDonald, P. 2006b, JCAP, 2006, 014
- Slosar, A., Font-Ribera, A., Pieri, M. M., et al. 2011, JCAP, 2011, 001
- Slosar, A., Iršič, V., Kirkby, D., et al. 2013, JCAP, 2013, 026
- Spergel, D. N., Verde, L., Peiris, H. V., et al. 2003, ApJS, 148, 175
- Springel, V. 2005, MNRAS, 364, 1105
- Takhtaganov, T., Lukić, Z., Müller, J., & Morozov, D. 2021, ApJ, 906, 74
- Verde, L., Peiris, H. V., Spergel, D. N., et al. 2003, ApJS, 148, 195
- Viel, M., & Haehnelt, M. G. 2006, MNRAS, 365, 231
- Viel, M., Weller, J., & Haehnelt, M. G. 2004, MNRAS, 355, L23
- Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013, Phys. Rev. D, 88, 043502
- Villaescusa-Navarro, F., Naess, S., Genel, S., et al. 2018, ApJ, 867, 137
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
- Walther, M., Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2019, ApJ, 872, 13
- Walther, M., Armengaud, E., Ravoux, C., et al. 2021, JCAP, 2021, 059
- Winkler, C., Worrall, D., Hoogeboom, E., & Welling, M. 2019, arXiv e-prints [arXiv:1912.00042]

Zaldarriaga, M., Hui, L., & Tegmark, M. 2001, ApJ, 557, 519

- <sup>1</sup> Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, 08193 Bellaterra (Barcelona), Spain
- <sup>2</sup> Port d'Informació Científica, Campus UAB, C. Albareda s/n, 08193 Bellaterra (Barcelona), Spain
- <sup>3</sup> Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA
- <sup>4</sup> Physics Dept., Boston University, 590 Commonwealth Avenue, Boston, MA 02215, USA
- <sup>5</sup> Dipartimento di Fisica "Aldo Pontremoli", Università degli Studi di Milano, Via Celoria 16, I-20133 Milano, Italy
- <sup>6</sup> Department of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, UK
- <sup>7</sup> Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK

- <sup>8</sup> Instituto de Física, Universidad Nacional Autónoma de México, Cd. de México C.P. 04510, Mexico
- <sup>9</sup> University of California, Berkeley, 110 Sproul Hall #5800, Berkeley, CA 94720, USA
- <sup>10</sup> Departamento de Física, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio Ip, CP 111711 Bogotá, Colombia
- <sup>11</sup> Observatorio Astronómico, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia
- <sup>12</sup> Institut d'Estudis Espacials de Catalunya (IEEC), 08034 Barcelona, Spain
- <sup>13</sup> Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth PO1 3FX, UK
- <sup>14</sup> Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, USA
- <sup>15</sup> Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA
- <sup>16</sup> Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA
- <sup>17</sup> The Ohio State University, Columbus 43210, OH, USA
- <sup>18</sup> Department of Physics, Southern Methodist University, 3215 Daniel Avenue, Dallas, TX 75275, USA
- <sup>19</sup> Department of Physics and Astronomy, University of California, Irvine 92697, USA
- <sup>20</sup> Departament de Física, Serra Húnter, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain
- <sup>21</sup> Department of Astronomy, The Ohio State University, 4055 McPherson Laboratory, 140 W 18th Avenue, Columbus, OH 43210, USA
- <sup>22</sup> Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain
- <sup>23</sup> Departamento de Física, Universidad de Guanajuato DCI, C.P. 37150 Leon, Guanajuato, Mexico
- <sup>24</sup> Instituto Avanzado de Cosmología A. C., San Marcos 11 Atenas 202. Magdalena Contreras, 10720 Ciudad de México, Mexico
- <sup>25</sup> Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 10, 08930 Barcelona, Spain
- <sup>26</sup> Department of Physics and Astronomy, Sejong University, Seoul 143-747, Korea
- <sup>27</sup> CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain
- <sup>28</sup> Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA
- <sup>29</sup> NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA



**Fig. A.1.** Impact of cosmic variance on  $P_{3D}$  (top panel) and  $P_{1D}$  (bottom panel) measurements from our simulations at z = 3. Lines show the difference between measurements from the CENTRAL and SEED simulations, which only differ on their initial distribution of Fourier phases, divided by  $\sqrt{2}$  times their average. Cosmic variance induces errors as large as 10% on  $P_{3D}$  for  $k \approx 0.3$  Mpc<sup>-1</sup>, while these are on the order of 1% for  $P_{1D}$ .

## Appendix A: Cosmic variance

Throughout this work, we trained and tested FORESTFLOW using simulations run employing the "fixed-and-paired" technique (Angulo & Pontzen 2016; Pontzen et al. 2016), which significantly reduces cosmic variance for the clustering of the Lyman- $\alpha$  forest (Anderson et al. 2019). We could further mitigate the impact of cosmic variance by using control variates (Kokron et al. 2022), but this is outside the scope of the current work. The impact of cosmic variance on fixed-and-paired simulations is not straightforward (Maion et al. 2022), and thus we would ideally use multiple fixed-and-paired simulations with different initial distributions of Fourier phases to estimate the precision of measurements from our simulations. However, we only have two simulations with these properties: CENTRAL and SEED. In this section, we use these two simulations to estimate the impact of cosmic variance on simulation measurements and best-fitting models. It is crucial to acknowledge that our findings are subject to significant noise because we only have access to two independent realizations.

In Fig. A.1, we show the difference between measurements from the CENTRAL and SEED simulations at z = 3, normalized by  $\sqrt{2}$  times their average<sup>12</sup>. The CENTRAL and SEED simulations differ only in their initial Fourier phase distributions, allowing their difference to isolate the effects of cosmic variance. Unlike traditional simulations, where cosmic variance for  $P_{3D}$ scales inversely with the square root of the number of modes, this uncertainty peaks at  $\simeq 10\%$  around  $k \simeq 0.3 \,\mathrm{Mpc}^{-1}$  and decreases at both larger and smaller scales. This behavior can be explained as follows: at the largest scales, the fix-and-paired technique cancels cosmic variance for linear density modes, reducing variance. At intermediate scales, nonlinear evolution, particularly mode coupling, reintroduces cosmic variance, increasing the uncertainty. At smaller scales, the increasing number of modes leads to a decrease in cosmic variance, similar to trends observed in traditional simulations.



**Fig. A.2.** Impact of cosmic variance on predictions from the parametric model. Lines show the difference between the best-fitting models to  $P_{3D}$  and  $P_{1D}$  measurements from the CENTRAL and SEED simulations, divided by  $\sqrt{2}$  times the best-fitting model to their average. The top panel shows the results for the Lyman- $\alpha$  linear biases ( $b_{\delta}$  and  $b_{\eta}$ ), while the middle and bottom panels display the results for  $P_{3D}$  and  $P_{1D}$  at z = 3, respectively. The impact of cosmic variance on model predictions is approximately an order of magnitude smaller than on simulation measurements (see Fig. A.1).

Thus, cosmic variance limits our ability to accurately evaluate both the parametric model and FORESTFLOW using our simulations. To mitigate its impact, we restricted the analysis of  $P_{3D}$  performance to scales  $k > 0.5 \text{ Mpc}^{-1}$  in the main results. In contrast, cosmic variance has a much smaller effect on  $P_{1D}$ , contributing only about 1.5% uncertainty at  $k_{\parallel} < 2 \text{ Mpc}^{-1}$ . This smaller effect allowed us to include all scales in  $P_{1D}$  tests without concern. To more precisely quantify the impact of cosmic variance on simulation measurements, we computed the standard deviation of the results shown in Fig. A.1 across redshift. We did it within the intervals  $0.5 < k[\text{ Mpc}^{-1}] < 5$  for  $P_{3D}$  and  $0.09 < k_{\parallel}[\text{ Mpc}^{-1}] < 4$  for  $P_{1D}$ , based on the scales discussed earlier and those used for fitting the  $P_{3D}$  model in Sect. 3.3. We found that the average impact of cosmic variance is 1.3% for  $P_{3D}$ and 0.5% for  $P_{1D}$ .

We expect the impact of cosmic variance on the best-fitting model to  $P_{3D}$  and  $P_{1D}$  measurements to be weaker than on simulation measurements, as multiple  $P_{3D}$  and  $P_{1D}$  bins collectively constrain the eight parameters of the model. In Fig. A.2, we show the difference between the best-fitting models to the CENTRAL and SEED simulations, normalized by  $\sqrt{2}$  times the best-fitting model to their average. The top panel displays results for the two Lyman- $\alpha$  linear biases,  $b_{\delta}$  and  $b_{\eta}$ . The standard deviation of the differences is 0.6% and 1.8% for  $b_{\delta}$  and  $b_{\eta}$ , respectively, demonstrating that we can measure the two Lyman- $\alpha$  linear biases with percent-level accuracy using our simulations. This precision is achievable because the combination of small and large scales in

 $<sup>^{12}</sup>$  The factor  $\sqrt{2}$  accounts for the noise estimate from a single simulation.



**Fig. B.1.** PIT distribution for  $P_{1D}$  (blue) and  $P_{3D}$  (red). This plot validates the uncertainties predicted by FORESTFLOW across TRAINING simulations via a leave-simulation-out approach. The PIT distribution is approximately uniform, indicating well-calibrated uncertainties for most samples, while the peaks at the edges indicate underestimated uncertainties for some samples.

the fits helps break degeneracies between the two linear biases and the other six model parameters.

By propagating these uncertainties to the behavior of  $P_{3D}$  on linear scales, we find that the impact of cosmic variance on perpendicular and parallel modes is 1.2% and 1.8%, respectively. In the middle and bottom panels of Fig. A.2, we analyze the influence of cosmic variance on model predictions for  $P_{3D}$  and  $P_{1D}$ . The overall impact of cosmic variance on  $P_{3D}$  and  $P_{1D}$  predictions is 0.8% and 0.1%, respectively, confirming that the bestfitting model is less sensitive to cosmic variance than individual simulation measurements. Consequently, FORESTFLOW is more robust against cosmic variance than emulators that predict the power spectrum at fixed *k*-bins.

# Appendix B: Validation of uncertainties predicted for $P_{3D}$ and $P_{1D}$

Normalizing flows predict the full posterior distribution of the target data rather than only their mean like fully connected neural networks or their mean and width like Mixture Density Networks (see Ramachandra et al. 2022; Cabayol-Garcia et al. 2023, for some applications in cosmology). This is achieved through multiple sampling iterations from the target latent distribution, an eight-dimensional Gaussian in our case. In FOREST-FLOW, each sampled realization of the  $P_{3D}$  model parameters is propagated to generate predictions for  $P_{3D}$  and  $P_{1D}$  (see Sect. 4.1), producing a covariance matrix for these statistics. In this appendix, we validate its diagonal elements. Note that well-calibrated uncertainties are critical for future uses of FOREST-FLOW such as cosmology inference.

We validated the uncertainty in  $P_{3D}$  and  $P_{1D}$  predictions using the Probability Integral Transform test (PIT), which is the value of the cumulative distribution function (CDF) of a distribution evaluated at the ground-truth value  $z_t$ 

PIT = CDF[p, 
$$z_t$$
] =  $\int_{-\infty}^{z_t} p(z)dz$ , (B.1)  
A187, page 18 of 18

where p is in our case the distribution of FORESTFLOW predictions for  $P_{3D}$  or  $P_{1D}$  and  $z_t$  stands for measurements of these statistics from the simulations. A model that displays a wellcalibrated uncertainty distribution yields PIT values that are uniformly distributed between zero and one. This indicates that the observed outcomes have an equal likelihood of falling at any point along the predicted CDF. In contrast, an excess of values close to zero or one indicates that the width of the distribution is underestimated.

In Fig. B.1, we display a PIT test produced using all the TRAINING simulations via a leave-simulation-out approach (see Sect. 5.1). This process validates average predictions and uncertainties against simulations excluded in the training process. The red and blue lines display the results for  $P_{3D}$  and  $P_{1D}$ , respectively, which were generated by combining results from different scales and redshifts. As we can see, the PIT distribution is approximately uniform for the two statistics but it presents peaks at the low and high ends, indicating underestimated uncertainties for some samples. The cause behind this feature is unclear and it demands further investigation beyond the scope of this project.