

Lawrence Berkeley National Laboratory

LBL Publications

Title

An emulator for the Lyman- α forest in beyond- Λ CDM cosmologies

Permalink

<https://escholarship.org/uc/item/96d469w8>

Journal

Journal of Cosmology and Astroparticle Physics, 2021(05)

ISSN

1475-7516

Authors

Pedersen, Christian
Font-Ribera, Andreu
Rogers, Keir K
[et al.](#)

Publication Date

2021-05-01

DOI

10.1088/1475-7516/2021/05/033


Peer reviewed

PAPER • OPEN ACCESS

An emulator for the Lyman- α forest in beyond- Λ CDM cosmologies

To cite this article: Christian Pedersen *et al* JCAP05(2021)033

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



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

An emulator for the Lyman- α forest in beyond- Λ CDM cosmologies

Christian Pedersen,^{a,1} Andreu Font-Ribera,^{b,a} Keir K. Rogers,^{c,d}
Patrick McDonald,^e Hiranya V. Peiris,^{a,c} Andrew Pontzen^a
and Anže Slosar^f

^aDepartment of Physics and Astronomy, University College London,
Gower Street, London, United Kingdom

^bInstitut de Física d'Altes Energies (IFAE),
The Barcelona Institute of Science and Technology,
08193 Bellaterra (Barcelona), Spain

^cOskar Klein Centre for Cosmoparticle Physics, Department of Physics,
Stockholm University,
AlbaNova University Center, Stockholm 10691, Sweden

^dDunlap Institute for Astronomy & Astrophysics, University of Toronto,
50 St. George Street, Toronto, ON M5S 3H4, Canada

^eLawrence Berkeley National Laboratory,
One Cyclotron Road, Berkeley, CA 94720, U.S.A.

^fPhysics Department, Brookhaven National Laboratory,
Upton, NY 11973, U.S.A.

E-mail: christian.pedersen.17@ucl.ac.uk, afont@ifae.es, keir.rogers@utoronto.ca,
pvmcdonald@lbl.gov, h.peiris@ucl.ac.uk, a.pontzen@ucl.ac.uk, anze@bnl.gov

Received December 15, 2020

Revised March 24, 2021

Accepted April 14, 2021

Published May 13, 2021

Abstract. Interpreting observations of the Lyman- α forest flux power spectrum requires interpolation between a small number of expensive simulations. We present a Gaussian process emulator modelling the 1D flux power spectrum as a function of the amplitude and slope of the small-scale linear matter power spectrum, and the state of the intergalactic medium at the epoch of interest ($2 < z < 4$). This parameterisation enables the prediction of the flux power spectrum in extended cosmological models that are not explicitly included in the training set, eliminating the need to construct bespoke emulators for a number of extensions to Λ CDM. Our emulator is appropriate for cosmologies in which the linear matter power spectrum is described to percent level accuracy by just an amplitude and slope across the epoch of interest, and in the regime probed by eBOSS/DESI data. We demonstrate this for

¹Corresponding author.



massive neutrino cosmologies, where the emulator is able to predict the flux power spectrum in a $\Sigma m_\nu = 0.3 \text{ eV}$ neutrino cosmology to sub-percent accuracy, without including massive neutrinos in the training simulations. Further parameters would be required to describe models with sharp features in the linear power, such as warm or light axion dark matter. This work will facilitate the combination of upcoming DESI data with observations of the cosmic microwave background, to obtain constraints on neutrino mass and other extensions to ΛCDM cosmology.

Keywords: cosmological parameters from LSS, cosmological simulations, Lyman alpha forest, neutrino masses from cosmology

ArXiv ePrint: [2011.15127](https://arxiv.org/abs/2011.15127)

Contents

1	Introduction	1
2	Parameterisation	3
2.1	Linear matter power spectrum	3
2.2	Intergalactic medium	5
3	Simulations	6
4	Gaussian process emulator	8
4.1	Emulator validation	9
4.2	Predictions at a redshift without training data	11
4.3	Predicting extended models	12
5	Parameter constraints	13
5.1	Likelihood parameters	14
5.2	Likelihood function	15
5.3	MCMC results	15
5.4	Extended models	15
6	Discussion	16
A	Parameter dependence of the P_{1D}	19
B	IGM evolution	19
C	Full posterior with IGM parameters	21

1 Introduction

The Lyman- α forest (Ly α forest) is a series of absorption features in the spectra of high-redshift quasars, which is caused by the presence of intervening neutral hydrogen. During the last decade, the Baryon Oscillation Spectroscopic Survey (BOSS, [1]) and its extension eBOSS [2] have obtained over 200,000 quasar spectra. These surveys were able to measure the 3D correlation of absorption features and detect baryon acoustic oscillations (BAO) around $z = 2.3$ [3]. In addition, the same dataset has been used to provide precise measurements of the flux power spectrum along the line of sight, P_{1D} [4, 5]. In this paper we present a framework to interpret these P_{1D} measurements and obtain robust constraints on cosmological parameters. Starting in 2021, the Dark Energy Spectroscopic Instrument (DESI, [6]) will observe 700,000 Ly α forest quasar spectra, and will provide the most precise measurement of the P_{1D} to date. P_{1D} measurements are a unique probe of the clustering of matter on megaparsec scales [7–13]. In particular, measurements from BOSS, eBOSS and DESI¹ are sensitive to the clustering of matter on scales of $0.2 \text{ Mpc}^{-1} < k < 3 \text{ Mpc}^{-1}$ [22]. This

¹ P_{1D} measurements from high-resolution spectra of $z > 4$ quasars have been used to study the reionisation epoch [14–16] and to constrain dark matter models [17–21]. We do not discuss these measurements here.

information is highly complementary to other probes, such as observations of the cosmic microwave background (CMB), and combined analyses have provided some of the tightest constraints on the shape of the primordial power spectrum and upper limits on the neutrino mass scale [23–32].

Obtaining cosmological constraints from P_{1D} measurements requires theoretical modelling. Perturbative approaches are not sufficiently accurate on the small scales to which the P_{1D} is sensitive. Moreover, the absorption features are also sensitive to the state of the intergalactic medium (IGM), which is dependent on baryonic physics and the ultra-violet background [33]. Modelling these effects simultaneously requires expensive hydrodynamical simulations, each costing $\sim 10^5$ CPU hours, which means, in practice, that only a small number ($\lesssim 100$) of simulations can be run for a given analysis. This is far fewer than the $\sim 10^6$ likelihood evaluations necessary for common sampling techniques, e.g., Markov chain Monte Carlo methods (MCMC), to provide robust statistical constraints on parameters. A solution to this problem has been to construct frameworks that interpolate between simulations. Similar problems exist in other areas of cosmology, such as in analysis of galaxy surveys, where Gaussian processes have been used to interpolate probabilistically (or *emulate*) the galaxy power spectrum [34] and correlation function [35] as functions of cosmology from a small number of expensive simulations. Analysis of the P_{1D} over the last decade most commonly used quadratic polynomial interpolation [36, 37]; however recent work has demonstrated the advantages of using a Gaussian process emulator to perform the interpolation [15, 20, 21, 38–41].

Historically, approaches to combining small-scale clustering information from the Ly α forest with CMB observations fall into two categories. Early work on the Ly α forest recognised that the Ly α forest probes the universe at an epoch that is close to an Einstein de-Sitter (EdS) model, and therefore considered the cosmological information in the Ly α forest to be concentrated in the linear matter power spectrum [7, 9–11]. The procedure was then to constrain the $z \sim 3$, linear matter power spectrum on megaparsec scales from Ly α forest observations [8, 9, 12, 13], and combine these measurements with the CMB to constrain cosmological parameters [23–28]. More recent analyses [30–32, 36, 37, 42] have instead modelled the P_{1D} as a function of Λ CDM parameters, plus neutrino mass in the case of studies into massive neutrinos.

The work we present here returns to the original approach. The motivation for this is twofold. First, the effects of Λ CDM parameters on the Ly α forest are strongly degenerate with one another. Minimising parameter degeneracies is important at all stages of the analysis. Degeneracies can be artificially broken by interpolation errors, and lower dimensional parameter spaces allow for a more efficient distribution of simulations. Second, there are several well-motivated extensions to Λ CDM that will further characterise physical properties of the universe, such as the number and mass scale of the neutrinos, or running of the spectral index (see section 7 of ref. [43] for further examples). When combined with CMB observations, the small-scale clustering information provided by the P_{1D} provides constraints on these and many other extensions to Λ CDM. However, given the large number of prospective parameter spaces that one would want to explore, it is impractical to construct an emulator for each. A potential solution is to emulate the P_{1D} in a parameter space that can simultaneously describe the effect of these extended models on the flux power spectrum.

In this work we investigate this idea by constructing a Gaussian process emulator modelling the P_{1D} with just two parameters in the cosmological sector: the amplitude and slope of the linear matter power spectrum at the epoch of interest. We demonstrate that this

emulator can accurately predict the P_{1D} in massive neutrino cosmologies without including massive neutrinos in the training simulations.

As we will show in section 2, parameterising the linear power spectrum in terms of just an amplitude and slope well describes two commonly investigated extensions to Λ CDM, neutrino mass [44, 45] and curved cosmologies [43], over the relevant range of length scales. This is because these cosmologies induce smooth variations in the small-scale linear power spectrum. Many proposed alternatives to cold dark matter (such as light axion [46, 47], self-interacting [48] or sterile neutrino [49, 50] dark matter) feature a sharp suppression of the linear matter power spectrum on small scales, which this parameterisation does not describe. High resolution P_{1D} measurements [16, 51] have been used to put bounds on properties of these models by constraining this cutoff scale [17–19, 21, 32, 38]. Recent results bound the cutoff scale at wavenumbers higher than those probed by the P_{1D} measured from BOSS/eBOSS/DESI, which is the focus of this paper. Therefore, while in principle our parameterisation of the linear power spectrum can be extended to include a cutoff scale, we leave this investigation for future work.

The outline of this paper is as follows: in section 2, we describe the parameterisation of the emulator in terms of the linear matter power spectrum and state of the IGM. Section 3 describes the suite of simulations, and the post-processing steps. In section 4, we detail the construction of the Gaussian process emulator and validate the predictions on test data. In section 5, we run MCMC chains to obtain constraints on the primordial power spectrum from simulated mock data. We discuss our results in section 6.

2 Parameterisation

In this section we describe the parameterisation of our emulator. The absorption features that constitute the Ly α forest are caused by the presence of intervening neutral hydrogen along the lines of sight to quasars. The density of this neutral hydrogen is determined both by fluctuations in the baryon density which trace the underlying dark matter field, and the thermal and ionisation state of the IGM. We use a total of six parameters to describe these effects: two to describe the clustering of matter and four to characterise uncertainty in the state of the IGM. These are discussed in sections 2.1 and 2.2 respectively.

2.1 Linear matter power spectrum

The Ly α forest is sensitive to the matter density field. We parameterise the P_{1D} with the amplitude $\Delta_p^2(z)$ and slope $n_p(z)$ of the combined cold dark matter + baryon linear power spectrum,² P_{lin} , defined as:

$$\Delta_p^2(z) = k^3 P_{\text{lin}}(k, z)|_{k=k_p}, \quad (2.1)$$

$$n_p(z) = \text{dln}P_{\text{lin}}(k, z)/\text{dln}k|_{k=k_p}, \quad (2.2)$$

where we use a pivot scale $k_p = 0.7 \text{ Mpc}^{-1}$ which was used in the measurement of the small-scale linear matter power spectrum of ref. [13]. To numerically evaluate Δ_p^2 and n_p , we use CAMB to generate a linear matter power spectrum and fit a second order log polynomial over the range $0.5k_p < k < 2k_p$. We do not include parameters to describe changes to the growth rate or expansion history. This choice is motivated by the fact that in the redshift range

²Note that in the case of cosmologies with massive neutrinos, we do not include the clustering of the neutrinos in this quantity.




Figure 1. Information compression applied to two extended- Λ CDM parameter spaces. For 100 randomly sampled points in the *Planck* chain, we plot the ratio of the $z = 3$ linear matter power spectrum for each model with the baseline Λ CDM best fit model [43]. In panel (a), we sample from a chain with free Σm_ν . In panel (b), for each of these models we rescale the amplitude and slope of the linear matter power spectrum to match at $k_p = 0.7 \text{ Mpc}^{-1}$, shown in the vertical dashed line. The colour of each line represents the value of Σm_ν in that model. In the bottom panels we repeat the procedure for a chain with free Ω_k . Panel (c) shows the ratio for 100 randomly chosen samples, and in panel (d) we have rescaled the slope and amplitude to match at the vertical dashed line. The vertical shaded area represents the length scales to which the P_{1D} is sensitive in BOSS data [22], and the horizontal gray band represents the region of 1% agreement.

that gives rise to the Ly α forest ($5 > z > 2$), the universe is very close to Einstein-de Sitter (EdS) with $\Omega_m(z = 2) \approx 0.91$ and $\Omega_m(z = 5) \approx 0.99$ for typical cosmological models allowed by the CMB. In an EdS universe, the expansion rate evolves $H(z) = \dot{a}/a \propto (1+z)^{3/2}$, and the growth factor evolves proportional to a . In particular, for the expansion history, there is a 1.2% difference in $H(z = 2)$ between flat cosmologies with $h = 0.67$ and $h = 0.74$ at fixed physical density $\omega_m = \Omega_m h^2$, which is the range of the current tension between early and late universe constraints on H_0 [52]. This difference drops to 0.5% at $z = 4$ as we move deeper into the EdS regime. Similarly, when we compare the logarithmic growth rates for cosmologies with $\Omega_m = 0.25$ and $\Omega_m = 0.35$, we see a 2.1% difference at $z = 2$ and a 0.4% difference at $z = 4$. We construct our emulator under the assumption that the relationship between the matter power spectrum and the P_{1D} is insensitive to these slight changes in expansion history and growth rate, and we consider the results shown in sections 4 and 5 a test of this assumption.

Figure 1 demonstrates why we consider just two parameters describing a slope and an amplitude to be sufficient to parameterise the linear matter power spectrum within this regime. In panel (a), we randomly sample 100 points from the *Planck* chains³ [43] with free Σm_ν , and plot the ratio of the linear matter power spectrum at $z = 3$ with the baseline Λ CDM

³https://wiki.cosmos.esa.int/planck-legacy-archive/index.php/Cosmological_Parameters.

best fit model, which we denote as $P_{\text{lin}}^{\text{best fit}}$. The colour of each line indicates the neutrino mass value for each model. We see the expected behaviour where cosmologies with a higher neutrino mass tend to lead to a suppression of the power spectrum on small scales at late times. In panel (b), we take the same 100 samples and rescale the amplitude and slope of the power spectrum in order to match the small-scale linear power in $P_{\text{lin}}^{\text{best fit}}$ as parameterised by equations (2.1) and (2.2). The vertical shaded region shows the approximate length scales to which the P_{1D} is sensitive in BOSS data [22]. The upcoming DESI results will provide an improved measurement of the P_{1D} over the same range of wavenumbers, so the arguments presented in this paper will continue to apply. The horizontal band shows the region of 1% agreement. The effects on the linear matter power spectrum of variations in the 7 free parameters are captured by changes to the slope and amplitude to within 1% (and often much better) when considering only the length scales within the shaded region. This is compared with the most recent constraints on the amplitude of the small-scale linear power spectrum from BOSS which have an uncertainty of 6% [5]. In the bottom panels (c) and (d) we repeat the procedure for a chain with free curvature parameter Ω_k instead of Σm_ν . Again we see that the all the variations in cosmology are degenerate with changing just the slope and amplitude of the power spectrum on the scales of interest.

A natural extension to Δ_p^2 and n_p would be to include a parameter describing the second order derivative in the linear power around k_p (a running of the small-scale linear power spectrum). The necessity of this parameter would be determined by the variation in the running within the prior range of the cosmological models under investigation, and the k range and precision of the observational dataset. It would also be possible to include an extra parameter describing the growth rate if deemed necessary. We leave such extensions to future work.

2.2 Intergalactic medium

The neutral hydrogen density is determined both by the density of hydrogen (dependent upon the clustering as described in the previous section, as well as the mean baryon density ω_b) and the ionisation state, which is itself affected by several astrophysical factors. A primary factor is the intensity of the ionising metagalactic UV background (UVB), which is highly uncertain. Assuming ionisation equilibrium this is balanced with the rate of recombination, which is dependent on the gas temperature. We parameterise these effects in terms of an effective optical depth, τ_{eff} , which is related to the mean transmitted flux fraction $\langle F \rangle$ by $\tau_{\text{eff}} = -\ln\langle F \rangle$.

The thermal state of the gas around the mean cosmic density regions that are probed by the Ly α forest is well approximated by a power-law:

$$T(\Delta_b) = T_0 \Delta_b^{\gamma-1}, \quad (2.3)$$

where $\Delta_b = \rho_b/\bar{\rho}_b$ is the baryon overdensity, T_0 is the gas temperature in K at mean density [53] and γ is a free parameter. The effect of T_0 on the recombination rate is captured by $\langle F \rangle$, however thermal motion of the gas particles also induces a Doppler broadening of the absorption features. This has the effect of smoothing the P_{1D} on small scales. We describe this with a thermal broadening scale

$$\sigma_T = 9.1 \sqrt{\frac{T_0}{10000} \frac{1+z}{H(z)}}, \quad (2.4)$$

where the factor $(1+z)/H(z)$ is required to convert the broadening scale from km/s to Mpc (to match the units in which we store the P_{1D}). We include the parameters σ_T and γ in order to characterise the effects of gas temperature on the P_{1D} .

On small scales the baryons are supported by gas pressure, leading to a small-scale suppression in the baryon power spectrum which therefore affects the Ly α forest. Unlike the thermal broadening, this is a 3D effect that also affects correlations perpendicular to the line of sight. Following ref. [54] we parameterise this effect with a pressure smoothing scale, k_F , in units of Mpc^{-1} .

3 Simulations

To create a set of training data for constructing the emulator we present in section 4, we run a suite of 30 hydrodynamical simulations using the `TreeSPH` code `MP-Gadget`⁴ [55], a variant of the widely used `Gadget-2` [56] with improved parallelisation and gravitational force solver. The speed of these simulations is increased using the `QuickLymanAlpha` option which turns regions of baryon overdensity exceeding 10^3 and with temperatures $T < 10^5$ K into stars. These highly non-linear regions are computationally expensive to evolve but do not contribute to the Ly α forest [26]. Our simulation box size is $L = 67.5$ Mpc with 768^3 CDM and baryon particles. These values were chosen as a compromise between modelling the low k_{\parallel} modes that play an important role in cosmological constraints from BOSS data, and modelling the small-scale effects of pressure smoothing of the IGM within computational constraints.

The simulations start at $z = 99$ from initial conditions created using `MP-GenIC`, which computes initial displacements using the Zel’dovich approximation. The baryons and dark matter are initialised on an offset grid using species-specific transfer functions calculated in `CLASS` [57]. Following the simulations presented in ref. [58], in order to reduce the effects of cosmic variance we adopt the ‘paired-and-fixed’ approach from refs. [59–61]. In this method, the initial amplitudes of the Fourier modes are fixed, with random phases. The same random seed is used for all simulations, meaning that the residual cosmic variance after applying the paired-fixed approach is the same in each simulation. We do not include any modelling of AGN feedback. We output snapshots at $4 > z > 2$ at intervals of $\Delta z = 0.25$. Note that this redshift information is not included in the emulator. For each simulation snapshot we calculate the transmitted flux fraction along 500^2 lines of sight using the postprocessing code “`fake_spectra`”⁵ [62]. We take the Fourier transform of the flux fraction along each line of sight, and find the average of each Fourier mode to calculate the P_{1D} . The line-of-sight resolution is set to 0.05 Mpc, which is high enough to resolve the thermal broadening scale σ_T whose distribution is shown in figure 2. Box sizes and flux line-of-sight resolutions are chosen such that our mock P_{1D} have the same spacing in k_{\parallel} for all simulations and snapshots.

In order to ensure that our training data are well distributed within the range of physically interesting models, the suite was set up using a Latin hypercube. The Latin hypercube spans the parameters $[\Delta_p^2(z=3), n_p(z=3), z_{\text{rei}}, H_A, H_S]$, with the ranges shown in the left column of table 1. We use a flat Λ CDM cosmology, and vary only the primordial power spectrum in order to sample Δ_p^2 and n_p . This will allow us to test later our hypotheses from section 2. We fix $\omega_c = 0.12$, $\omega_b = 0.022$ throughout this paper, since they are very well constrained by the CMB and otherwise poorly constrained by the Ly α forest. We also explore a range of IGM histories so that the astrophysical parameters in section 2.2 are well

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

⁵https://github.com/sbird/fake_spectra.

	Training simulations	<i>central sim</i>	<i>h sim</i>	ν <i>sim</i>
$A_s (\times 10^{-9})$	[1.35–2.71]	2.00	2.01	2.25
n_s	[0.92–1.02]	0.97	0.97	0.97
h	0.67	0.67	0.74	0.67
Ω_m	0.316	0.316	0.259	0.324
Σm_ν (eV)	0.0	0.0	0.0	0.3
$\Delta_p^2(z=3)$	[0.25–0.45]	0.35		
$n_p(z=3)$	–[2.35–2.25]	–2.30		
z_{rei}	[5.5–15]	10.5		
H_A	[0.5–1.5]	1.0		
H_S	[0.5–1.5]	1.0		

Table 1. Cosmological and astrophysical parameters for the training and test simulations. The limits of the Latin hypercube for the training simulations are shown in the left column, where only the primordial power spectrum and astrophysical parameters are varied. The primordial parameters A_s and n_s here are defined at the CMB pivot scale of $k = 0.05 \text{ Mpc}^{-1}$. The *central sim*, *h sim* and ν *sim* test simulations are constructed such that they have the same small scale linear matter power spectrum (Δ_p^2 and n_p) at $z = 3$. For all simulations, we fix $\omega_c = 0.12$, $\omega_b = 0.022$.

sampled. This is done using modifications [63] around a fiducial reionisation model presented in ref. [64], using a spatially uniform UVB. We define z_{rei} as the midpoint of HI reionisation, i.e. when the ionisation fraction of HI is 0.5, and modify the UVB rates such that this point ranges between $z_{\text{rei}} \in [5.5, 15]$. Following ref. [65], the thermal history is varied using parameters H_A and H_S , which respectively rescale the amplitude of the gas heating rates, and the slope of their density dependence, $\epsilon = H_A \epsilon_0 \Delta_b^{H_S}$. These translate most directly into changes to T_0 (and therefore σ_T) and γ respectively.

In addition to the training simulations, we run three more simulations that we use for various tests, shown in the right side of table 1. The first is a simulation at the centre of the Latin hypercube of simulation parameters in table 1. We refer to this as the *central sim*. Second we run a Λ CDM simulation with a different background expansion and growth, labelled *h sim*, with a larger value of $h = 0.74$ at fixed physical densities (ω_c , ω_b). Finally we run a simulation with 0.3 eV massive neutrinos, labelled ν *sim*, where the neutrinos are implemented using the linear response approximation from ref. [66]. Again we keep ω_c and ω_b fixed. The values of A_s in these test simulations are all set such that $\Delta_p^2(z=3) = 0.35$, which is the midpoint of the Latin hypercube of the training data. We use the same reionisation model in all test simulations. These simulations are therefore constructed in such a way that they lie in the same central part of emulator parameter space. In the case of the ν *sim*, this leads to an increase in A_s to compensate for the suppression of the linear power on small scales caused by massive neutrino free streaming [58].

The 30 training simulations, each with 9 snapshots, produce a total of 270 mock P_{1D} measurements. The Latin hypercube of simulation parameters in table 1 does not propagate into a Latin hypercube in the parameters described in section 2, but it does ensure that the physically motivated regions of this parameter space are well populated with training points. In figure 2 we show the distribution of these training points in three projections of our six dimensional emulator parameter space. The points are coloured corresponding to the redshift




Figure 2. Distribution of the 270 training points in three example 2D projections. We colour training points by redshift purely for illustration, since we do not include redshift information in our emulator. The $z = 3$ training points are highlighted in red for use in figure 4. The Latin hypercube in simulation parameters shown in table 1 does not propagate into a Latin hypercube in emulator parameters, but ensures that the physically interesting parts of parameter space are populated with training points.

of each training point, although we show this purely for illustrative purposes. We assume that these 6 parameters are sufficient to describe the P_{1D} without including explicit redshift information, and so we construct a single training set combining points from all redshifts. We highlight training points at $z = 3$ in red for discussion in section 4.

The trajectories of individual simulations can be identified most clearly in the left panel which shows the thermal parameters σ_T and γ . The middle panel shows points in the $\langle F \rangle - \Delta_p^2$ plane. Both of these parameters are strongly correlated with redshift, and so large parts of this plane are unphysical, and therefore not populated with training points. Although we do not explicitly include redshift information in the emulator, we observe the redshift evolution of P_{1D} through the evolution of Δ_p^2 and $\langle F \rangle$. In the right panel we show the distribution of training points in k_F and n_p . There is no redshift evolution of n_p , and so once again the trajectories of individual simulations are identifiable. The pressure smoothing scale evolves monotonically with redshift as the gas accumulates thermal energy from the UV background.

4 Gaussian process emulator

The purpose of this paper is to model the P_{1D} as a function of the parameters described in section 2 using the simulations presented in section 3. To do this we use a Gaussian process (GP) [67] implemented in the Python package GPy [68]. Our simulations are constructed in such a way that every snapshot provides a measurement of the P_{1D} in 85 linearly spaced k_{\parallel} bins covering the range $[0.0931 < k_{\parallel} < 7.91]$ Mpc^{-1} . Therefore we have a set of training points for each k_{\parallel} bin, and we drop the k_{\parallel} notation for the following discussion. We denote a point in parameter space as $\boldsymbol{\theta} = [\Delta_p^2, n_p, \langle F \rangle, \sigma_T, \gamma, k_F]$. A GP models the function one wants to learn, in this case $P_{1D}(\boldsymbol{\theta})$, as a collection of random variables that form a joint Gaussian distribution, i.e. $P_{1D}(\boldsymbol{\theta}) \sim \mathcal{N}(0, K(\boldsymbol{\theta}, \boldsymbol{\theta}_i))$. Here $\boldsymbol{\theta}_i$ are the locations of the training points where we have modelled the P_{1D} using simulations, and $K(\boldsymbol{\theta}, \boldsymbol{\theta}')$ is a covariance kernel whose precise form will be discussed shortly. In order to approximate the Gaussian process prior of zero mean, we normalise the P_{1D} by the median value of the training data. At a test point $\boldsymbol{\theta}_*$, the joint distribution of the test data $P_{1D}(\boldsymbol{\theta}_*)$ and training data $P_{1D}(\boldsymbol{\theta}_i)$ can be written as

$$\begin{bmatrix} P_{1D}(\boldsymbol{\theta}_i) \\ P_{1D}(\boldsymbol{\theta}_*) \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_i) + \sigma_n^2 \mathbf{I} & K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_*) \\ K(\boldsymbol{\theta}_*, \boldsymbol{\theta}_i) & K(\boldsymbol{\theta}_*, \boldsymbol{\theta}_*) \end{bmatrix}\right) \quad (4.1)$$

where σ_n^2 is a hyperparameter representing Gaussian noise on the training data. The mean and variance of the posterior predictive distribution given the training information at $\boldsymbol{\theta}_i$ are then taken as our estimates of the value and interpolation uncertainty for the $P_{1D}(\boldsymbol{\theta}_\star)$:

$$\mu = K(\boldsymbol{\theta}_\star, \boldsymbol{\theta}_i)[K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_i) + \sigma_n^2 \mathbf{I}]^{-1} P_{1D}(\boldsymbol{\theta}_i) \quad (4.2)$$

$$\sigma^2 = K(\boldsymbol{\theta}_\star, \boldsymbol{\theta}_\star) - K(\boldsymbol{\theta}_\star, \boldsymbol{\theta}_i)[K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_i) + \sigma_n^2 \mathbf{I}]^{-1} K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_\star). \quad (4.3)$$

There is significant freedom in the choice of covariance kernel, which represents a weak prior on the behaviour of the underlying function one wants to model. We use a squared exponential kernel, also known as a radial basis function

$$K(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_0^2 \exp\left(-\sum_j \frac{(\theta_j - \theta'_j)^2}{2l_j^2}\right), \quad (4.4)$$

with an independent correlation length l_i in each parameter direction. It follows that our GP has a total of 8 hyperparameters: 6 correlation lengths, σ_0^2 and σ_n^2 . The hyperparameters are optimised by maximising the marginal log likelihood of the training data [67]. This is done across all k_{\parallel} bins simultaneously, so we use the same set of hyperparameters for the full bandpower description of the P_{1D} . To avoid numerical issues, all of our parameters $\boldsymbol{\theta}$ are rescaled to vary within a unit volume before being passed to the emulator.

We briefly highlight the differences in this setup compared to previous applications of GPs to emulation of the P_{1D} [20, 39, 40]. These emulators trained an independent GP on each redshift bin, whereas we combine training points from all redshifts in one GP. These works also used a different covariance kernel, combining the squared exponential in equation (4.4) with a linear term, and used an isotropic correlation length. The parameterisation used in these works was such that the P_{1D} varied by a similar order of magnitude within the prior volume in each parameter direction, meaning that an isotropic kernel performed well. A major difference in our setup is that the sensitivity of the P_{1D} to each parameter varies strongly. This is reflected in the optimised correlation lengths, which range from 0.47 for $\langle F \rangle$ to 11.5 for n_p . Since we combine all training points in a single emulator and we use larger simulations with more k_{\parallel} bins, there is more information in the training set with which to optimise a larger number of hyperparameters in an anisotropic kernel.

4.1 Emulator validation

In this section we validate the emulator predictions on simulation data. In figure 3 we show the results of a leave-one-out validation test iterated across the full suite of 30 training simulations.⁶ We compare the P_{1D} in the nine snapshots of each validation simulation with the predicted P_{1D} from an emulator trained on the rest of the suite. The solid red lines and shaded red regions respectively show the mean and variance standard deviation of the percentage error on these 30 predictions. The black solid lines show the \pm median $1 - \sigma$ uncertainty on the predictions as estimated by the GP using equation (4.3), and the gray shaded area shows the region of 1% error. We show the accuracy of the emulator predictions at different redshifts to compare it with the uncertainty on BOSS measurements of the P_{1D} from ref. [5], represented by blue shaded regions.⁷

⁶This is in fact leave-*nine*-out cross-validation, since we remove each time from the training set all nine snapshots from a particular simulation.

⁷For each panel we use the nearest redshift bin actually measured in ref. [5].




Figure 3. Results of a ‘leave-one-out’ validation test on the GP emulator. For each of the 30 training simulations, we train the GP on the snapshots from 29 simulations and compare the GP prediction with the truth in each of the nine snapshots of the remaining simulation. The solid red line and red shaded region respectively show the mean and standard deviation of the empirical error (prediction – truth) from the 30 validation tests. The black solid lines around zero show the \pm median of the theoretical uncertainty on the emulator predictions (computed from equation (4.3)). The gray shaded area shows the region of 1% error. We show the performance of the emulator at several redshifts to compare it with the observational uncertainties from the latest analysis of BOSS/eBOSS data (blue shaded regions [5]). We expect that P_{1D} measurements from DESI will cover a similar range of wavenumbers, but the systematic contribution to the data covariance is currently difficult to predict.

The emulator predictions are unbiased, with a mean error of $< 1\%$ at the redshifts and wavenumbers measured by BOSS, eBOSS and DESI. The variance indicated by the shaded red region can be interpreted as an empirical estimate of the average error on emulator predictions. In all regimes probed by BOSS data this quantity is $< 0.5\%$, considerably lower than observational uncertainties. The median theoretical error is slightly larger implying that emulator uncertainties are slightly overestimated, which was also the case in refs. [20, 39, 40]. Nevertheless, these uncertainties are still also well below those of the BOSS data. We do not include lines for DESI here as the exact increase in sensitivity is still uncertain. However it is worth noting that in the case where DESI errors become comparable to these

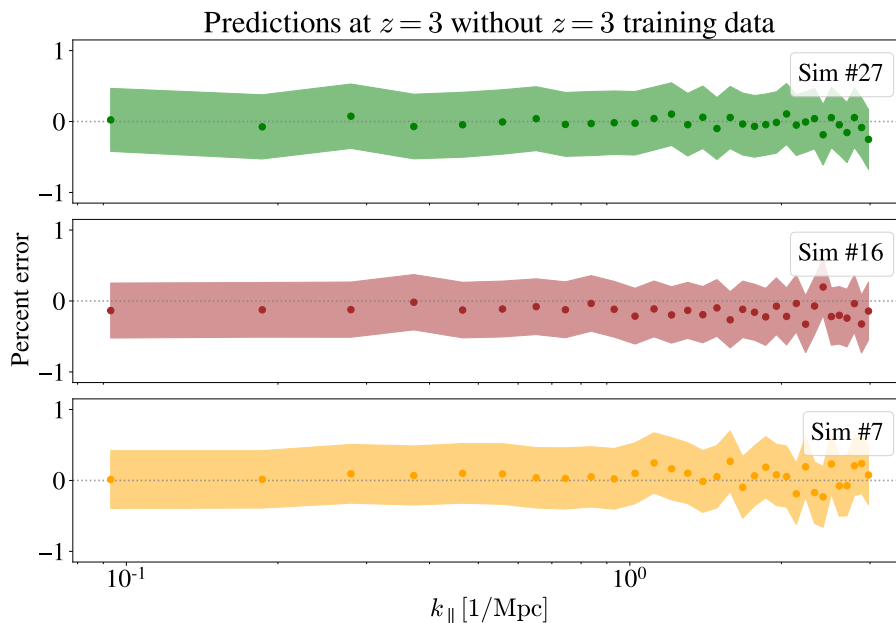


Figure 4. Emulator predictions for the P_{1D} at $z = 3$ when using no $z = 3$ training points. We leave the red points from figure 2 out of the training set and reoptimise the GP hyperparameters. The three panels show the error on the emulated P_{1D} for three randomly chosen simulations. The circles show the empirical error, and the coloured shaded region shows the $1 - \sigma$ theoretical emulator uncertainty.

uncertainties, there are several avenues to be explored such as changes to the covariance kernel to improve emulator performance or the addition of refinement simulations using Bayesian optimisation [20, 40].

Besides the measurements from BOSS and eBOSS, the Ly α forest P_{1D} has also been measured from small samples of very high-resolution spectra [15–17, 51, 69]. These measurements extend to higher redshift and smaller scales, and e.g. can be useful in breaking some of the parameter degeneracies (see appendix A). However, our simulations do not have the resolution required to resolve the pressure scale for some of the colder IGM models at $z > 4$ (see right panel of figure 2).

4.2 Predictions at a redshift without training data

In constructing a single GP using training points from all redshifts, we are assuming that the P_{1D} is sufficiently well described by the six parameters from section 2, without the need to explicitly include redshift information in the emulator. Under this assumption, if one were to construct two simulations that had the same set of θ parameters at different redshifts, the P_{1D} at these two different redshifts would be the same. In practice, engineering such a simulation is non-trivial due to the complex relationship between simulation parameters that model the IGM and the astrophysical emulator parameters. Therefore, in order to test this hypothesis, we construct an emulator dropping all training points at $z = 3$ (the points shown in red in figure 2). We then use this emulator to predict the P_{1D} in the $z = 3$ snapshots for 3 randomly chosen simulations, i.e. we are using training points at $z \leq 2.75$ and $3.25 \leq z$ to predict the P_{1D} at $z = 3$.

The accuracy of these predictions is shown in figure 4, with the theoretical uncertainty (computed using equation (4.3)) on each prediction shown in the shaded coloured region. In

each case the predictions have sub-percent accuracy over the full k_{\parallel} range. This is a novel result that has not been possible with previous Ly α forest emulators, which validates our choice to combine training data from all redshifts and omit redshift information. The ability to emulate the P_{1D} at arbitrary redshifts also confers several advantages. In particular, our emulator could be used to analyse DESI measurements of P_{1D} , regardless of the particular redshift binning of the data, which is currently unknown. Additionally, it is possible to combine training data from simulation suites with different output redshifts, making the construction of training sets more flexible.

4.3 Predicting extended models

Finally, we verify the emulator predictions on two test simulations: the h *sim* and ν *sim*. The suite of training simulations has a fixed expansion history and growth rate, with only the primordial power spectrum varied in order to sample Δ_p^2 and n_p . By fixing the growth rate in all training simulations, the emulator is learning the transfer function between the linear matter power spectrum and the P_{1D} for a fixed relationship between the density power spectrum and the velocity power spectrum. The velocity power spectrum will affect the P_{1D} through redshift space distortions (RSDs). Therefore it is important to test the accuracy of the emulator predictions in cosmologies with different growth rates to the training set. This will act as a test of our hypothesis that the P_{1D} is sensitive primarily to the linear matter power spectrum and insensitive to minor changes in the background expansion.

On the left side of figure 5, we show the accuracy of the emulated P_{1D} when compared with the truth in the h *sim*. The vertical dotted line shows the highest wavenumbers in BOSS measurements [5], and the gray shaded area shows the region of 1% accuracy. One way to consider this difference in h within flat cosmologies is as a variation in the onset of dark energy. We are therefore changing the growth factor in the redshift range $4 > z > 2$, which will have two clear effects. First, the redshift evolution of Δ_p^2 will change, and this effect is captured by our parameterisation. Second, the velocity power spectrum and therefore the RSDs will be different to the simulations in the training set. We are interested in determining whether this effect, which is not captured by any of our parameters, will impact the accuracy of the emulator predictions to a statistically-significant degree. Similarly to the tests shown in figure 3, the predictions have sub-percent accuracy in the regime probed by BOSS and DESI data. Therefore we can conclude that the effect on the growth rate for this scale of change in h is negligible.

In ref. [58] it was shown that the effects of massive neutrinos on the P_{1D} are degenerate with changes to the amplitude of the linear power spectrum to within 1%. In that work it was difficult to disentangle effects of massive neutrinos on the IGM, that were partially accounted for by changing $\langle F \rangle$ and σ_T in postprocessing. The emulator presented in this paper allows for this degeneracy to be tested more clearly since the effects of variations in the IGM and linear power spectrum can be predicted independently. The right side of figure 5 shows the accuracy of predictions compared with the truth in the ν *sim*, which has $\Sigma m_\nu = 0.3$ eV. When adding massive neutrinos, we have kept ω_b and ω_c fixed. Therefore we are modifying the late-time growth and expansion rate governed by Ω_m , which will have similar effects to the previous case on the IGM, the growth of structure, and baryon velocities. Another effect is the suppression of late-time, small-scale clustering due to free-streaming. This suppression is fairly large — approximately 8% in the case of $\Sigma m_\nu = 0.3$ eV, and it is this suppression that makes the Ly α forest a valuable probe of neutrino mass. Again the effect on the amplitude of clustering is captured in our setup by changes to Δ_p^2 . We see similar results to the previous




Figure 5. Accuracy of emulator predictions when tested on two test simulations from table 1. We test the emulator on different cosmological models to the training set, in which only the primordial power spectrum was varied. The red solid line shows the empirical accuracy of the emulator predictions, and the theoretical emulator uncertainty from equation (4.3) is shown in the red shaded region. In the left panel we show predictions for the h *sim*, with $h = 0.74$, where we have trained the emulator only on simulations with $h = 0.67$. In the right panel, we show the accuracy of predictions for the ν *sim* which includes $\Sigma m_\nu = 0.3$ eV massive neutrinos. The vertical dotted line shows the highest k_\parallel bin of BOSS data, and the horizontal gray band shows the region of 1% error.

test, in that the emulator predictions are accurate to less than 1 %, with the best performance at high z . Note that this neutrino mass is larger than current bounds $\Sigma m_\nu < 0.12$ (95% credibility) from *Planck* + baryon acoustic oscillations (BAO) [43]. In this extreme case the emulator still performs well; the methodological errors for smaller neutrino masses will be even smaller. These results demonstrate that the relationship between the linear matter power spectrum and the P_{1D} is insensitive to the specific background expansion used in training the emulator. The test simulations were set up such that each emulator parameter is the same in both simulations to within 2%. This choice was made so that the accuracy of the predictions shown in figure 5 serves as a test of our parameterisation rather than as a test of the interpolation.

5 Parameter constraints

The ability of Ly α forest information alone to constrain cosmological parameters is limited by parameter degeneracies and insensitivities. The constraining power of the P_{1D} on cosmological parameters comes when combined with CMB information, and we leave a discussion on this for future work. However, in order to test the robustness of our emulator, we present parameter constraints from a Ly α forest-only likelihood on the amplitude and slope of the primordial power spectrum, A_s and n_s when using simulated mock data. In order to min-

imise degeneracies, these quantities are defined at a scale $k = 0.7 \text{ Mpc}^{-1}$. To further test the assumptions of fixing the growth and expansion rate, we also explore the sensitivity of the P_{1D} to H_0 . Finally, we show constraints on Σm_ν in order to investigate the degeneracy between Σm_ν and A_s [42, 58]. Given that constraints on Σm_ν are a major motivation in P_{1D} analysis, understanding this degeneracy will be important in the analysis of DESI data.

5.1 Likelihood parameters

Given that our emulator parameters are not associated with any redshift, we do not attempt to constrain the emulator parameters themselves. For our likelihood parameter vector $\tilde{\theta}$ we consider two classes of parameters — cosmological and IGM parameters. In section 5.3 the cosmological parameters we vary are A_s and n_s , and we use uniform priors covering the ranges $[1.1, 2.5] \times 10^{-9}$ and $[0.89, 1.05]$ respectively. In section 5.4 we also vary Σm_ν and H_0 , with uniform priors covering the ranges $[0, 1] \text{ eV}$ and $[55, 90] \text{ km s}^{-1} \text{ Mpc}^{-1}$ respectively. We keep the physical baryon and cold dark matter densities, ω_c and ω_b , fixed to the values set in section 3. For each combination of these parameters we use CAMB to calculate the linear matter power spectra at each redshift where we have data, and we compute the corresponding values of $\Delta_p^2(z)$ and $n_p(z)$. Therefore each likelihood evaluation results in nine calls to the emulator.

To model the redshift evolution of the IGM parameters we use the redshift evolution of each parameter in the *central sim* described in table 1 as a fiducial model. The evolution of the IGM for this simulation and all 30 training simulations is shown in appendix B. To vary the IGM, we use a multiplicative rescaling $m(z)$ of this fiducial model using a power law around a pivot redshift:

$$\ln m(z) = X_0 + X_1 \ln \frac{1+z}{1+z_*}, \quad (5.1)$$

where we set $z_* = 3$ as it is the midpoint of our redshift outputs. X_0 controls the amplitude of the rescaling, and X_1 accommodates changes in the redshift evolution of a given parameter. The same approach is used for all IGM parameters, meaning we have a total of eight free parameters to model variations in the IGM. This model will not perfectly describe the evolution of the IGM in any simulation except the *central sim* and so we are not interested in constraints on these parameters *per se*, but instead interested in ensuring that we allow for enough variation in the IGM for a thorough marginalisation. The prior range for each parameter is $\in [-0.4, 0.4]$, and we have verified that marginalised constraints on cosmological parameters are not impacted by expanding this range.

An advantage of omitting redshift information from the emulator is the increased flexibility in modelling the redshift evolution of the IGM. Since our emulator parameters are decoupled from simulation and likelihood parameters, the redshift dependence of the IGM is not fixed at the point of running the simulations as is the case in many previous setups. Whilst we use equation (5.1) to vary the IGM in this analysis, in principle any form of redshift evolution can be chosen provided it does not involve calls to the emulator that lie outside the convex hull of the training set.

Our baseline likelihood parameter vector is therefore

$$\tilde{\theta} = [A_s, n_s, \ln \tau_0, \ln \tau_1, \ln \sigma_0^T, \ln \sigma_1^T, \ln \gamma_0, \ln \gamma_1, \ln k_0^F, \ln k_1^F] \quad (5.2)$$

where τ_i represents rescalings to the effective optical depth, τ_{eff} , σ_i^T to σ_T , γ_i to γ , k_i^F to k_F . The quantities that relate to physical length scales, σ_T and k_F are evaluated in velocity units.

5.2 Likelihood function

To more closely resemble a real analysis we perform our likelihood evaluations in velocity units. Since our emulator is constructed in comoving units, a cosmological model is required to convert between these using the relation $k = qH(z)/(1+z)$, where k and q are wavenumbers in comoving and velocity units respectively, and denote a flux power spectrum in velocity units as $\tilde{P}_{1D}(q_{\parallel})$. We use the same q_{\parallel} bins as the BOSS data [5], and use a cubic interpolation to rebin our mock and emulated flux power spectra to match the BOSS bins. We use a Gaussian likelihood function:

$$\ln\mathcal{L}(\tilde{\theta}) = \sum_z \left[-\frac{1}{2}\Delta_i(q_i, z, \tilde{\theta})C_{ij}(z, \tilde{\theta})^{-1}\Delta_j(q_j, z, \tilde{\theta}) - \frac{1}{2}\log\det C_{ij}(z, \tilde{\theta}) \right] \quad (5.3)$$

with

$$\Delta_i(q_i, z, \tilde{\theta}) = \tilde{P}_{1D}^{\text{emu}}(q_i, z, \tilde{\theta}) - \tilde{P}_{1D}^{\text{data}}(q_i, z) \quad (5.4)$$

where $\tilde{P}_{1D}^{\text{emu}}(q_i, z, \tilde{\theta})$ is the theoretical prediction from the emulator for a given set of parameters $\tilde{\theta}$, and z is the redshift at which we are making the calls to the emulator. The covariance matrix C_{ij} is a combination of the data and emulator covariance:

$$C_{ij}(z, \tilde{\theta}) = C_{ij}^{\text{data}}(z) + C_{ij}^{\text{emu}}(z, \tilde{\theta}). \quad (5.5)$$

In order to obtain results that resemble an analysis using current datasets, for the data covariance $C_{ij}^{\text{data}}(z)$ we use the covariance matrices from the most recent measurement of the $\tilde{P}_{1D}(q_{\parallel})$ from BOSS [5]. Our simulation snapshot outputs are different to the redshift binning of the BOSS measurements, so for each z where we have mock data, we simply use the covariance matrix from the closest BOSS redshift bin, prioritising the lower z bin in the case where the bins are equidistant. The emulator covariance matrix is constructed from the theoretical uncertainty on each emulator call, which is a function of the emulator hyperparameters and the local density of the training points. Given that both these properties are common to all q bins, we assume the errors are maximally correlated and build $C_{ij}^{\text{emu}}(z, \tilde{\theta})$ under this assumption. Since we use the same random seed in all simulations, any residual cosmic variance that remains after applying the ‘paired-and-fixed’ approach is the same in all simulations. We do not add noise realisations to the mock data drawn from the mock data covariance. The likelihood function is sampled using MCMC sampling implemented in `emcee` [70], and we check for convergence of the integrated autocorrelation time.

5.3 MCMC results

In figure 6 we show the accuracy of marginalised constraints on A_s and n_s for each of the 30 training simulations when running with 10 free parameters: A_s and n_s and the 8 IGM parameters described earlier. We fix h and Σm_{ν} to the true values of $h = 0.67$ and $\Sigma m_{\nu} = 0$ eV. As in section 4.1, for each constraint we leave the validation simulation out of the emulator training set. The error bars show the $1 - \sigma$ uncertainty. In all but 3 cases we recover the correct cosmology to within $1 - \sigma$.

5.4 Extended models

In figure 7 we show constraints on A_s , n_s and H_0 after marginalising over the IGM parameters, using the *h sim* as mock data. We set a wide prior range of $H_0 \in [55, 90] \text{ km s}^{-1} \text{ Mpc}^{-1}$. Since the Ly α forest is observed in velocity units, our emulator predictions undergo a conversion




Figure 6. Accuracy of marginalised constraints on A_s and n_s for each of the 30 training simulations, with $1 - \sigma$ credible bounds.

from comoving to velocity units using the factor $H(z)/(1+z)$. Therefore we expect the constraints on the amplitude of the power spectrum to be degenerate with $H(z)$. This is evident in the lower left contour of figure 7. The degeneracy is weakened by the fact that we are probing a high redshift where the differences in $H(z)$ are smaller. For example, across the wide prior range, there is only a 6% difference in $H(z=2)$, and a 1.5% difference at $H(z=4)$. The constraints on A_s and n_s are accurate. This is an important demonstration, as we are correctly recovering parameters in a cosmology with a different background evolution to the one used in constructing the emulator.

In figure 8 we show marginalised constraints when using the ν *sim* as mock data, allowing neutrino mass to vary in the range $\Sigma m_\nu \in [0.0, 1.0]$ eV. The effects of massive neutrinos on cosmology are to vary both the expansion history and the growth rate, and to suppress the amplitude of matter clustering on small scales. Given the results of the previous figure, we expect the effects of massive neutrinos on the expansion and growth to have negligible contributions to the constraints. We use an extreme value of $\Sigma m_\nu = 0.3$ eV that is significantly higher than the current bound $\Sigma m_\nu < 0.12$ eV (95% credibility) from joint *Planck* + BAO data [43]. The neutrino mass is strongly degenerate with A_s when using the Ly α forest alone [42, 58]; stronger bounds arise from a combination with CMB data by breaking the degeneracy (e.g [31, 32]).

6 Discussion

We have presented a Gaussian process emulator that predicts the P_{1D} of the Ly α forest as a function of the amplitude and slope of the small-scale linear matter power spectrum and the state of the intergalactic medium. This work is an application of modern simulation interpolation techniques to a modelling approach that was more prevalent in early analysis of the Ly α forest, based on the fact that the cosmological information in the Ly α forest is concentrated in the linear matter power spectrum at the observed redshift. This is due to

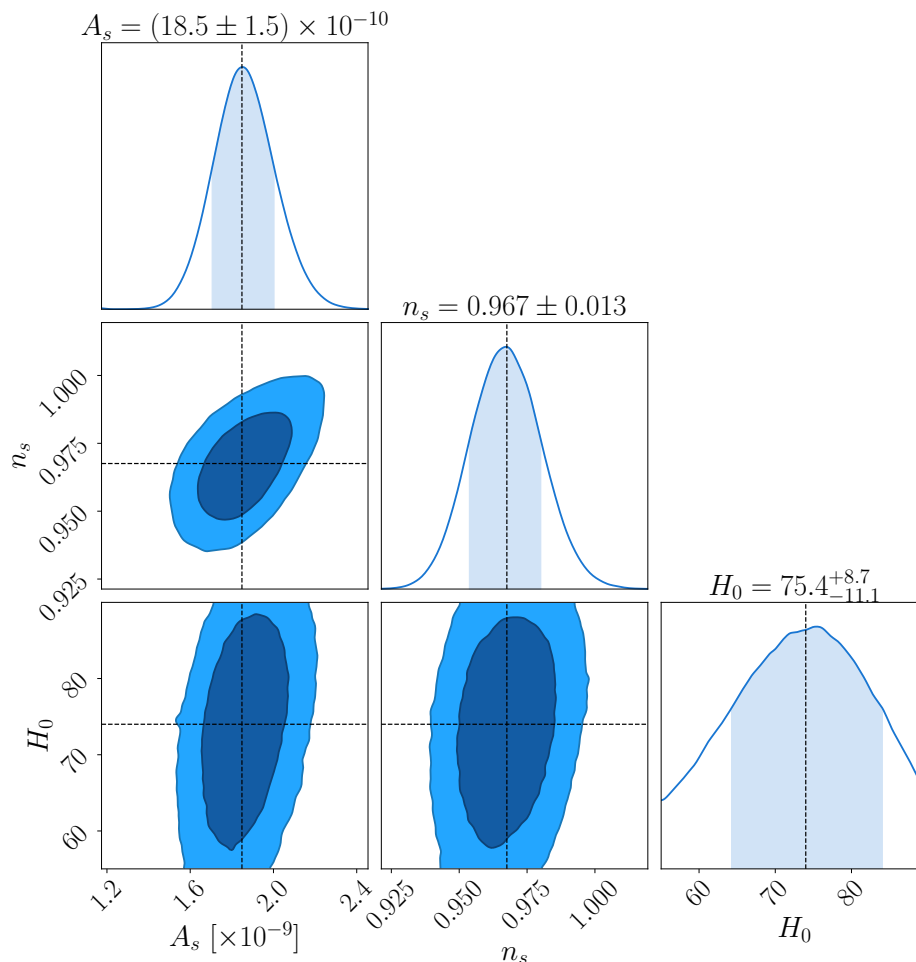


Figure 7. Marginalised 1D and 2D posterior distributions when using the *h sim* as mock data. The true values are indicated by the dashed lines, with the contours showing the 68% and 95% credible regions. A_s and n_s are defined at a pivot scale of $k = 0.7 \text{Mpc}^{-1}$. The shaded area of the 1D posteriors shows the 68% credible region. Posteriors for the remaining parameters are shown in appendix C.

a unique property of the Ly α forest, in that it probes the universe at a period very close to an Einstein de-Sitter model. In this regime, variations in the P_{1D} due to differences in the growth rate and expansion history (allowed by current data) are negligible. For the first time we test the accuracy of this approximation at the level required by current datasets, and show that it holds. In particular, we showed that even in models with massive neutrinos, the cosmological information in P_{1D} analysis is contained in the slope and amplitude of the linear power spectrum within current observational uncertainty.

In addition to being lower dimensional in the cosmological sector, our emulator affords other practical advantages. By emulating the flux power spectrum independently of redshift, predictions can be made to sub-percent accuracy at redshifts not necessarily included in the training simulations. This affords flexibility with regards to the analysis of future data. Moreover, the redshift evolution of the IGM can be modelled at the point of constructing a likelihood function, as opposed to being determined at the point of constructing the emulator as in many previous frameworks.

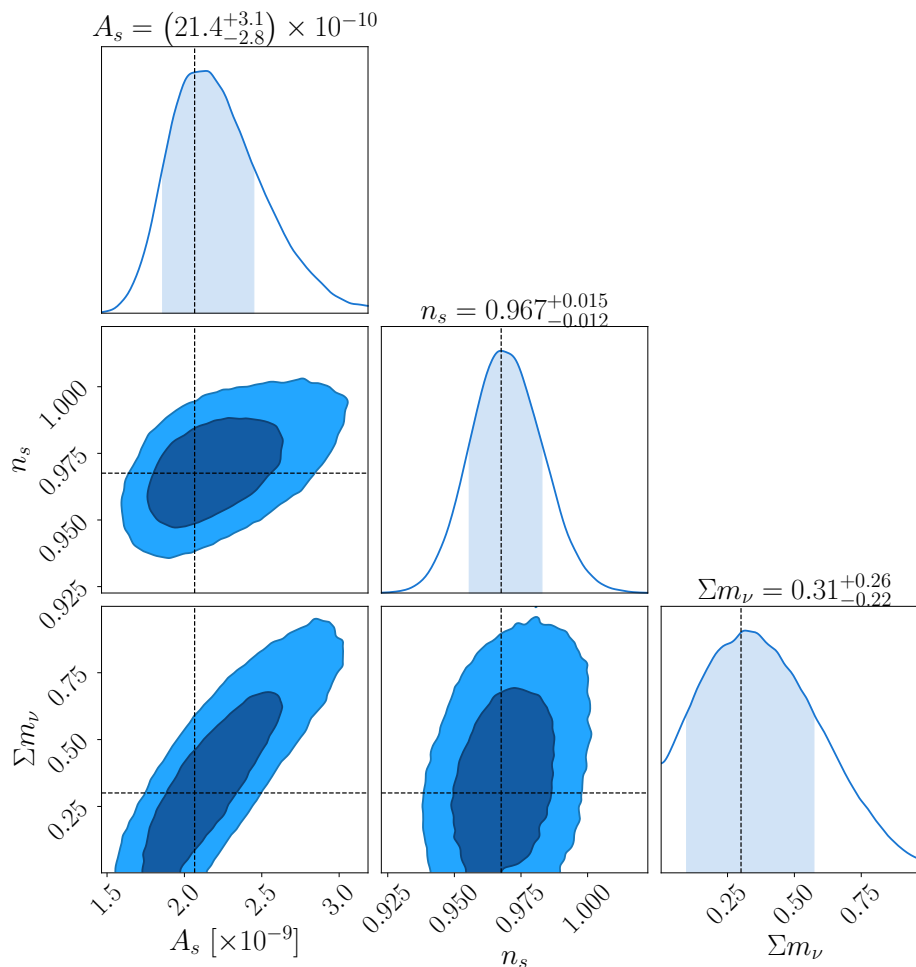


Figure 8. Marginalised 1D and 2D posterior distributions when using the ν *sim* as mock data, with Σm_ν as a free parameter. The true values are indicated by the dashed lines, with the contours showing the 68% and 95% credible regions. A_s and n_s are defined at a pivot scale of $k = 0.7 \text{ Mpc}^{-1}$. The shaded areas of the 1D posteriors show the 68% credible region. For clarity in the figure, we do not show posteriors for the 8 free IGM parameters in this chain.

Whilst we demonstrated that this emulator can be applied to massive neutrino cosmologies without massive neutrinos used in the training set, the same principle can be applied to a wider range of extended cosmological models, e.g. running of the spectral index in the primordial power spectrum or curved universes [43]. There are two conditions under which our implementation would not be valid. Models with sharp features in the small-scale matter power spectrum (such as warm dark matter models [17, 71] or ultra-light axion dark matter [46, 47]), or models with a non-standard relationship between the baryon density and velocities (such as modified gravity models [72]). However, simple extensions to the emulator parameter space could incorporate some of these models too (e.g. [20, 38]), and the motivations for parameterising the P_{1D} in terms of the linear power spectrum still apply.

The joint analysis of CMB observations with measurements of the small-scale linear matter power spectrum from the Ly α forest has long been a powerful combination in constraining extensions to Λ CDM [23–28, 73–75]. We presented an emulator that will facilitate

the combination of current (BOSS/eBOSS) and upcoming DESI measurements of the P_{1D} with CMB observations, to provide competitive constraints on beyond- Λ CDM cosmologies.

Acknowledgments

The authors thank Jose Oñorbe for sharing his code for the reionisation model of ref. [63] and for valuable discussions. This work was partially enabled by funding from the UCL Cosmoparticle Initiative. This work used computing equipment funded by the Research Capital Investment Fund (RCIF) provided by UK Research and Innovation (UKRI), and partially funded by the UCL Cosmoparticle Initiative. This work was supported by collaborative visits funded by the Cosmology and Astroparticle Student and Postdoc Exchange Network (CASPEN). This work was performed using the Cambridge Service for Data Driven Discovery (CSD3), part of which is operated by the University of Cambridge Research Computing on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The DiRAC component of CSD3 was funded by BEIS capital funding via STFC capital grants ST/P002307/1 and ST/R002452/1 and STFC operations grant ST/R00689X/1. DiRAC is part of the National e-Infrastructure. AFR acknowledges support by an STFC Ernest Rutherford Fellowship, grant reference ST/N003853/1, and by FSE funds through the program Ramon y Cajal (RYC-2018-025210) of the Spanish Ministry of Science and Innovation. AFR, HVP and AP were partially supported by the Science and Technology Facilities Council (STFC) Consolidated Grant number ST/R000476/1. KKR was supported by funding from the Science Research Council (VR) of Sweden; and the Dunlap Institute for Astronomy & Astrophysics, University of Toronto. PM was supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Contract no. DE-AC02-05CH11231. HVP was additionally supported by the research project grant “Fundamental Physics from Cosmological Surveys” funded by the Swedish Research Council (VR) under Dnr 2017-04212.

A Parameter dependence of the P_{1D}

In figure 9 we use the GP presented in section 4 to show the fractional change of the P_{1D} in response to changes of each of the parameters described in section 2. We show the fractional change around a central point in parameter space set by the midpoint of all training points at $z = 3$, i.e. the points shown in red in figure 2. We denote this point as $P_{1D}^0(k_{\parallel})$. We show the change in the P_{1D} with respect to $P_{1D}^0(k_{\parallel})$ when moving from edge to edge of the convex hull of the full training set, to ensure that we are not extrapolating into regions of parameter space where the GP cannot reliably predict the P_{1D} . The blue dotted lines show the highest k_{\parallel} probed by BOSS [5]. The effects of IGM parameters are most dominant at higher wavenumbers than those probed by BOSS data, with a degeneracy between $\langle F \rangle$ and γ when looking only at the lower k_{\parallel} modes.

B IGM evolution

We plot the redshift evolution of each IGM parameter for all 30 training simulations in gray lines in figure 10, with the central simulation shown in red. The values in red are used to provide a fiducial model for the marginalisation over the IGM described in section 5. The mean flux, $\langle F \rangle$ is found by averaging over the flux transmission fraction in every line-of-sight pixel of a given snapshot. To calculate the thermal parameters, σ_T and γ , the particle snapshot data is binned in the $\log_{10}T - \log_{10}\Delta_b$ plane, where T is in units of K. We then




Figure 9. Fractional effect on the P_{1D} of changing each of the six emulator parameters. We chose as a reference model, $P_{1D}^0(k_{||})$ the midpoint of the $z = 3$ training points, and show the effect on the P_{1D} when moving from edge to edge of the convex hull of the training set. Lowest values for each parameter are coloured blue, transitioning to the highest values in red. The blue dotted line represents the highest $k_{||}$ bin in BOSS data [5].

perform a linear fit to the mode within the range $-1.5 < \log_{10}\Delta_b < 0.5$ to obtain values for T_0 and γ from equation (2.3). T_0 is then converted into a value for σ_T using equation (2.4). In order to calculate the pressure smoothing scale, k_F , we follow the approach in ref. [54]. We calculate the real-space Lyman- α forest 3D flux power spectrum, $\Delta_F^2(k)$, and fit an exponential cutoff of the following form:

$$\Delta_F^2(k) = Ak^n \exp\left(-\frac{k^2}{k_F^2}\right), \quad (\text{B.1})$$

where A and n are free parameters.




Figure 10. Redshift evolution of the IGM parameters for each of our 30 training simulations in gray. In red we show the IGM parameters for the *central sim* which is used as a fiducial model in section 5.

C Full posterior with IGM parameters

In section 5 we presented MCMC posteriors for several analyses on mock data, after marginalising over the 8 IGM parameters. In figure 11 we show the full posterior for one of the analyses, including cosmological and IGM parameters.




Figure 11. Marginalised 1D and 2D posterior distributions for all free parameters when using the h *sim* as mock data (the same as in figure 7).

References

- [1] BOSS collaboration, *The Baryon Oscillation Spectroscopic Survey of SDSS-III*, *Astron. J.* **145** (2013) 10 [[arXiv:1208.0022](#)] [[INSPIRE](#)].
- [2] K.S. Dawson et al., *The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data*, *Astron. J.* **151** (2016) 44 [[arXiv:1508.04473](#)] [[INSPIRE](#)].
- [3] H. du Mas des Bourboux et al., *The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Baryon acoustic oscillations with Lyman- α forests*, *Astrophys. J.* **901** (2020) 153 [[arXiv:2007.08995](#)] [[INSPIRE](#)].
- [4] N. Palanque-Delabrouille et al., *The one-dimensional Ly α forest power spectrum from BOSS*, *Astron. Astrophys.* **559** (2013) A85 [[arXiv:1306.5896](#)] [[INSPIRE](#)].
- [5] S. Chabanier et al., *The one-dimensional power spectrum from the SDSS DR14 Ly α forests*, *JCAP* **07** (2019) 017 [[arXiv:1812.03554](#)] [[INSPIRE](#)].
- [6] DESI collaboration, *The DESI Experiment Part I: Science, Targeting, and Survey Design*, [arXiv:1611.00036](#) [[INSPIRE](#)].

- [7] R.A.C. Croft, D.H. Weinberg, N. Katz and L. Hernquist, *Recovery of the Power Spectrum of Mass Fluctuations from Observations of the Ly α Forest*, *Astrophys. J.* **495** (1998) 44 [[astro-ph/9708018](#)] [[INSPIRE](#)].
- [8] R.A.C. Croft, D.H. Weinberg, M. Pettini, L. Hernquist and N. Katz, *The Power Spectrum of Mass Fluctuations Measured from the Ly α Forest at Redshift $z = 2.5$* , *Astrophys. J.* **520** (1999) 1 [[astro-ph/9809401](#)] [[INSPIRE](#)].
- [9] P. McDonald et al., *The Observed Probability Distribution Function, Power Spectrum, and Correlation Function of the Transmitted Flux in the Ly α Forest*, *Astrophys. J.* **543** (2000) 1 [[astro-ph/9911196](#)] [[INSPIRE](#)].
- [10] N.Y. Gnedin and A.J.S. Hamilton, *Matter power spectrum from the Lyman-alpha forest: Myth or reality?*, *Mon. Not. Roy. Astron. Soc.* **334** (2002) 107 [[astro-ph/0111194](#)] [[INSPIRE](#)].
- [11] R.A.C. Croft et al., *Toward a Precise Measurement of Matter Clustering: Ly α Forest Data at Redshifts 2–4*, *Astrophys. J.* **581** (2002) 20 [[astro-ph/0012324](#)] [[INSPIRE](#)].
- [12] M. Viel, M.G. Haehnelt and V. Springel, *Inferring the dark matter power spectrum from the Lyman α forest in high-resolution QSO absorption spectra*, *Mon. Not. Roy. Astron. Soc.* **354** (2004) 684 [[astro-ph/0404600](#)] [[INSPIRE](#)].
- [13] SDSS collaboration, *The Linear Theory Power Spectrum from the Ly α Forest in the Sloan Digital Sky Survey*, *Astrophys. J.* **635** (2005) 761 [[astro-ph/0407377](#)] [[INSPIRE](#)].
- [14] J. Oñorbe, J.F. Hennawi, Z. Lukić and M. Walther, *Constraining Reionization with the $z \sim 5 - 6$ Lyman- α Forest Power Spectrum: the Outlook after Planck*, *Astrophys. J.* **847** (2017) 63 [[arXiv:1703.08633](#)] [[INSPIRE](#)].
- [15] M. Walther, J. Oñorbe, J.F. Hennawi and Z. Lukić, *New Constraints on IGM Thermal Evolution from the Ly α Forest Power Spectrum*, *Astrophys. J.* **872** (2019) 13 [[arXiv:1808.04367](#)] [[INSPIRE](#)].
- [16] E. Boera, G.D. Becker, J.S. Bolton and F. Nasir, *Revealing Reionization with the Thermal History of the Intergalactic Medium: New Constraints from the Ly α Flux Power Spectrum*, *Astrophys. J.* **872** (2019) 101 [[arXiv:1809.06980](#)] [[INSPIRE](#)].
- [17] M. Viel, G.D. Becker, J.S. Bolton and M.G. Haehnelt, *Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data*, *Phys. Rev. D* **88** (2013) 043502 [[arXiv:1306.2314](#)] [[INSPIRE](#)].
- [18] V. Iršič et al., *New Constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman- α forest data*, *Phys. Rev. D* **96** (2017) 023522 [[arXiv:1702.01764](#)] [[INSPIRE](#)].
- [19] V. Iršič, M. Viel, M.G. Haehnelt, J.S. Bolton and G.D. Becker, *First constraints on fuzzy dark matter from Lyman- α forest data and hydrodynamical simulations*, *Phys. Rev. Lett.* **119** (2017) 031302 [[arXiv:1703.04683](#)] [[INSPIRE](#)].
- [20] K.K. Rogers and H.V. Peiris, *General framework for cosmological dark matter bounds using N -body simulations*, *Phys. Rev. D* **103** (2021) 043526 [[arXiv:2007.13751](#)] [[INSPIRE](#)].
- [21] K.K. Rogers and H.V. Peiris, *Strong bound on canonical ultra-light axion dark matter from the Lyman-alpha forest*, *Phys. Rev. Lett.* **126** (2021) 071302 [[arXiv:2007.12705](#)] [[INSPIRE](#)].
- [22] S. Chabanier, M. Millea and N. Palanque-Delabrouille, *Matter power spectrum: from Ly α forest to CMB scales*, *Mon. Not. Roy. Astron. Soc.* **489** (2019) 2247 [[arXiv:1905.08103](#)] [[INSPIRE](#)].
- [23] J. Phillips, D.H. Weinberg, R.A.C. Croft, L. Hernquist, N. Katz and M. Pettini, *Constraints on Cosmological Parameters from the Ly α Forest Power Spectrum and COBE DMR*, *Astrophys. J.* **560** (2001) 15 [[astro-ph/0001089](#)] [[INSPIRE](#)].

- [24] WMAP collaboration, *First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Parameter Estimation Methodology*, *Astrophys. J. Suppl.* **148** (2003) 195 [[astro-ph/0302218](#)] [[INSPIRE](#)].
- [25] WMAP collaboration, *First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters*, *Astrophys. J. Suppl.* **148** (2003) 175 [[astro-ph/0302209](#)] [[INSPIRE](#)].
- [26] M. Viel, J. Weller and M. Haehnelt, *Constraints on the primordial power spectrum from high-resolution Lyman α forest spectra and WMAP*, *Mon. Not. Roy. Astron. Soc.* **355** (2004) L23 [[astro-ph/0407294](#)] [[INSPIRE](#)].
- [27] SDSS collaboration, *Cosmological parameter analysis including SDSS Ly α forest and galaxy bias: Constraints on the primordial spectrum of fluctuations, neutrino mass, and dark energy*, *Phys. Rev. D* **71** (2005) 103515 [[astro-ph/0407372](#)] [[INSPIRE](#)].
- [28] U. Seljak, A. Slosar and P. McDonald, *Cosmological parameters from combining the Lyman- α forest with CMB, galaxy clustering and SN constraints*, *JCAP* **10** (2006) 014 [[astro-ph/0604335](#)] [[INSPIRE](#)].
- [29] S. Bird, H.V. Peiris, M. Viel and L. Verde, *Minimally parametric power spectrum reconstruction from the Lyman α forest*, *Mon. Not. Roy. Astron. Soc.* **413** (2011) 1717 [[arXiv:1010.1519](#)] [[INSPIRE](#)].
- [30] N. Palanque-Delabrouille et al., *Constraint on neutrino masses from SDSS-III/BOSS Ly α forest and other cosmological probes*, *JCAP* **02** (2015) 045 [[arXiv:1410.7244](#)] [[INSPIRE](#)].
- [31] N. Palanque-Delabrouille et al., *Neutrino masses and cosmology with Lyman-alpha forest power spectrum*, *JCAP* **11** (2015) 011 [[arXiv:1506.05976](#)] [[INSPIRE](#)].
- [32] N. Palanque-Delabrouille et al., *Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman- α and Planck full-survey data*, *JCAP* **04** (2020) 038 [[arXiv:1911.09073](#)] [[INSPIRE](#)].
- [33] M. McQuinn, *The Evolution of the Intergalactic Medium*, *Ann. Rev. Astron. Astrophys.* **54** (2016) 313 [[arXiv:1512.00086](#)] [[INSPIRE](#)].
- [34] J. Kwan et al., *Cosmic Emulation: Fast Predictions for the Galaxy Power Spectrum*, *Astrophys. J.* **810** (2015) 35 [[arXiv:1311.6444](#)] [[INSPIRE](#)].
- [35] Z. Zhai et al., *The Aemulus Project. III. Emulation of the Galaxy Correlation Function*, *Astrophys. J.* **874** (2019) 95 [[arXiv:1804.05867](#)] [[INSPIRE](#)].
- [36] M. Viel and M.G. Haehnelt, *Cosmological and astrophysical parameters from the Sloan Digital Sky Survey flux power spectrum and hydrodynamical simulations of the Lyman α forest*, *Mon. Not. Roy. Astron. Soc.* **365** (2006) 231 [[astro-ph/0508177](#)] [[INSPIRE](#)].
- [37] A. Borde et al., *New approach for precise computation of Lyman- α forest power spectrum with hydrodynamical simulations*, *JCAP* **07** (2014) 005 [[arXiv:1401.6472](#)] [[INSPIRE](#)].
- [38] R. Murgia, V. Iršič and M. Viel, *Novel constraints on noncold, nonthermal dark matter from Lyman- α forest data*, *Phys. Rev. D* **98** (2018) 083540 [[arXiv:1806.08371](#)] [[INSPIRE](#)].
- [39] S. Bird, K.K. Rogers, H.V. Peiris, L. Verde, A. Font-Ribera and A. Pontzen, *An emulator for the Lyman- α forest*, *JCAP* **02** (2019) 050 [[arXiv:1812.04654](#)] [[INSPIRE](#)].
- [40] K.K. Rogers, H.V. Peiris, A. Pontzen, S. Bird, L. Verde and A. Font-Ribera, *Bayesian emulator optimisation for cosmology: application to the Lyman-alpha forest*, *JCAP* **02** (2019) 031 [[arXiv:1812.04631](#)] [[INSPIRE](#)].
- [41] T. Takhtaganov, Z. Lukić, J. Müller and D. Morozov, *Cosmic Inference: Constraining Parameters with Observations and a Highly Limited Number of Simulations*, *Astrophys. J.* **906** (2021) 74 [[arXiv:1905.07410](#)] [[INSPIRE](#)].

- [42] M. Viel, M.G. Haehnelt and V. Springel, *The effect of neutrinos on the matter distribution as probed by the intergalactic medium*, *JCAP* **06** (2010) 015 [[arXiv:1003.2422](#)] [[INSPIRE](#)].
- [43] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [44] J. Lesgourgues and S. Pastor, *Massive neutrinos and cosmology*, *Phys. Rept.* **429** (2006) 307 [[astro-ph/0603494](#)] [[INSPIRE](#)].
- [45] J. Lesgourgues and S. Pastor, *Neutrino cosmology and Planck*, *New J. Phys.* **16** (2014) 065002 [[arXiv:1404.1740](#)] [[INSPIRE](#)].
- [46] W. Hu, R. Barkana and A. Gruzinov, *Fuzzy Cold Dark Matter: The Wave Properties of Ultralight Particles*, *Phys. Rev. Lett.* **85** (2000) 1158 [[astro-ph/0003365](#)] [[INSPIRE](#)].
- [47] L. Hui, J.P. Ostriker, S. Tremaine and E. Witten, *Ultralight scalars as cosmological dark matter*, *Phys. Rev. D* **95** (2017) 043541 [[arXiv:1610.08297](#)] [[INSPIRE](#)].
- [48] M. Vogelsberger, J. Zavala, F.-Y. Cyr-Racine, C. Pfrommer, T. Bringmann and K. Sigurdson, *ETHOS — an effective theory of structure formation: dark matter physics as a possible explanation of the small-scale CDM problems*, *Mon. Not. Roy. Astron. Soc.* **460** (2016) 1399 [[arXiv:1512.05349](#)] [[INSPIRE](#)].
- [49] J. König, A. Merle and M. Totzauer, *keV sterile neutrino dark matter from singlet scalar decays: the most general case*, *JCAP* **11** (2016) 038 [[arXiv:1609.01289](#)] [[INSPIRE](#)].
- [50] M. Drewes et al., *A White Paper on keV sterile neutrino Dark Matter*, *JCAP* **01** (2017) 025 [[arXiv:1602.04816](#)] [[INSPIRE](#)].
- [51] V. Iršič et al., *The Lyman α forest power spectrum from the XQ-100 Legacy Survey*, *Mon. Not. Roy. Astron. Soc.* **466** (2017) 4332 [[arXiv:1702.01761](#)] [[INSPIRE](#)].
- [52] A.G. Riess et al., *A 2.4% Determination of the Local Value of the Hubble Constant*, *Astrophys. J.* **826** (2016) 56 [[arXiv:1604.01424](#)] [[INSPIRE](#)].
- [53] Z. Lukić, C.W. Stark, P. Nugent, M. White, A.A. Meiksin and A. Almgren, *The Lyman- α forest in optically thin hydrodynamical simulations*, *Mon. Not. Roy. Astron. Soc.* **446** (2015) 3697 [[arXiv:1406.6361](#)] [[INSPIRE](#)].
- [54] G. Kulkarni, J.F. Hennawi, J. Oñorbe, A. Rorai and V. Springel, *Characterizing the Pressure Smoothing Scale of the Intergalactic Medium*, *Astrophys. J.* **812** (2015) 30 [[arXiv:1504.00366](#)] [[INSPIRE](#)].
- [55] Y. Feng, S. Bird, L. Anderson, A. Font-Ribera and C. Pedersen, *Mp-gadget/mp-gadget: A tag for getting a doi*, (Oct. 2018) [[DOI](#)].
- [56] V. Springel, *The cosmological simulation code GADGET-2*, *Mon. Not. Roy. Astron. Soc.* **364** (2005) 1105 [[astro-ph/0505010](#)] [[INSPIRE](#)].
- [57] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, [[arXiv:1104.2932](#)] [[INSPIRE](#)].
- [58] C. Pedersen et al., *Massive neutrinos and degeneracies in Lyman-alpha forest simulations*, *JCAP* **04** (2020) 025 [[arXiv:1911.09596](#)] [[INSPIRE](#)].
- [59] R.E. Angulo and A. Pontzen, *Cosmological N-body simulations with suppressed variance*, *Mon. Not. Roy. Astron. Soc.* **462** (2016) L1 [[arXiv:1603.05253](#)] [[INSPIRE](#)].
- [60] L. Anderson, A. Pontzen, A. Font-Ribera, F. Villaescusa-Navarro, K.K. Rogers and S. Genel, *Cosmological Hydrodynamic Simulations with Suppressed Variance in the Lyman- α Forest Power Spectrum*, *Astrophys. J.* **871** (2019) 144 [[arXiv:1811.00043](#)] [[INSPIRE](#)].
- [61] F. Villaescusa-Navarro et al., *Statistical properties of paired fixed fields*, *Astrophys. J.* **867** (2018) 137 [[arXiv:1806.01871](#)] [[INSPIRE](#)].

- [62] S. Bird, *FSFE: Fake Spectra Flux Extractor*, Astrophysics Source Code Library, (Oct. 2017).
- [63] J. Oñorbe, J.F. Hennawi and Z. Lukić, *Self-Consistent Modeling of Reionization in Cosmological Hydrodynamical Simulations*, *Astrophys. J.* **837** (2017) 106 [[arXiv:1607.04218](#)] [[INSPIRE](#)].
- [64] F. Haardt and P. Madau, *Radiative Transfer in a Clumpy Universe. IV. New Synthesis Models of the Cosmic UV/X-Ray Background*, *Astrophys. J.* **746** (2012) 125 [[arXiv:1105.2039](#)] [[INSPIRE](#)].
- [65] J.S. Bolton, M. Viel, T.S. Kim, M.G. Haehnelt and R.F. Carswell, *Possible evidence for an inverted temperature-density relation in the intergalactic medium from the flux distribution of the Ly α forest*, *Mon. Not. Roy. Astron. Soc.* **386** (2008) 1131 [[arXiv:0711.2064](#)] [[INSPIRE](#)].
- [66] Y. Ali-Haïmoud and S. Bird, *An efficient implementation of massive neutrinos in non-linear structure formation simulations*, *Mon. Not. Roy. Astron. Soc.* **428** (2012) 3375 [[arXiv:1209.0461](#)] [[INSPIRE](#)].
- [67] C.E. Rasmussen and C.K.I. Williams, *Gaussian Processes for Machine Learning*, MIT Press, (2006).
- [68] GPy, *GPy: a gaussian process framework in python*, <http://github.com/SheffieldML/GPy>, since (2012).
- [69] C. Yèche, N. Palanque-Delabrouille, J. Baur and H. du Mas des Bourboux, *Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100*, *JCAP* **06** (2017) 047 [[arXiv:1702.03314](#)] [[INSPIRE](#)].
- [70] D. Foreman-Mackey, D.W. Hogg, D. Lang and J. Goodman, *emcee: The MCMC Hammer*, *Publ. Astron. Soc. Pac.* **125** (2013) 306 [[arXiv:1202.3665](#)] [[INSPIRE](#)].
- [71] M. Viel, J. Lesgourgues, M.G. Haehnelt, S. Matarrese and A. Riotto, *Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman- α forest*, *Phys. Rev. D* **71** (2005) 063534 [[astro-ph/0501562](#)] [[INSPIRE](#)].
- [72] M. Ishak, *Testing general relativity in cosmology*, *Living Rev. Rel.* **22** (2019) 1 [[arXiv:1806.10122](#)] [[INSPIRE](#)].
- [73] C. Dvorkin, K. Blum and M. Kamionkowski, *Constraining dark matter-baryon scattering with linear cosmology*, *Phys. Rev. D* **89** (2014) 023519 [[arXiv:1311.2937](#)] [[INSPIRE](#)].
- [74] R. Krall, F.-Y. Cyr-Racine and C. Dvorkin, *Wandering in the Lyman-alpha Forest: A Study of Dark Matter-Dark Radiation Interactions*, *JCAP* **09** (2017) 003 [[arXiv:1705.08894](#)] [[INSPIRE](#)].
- [75] M. Archidiacono, D.C. Hooper, R. Murgia, S. Bohr, J. Lesgourgues and M. Viel, *Constraining Dark Matter-Dark Radiation interactions with CMB, BAO, and Lyman- α* , *JCAP* **10** (2019) 055 [[arXiv:1907.01496](#)] [[INSPIRE](#)].