

PAPER • OPEN ACCESS

Emulating the Lyman-Alpha forest 1D power spectrum from cosmological simulations: new models and constraints from the eBOSS measurement

To cite this article: Michael Walther *et al* JCAP05(2025)099

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

You may also like

- [Estimation of the Lyman- signal of the EFIL diagnostic under static or radiofrequency electric field in vacuum](#)
Carlo POGGI, Théo GUILLAUME, Fabrice DOVEIL et al.
- [Quantum state preparation of hydrogen atoms by hyperfine quenching](#)
Amanda Alencar, I Prazeres, C R de Carvalho et al.
- [High-precision x-ray spectroscopy of highly charged ions with microcalorimeters](#)
S Kraft-Bermuth, V Andrianov, A Bleile et al.

RECEIVED: December 20, 2024

REVISED: April 24, 2025

ACCEPTED: April 27, 2025

PUBLISHED: May 29, 2025

Emulating the Lyman-Alpha forest 1D power spectrum from cosmological simulations: new models and constraints from the eBOSS measurement

Michael Walther^{a,b,*}, Nils Schöneberg^{a,b,*}, Solène Chabanier^c, Eric Armengaud^d, Jean Sexton^c, Christophe Yèche^d, Julien Lesgourgues^e, Markus R. Mosbech^{e,f}, Corentin Ravoux^{g,d}, Nathalie Palanque-Delabrouille^{c,d} and Zarija Lukić^c

^aUniversity Observatory, Faculty of Physics, Ludwig-Maximilians-Universität, Scheinerstr. 1, 81677 Munich, Germany

^bExcellence Cluster ORIGINS, Boltzmannstr. 2, 85748 Garching, Germany

^cLawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720, U.S.A.

^dUniversité Paris-Saclay, CEA, IRFU, 91191, Gif-sur-Yvette, France

^eInstitute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University, 52056 Aachen, Germany

^fInstitute for Theoretical Particle Physics (TTP), Karlsruhe Institute of Technology (KIT), 76128 Karlsruhe, Germany

^gUniversité Clermont-Auvergne, CNRS, LPCA, 63000 Clermont-Ferrand, France

E-mail: michael.walther@lmu.de, nils.science@gmail.com, chabaniersolene@gmail.com, eric.armengaud@cea.fr, jmsexton@lbl.gov, christophe.yeche@cea.fr, lesgourg@physik.rwth-aachen.de, mosbech@physik.rwth-aachen.de, corentin.ravoux@clermont.in2p3.fr, npalanque-delabrouille@lbl.gov, zarija@lbl.gov

ABSTRACT: We present the Lyssa suite of high-resolution cosmological simulations of the Lyman- α forest designed for cosmological analyses. These 18 simulations have been run using the Nyx code with 4096^3 hydrodynamical cells in a 120 Mpc (~ 81 Mpc/ h) comoving box and individually provide sub-percent level convergence of the Lyman- α forest 1d flux power spectrum. We build a Gaussian process emulator for the Lyssa simulations in the `lym1d` likelihood framework to interpolate the power spectrum at arbitrary parameter values. We validate this emulator based on leave-one-out tests and based on the parameter constraints for simulations outside of the training set. We also perform comparisons with a previous

*Corresponding author.

emulator, showing a percent level accuracy and a good recovery of the expected cosmological parameters. Using this emulator we derive constraints on the linear matter power spectrum amplitude and slope parameters $A_{\text{Ly}\alpha}$ and $n_{\text{Ly}\alpha}$. While the best-fit Planck Λ CDM model has $A_{\text{Ly}\alpha} = 8.79$ and $n_{\text{Ly}\alpha} = -2.363$, from DR14 eBOSS data we find that $A_{\text{Ly}\alpha} < 7.6$ (95% CI) and $n_{\text{Ly}\alpha} = -2.369 \pm 0.008$. The low value of $A_{\text{Ly}\alpha}$, in tension with Planck, is driven by the correlation of this parameter with the mean transmission of the Lyman- α forest. This tension disappears when imposing a well-motivated external prior on this mean transmission, in which case we find $A_{\text{Ly}\alpha} = 9.8 \pm 1.1$ in accordance with Planck.

KEYWORDS: Lyman alpha forest, cosmological parameters from LSS

ARXIV EPRINT: [2412.05372](https://arxiv.org/abs/2412.05372)

Contents

1	Introduction	1
2	Lyssa simulations and inference pipeline	4
2.1	Individual simulation design	4
2.2	Cosmological parameterization and suite design	6
2.3	Emulator	8
2.4	Likelihood and inference procedure	10
3	Emulator validation	10
3.1	Leave-one-out power spectra comparison	10
3.2	Fitting the fiducial model	11
4	Comparison with pre-existing simulations	13
4.1	A caveat and its correction: initial conditions	14
4.2	Fitting previous B13 simulations	15
5	Fitting observational data	17
5.1	Additional analytical corrections for observational nuisances	17
5.2	Application to eBOSS data — baseline analysis	18
5.3	Variations with more informative priors	20
6	Conclusions	22
A	Description of nuisance parameters	25
B	Variations of the baseline analysis	26
C	Best-fitting power spectra	27
D	Additional plots of constraints on nuisance parameters	27

1 Introduction

The Lyman- α forest consists of a densely spaced series of absorption features in the spectra of distant quasars. These features track the distribution of neutral hydrogen and can consequently be used to probe cosmological structures. As the Lyman- α forest is sensitive to high redshift ($2 < z < 5$) structures at small scales up to around $k_{\text{max}} \sim 5/\text{Mpc}$, it probes structure formation at the smallest scales in the mildly nonlinear regime. However, the Lyman- α forest is also very sensitive to the thermal state of the intergalactic medium (IGM), such as its mean temperature or the pressure smoothing scale, and thus computationally expensive hydrodynamical simulations are typically required to model this observable with high accuracy. For a review, see [1, 2].

The 1d power spectrum of the Lyman- α forest ($P_{F,1d}$) is the power spectrum of the so-called transmission contrast $\delta_F = F/\bar{F} - 1$ where F is the flux transmission of the Lyman- α forest and \bar{F} is its mean. Recent interest for $P_{F,1d}$ has grown due to it being an excellent probe of many models that predict a small-scale suppression or enhancement of power, as well as from claims of a high tension between the power measured with the Lyman- α forest and that expected from the CMB anisotropies [3–5] (see also below). In particular, due to the high reach in wavenumbers and redshift, high precision measurements of the $P_{F,1d}$ can be used to put precise constraints on many dark matter candidates (warm dark matter, interacting dark matter, fuzzy dark matter), neutrino masses, inflationary scenarios predicting a running of the power spectrum, and other models that impact the small-scale power spectrum [3–5].

The Lyman- α forest has been measured to high precision by the extended Baryon Oscillation Sky Survey (eBOSS), which is part of the Sloan Digital Sky Survey, using data release 14 (SDSS DR14), where around 44 000 quasars provide a near-percent level statistical error on $P_{F,1d}$ [6]. Furthermore, the ongoing Dark Energy Spectroscopic Instrument (DESI) survey [7] will also measure $P_{F,1d}$ at excellent precision, with improved statistics and spectrograph resolution [8]. Additionally, there are various datasets containing smaller amounts of high-resolution quasar spectra allowing to extend the survey measurements to smaller scales, albeit at lower precision (see [9]). These also require a high degree of accuracy in the theoretical predictions. As such, there is a great need for highly converged, accurate simulations of the Lyman- α forest.

There have been a number of recent approaches to simulating $P_{F,1d}$ with different simulation tools and parametrizations, which we summarize in table 1. Notably, about a decade ago the authors of [10] have created a suite of simulations for the purpose of analyzing BOSS data based on the **gadget** code (hereafter B13), splicing together low- and high-resolution simulations (see table 1). The latest analysis of eBOSS data with this set of simulations found a tension in the cosmological parameters compared to those derived from CMB anisotropies within the Λ CDM framework. In the original work by [3], based on the full likelihood modeling of the Lyman- α data, this tension was found at the 3.6σ level. More recent studies, focusing on the slope and amplitude of the matter power spectrum at the Lyman- α scale, made use of the likelihood contours in this reduced parameter space (previously published in [6]), and found a higher tension at the level of 4.9σ [4, 5].

Recent simulation suites listed in table 1 have typically not been applied to eBOSS data (due to insufficient resolution, lack of convergence, or not varying cosmological parameters), with the exception of the PRIYA suite described in [11, 12]. This suite involves a number of low- and high-resolution simulations that are used to build a multi-fidelity emulator.

In this paper, we present a new suite of simulations, which we dub “Lyman- α simulation suite in absorption” (Lyssa), designed to provide high precision predictions for the one-dimensional flux power spectrum of the Lyman- α forest at both small and large scales with a box of size $L = 120$ Mpc and resolution 29.3 kpc, defined as $L/N^{1/3}$ where N is the total particle number. As such, each simulation is covering modes from $k_{\min} = 0.052 \text{ Mpc}^{-1}$ ($= 2\pi/L$) to the Nyquist frequency $k_{\text{Nyq}} = 107 \text{ Mpc}^{-1}$. However the actual maximum wavenumber used in this work is $k_{\max} = 5 \text{ Mpc}^{-1}$, to ensure proper convergence to the sub-% level in the wavenumber range corresponding to current Lyman- α forest data. Using these

Simulation suite	Box size	Resolution	Code
B13* [10]	25 Mpc/h \sim 37.0 Mpc	$130.2 h^{-1}$ kpc \sim 193 kpc	GADGET-3
	100 Mpc/h \sim 148 Mpc	$130.2 h^{-1}$ kpc \sim 193 kpc	
	100 Mpc/h \sim 148 Mpc	$32.55 h^{-1}$ kpc \sim 48.2 kpc	
B19* [18]	40 Mpc/h \sim 57.1 Mpc	$156.3 h^{-1}$ kpc \sim 223 kpc	MP-GADGET
THERMAL [19]	20 Mpc/h \sim 29.8 Mpc	$19.5 h^{-1}$ kpc \sim 29.1 kpc	Nyx
P21* [20–22]	\sim 47.3 Mpc/h 67.5 Mpc	\sim 61.4 h^{-1} kpc 87.7 kpc	MP-GADGET
SHERWOOD ⁺ [23, 24]	40 Mpc/h \sim 59.0 Mpc	$19.5 h^{-1}$ kpc \sim 28.8 kpc	P-Gadget3
PRIYA [11, 12]	120 Mpc/h \sim 171 Mpc	$39.1 h^{-1}$ kpc \sim 55.9 kpc	MP-GADGET
	120 Mpc/h \sim 171 Mpc	$78.1 h^{-1}$ kpc \sim 111 kpc	
Lyssa (this work)	\sim 80.8 Mpc/h 120 Mpc	\sim 19.7 h^{-1} kpc 29.3 kpc	Nyx

Table 1. Summary of different simulation suites for $P_{F,1d}$ Lyman- α analyses. The conversions have been performed with the given central values of h or the mean of the ranges, and are marked with “ \sim ” as different simulations do not necessarily share the same box size in those units. Note that for SPH codes, mass resolution is fixed and the “Resolution” column shows the value for gas at mean density. In over-/underdense regions the SPH resolution is accordingly better/worse.

* These simulation suites are not named. They are described in detail in references [10] (B13), [18] (B19) and [20] (P21), respectively. Note that the B13 suite has been used to generate the Taylor emulator we compare to in this work.

⁺ In this case we are only quoting the simulation with the best resolution among those with the highest number of particles.

simulations, we build an emulator that can be used to obtain constraints on the small scale power spectrum amplitude and slope. This information can then, in turn, be compared to the predictions of various cosmological models, in particular those claiming to ease the σ_8 tension (see for example [13, 14]). This new simulation suite is targeted for analyses of the $P_{F,1d}$ in eBOSS and DESI, though it might also be useful in future analysis of the three-dimensional flux transmission power spectrum.

Comparing the Lyssa simulations to the high-fidelity simulations of [11], we see that Lyssa has 18 compared to 3 high-resolution simulations. Instead, Lyssa is missing the 48 low fidelity simulations (which are used in [11] to construct a multi-fidelity emulator) and therefore covers cosmological parameter space with a lower density. Compared to the high-fidelity simulations, Lyssa has a smaller box, but higher resolution at mean density. However, the resolution at mean density does not necessarily provide a complete description of convergence properties. For example, [15] compared Nyx to CRK-HACC showing that Nyx reaches similar $P_{F,1d}$ accuracy with half the resolution at mean density at $z = 3$ if SPH-based density estimation is used for CRK-HACC (and better accuracy at $z = 5$ under the same conditions). Given the smaller box size, we caution about cosmic variance concerns, see [10, 11, 16]. More detailed comparisons of the fidelity of the simulations between [11, 15], and this work as well as more detailed convergence studies overall are left for future work. For resolution convergence tests of our setup based on smaller simulation boxes, see [17].

This paper is structured as follows. In section 2 we describe the new high-precision suite of simulations “Lyssa”, the choice of the thermal and cosmological parameter basis, as well as the overall likelihood pipeline. In section 3 we validate the emulator using leave-one-out




Figure 1. A rendering of a $21 \text{ Mpc} \times 21 \text{ Mpc} \times 120 \text{ Mpc}$ cutout from one of the simulation boxes. Brighter colors correspond to larger baryon density.

tests and cosmological parameter fits. In section 4 we compare with simulations from B13 and check the agreement with the results of [3]. In section 5 we then obtain constraints from the eBOSS data using the emulator, and we conclude in section 6.

2 Lyssa simulations and inference pipeline

In this section, we describe the set of simulations and the corresponding emulator that has been built in preparation for the upcoming DESI data. We highlight the similarities and differences of Lyssa with respect to other similar simulation suites, and in particular with the simulation set from [10] for which an emulator has been designed based on a Taylor approximation around a central simulation. We refer to this emulator as the *Taylor emulator*.

2.1 Individual simulation design

To generate an individual model of the Lyman- α forest power spectrum, we proceed as follows. First, we run a cosmological hydrodynamical simulation with the *Nyx* code¹ [25–27]. *Nyx* evolves a hydrodynamical fluid on an Eulerian grid, which is gravitationally coupled to N-body particles tracing the dark matter, all deposited on the same grid. While adaptive mesh refinement techniques are available in *Nyx*, they are inefficient for Lyman- α forest simulations as the gas of interest spans most of the simulation volume. The Euler equations are solved with a second-order accurate scheme, and key physical processes affecting the intergalactic medium — such as photoionization, collisional ionization, recombination, heating, and atomic cooling — are included in the simulation. We show an example cutout of the simulation box in figure 1.

¹<https://github.com/AMReX-Astro/Nyx>.

Based on the convergence considerations of [15, 17, 26], each simulation run is using 4096^3 hydrodynamical cells (plus an equal amount of dark matter particles) in a 120 Mpc comoving box.² Notably, the simulations cover the whole range of scales probed by the $P_{F,1d}$ measurement of DESI, without resorting to the splicing method used in [10] for the Taylor emulator. We initialize the simulations at $z_{\text{ini}} = 99$ and take snapshots of the evolved fluid properties and dark matter densities at thirteen redshifts in intervals of 0.2 between $z = 2.2$ and 4.6, as well as two additional high-redshift snapshots at redshifts 5.0 and 5.4.

To generate initial conditions for our simulations, we use `2lpt-ic` [28, 29] and produce particles following the total matter transfer function at z_{ini} (see section 4.1 for a comparison with the individual transfer functions for both fluids and a different IC code: `MUSIC2-monofonIC`). The baryonic part of the particle mass is then subtracted from the particles and deposited on a regular mesh using a cloud-in-cell scheme.

The effect of a time-varying, spatially uniform ultraviolet background (UVB) on photoionization and photo-heating is modeled by the code based on a list of redshift-dependent rates [30]. These rates were derived from the late reionization model of [31], and a warm He II reionization scenario with $\Delta T = 20000K$, whose resulting thermal evolution is in good agreement with the measurement of [19]. To allow different thermal evolutions, we rescale the overall heating rates by an additional density- and redshift-independent factor A_{UVB} . Considering two simulations with different A_{UVB} , even after rescaling their snapshot's temperature at a given redshift to a common value (as described below), they will still have different gas pressure smoothing scales λ_P which depends on the full thermal history of the gas within the simulation (see for example [32]). Therefore, within this approach, the A_{UVB} parameter allows us to vary the smoothing scale. By interpolating between simulations with different A_{UVB} values for different redshifts, an arbitrary $\lambda_P(z)$ history can be emulated.

After the simulations have been run, we generate Lyman- α forest spectra from the snapshots and compute 1d- and 3d- power spectra of different quantities (baryons, dark matter, matter, Lyman- α forest flux transmission), making use of the `gimlet` postprocessing suite [33]. This software package allows for efficient, MPI-parallel computation of typical statistical properties in a `Nyx` simulation based on the hydrodynamical grid and cloud-in-cell depositions of the dark matter particles. The extraction of flux skewers is performed along each of the principal axes of the box and statistical properties are averaged over all 3 axes to reduce cosmic variance. As described in [17], to allow generating a variety of thermal states at each redshift we additionally rescale the output temperatures in all pixels by a power law in $(1 + \delta) = \rho_b / \bar{\rho}_b$, where ρ_b is the baryon density and $\bar{\rho}_b$ is its spatial average. This allows us to emulate, with arbitrary redshift dependence, the coefficients T_0 and γ of the usual temperature-density relation

$$T(\delta, z) = T_0(z)(1 + \delta)^{\gamma(z)-1} \quad (2.1)$$

Explicitly, we do not set all pixel temperatures equal to the power law, but instead we keep the scatter of the temperature-density relation intact. According to [34] this allows for

²Note that the Lyssa fiducial simulation is in fact the largest simulation run in [17] (labeled 120 Mpc 4096^3). Both resolution and box size convergence have been studied there. For an internal cross-check on cosmic variance, we compared extractions taken along the different simulation axes. Further tests using initial conditions with different seeds are planned for future work.

computing $P_{\text{F,1d}}$ with sub-percent accuracy for $k < 0.03 \text{ s km}^{-1}$ (at $z = 3.0$) compared to running a new simulation with a rescaled thermal history. Finally, to match a range of mean transmissions, we rescale the resulting optical depths with a global factor f_τ .

We evaluate λ_P by fitting a cutoff power law $Ak^n \exp(-k^2 \lambda_P^2)$ to the “real space” three-dimensional flux power spectrum, which is estimated similarly to the (redshift-space) flux power spectrum, but ignoring line broadening and peculiar velocity shifts during skewer generation. When computing the real space flux field, we rescale the real space optical depth of the simulation (τ_{real}) by the same factor f_τ which is needed for the redshift space transmission to match the fiducial mean transmission (\bar{F}) for a given redshift.

We acknowledge that some physical effects are not included in our simulation suite. These will be included as additional corrections during the likelihood fit, see section 5.1. While we already account for SN and AGN feedback in this work (similar to the previous eBOSS analysis by [3]), further corrections of e.g. inhomogeneous He II-reionization are left for future work. Therefore, the constraints on cosmological parameters shown in section 5 might change once these physical effects are marginalized over (see e.g. [11, 35, 36] for the effects of He II reionization on the $P_{\text{F,1d}}$).

2.2 Cosmological parameterization and suite design

For the purpose of the Lyman- α $P_{\text{F,1d}}$ analysis, we would ideally sample simulation parameters from a broad region based on the (expected) constraints of current (and upcoming) data.

Some simulation suites dedicated to the cosmological analysis of the flux power spectrum, including in particular the B13 suite and its Taylor emulator [3, 10] have used cosmological base parameters widely used within the cosmological community (especially n_s and A_s/σ_8), which were commonly used in CMB anisotropy measurements and low-redshift galaxy surveys. However, these cannot be connected in a straightforward way to the Lyman- α forest data probing both smaller scales than the CMB and higher redshifts than most galaxy surveys.

Other recent works such as [18, 21] have adopted a different basis of cosmological parameters, more adapted to the actual sensitivity of $P_{\text{F,1d}}$ measurements. In particular, as already found by [37], $P_{\text{F,1d}}$ best constrains the amplitude and slope of the linear matter power spectrum computed at the redshift and pivot wavenumber of the Lyman- α data. Following this approach, we define the parameters $A_{\text{Ly}\alpha}$ and $n_{\text{Ly}\alpha}$ as the amplitude and logarithmic slope of the linear matter power spectrum $P_{\text{lin}}(k_p, z_p)$ at a pivot redshift of $z_p = 3$ and a pivot scale of $k_p = 1 \text{ Mpc}^{-1}$.³ For comparison, [21] makes use of the parameters Δ_\star and n_\star at $k_p = 0.009 \text{ s km}^{-1}$ (about 0.7 Mpc^{-1}), while [11] defines A_p as the amplitude and n_p as the slope of the primordial power spectrum at $k_p = 0.78 \text{ Mpc}^{-1}$. Note that conversions between Δ_\star^2, n_\star and $A_{\text{Ly}\alpha}, n_{\text{Ly}\alpha}$ can be approximated as

$$\Delta_\star^2 \approx A_{\text{Ly}\alpha} \frac{k_\star^3}{2\pi^2} \frac{H(z)}{(1+z)} \left[\frac{k_\star}{k_{\text{Ly}\alpha}} \right]^{n_{\text{Ly}\alpha}}, \quad (2.2)$$

$$n_\star \approx n_{\text{Ly}\alpha} \quad (2.3)$$

³In practice we derive these parameters by convolving the power spectrum with a narrow Gaussian of width $\Delta \ln k = 0.1$ (for the slope we use the derivative of a Gaussian). This allows us to reduce numerical errors from interpolation. We have checked that the results are consistent with similar definitions. Whenever we show $A_{\text{Ly}\alpha}$ it is implicitly assumed to be in units of Mpc^3 .

or otherwise converted, for example by building an emulation scheme for their mapping (which we employ in figure 10 and table 3).

Within the flat Λ CDM model, the other parameters we use for the suite are the physical matter density $\Omega_m h^2$ and h (defined as $H_0 = h \cdot 100 \text{ km/s/Mpc}$). Indeed, as highlighted in [16, 21], the sensitivity of $P_{F,1d}$ measurements to Ω_m and H_0 is driven, at first order, by the conversion factor between inverse velocity units (s km^{-1}) from spectroscopic observations and physical units in Mpc^{-1} , which is:

$$\frac{H(z)}{1+z} = \frac{100 \text{ km/s/Mpc}}{1+z} \sqrt{\Omega_m h^2 [(1+z)^3 - 1] + h^2} \quad (2.4)$$

At redshift $z > 2$, the effect of dark energy on the Hubble rate is subdominant, so that to a good approximation this conversion factor scales like $\sqrt{\Omega_m h^2 (1+z)}$. To avoid artificial parameter dependencies, we therefore use $\Omega_m h^2$ as a design parameter. While our suite also varies h as a cosmological parameter, we confirm in the following section that, as expected according to [16, 21], realistic $P_{F,1d}$ measurements are not expected to provide a significant sensitivity to this additional parameter.

We explicitly keep $\Omega_b h^2$ fixed to 0.02233 based on [38] and our simulations are run with massless neutrinos. As a matter of fact, the flux power spectrum measured from simulations with either massive or massless neutrinos are nearly indistinguishable from each other provided that one performs a re-mapping of redshifts within the massless neutrino simulation such that the linear matter power spectrum amplitude evolution $A(z)$ is equivalent to that of the massive neutrino simulation. Therefore, the sum of neutrino masses can still be constrained using this simulation suite, as was demonstrated in [16, 21]. To summarize, in total our simulations are characterized by a set of five parameters $\{A_{Ly\alpha}, n_{Ly\alpha}, \Omega_m h^2, h, A_{UVB}\}$. The UV background scaling parameter A_{UVB} was discussed in section 2.1.

The original layout of the suite, illustrated in figure 2, is a Latin hypercube design in the varied parameters using a total of 14 simulations. It also includes a “fiducial” simulation at the approximate center of parameter ranges, which was already described in [17]. For historic reasons, this simulation features a slightly different box length (see table 2) and was initialized with `camb` instead of `class`. Later, we also added 4 refinement simulations to increase the coverage of lower values of $\Omega_m h^2$. The cosmological parameters of each of the 18 simulations are given in table 2.

In addition to the parameters mentioned above, at each redshift and for each simulation, we generate a wide latin hypercube design in the parameters describing the thermal model, that is, the mean transmission $\bar{F}(z)$ and the temperature-density relation of eq. (2.1), using the approach described in section 2.1. We use a combination of two different thermal model designs:

- In the main suite, we sample the ranges $3800 \text{ K} < T_0 < 27800 \text{ K}$, $1.0 < \gamma < 1.9$, and $|\bar{F} - \exp(-0.0025(1+z)^{3.7})| < 0.15$. For each of the 18 simulations we post-process 4 thermal samples for a total of 72 samples per redshift, see the purple points of figure 2.

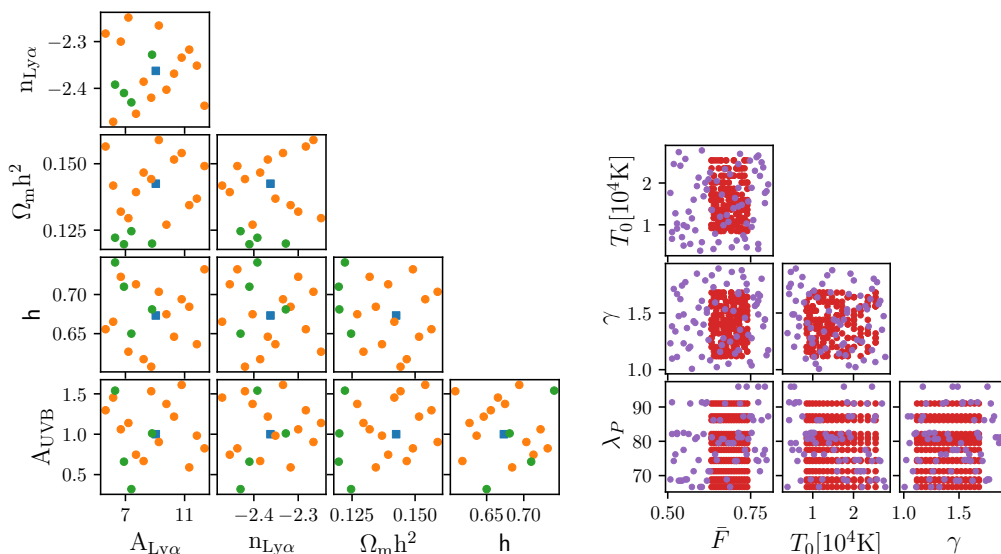


Figure 2. The parameter designs sampled by our hydrodynamical simulations. Left: cosmological parameter space. Here the orange points are from the base latin hypercube design, the green points are from the refinement simulations and the blue point is from the fiducial. Right: thermal parameter space at $z = 3$ (for the redshift evolution, see figure 3). In purple we show the main suite, and in red the auxilliary suite. The sparser clustering of λ_P around specific values arises from strong anti-correlation with A_{UVB} (as it is not explicitly modified in post-processing). The values for λ_P (measured here in kpc) compare well to existing measurements like [39].

- In an auxilliary suite we use a narrower Latin hypercube design of 15 thermal models per cosmological parameter point of the base suite ($15 \times 14 = 210$) (red points in figure 2) forming a Latin hypercube design for each simulation.⁴

We show the evolution of the covered thermal parameter ranges in figure 3.

2.3 Emulator

We use the previously described set of simulations with a Gaussian process emulator based on the `george` package [42]. The emulator code is publicly available as part of the `lym1d` package.⁵ The emulator code builds a separate Gaussian process at each redshift that predicts all wavenumber bins of the one-dimensional flux power spectrum (in Mpc units) as a function of the cosmological and thermal parameters.

The total set of 8 input parameters for the emulator is thus

$$\{A_{Ly\alpha}, n_{Ly\alpha}, \Omega_m h^2, h, T_0(z), \gamma(z), \bar{F}(z), \lambda_P(z)\}. \quad (2.5)$$

⁴Those were generated by exploring the ranges covered by the 3 parameters at each redshift for bounds defined via, firstly, $\bar{F} = \exp(-0.0025 A_F (1+z)^{B_F+3.7} (1+z_p)^{-B_F})$ for $0.7 < A_F < 1.1$ and $|B_F| < 0.25$; secondly, a broken power-law in $T_0 = T_p((1+z)/(1+z_p))^{\eta(z)}$ for $12000 < T_p/K < 18000$, $|\eta(z < z_p) - 2.1| < 0.05$, and $|\eta(z > z_p) - 3.4| < 0.05$; and thirdly, a power law in $\gamma = \gamma_p((1+z)/(1+z_p))^\xi$ with $1.3 < \gamma_p < 1.7$ and $|\xi| < 0.1$. We used $z_p = 3$ for all parameters.

⁵The `lym1d` python package is freely available at <https://github.com/schoeneberg/lym1d>. It includes the Gaussian Process emulator presented here, a version of the Taylor expansion used in [3], and the likelihood framework used throughout this work. It does not include the Lyssa simulations because of the large file size.

Simulation	$\Omega_m h^2$	h	$A_{\text{Ly}\alpha}$	$n_{\text{Ly}\alpha}$	A_{UVB}
fiducial*	0.1425	0.6732	9.0514	-2.3626	1.0000
0	0.1516	0.6461	10.2857	-2.3686	1.2200
1	0.1589	0.7035	9.2571	-2.2657	0.9036
2	0.1271	0.6748	9.7714	-2.4029	1.3750
3	0.1344	0.6844	11.3143	-2.3171	0.5893
4	0.1442	0.6078	8.7429	-2.4200	1.5321
5	0.1369	0.6365	11.8286	-2.3514	0.9821
6	0.1540	0.6939	10.8000	-2.3343	1.6107
7	0.1418	0.6652	6.1714	-2.4714	1.4535
8	0.1320	0.7226	6.6857	-2.3000	1.0607
9	0.1393	0.7131	7.7143	-2.4543	0.7464
10	0.1565	0.6556	5.6571	-2.2829	1.2964
11	0.1467	0.6173	8.2286	-2.3857	0.6679
12	0.1295	0.6267	7.2000	-2.2486	1.1393
13	0.1491	0.7322	12.3429	-2.4371	0.8250
14	0.1197	0.7100	6.9000	-2.4100	0.6600
15	0.1222	0.7400	6.3000	-2.3920	1.5400
16	0.1246	0.7100	7.4000	-2.4300	0.3200
17	0.1200	0.6810	8.8000	-2.3280	1.0100

Table 2. Cosmological parameters of the simulations in our suite. The baryon density is fixed to $\Omega_b h^2 = 0.02233$. See also figure 2.

* This simulation has been run using transfer functions from CAMB with $A_s = 2.1 \times 10^{-9}$ and $n_s = 0.966$, but the corresponding values of $A_{\text{Ly}\alpha}$ and $n_{\text{Ly}\alpha}$ were derived from those parameters using CLASS. It also uses a slightly different box size of $80 \text{ Mpc}/h \approx 118.8 \text{ Mpc}$ whereas the others use 120 Mpc .




Figure 3. Redshift evolution of the thermal parameters in the different post-processings of the simulations. We show in purple/red the entire covered region (minimum and maximum covered values) of the main/auxilliary suite, and in black the fiducial thermal history. The black points represent data from [40] (T_0 and γ) and [41] (\bar{F}) just for reference.

In addition to this default set of basis parameters, a few alternative parameterizations are also implemented. First, it is possible to use $\{\Delta_*, n_*\}$ or $\{\sigma_8, n_s\}$ instead of $\{A_{\text{Ly}\alpha}, n_{\text{Ly}\alpha}\}$. Second, the native parameter A_{UVB} may be used in place of λ_P . For the Gaussian process we use a stationary Matern-5/2 covariance function with an independent length scale l_i for

each of the eight emulation parameters i and an overall scaling σ_0 , which were optimized for each wavenumber and redshift bin independently. The typical length scales we obtain fall between 3 and 20 times the range covered by simulations for each parameter. For the hyperparameter optimization, we use all thermal re-processings for all simulations, leading to a total of 297 points per redshift. More details about how the emulator is generated can be found in [17].

When using the emulator one can either specify the thermal parameters $\{T_0, \gamma, \bar{F}, \lambda_P\}$ at each redshift independently or describe them by (broken) power laws. For simplicity, and in order to better compare with previous fits of eBOSS data, within this work we use a parameterization equivalent to the one assumed in [3] (except for the validation in section 3.1): broken power laws with a pivot redshift $z_p = 3$ for the functions $T_0(z)$ and $\lambda_P(z)$, and power laws for the functions $\gamma(z)$ and $\tau_{\text{eff}}(z) = -\ln \bar{F}(z)$. For instance, the temperature at mean density and the effective optical depth are parameterized explicitly as

$$T_0(z) = T_0(z_p) \left[\frac{1+z}{1+z_p} \right]^{\eta_T (+\Delta\eta_T \text{ if } z > z_p)}, \quad \tau_{\text{eff}}(z) = A_\tau \left[\frac{1+z}{1+z_p} \right]^{\eta_\tau}. \quad (2.6)$$

The cosmological parameters are of course not varied across redshift. For the parameters and priors related to this parameterization, see appendix A.

2.4 Likelihood and inference procedure

The likelihood is based on a simple Gaussian assumption with the covariance matrix taken from the original analysis performed in [3]. The computation of the theoretical flux power spectrum is performed using the emulator outlined in section 2.3 with possible nuisance corrections (as for example required for section 5.1) added on top. The likelihood pipeline code is freely available in the `lym1d` package.

In order to perform the MCMC sampling we use the `MontePython v3.0` package [43]. The corner plots are produced with `liquidcosmo`, which is a free chain analysis and plotting utility available at <https://github.com/schoeneberg/liquidcosmo>.

3 Emulator validation

In this section we present the results of the validation of the emulator described in section 2.3 applied to the $P_{\text{F},1d}$ inference. We first perform a number of leave-one-out tests to gauge the inherent systematic uncertainty of the emulator and subsequently fit the $P_{\text{F},1d}$ of the fiducial simulation to check if we correctly recover the cosmological and astrophysical parameters.

3.1 Leave-one-out power spectra comparison

In this first test we check that the emulator correctly reproduces the power spectrum of a given simulation when only the other remaining simulations are used to train it. Explicitly, all astrophysical post-processings of one simulation are removed from the training set while all post-processings of other simulations are kept. We then attempt to emulate the intrinsic astrophysical evolution of the left out simulation. This allows us to estimate the accuracy of the emulator's predictions for the transmission power spectrum.

Note that this result actually quantifies the joint residual from two effects: a) the accuracy of the astrophysical post-processing when generating different realizations of $\{T_0, \gamma, \bar{F}\}$ and b) the interpolation accuracy of the emulator in different parts of the cosmological parameter space⁶ in the worst-case scenario of reduced model density at the point of interest.

Figure 4 shows the result of the leave-one-out test for redshifts between 2.2 and 4.6. The accuracy of the emulation procedure is at the percent level on all relevant wavenumbers for the considered redshift bins as long as the left out simulation is not on the edge of the cosmological parameter space used for training. We quantify whether a left out simulation is “inside” or “on the edge” of the training volume by checking whether the smallest hypercube encompassing the parameters of all remaining simulations contains the parameters of the left out simulation or not.

At $z \leq 2.6$, we find an emulation error of up to around 2% for all models “inside” the training volume. The accuracy is better than 1% (corresponding to the grey bands) at higher redshift. There is a notable exception at $z = 3.6$, where a critical computing issue in one of the hydrodynamical simulations resulted in a corrupted snapshot at this redshift. This prevented the generation of the main suite of thermal models, such that only the auxiliary suite was available in this case. Attempts to re-simulate the missing snapshot are ongoing.

Our emulator is thus able to achieve good accuracy. We will now check whether this accuracy is sufficient for inferring cosmological parameters from mock data without significant bias.

3.2 Fitting the fiducial model

In this section, we check whether the emulator can recover the cosmological and thermal parameters in a full inference run, using some mock data created from a simulation that was not used to train the emulator. For this purpose, we use the fiducial simulation mentioned in section 2.2. This test does not only check that the fiducial power spectrum is correctly recovered at fiducial parameter values, but also that there are no other parameter combinations in a different but nearby region of parameter space that predict the same power spectrum.

To create mock data, we measure the flux power spectrum in the fiducial simulation at the wavenumbers used in eBOSS measurements (in units of s/km). We do not add astrophysical (DLA, metals, ...) or instrumental corrections to this mock data. In order to obtain some realistic uncertainties, we rely on the eBOSS data analysis. Thus, our mock likelihood incorporates the same covariance matrix as in [3].

The results of this test are shown in figure 5. The emulator recovers very well the cosmological parameters of the fiducial simulation, and in particular the amplitude $A_{\text{Ly}\alpha}$ and tilt $n_{\text{Ly}\alpha}$ to which Ly α data are most sensitive. We even find some sensitivity to $\Omega_{\text{m}} h^2$, while the posteriors on H_0 are almost flat, as expected given the arguments of [16, 21]. The thermal nuisance parameters are also nicely recovered, as shown in figure 14.

Since the emulator does not provide significant constraints on H_0 or Ω_{m} individually, we also check whether setting priors in the $\{\Omega_{\text{m}}, h\}$ plane biases the results on $A_{\text{Ly}\alpha}$ or $n_{\text{Ly}\alpha}$. This check is performed in view of combining the Lyman- α data with data from CMB experiments such as Planck. We use a prior with a covariance matrix in $\{\Omega_{\text{m}}, h\}$ derived from the Planck 2018 TTTEEE+lensing results [38], centered at the fiducial simulation.

⁶This also includes variations of the heating rate and thus emulation in λ_P .




Figure 4. Results of leave-one-out tests. In each panel (corresponding to different redshift bins) we display the relative difference between the emulated one-dimensional flux power spectrum and that of the original simulation. The color of the line specifies whether the parameters of the left out model are on the edge of (red) or inside (blue) the cosmological parameter hypercube spanned by the other simulations (see text for explanations). The light grey horizontal bands show the 1% error range. The dark grey band on the right is not accessible with eBOSS, but can be partially accessed with DESI.

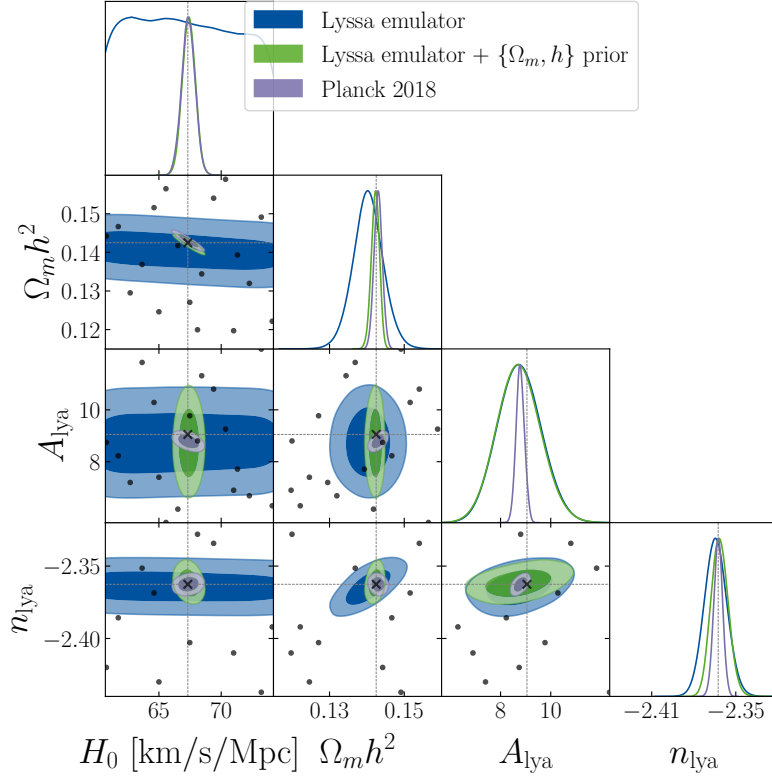


Figure 5. 68% and 95% CL credible regions and marginalised posteriors when fitting mock data extracted from the fiducial simulation with the new Lyssa-based emulator. In blue, we show the posteriors obtained with our baseline priors, summarized in table 4. The green contours and posteriors include additionally a Gaussian prior on $\{\Omega_m, h\}$ derived from Planck (see text). The purple contours and posteriors, post-processed from the Planck chains, are shown for comparison. The black points represent the parameters of the simulation suite and the black crosses the parameters of the fiducial simulation (which are very close, but not exactly equal to the Planck mean values). The grey lines show the parameter values of the fiducial simulation. The fit is able to reconstruct these values without any significant bias. For correlations with nuisance parameters, see figure 14.

Figure 5 shows that with such priors we recover almost the same constraints, with improved precision by 20% on $n_{\text{Ly}\alpha}$ due to its correlation with $\Omega_m h^2$. We stress that such a correlation between $\Omega_m h^2$ and $n_{\text{Ly}\alpha}$ is present in all of our results (see below), motivating a potential optimization of the parametrization in future studies.

4 Comparison with pre-existing simulations

In this section, we compare the Lyssa emulator to the existing B13-based Taylor emulator and check how compatible they are within our likelihood framework. We start with introducing an additional caveat due to the approach we took in generating initial conditions, and develop a correction for this effect which has to be applied when comparing to external simulations or observational measurements. Afterward, we perform inference on the fiducial B13 model using both emulators.

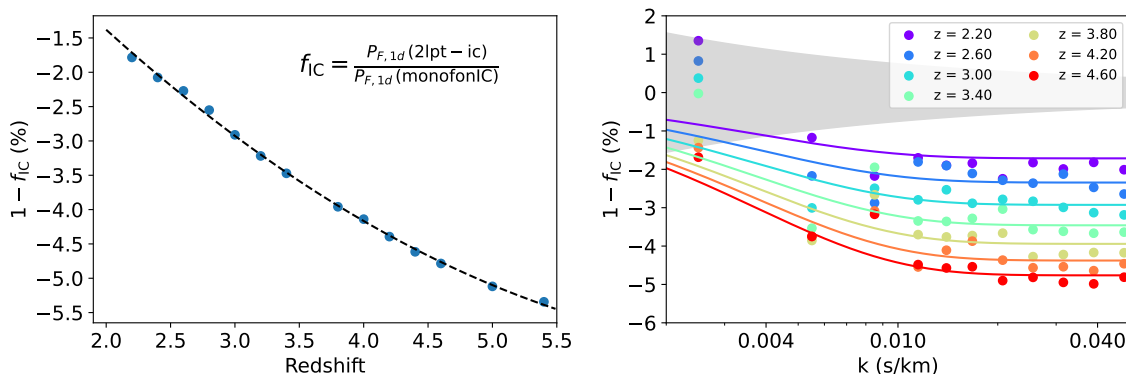


Figure 6. Difference in the $P_{F,1d}$ caused by using different initial conditions (one-particle transfer function vs split between baryons and dark matter). The corrections represent the relative difference in % between the `monofonIC` run and the corresponding `2lpt-ic` simulation. Left: mean suppression as a function of redshift for $0.01 < k < 0.05$ s/km, where points are from simulations and the dashed line is the correction factor as parameterized by eq. (4.1). Right: suppression as a function of k , plotted for a subset of redshifts, with the continuous lines from eq. (4.1). The shaded band represents an estimation of the 1σ cosmic variance associated to the finite box size.

4.1 A caveat and its correction: initial conditions

As described in section 2.1, initial conditions were generated with the `2lpt-ic` code, using a single particle species with the total matter transfer function at $z_{\text{ini}} = 99$, instead of two species, with the transfer functions of baryons and dark matter respectively.⁷ To quantify the impact of the choice of initial condition code and setting, we ran an additional simulation using the `MUSIC2-monofonIC` code [45, 46], generating two-fluid initial conditions, while using the same seed and random number generator as our `2lpt-ic` runs.

We find that, after rescaling mean optical depths to a common value, the predicted flux power spectra differ by a few percent, as shown in figure 6. The main difference is an overall suppression of the power, with a smooth evolution as a function of redshift: this suppression reaches 5% at high redshift, but it is 2% for $z \simeq 2.4$, where the precision of $P_{F,1d}$ measurements is highest. As expected, the impact of separating out baryons and dark matter initial conditions is larger at high redshift and reduces at later times. We checked that this suppression does not vary significantly when post-processing the simulations to emulate different thermal histories. The suppression does not vary as a function of the wavenumber k for $0.01 < k < 0.05$ s/km. The results indicate that the amplitude of the suppression is modified for smaller values of k , as one may expect, but the exact dependence as a function of wavenumber on large scales is difficult to quantify due to the finite box size effect, and simulation-related noise that manifests as oscillations of the flux power spectrum.

We correct for this effect using an analytical correction as a multiplicative factor f_{IC} to the emulator’s predicted $P_{F,1d}$:

$$f_{\text{IC}} = 1 - 0.01 \times [Az^2 + Bz + C][1 - \exp(-k/k_{\text{IC}})] \quad (4.1)$$

⁷These differ due to non-adiabatic perturbations seeded by the pre-recombination acoustic oscillations [44]. The species differ in both their over- and underdensities and velocities.

with $k_{\text{IC}} = 0.0037 \text{ s/km}$, $A = 0.15$, $B = -2.3$, $C = 2.6$. The main outcome of this correction is a redshift-dependent offset relevant at all wavenumbers as well as a subdominant exponential part that is relevant for $k \lesssim 0.01 \text{ s/km}$. Note that the modeled range in k is sufficient for correcting all measured scales in BOSS ($k < 0.02 \text{ s km}^{-1}$) and DESI ($k \lesssim 0.05 \text{ s km}^{-1}$).

We do not attempt to model the noise-related oscillations. We enable the correction f_{IC} when comparing the emulator with B13 simulations in section 4.2, and in the application to data shown in section 5. We also explicitly test the impact of not using the correction (see figure 8, right panel). In [47] the authors investigate the impact of using a single-species or two-species initial conditions on the `gadget`-based Lyman- α forest. A more detailed investigation into the effect of initial conditions generated from different codes on the flux power spectrum in the context of Nyx is left for future work.

4.2 Fitting previous B13 simulations

We additionally wish to test whether our new emulator is able to provide a reasonable fit to the power spectra extracted from one of the B13 simulations presented in [3, 10]. For this test, we choose the central simulation that was used as the expansion point of the old Taylor emulator. Our purpose is to assess whether the improved modeling of the flux power spectrum presented in this work is likely to change the cosmological parameter constraints derived from previous analyses.

In principle, there could be many reasons for which the new emulator would fail to extract the correct parameters from the old B13 simulation. First and foremost, the new simulations are run with much higher dynamic range, using 4096^3 dark matter particles and hydrodynamical cells each, whereas in [10] simulations were spliced together from three simulations, two with 2×768^3 particles with a large and a small box, and one with 2×192^3 particles in a small box at the same resolution as the large one. This approach required a splicing correction which was marginalized over in previous analyses, and is not needed anymore when using a larger dynamic range simulation as in the Lyssa suite. Additionally, fundamentally different codes were used for the simulations (SPH-based `gadget` compared to grid-based Nyx). The difference between these simulations is well described in [17], showing that there is a growing deviation in the flux power spectrum towards smaller scales. Across the eBOSS scales adopted for this test ($k < 0.02 \text{ s/km}$), this deviation can be approximated as a constant offset plus a difference in slope for higher redshifts. Thus, we may find a small reconstruction bias for the amplitude parameter $A_{\text{Ly}\alpha}$ and/or the mean transmission, $\bar{F}(z)$.

We also check here whether adding a Planck prior on $\{\Omega_{\text{m}}, h\}$ — centered on the B13 central values — leads to better constraints or biases the results. This is to estimate if the cosmological constraints would be significantly changed by combining with CMB data.

Finally, we compare the contours obtained with our new Lyssa-based emulator with those from the old Taylor emulator for this mock data, post-processed into the new parameter space. Figure 7 shows that, despite different codes and models of the power spectrum, the constraints obtained with the Lyssa-based emulator are broadly consistent with the fiducial values assumed in the B13 simulation.

There are a few notable differences compared to before:

1. The constraint obtained from the mock data with the Lyssa-based emulator is now somewhat degraded compared to the results of the previous section. The largest change

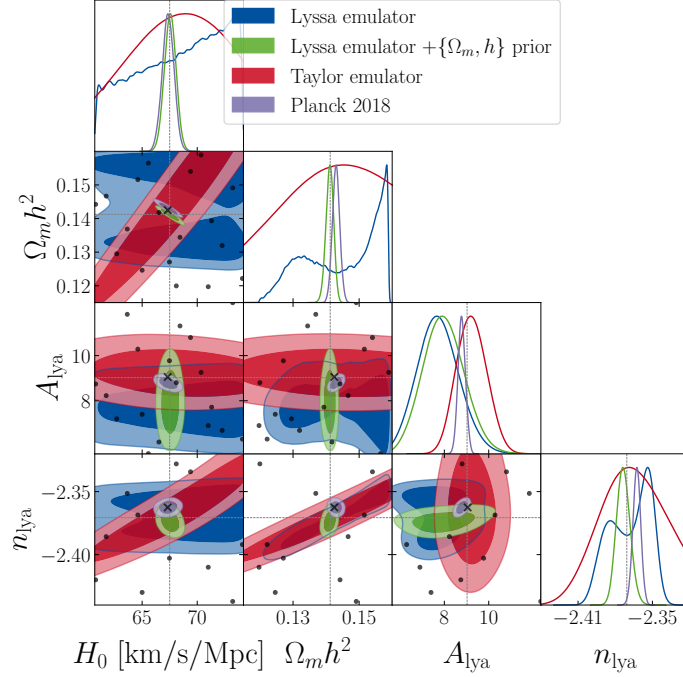


Figure 7. Same as in figure 5, but fitting instead the old B13 central simulation, whose parameters are represented by the grey lines. The fits with the Lyssa emulator are shown in blue and green. We also include results using the Taylor emulator of [3] as a reference in red.

is that the posterior on $\Omega_m h^2$ is flat and wide, with a very mild bimodality and an extremely mild preference for large values of $\Omega_m h^2$ towards the edge of the design. This also leads to a degraded sensitivity on $n_{\text{Ly}\alpha}$ due to their correlation. The constraints orthogonal to the $\Omega_m h^2 - n_{\text{Ly}\alpha}$ correlation direction or for $A_{\text{Ly}\alpha}$ do not really degrade.

2. The constraints on the $A_{\text{Ly}\alpha}$ and $n_{\text{Ly}\alpha}$ parameters are consistently recovered with a mild -1.4σ downward bias on $A_{\text{Ly}\alpha}$ ($n_{\text{Ly}\alpha}$ is recovered with only a $+0.15\sigma$ bias).
3. Since the prior on $\{\Omega_m, h\}$ breaks the $\Omega_m h^2 - n_{\text{Ly}\alpha}$ degeneracy, the $n_{\text{Ly}\alpha}$ constraints become significantly tighter once the prior is added. The slope remains almost unbiased with the prior (-0.3σ bias), while the bias of the amplitude is reduced (-0.9σ bias).
4. The degeneracy between the $n_{\text{Ly}\alpha} - h$ parameters as well as the $\Omega_m h^2 - h$ parameters is somewhat different in the Taylor emulator case, which we attribute to the finite accuracy of the employed method of using single cross-derivatives of the Taylor emulator and to the fact that the Taylor expansion was limited to second order, albeit a more detailed investigation is left for future work.

Overall, we notice that despite the very different codes the cosmology is recovered very consistently (apart from a slight $\sim 1\sigma$ bias in the amplitude). For this reason, we do not expect a strong difference between the constraints on cosmological parameters obtained with the Taylor and Lyssa-based emulator when applied to real data. We also stress that this $\sim 1\sigma$ bias on $A_{\text{Ly}\alpha}$ is reduced when adding a prior on $\{\Omega_m, h\}$ from Planck.

5 Fitting observational data

In this section we aim to fit the observational data for $P_{\text{F,1d}}$ of SDSS eBOSS DR14 presented in [3]. The goal is both to compare the inferred parameter values between the old Taylor emulator and the new Lyssa-based emulator, as well as to investigate a variety of settings and analysis choices with regard to the impact of such choices on a real data analysis. We first discuss our treatment of nuisance parameters, summarize a baseline analysis including tests of using different parametrization and corrections, and finally show how additional informative priors (on flux and/or cosmological parameters) change the given results.

5.1 Additional analytical corrections for observational nuisances

In order to fit eBOSS data, we have to include a number of nuisance parameters meant to marginalize over astrophysical and instrumental effects. These corrections are purposefully kept identical to those of [3] to ease the comparison, while ones more suited for DESI will be developed in future work. We also apply the correction discussed in section 4.1. The final power spectrum incorporates all corrections as

$$P_{\text{F,1d}} = f_{\text{metals}} \frac{(P_{\text{F,emu}} + f_{\text{noise}} P_{\text{noise}})}{f_{\text{DLA}} f_{\text{SN}} f_{\text{AGN}} f_{\text{IC}}} \quad (5.1)$$

with all quantities depending on both k and z . We use the following corrections:

1. The $P_{\text{F,1d}}$ measurement relies on the subtraction of a noise component P_{noise} inferred from the spectrograph noise. We add a rescaled (by a factor f_{noise}) version of this component, accounting for incomplete or overzealous subtraction. The wavenumber independent prefactor is left free for each redshift bin independently, and marginalized over.
2. The simulations do not model the presence of Damped Lyman- α Absorbers (DLAs), that are associated to the circumgalactic medium of (proto)-galaxies close to the line-of sight. While a fraction of those DLAs could be removed from the SDSS data used to estimate $P_{\text{F,1d}}$, there is a residual contamination. Here we use a correction factor (originally based on [48] and [49])

$$f_{\text{DLA}} = 1 - \left(\frac{1}{15\,000\,k - 8.9} + 0.018 \right) 0.2\,A_{\text{DLA}}, \quad (5.2)$$

where we marginalize over a single parameter A_{DLA} .

3. Although most of the signal in the Lyman- α forest originates from the pristine IGM, feedback processes from galaxies (such as supernovae, AGN winds, etc.) also contaminate the flux power spectrum. We parameterize this effect with the updated correction factor

$$f_{\text{SN}} = 1 + A_{\text{SN}} \left[d_0 + (d_1 - d_0) \frac{k - k_0}{k_1 - k_0} \right], \quad (5.3)$$

with $k_0 = 0.001\,\text{s/km}$, $k_1 = 0.02\,\text{s/km}$, $d_0 \in \{-0.06, -0.04, -0.02\}$ (for $z < 2.5$, $2.5 < z < 3.5$ and $z > 3.5$), $d_1 = -0.01$, and we marginalize over a single parameter A_{SN} [50].

4. We define an analytical correction for the presence of AGN (also discussed in [51], table 4) as

$$f_{\text{AGN}} = (1 + d(k, z) A_{\text{AGN}})^{-1}, \quad (5.4)$$

where $d(k, z)$ is a polynomial in redshift with a k -dependence of the form $a + b \exp(-c k)$. The parameters (a, b, c) were previously fitted to full Galaxy formation simulations including AGN and SN feedback prescriptions from the Horizon-AGN and Horizon-noAGN runs [52]. They have been determined independently at each redshift and are kept fixed in our analysis while a single parameter A_{AGN} is marginalized over.

5. We marginalize over a correction for silicon absorption lines,

$$f_{\text{metals}} = \prod_i \left(1 + a_i^2 + 2a_i \cos(k \Delta v_i) \right), \quad (5.5)$$

where $\Delta v_{\text{Si III}} = 2271 \text{ km/s}$, $\Delta v_{\text{Si II}} = 5577 \text{ km/s}$, the amplitude of each transition is given by $a_i = f_i / (1 - \bar{F})$,⁸ and a single parameter f_i is marginalized over (one for Si III, and one for Si II). Note that this correction is needed only for absorption lines with restframe wavelengths comparable to the Lyman- α transition as the respective metal forest then overlaps the probed Lyman- α forest leading to correlations in the absorption field. For metal transitions with restframe wavelengths $\lambda_{\text{rest}} > 1270 \text{ \AA}$ an estimation and subtraction of the metal power have been performed instead at the measurement level, specifically in the observed $P_{\text{F,1d}}$, see [6, 53, 54].

We expect many of these corrections to be insufficiently accurate for the analysis of upcoming survey data. However, they are well suited for performing an apple-to-apple comparison between the eBOSS fits performed with our new emulator and with the old Taylor emulator of [3] as well as the results of [4]. Future works will adopt more extensive corrections. In table 4 and appendix A, we summarize the set of nuisance parameters used in our analysis and the exact priors assumed on these parameters. These priors are nearly equivalent to those assumed in reference ([3], table 2).

The Taylor emulator of [3] includes additional corrections for splicing and background UV fluctuations. In the new emulator, the former are irrelevant since we no longer use splicing, and the latter we neglect — in the context of the present work — since the fit of [3] was compatible with no such fluctuations.

5.2 Application to eBOSS data — baseline analysis

We now use our Lyssa-based emulator to fit the current eBOSS data, which are described in [3, 6]. In this case, all the nuisance parameters described in section 5.1 and summarized in table 4 are marginalized over. For the Taylor emulator we also include a prior on

⁸In order to ease comparison with the previous approach of [3], in this particular formula, we fix the evolution of $\bar{F}(z)$ to the same *fiducial* evolution as in this reference, instead of matching it to the mean transmission function of the given model.

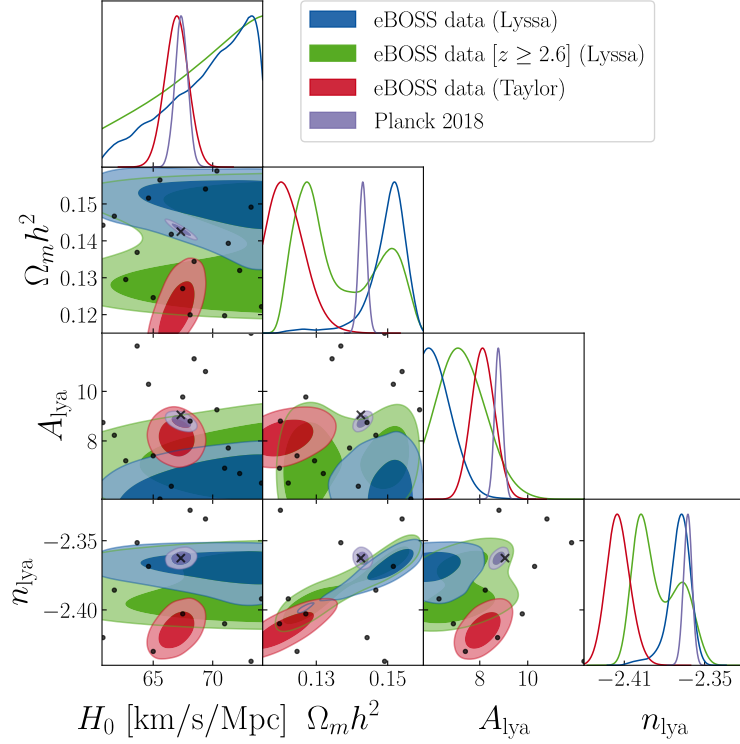


Figure 8. 68% and 95% confidence contours and marginalized posteriors on cosmological parameters, obtained by fitting eBOSS data with the Lyssa-based emulator. In blue we show the nominal run, in green a run with only data at $z \geq 2.6$, in red a run using the Taylor emulator of [3] (with an additional prior on H_0 as was used in the original chains), and in violet the constraints re-processed from Planck data assuming a flat Λ CDM cosmology. Note that, as we show in the right panel of figure 9, the results in $A_{\text{Ly}\alpha}$ and $n_{\text{Ly}\alpha}$ do not strongly change when a prior in H_0 is considered for the Lyssa-based emulator. The contours and posteriors involving nuisance parameters can be found in figures 15 and 16.

$H_0 = (67.3 \pm 1) \text{ km/s/Mpc}$ to make the results equivalent to those of [3] for the massless neutrino case.⁹

The results of the nominal analysis in which we use the emulator as described above are displayed in figure 8. These nominal results suggest that the data prefer a value of $A_{\text{Ly}\alpha}$ significantly lower (by -3.8σ) than the value extrapolated from Planck 2018 in the Λ CDM model. They also point at a consistent though somewhat higher value of $\Omega_m h^2$ (higher by 1.7σ) and at a value of $n_{\text{Ly}\alpha}$ fully consistent with Planck results [38] (lower by -0.7σ). Overall, this nominal analysis appears to point to a tension between the Planck

⁹We mention that for the Taylor emulator we have checked that we recover the constraints of [3] in the original parameter basis (and have indeed checked that the likelihoods agree at the $< 0.1\%$ level at a number of points). That $A_{\text{Ly}\alpha}$ in the constraints is lower than that of Planck 2018, while in [3] σ_8 was/is larger comes down to the fact that the scales at which the amplitudes are measured are very different. With very low n_s/Ω_m values (as in [3]), even a large σ_8 leads to a low $A_{\text{Ly}\alpha}$ at the corresponding scale. Similarly, while in [3] the discrepancy in n_s is only at the $1-2\sigma$ level, for $n_{\text{Ly}\alpha}$ it is clearly larger, again due to the different scale and the low Ω_m value found in [3]. These effects have also been pointed out in [4]. Explicitly, for $\sigma_8 = 0.82$, $n_s = 0.96$, $\Omega_m = 0.27$ (corresponding roughly to the mean of [3]), we find $A_{\text{Ly}\alpha} = 8.2$ and $n_{\text{Ly}\alpha} = -2.4$, in good agreement with figure 8.

CMB anisotropy data and the eBOSS Lyman- α data when a Λ CDM model is assumed. This is reminiscent of what has been found already in [3] and [4]. We notice that one major difference with respect to these previous analyses is that the tension appears in $A_{\text{Ly}\alpha}$, and not in $n_{\text{Ly}\alpha}$ (see figure 8, blue contour compared to red contour, compare also [4], figure 1). We also note that, somewhat similar to section 4.2, $\Omega_m h^2$ is poorly constrained by Ly α data alone (as expected from [12]). Indeed, in the constraints reported here and in section 5.3, we observe that the data do not exhibit a strong $\Omega_m h^2$ feature, leading to only very degraded or even bimodal constraints in $\Omega_m h^2$. For this reason we also check in section 5.3 if a prior on $\Omega_m h^2$ significantly impacts these conclusions.

To investigate the tension in $A_{\text{Ly}\alpha}$ further, we check whether it could be induced by the well motivated choices in the construction of our emulator, for example the assumption of a broken powerlaw redshift dependence of λ_P , see appendix B. In none of the test cases we find a relaxation of the tension in $A_{\text{Ly}\alpha}$. As such, it may seem that a tension similar to that reported in [3, 4] is confirmed. However, as we show below, there are several important caveats to such a conclusion.

As a first indication, we note that the removal of data below redshift $z < 2.6$ as advocated in [12] reduces the tension in $A_{\text{Ly}\alpha}$ significantly to 1.5σ (largely due to larger uncertainties and a slight shift in mean value by around $+1\sigma$). As such, the tension is disproportionately strongly driven by the lowest 2 out of 13 redshift bins, casting its physical nature into doubt. We do not conclude from our fits that there is a problem with the data at $z < 2.6$, as we discuss and show in section 5.3.

By investigating the correlations with nuisance parameters (see figure 15) we managed to identify that there is a strong correlation between the amplitude $A_{\text{Ly}\alpha}$ and the overall amplitude of the observed mean transmission, where in particular the A_τ parameter is highly degenerate with $A_{\text{Ly}\alpha}$ and appears to prefer values close to the lower boundary of the prior. This result puts the accuracy of emulation into question in this regime (as we have seen in section 3.1, the emulation accuracy degrades towards the edges of the design). As such, it is essential to check how the constraints change when imposing a well motivated prior on the mean transmission.

5.3 Variations with more informative priors

We have seen above that the critical degeneracies for cosmological analyses are the one between $A_{\text{Ly}\alpha} - A_\tau$ and between $n_{\text{Ly}\alpha} - \Omega_m h^2$. Luckily, for both of these degeneracies there are independent measurements that can be exploited. For A_τ there are independent measurements of the mean transmission, which we adopt in section 5.3.1. Instead, for $\Omega_m h^2$ we can add other cosmological information, such as from Planck CMB observations, as we do in section 5.3.2.

5.3.1 Mean transmission prior

To force the fit to remain in a regime that is well emulated, we impose a Gaussian prior on the observed effective optical depth $\tau_{\text{eff}}(z)$ (which is related to the mean transmission by $\bar{F}(z) = \exp[-\tau_{\text{eff}}(z)]$) according to the largely model independent results from [41] at each individual redshift. While this measurement was performed from a sample of $\sim 6\,000$ SDSS spectra which overlaps with that used in [6], it relies essentially on an anchor point from

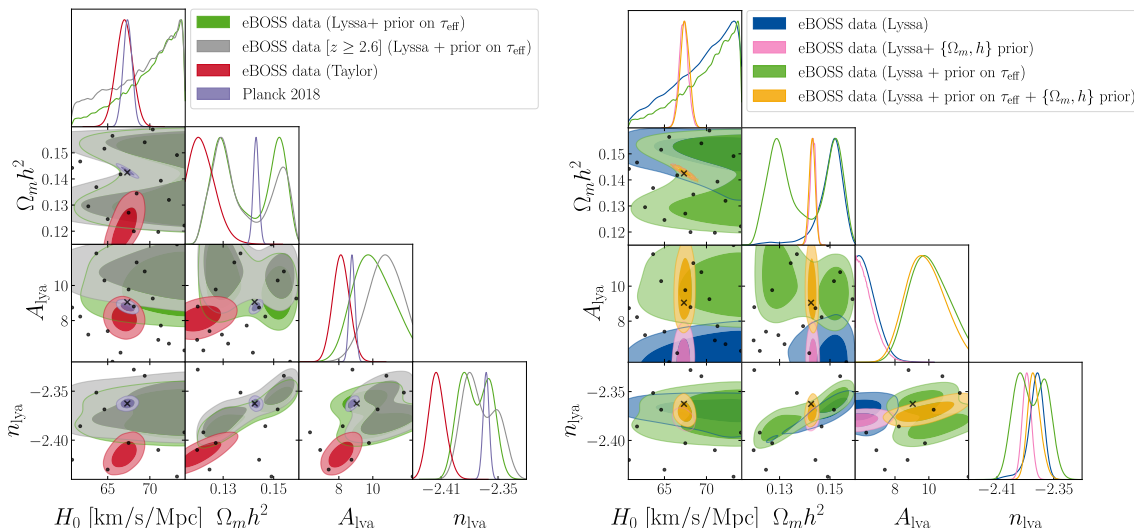


Figure 9. Same as figure 8, but including additional priors. Left: with an additional prior on the mean transmission from [41]. Right: with an additional prior on $\{\Omega_m, h\}$ from Planck, with or without the same additional mean transmission prior. Nuisance parameters can be found in figures 17 and 18. Note that as in figure 8 the red contours include a prior on H_0 .

high-resolution spectra [55], and on a method to estimate quasar continua fundamentally different from that used in [6]. Therefore, within the scope of this study we can consider that the measurement from [41] is largely independent of the eBOSS $P_{F,1d}$. Note that it is well known that the mean transmission strongly affects $P_{F,1d}$ (and is in particular degenerate with $A_{Ly\alpha}$). The idea to use external τ_{eff} measurements in order to ease the determination of the amplitude of matter power spectrum from $P_{F,1d}$ is not new at all, see for example [56], although several recent analysis like [3, 12] do not use such an external prior.

Figure 9 (left panel) displays our result when using this additional prior based on [41]. The tension in $A_{Ly\alpha}$ completely disappears in this case ($+0.3\sigma$) with all data included. However, subsets of the data still can show a tension (such as $+1.2\sigma$ when only including data at $z \geq 2.6$). We therefore conclude that, if the emulator is forced to remain within a regime of a mean transmission that agrees with independent measurements from [41] and thus lies away from the edge of its prior range, then the tension in $A_{Ly\alpha}$ completely disappears.

We also note that this same shift (in the $A_{Ly\alpha}$ - A_τ degeneracy direction) is what is observed in ([12], table 3) when cutting the data to $z \geq 2.6$. We argue that at least in the case of the Lyssa emulator applying well-motivated priors on the mean flux can restore concordance of the cosmological constraints between high- and low-redshift data. We note, however, that unlike [12] we do not see a big difference in the flux slope parameter or the primordial slope when fitting all redshift bins.

Given the differences between the analyses with and without a prior on the observed mean flux, we caution against strong interpretations of the baseline 3.8σ tension in $A_{Ly\alpha}/n_{Ly\alpha}$ (such as in [4, 5]), and instead argue to use the results when including a prior on the mean transmission from [41], until further investigations are performed.

A comprehensive comparison to [11] and an investigation into the degeneracy of the observed mean flux with the cosmological parameters and corresponding estimates of systematic uncertainties is left for future work. We also note that other mean transmission priors (such as based on the recent measurement by [57]) would lead to slightly different constraints on $A_{\text{Ly}\alpha}$ as well.

We also note that from figure 17 that we see a large anti-correlation between $A_{\text{Ly}\alpha}$ and T_0 , as well as a strong correlation between T_0 and its slope parameters η_T and $\Delta\eta_T$. In principle also priors on $T_0(z)$, $\gamma(z)$ or $\bar{F}(z)$ based on high-resolution Ly α forest observations such as [19, 40] could be applied. However, note that similarities in the modeling between the approaches taken in these measurements and here, e.g. concerning the approaches towards simulation post-processing, could lead to correlated methodological systematics. Also, parts of the BOSS dataset used in [19] alongside higher resolution data overlapped with the eBOSS data studied in here, and these analyses were performed using smaller 10–20 Mpc/ h simulation boxes and fixed initial conditions, i.e., fixed cosmological parameters and cosmic variance. Although such priors could be used for strengthening the constraints and reducing prior volume effects on our cosmological measurement, to be conservative, we leave such analysis for future work.

Similarly, one could also combine the medium resolution eBOSS data with higher resolution $P_{\text{F,1d}}$ data from X-SHOOTER, HIRES or UVES, such as in [9]. However, our post-processing and emulation technique has not yet been fully validated at $k > 0.05 \text{ s km}^{-1}$ (where the main information of those datasets beyond eBOSS and DESI lies). We therefore opted to not yet combine with those high-resolution products.

5.3.2 Cosmological prior

Since $\Omega_{\text{m}}h^2$ cannot be measured in the data at high significance, and H_0 is unconstrained, we also check the impact of imposing a prior on $\{\Omega_{\text{m}}, h\}$ from Planck 2018 data. This is shown in figure 9, right panel. This prior has a strong impact mostly on $n_{\text{Ly}\alpha}$ due to the correlation between $n_{\text{Ly}\alpha}$ and $\Omega_{\text{m}}h^2$ present in our constraints. Encouragingly, along the correlation between Ω_{m} and $n_{\text{Ly}\alpha}$ the cosmological prior forces the constraints into a regime consistent with Planck data. The $A_{\text{Ly}\alpha}$ constraint is otherwise unchanged, whether we also include a prior on the mean transmitted flux or not. As such, we conclude that an informative prior on $\{\Omega_{\text{m}}, h\}$ does not change any of our conclusions, and only slightly increases the constraining power in $n_{\text{Ly}\alpha}$.

When imposing both the mean transmission prior from [41] as well as the prior on $\Omega_{\text{m}}h^2$ from Planck, we find that we recover both an amplitude $A_{\text{Ly}\alpha}$ and a slope $n_{\text{Ly}\alpha}$ largely consistent with Planck data.

6 Conclusions

In this work, we have introduced the Lyssa suite of high resolution simulations performed with the Nyx hydrodynamical code. Each of these simulations followed 4096^3 dark matter particles and baryon mesh cells in a 120 Mpc comoving box ($\sim 81 \text{ Mpc}/h$). These simulations are used together with a Gaussian process emulation technique, and the resulting emulator of the 1D flux power spectrum is subjected to a variety of validation tests. We perform leave-one-out

Experiment	$A_{\text{Ly}\alpha}$	$n_{\text{Ly}\alpha}$	Δ_*	n_*
Planck (for reference)	8.79 ± 0.17	-2.3626 ± 0.0031	0.3476 ± 0.0068	-2.2899 ± 0.0032
Lyssa only	< 7.6	$-2.3693^{+0.0082}_{-0.0074}$	0.260 ± 0.024	$-2.2995^{+0.0068}_{-0.0064}$
Lyssa + $\tau_{\text{eff}}(z)$ + $\{\Omega_m, H_0\}$ priors	9.8 ± 1.1	-2.3705 ± 0.0068	0.388 ± 0.045	-2.2978 ± 0.0067

Table 3. Main results of this work: mean and 68% credible intervals (or 95% upper limit) for the parameters representing the amplitude and slope of the linear power spectrum at the scales probed by the Lyman- α forest. The constraints are illustrated in figure 10.

tests, where a simulation is left out from the training set and the emulator is tasked to reproduce its flux power spectrum using only the information from the other simulations. The emulation accuracy is around 0.5-2% in this test, depending on the redshift and wavenumber under consideration. We then check if fits using the full sampling pipeline consistently recovers a simulation that is not part of this training set (the “fiducial” simulation). The emulator nicely recovers the true cosmology without issues and this conclusion does not change when using an additional cosmological prior. We also fit the simulations described in [10] and show that despite large differences in the codes and the simulation setup, the new emulator can still consistently recover the $A_{\text{Ly}\alpha}$ and $n_{\text{Ly}\alpha}$ parameters (with no sensitivity to $\Omega_m h^2$) with only a mild bias on $A_{\text{Ly}\alpha}$ at the level of -1.4σ without a cosmological prior and -0.9σ with a cosmological prior.

We further apply the emulator and fitting pipeline to the 1D Lyman- α data from eBOSS DR14 derived in [6]. We find a tension in the amplitude $A_{\text{Ly}\alpha}$ compared to post-processed constraints from Planck 2018 (from [38]). This tension is reduced slightly when only data at $z \geq 2.6$ are considered (as advocated in [12]). More significantly, we find that such a tension is largely driven by a degeneracy between the amplitude $A_{\text{Ly}\alpha}$ and the mean transmission of the Lyman- α forest at $z = 3$ (A_τ), whose best-fit value lies towards the edge of the prior in the baseline analysis.

On this basis, we argue that the inclusion of a mean transmission prior from [41] is crucial to keep the Lyssa emulator in a well-behaved regime. With this well-motivated prior we find that our inferred value for $A_{\text{Ly}\alpha}$ becomes completely consistent with the results from Planck 2018, and our inferred parameters do not show any large changes when fitting only $z \geq 2.6$. We thus note that further investigations into systematic uncertainties based on the modeling of the mean flux transmission as well as into the differences between the emulation techniques will be necessary to improve the robustness of cosmological constraints from the 1D Lyman- α forest. We summarize our results in figure 10 and table 3.

Further comparisons with other simulation suites, such as those of [11] will establish the consistency of the conclusions. The simulations presented here may also be used in future $P_{\text{F,1d}}$ analyses in the context of DESI and for the analysis of the three-dimensional flux transmission power spectrum, though we expect that the employed emulation techniques will have to be adapted again for each case.

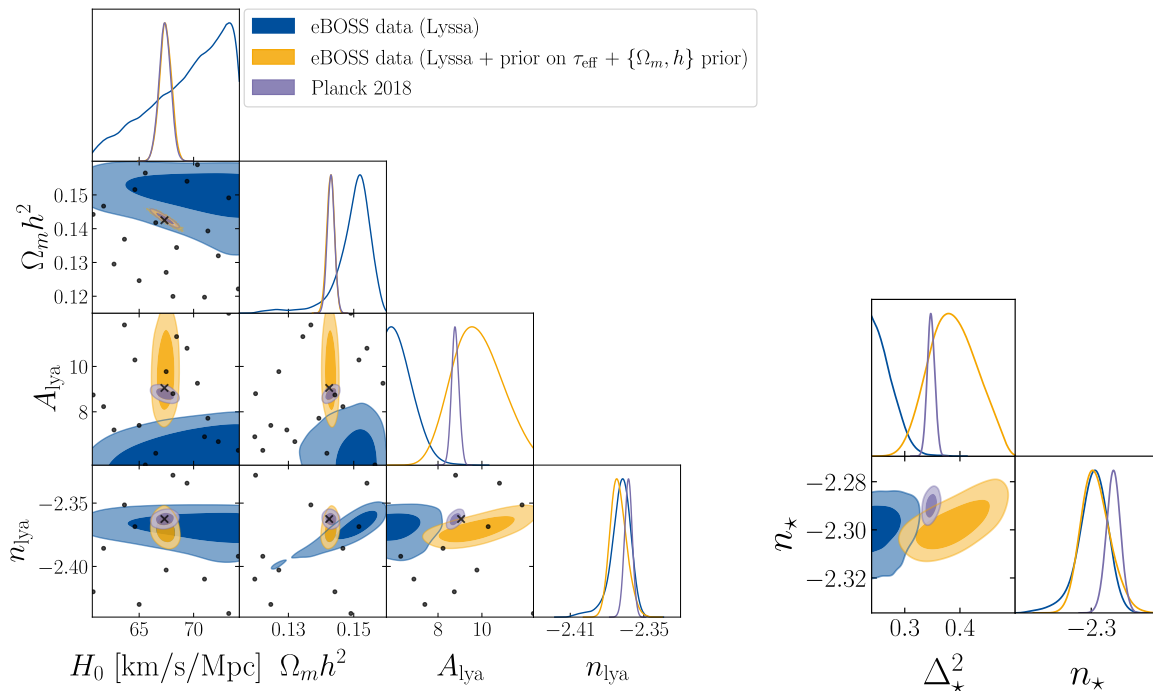


Figure 10. Same as the right panel of figure 9, but showing only the case without any additional informative priors and the case with a prior on $\{\Omega_m, h\}$ and $\tau_{\text{eff}}(z)$. Left: cosmological part of the parameter basis. Right: translation into the $\{\Delta_*, n_*\}$ parameter space (using an emulator technique). Notice the similarity to the $\{A_{\text{Ly}\alpha}, n_{\text{Ly}\alpha}\}$ parameter space. The numerical values are summarized in table 3.

Acknowledgments

The authors would like to thank Andreu Font-Ribera, Jonas Chaves-Montero, and Laura Cabayol-Garcia from the IFAE group for several useful discussions, as well as the DESI Ly α WG for their continued support and interesting ideas. We would like to thank Fajr Abdulrahman for his help with the conversion between $\{A_{\text{Ly}\alpha}, n_{\text{Ly}\alpha}\}$ and $\{\Delta_*^2, n_*\}$. The authors would also like to thank Felix Kahlhöfer for useful discussions. We finally thank Simeon Bird for comments on an earlier version of this paper.

We acknowledge PRACE for awarding us access to Joliot-Curie at TGCC, France via proposal 2019204900. We acknowledge the French national access to high-performance computing GENCI (Grand Equipement National de Calcul Intensif), for additional access to the HPC resources of TGCC under the allocations A0050410586, A0070410586 and AD01413571.

We also acknowledge the Oak Ridge Leadership Computing Facility, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This work used resources of National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. MW acknowledges support by the project AIM@LMU funded by the German Federal Ministry of Education and Research (BMBF) under the grant number 16DHBKI013. NS acknowledges

the support of the following Maria de Maetzu fellowship grant: Esta publicación es parte de la ayuda CEX2019-000918-M, financiada por MCIN/AEI/10.13039/501100011033. NS acknowledges support by MICINN grant number PID 2022-141125NB-I00. NS also was supported by a Fraunhofer-Schwarzschild Fellowship at Universitäts Sternwarte München. JL and MM acknowledge support from the DFG grant LE 3742/8-1. EA, CY and NPD acknowledge support from the ANR grant ANR-22-CE92-0037.

This work was partially supported by the DOE’s Office of Advanced Scientific Computing Research and Office of High Energy Physics through the Scientific Discovery through Advanced Computing (SciDAC) program. The development of Nyx as an AMReX application was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under contract number DE-AC02005CH11231, and by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.

A Description of nuisance parameters

In table 4 we show the priors on the parameters that we used during the analysis. Note that these priors were only used when the parameters were not fixed (as they were for the validation in sections 3.1, 3.2 and 4.2). Also that most of these parameters are equivalent to those described in [3], apart from a few exceptions. For example, instead of modeling the temperature slopes at $z < 3$ and $z > 3$ with different parameters, we instead put a linear prior on the difference $\Delta\eta^{T_0} = \eta^{T_0}(z > 3) - \eta^{T_0}(z < 3)$ between the slopes. Additionally, we use a

Parameter	Minimum	Maximum	Description
$T_0(z = 3)$	0	25000	Temperature at mean density at $z = 3$ in K [see eq. (2.6)]
η_T	-5	10	Logarithmic slope of T_0 at $z < 3$ [see eq. (2.6)]
$\Delta\eta_T$	-15	12	Difference in logarithmic slope between $z < 3$ and $z > 3$ [see eq. (2.6)]
$\gamma(z = 3)$	0.3	2.0	Logarithmic slope of temperature-density relation at $z = 3$
η_γ	-5	2	Logarithmic slope of γ at $z < 3$
A_τ	0	1.5	Effective optical depth at $z = 3$ [see eq. (2.6)]
η_τ	0	7.0	Logarithmic slope of effective optical depth τ [see eq. (2.6)]
A_{AGN}	0	3	AGN correction amplitude of eq. (5.4)
A_{SN}	0	3	SN correction amplitude of eq. (5.3)
$A_{n,i}$	0	3	Amplitude of noise power correction of bin i
f_{SiIII}	-0.2	0.2	Amplitude a of eq. (5.5) for SiIII scaled inversely with mean flux
f_{SiII}	-2.0	2.0	Amplitude a of eq. (5.5) for SiII scaled inversely with mean flux
λ_P	60	100	Pressure smoothing scale at $z = 3$
η_{λ_P}	-50	50	Logarithmic slope of λ_P at $z < 3$ in kpc
$\Delta\eta_{\lambda_P}$	-50	50	Difference in logarithmic slope of λ_P between $z < 3$ and $z > 3$
A_{UVfluct}	0	3	UV fluctuation amplitude

Table 4. Maximum allowed prior ranges on astrophysical/nuisance parameters. These are similar to those of [3], except for the changes explicitly mentioned in the text. There are 13 noise parameters $A_{n,i}$ with $i \in [1 \dots 13]$, one for each redshift bin. Additional Gaussian priors are described for the bold parameters in section 5.1. Note that there are additional effective priors from the finite extent of the Lyssa-based emulator, which are shown in figure 11.




Figure 11. Implicit priors set on thermal parameters for all constraints involving the Lyssa simulations with redshift evolution of parameters being described by (broken) power laws. Contours show the ranges reached when sampling only the priors (explicit from table 4 and implicit from the design of simulations) for all relevant 2d-subspaces. Priors are uniform inside the red areas for the parameters shown, but zero outside. Note that for fits with free \bar{F} (free λ_P) evolution, the implicit priors for A_τ , η_τ (λ_P , η_{λ_P} , $\Delta\eta_{\lambda_P}$) do not apply.

power law parameterization for the pressure smoothing scale λ_P instead of a simple overall rescaling of the heating rate. We put additional Gaussian priors on the noise level (0 ± 0.02 for each redshift bin), the SN correction amplitude (1 ± 0.3), the AGN correction amplitude (1 ± 0.3), and for the Taylor emulator also on the UV fluctuation amplitude (0 ± 0.3) and the splicing parameters as in [3] (offset 0.01 ± 0.05 and slope 0.9 ± 5.0).

Note that due to the finite extent of the Lyssa simulation thermal reprocessings, there are effective thermal priors on the powerlaw parameter basis, which are shown in figure 11.

B Variations of the baseline analysis

In this section we present variations of the baseline analysis (figure 8) in which we allow additional freedom to the fit (without imposing additional informative priors). See figure 12. The “free \bar{F} ” case allows a free mean transmission redshift dependence (with a nuisance parameter for each redshift bin and an implicit prior set by the range of our simulation grid at that redshift, see appendix A), whereas the “free λ_P ” allows a free λ_P redshift dependence (again limited by the extent of our model grid). Finally, the “no IC” case tests what happens if the IC correction of section 4.1 is disabled. In none of these cases we find a significant relaxation of the tension in A_{lya} .

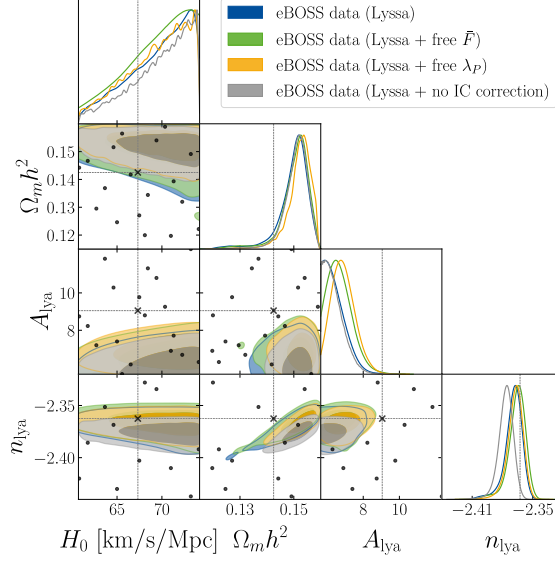


Figure 12. Results equivalent to figure 8 obtained when leaving the mean transmission (\bar{F}) redshift dependence free, when leaving the pressure smoothing scale (λ_P) redshift dependence free, or when turning off the IC correction.




Figure 13. Best-fitting models compared to the data for the two cases shown in figure 10. The data is the same as in [3], and the displayed theoretical predictions include the nuisance parameter corrections described in section 5.1.

C Best-fitting power spectra

We show the best-fitting power spectra for the main results of the eBOSS analysis from section 5.2 in figure 13. Both our baseline results and the results including a cosmological and effective mean transmission prior are presented. We note that the fit in both cases looks very good and are comparable to that of the bottom panel of figure 1 of [3].

D Additional plots of constraints on nuisance parameters

In this section we present the correlations of the cosmological parameter constraints with those of the astrophysical and nuisance parameters. These are shown in figure 14 for the fit to the fiducial simulation of section 3.2 and in figure 15 for the baseline fit to the eBOSS

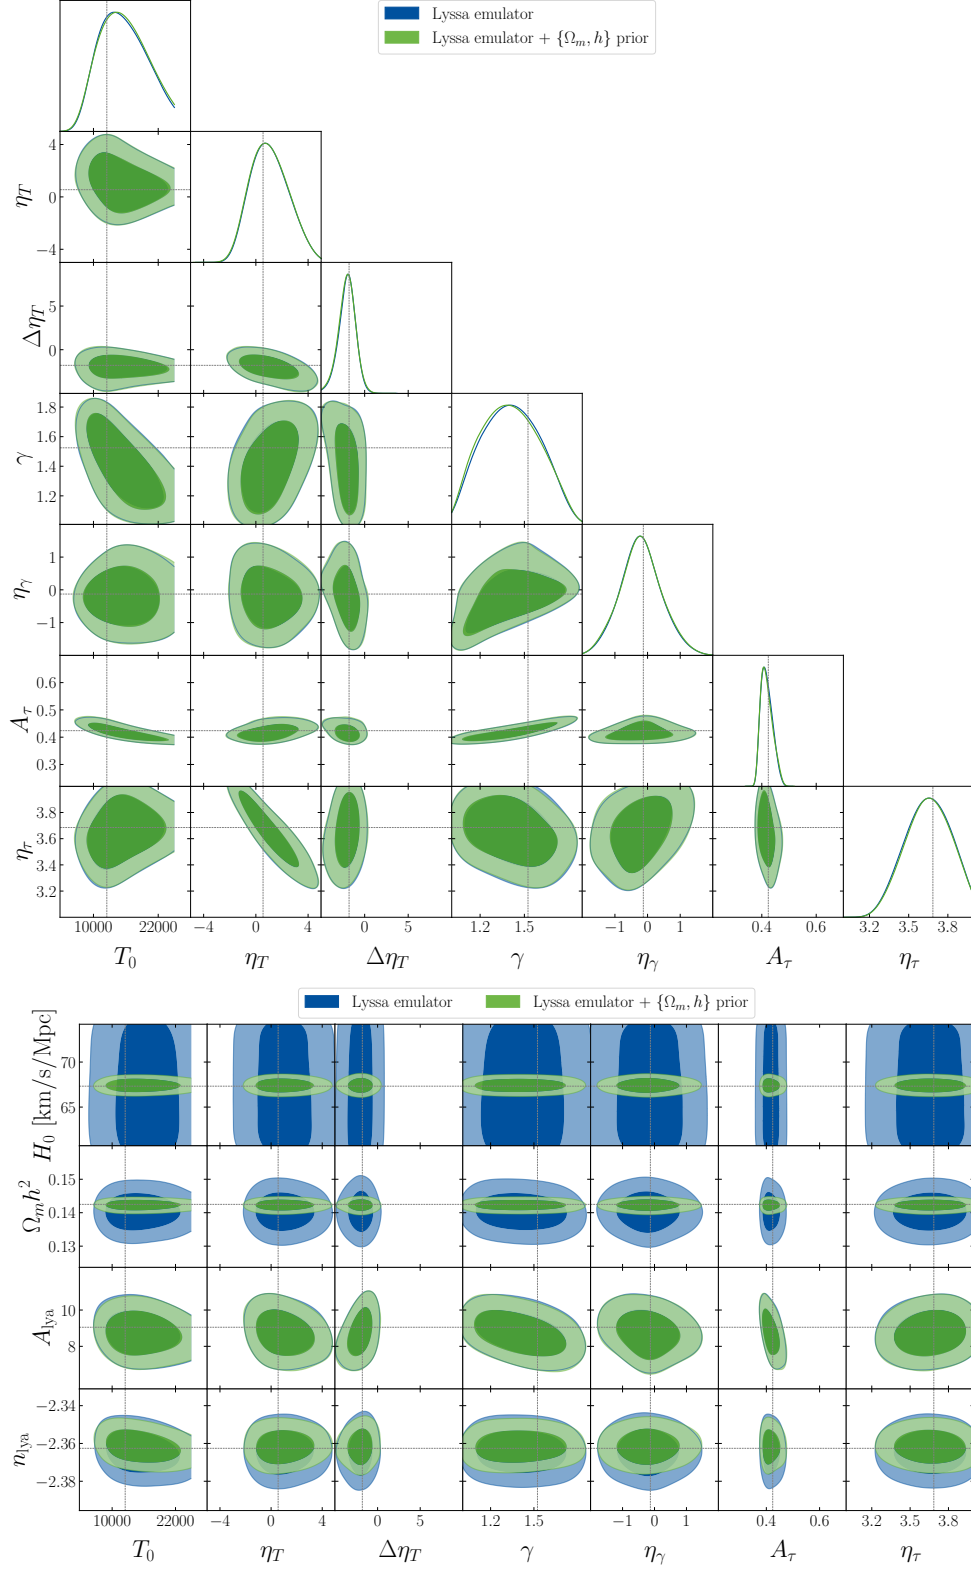


Figure 14. The 68% and 95% CL bounds for nuisance parameters when fitting the fiducial simulation with the Lyssa-based emulator, corresponding to figure 5.

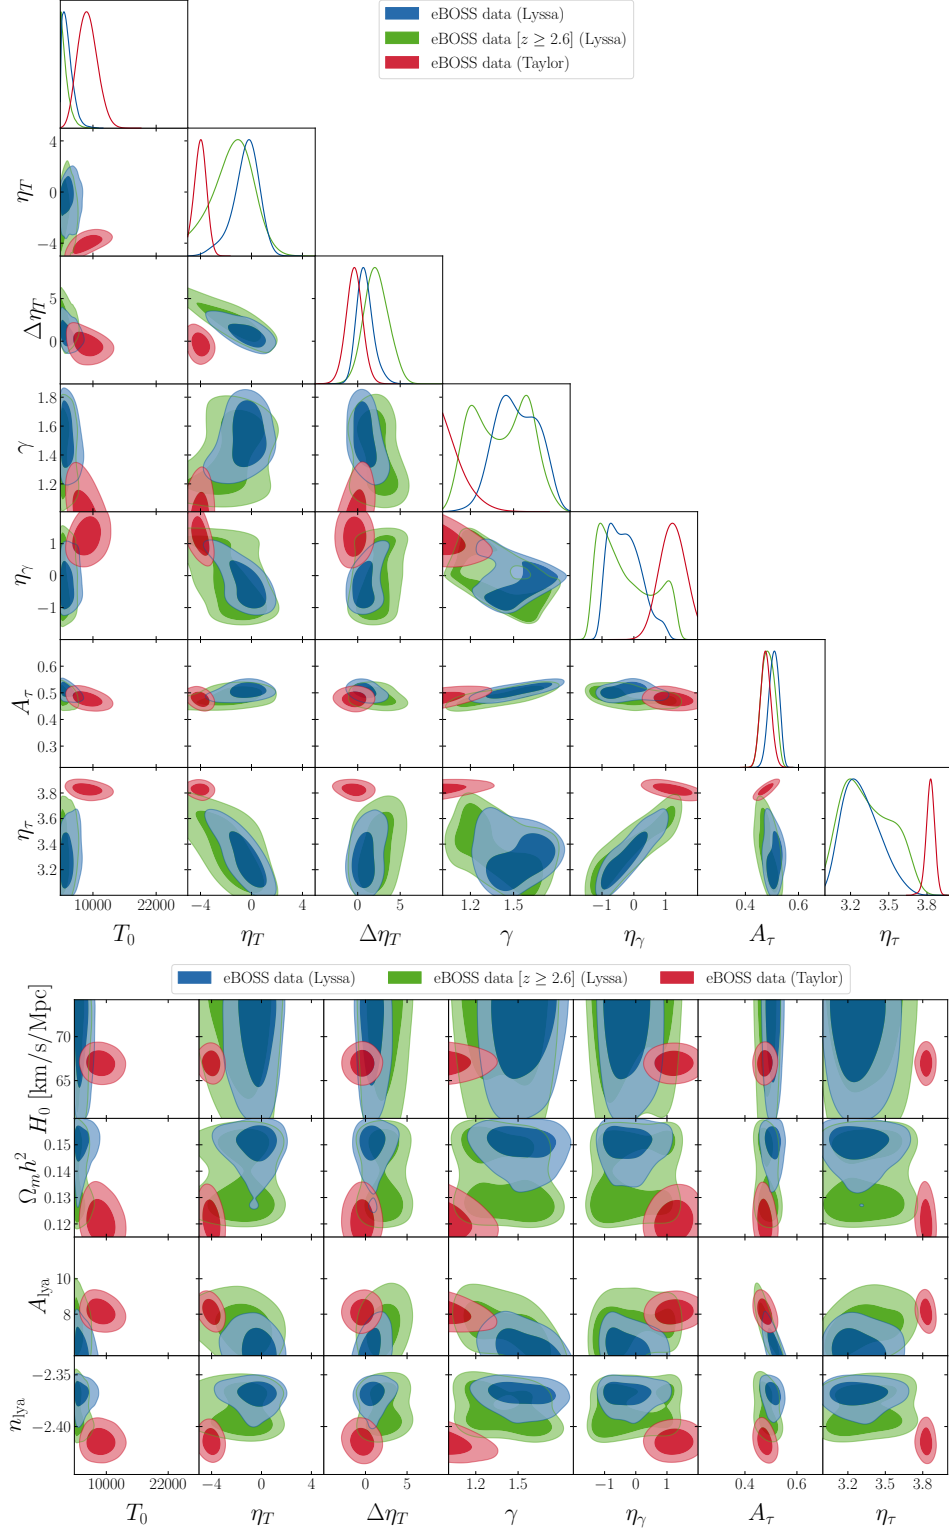


Figure 15. Same as figure 14, but instead corresponding to the run of figure 8 (left panel).

data. If lines are shown in the corresponding plots for fit to the fiducial simulation, they show the thermal parameters that would correspond to the best-fitting (broken) power law models corresponding to the true thermal history of the fiducial simulation as well as the

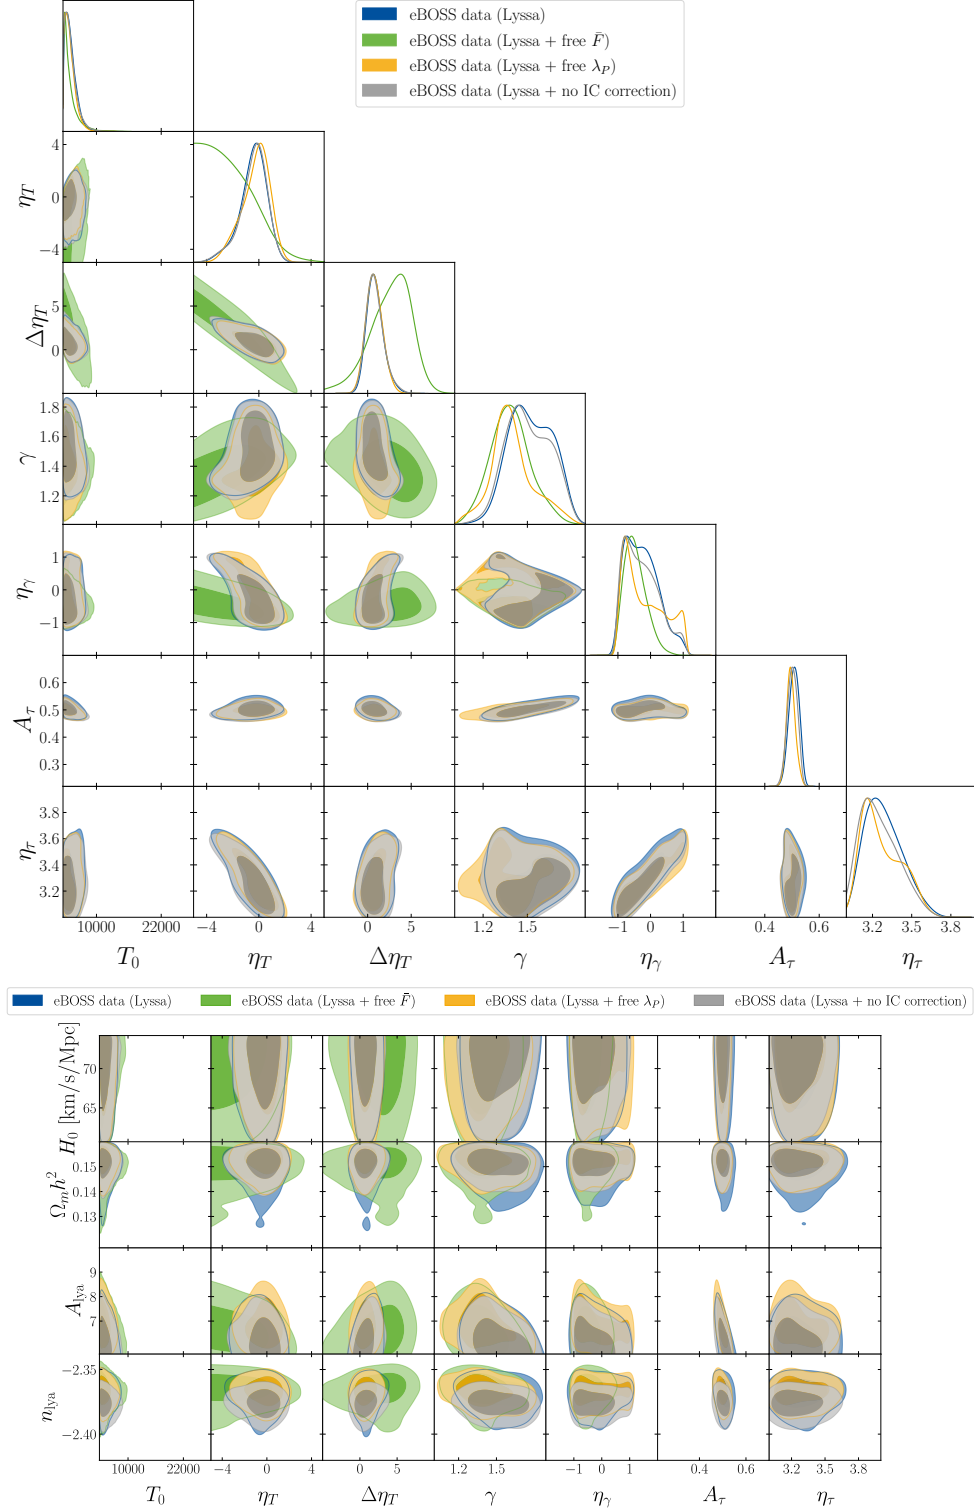


Figure 16. Same as figure 14, but instead corresponding to the run of figure 8 (right panel).

true cosmological parameters of that simulation. For example, we fit a broken powerlaw in to the actual $T_0(z)$ of the fiducial simulation and derive the corresponding $\{T_0, \eta_T, \Delta\eta_T\}$ parameter set.

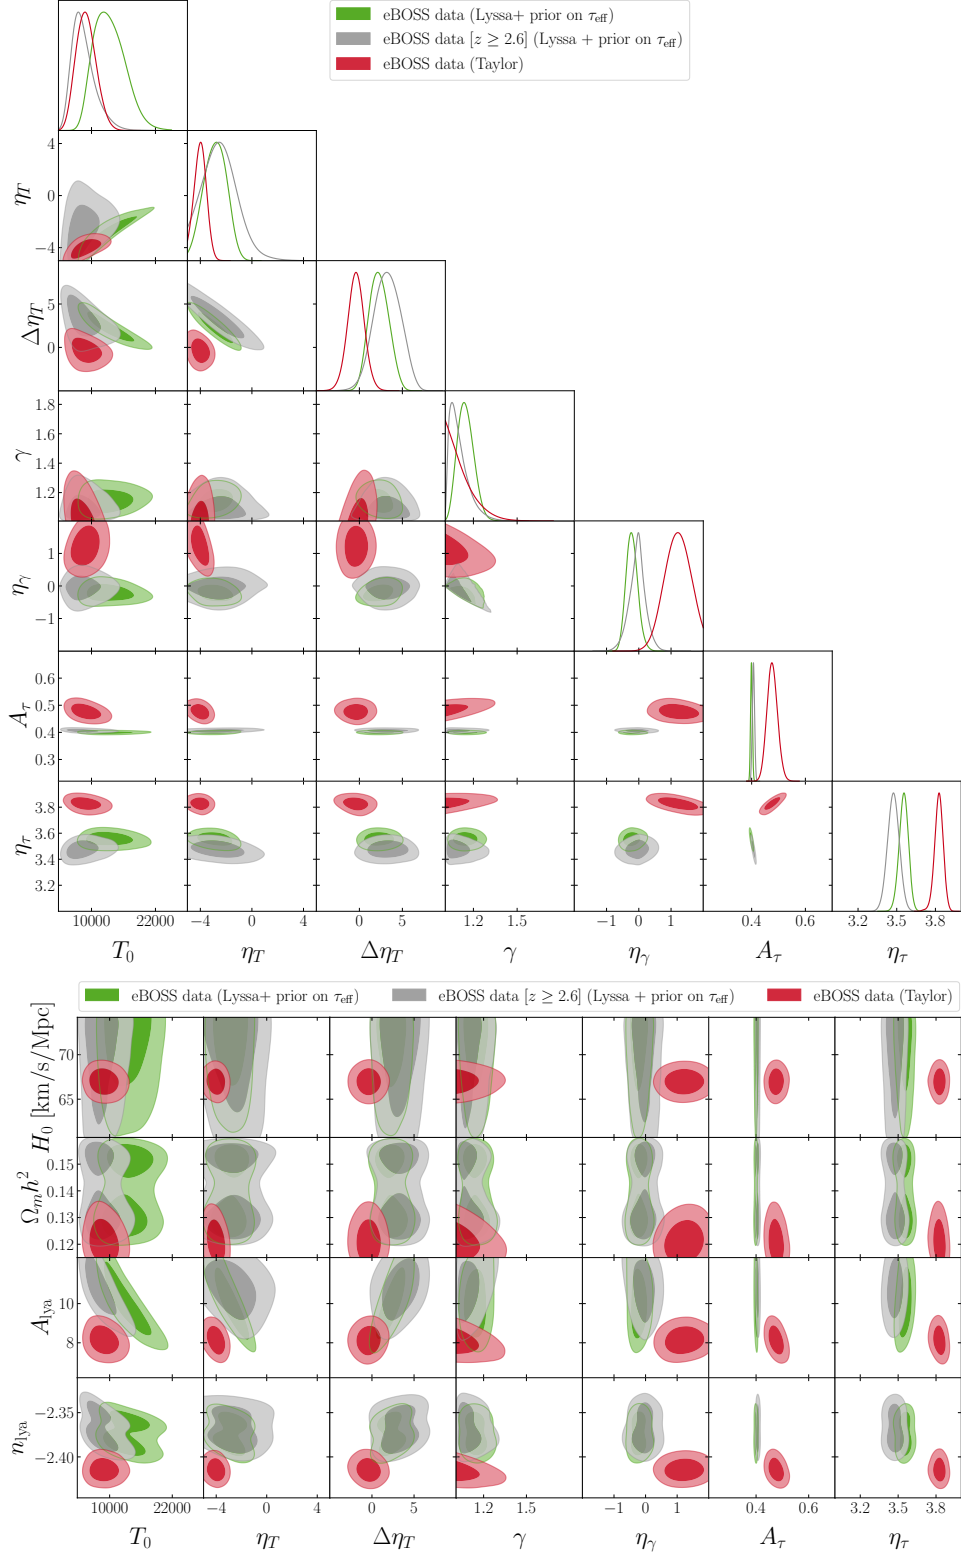


Figure 17. Same as figure 14, but instead corresponding to the run of figure 9 (left panel).

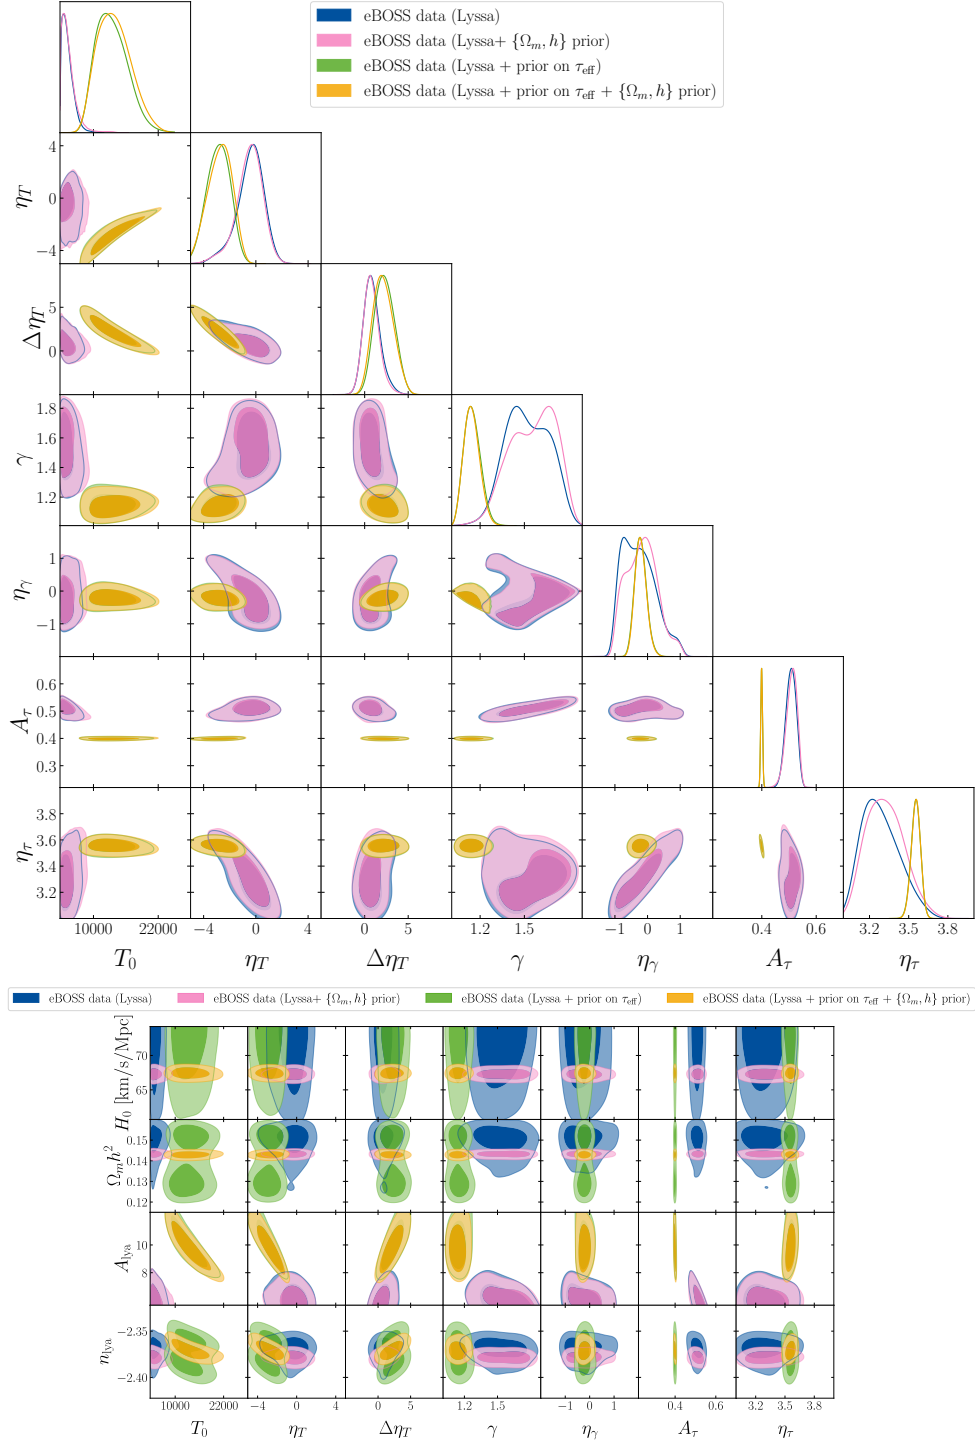


Figure 18. Same as figure 14, but instead corresponding to the run of figure 9 (right panel).

References

- [1] A.A. Meiksin, *The Physics of the Intergalactic Medium*, *Rev. Mod. Phys.* **81** (2009) 1405 [[arXiv:0711.3358](#)] [[INSPIRE](#)].
- [2] M. McQuinn, *The Evolution of the Intergalactic Medium*, *Ann. Rev. Astron. Astrophys.* **54** (2016) 313 [[arXiv:1512.00086](#)] [[INSPIRE](#)].
- [3] 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](#)].
- [4] K.K. Rogers and V. Poulin, *5 σ tension between Planck cosmic microwave background and eBOSS Lyman-alpha forest and constraints on physics beyond Λ CDM*, *Phys. Rev. Res.* **7** (2025) L012018 [[arXiv:2311.16377](#)] [[INSPIRE](#)].
- [5] S. Goldstein, J.C. Hill, V. Iršič and B.D. Sherwin, *Canonical Hubble-Tension-Resolving Early Dark Energy Cosmologies Are Inconsistent with the Lyman- α Forest*, *Phys. Rev. Lett.* **131** (2023) 201001 [[arXiv:2303.00746](#)] [[INSPIRE](#)].
- [6] EBOSS collaboration, *The one-dimensional power spectrum from the SDSS DR14 Ly α forests*, *JCAP* **07** (2019) 017 [[arXiv:1812.03554](#)] [[INSPIRE](#)].
- [7] DESI collaboration, *The Lyman- α forest catalog from the Dark Energy Spectroscopic Instrument Early Data Release*, *Mon. Not. Roy. Astron. Soc.* **528** (2024) 6666 [[arXiv:2306.06312](#)] [[INSPIRE](#)].
- [8] DESI collaboration, *Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **164** (2022) 207 [[arXiv:2205.10939](#)] [[INSPIRE](#)].
- [9] N.G. Karaçaylı et al., *Optimal 1D Ly α forest power spectrum estimation — II. KODIAQ, SQUAD, and XQ-100*, *Mon. Not. Roy. Astron. Soc.* **509** (2022) 2842 [[arXiv:2108.10870](#)] [[INSPIRE](#)].
- [10] 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](#)].
- [11] S. Bird et al., *PRIYA: a new suite of Lyman- α forest simulations for cosmology*, *JCAP* **10** (2023) 037 [[arXiv:2306.05471](#)] [[INSPIRE](#)].
- [12] M.A. Fernandez, S. Bird and M.-F. Ho, *Cosmological constraints from the eBOSS Lyman- α forest using the PRIYA simulations*, *JCAP* **07** (2024) 029 [[arXiv:2309.03943](#)] [[INSPIRE](#)].
- [13] D.C. Hooper and M. Lucca, *Hints of dark matter-neutrino interactions in Lyman- α data*, *Phys. Rev. D* **105** (2022) 103504 [[arXiv:2110.04024](#)] [[INSPIRE](#)].
- [14] D.C. Hooper et al., *One likelihood to bind them all: Lyman- α constraints on non-standard dark matter*, *JCAP* **10** (2022) 032 [[arXiv:2206.08188](#)] [[INSPIRE](#)].
- [15] S. Chabanier et al., *Modelling the Lyman- α forest with Eulerian and SPH hydrodynamical methods*, *Mon. Not. Roy. Astron. Soc.* **518** (2022) 3754 [[arXiv:2207.05023](#)] [[INSPIRE](#)].
- [16] C. Pedersen et al., *Massive neutrinos and degeneracies in Lyman-alpha forest simulations*, *JCAP* **04** (2020) 025 [[arXiv:1911.09596](#)] [[INSPIRE](#)].
- [17] M. Walther et al., *Simulating intergalactic gas for DESI-like small scale Lyman α forest observations*, *JCAP* **04** (2021) 059 [[arXiv:2012.04008](#)] [[INSPIRE](#)].
- [18] S. Bird et al., *An Emulator for the Lyman-alpha Forest*, *JCAP* **02** (2019) 050 [[arXiv:1812.04654](#)] [[INSPIRE](#)].

- [19] 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](#)].
- [20] C. Pedersen et al., *An emulator for the Lyman- α forest in beyond- Λ CDM cosmologies*, *JCAP* **05** (2021) 033 [[arXiv:2011.15127](#)] [[INSPIRE](#)].
- [21] C. Pedersen, A. Font-Ribera and N.Y. Gnedin, *Compressing the Cosmological Information in One-dimensional Correlations of the Ly α Forest*, *Astrophys. J.* **944** (2023) 223 [[arXiv:2209.09895](#)] [[INSPIRE](#)].
- [22] L. Cabayol-Garcia, J. Chaves-Montero, A. Font-Ribera and C. Pedersen, *A neural network emulator for the Lyman- α forest 1D flux power spectrum*, *Mon. Not. Roy. Astron. Soc.* **525** (2023) 3499 [[arXiv:2305.19064](#)] [[INSPIRE](#)].
- [23] J.S. Bolton et al., *The Sherwood simulation suite: overview and data comparisons with the Lyman α forest at redshifts $2 \leq z \leq 5$* , *Mon. Not. Roy. Astron. Soc.* **464** (2017) 897 [[arXiv:1605.03462](#)] [[INSPIRE](#)].
- [24] E. Puchwein et al., *The Sherwood-Relics simulations: overview and impact of patchy reionization and pressure smoothing on the intergalactic medium*, *Mon. Not. Roy. Astron. Soc.* **519** (2023) 6162 [[arXiv:2207.13098](#)] [[INSPIRE](#)].
- [25] A. Almgren et al., *Nyx: A Massively Parallel AMR Code for Computational Cosmology*, *Astrophys. J.* **765** (2013) 39 [[arXiv:1301.4498](#)] [[INSPIRE](#)].
- [26] Z. Lukić et al., *The Lyman- α forest in optically thin hydrodynamical simulations*, *Mon. Not. Roy. Astron. Soc.* **446** (2015) 3697 [[arXiv:1406.6361](#)] [[INSPIRE](#)].
- [27] J. Sexton et al., *Nyx: A Massively Parallel AMR Code for Computational Cosmology*, *J. Open Source Softw.* **6** (2021) 3068 [[INSPIRE](#)].
- [28] M. Crocce, S. Pueblas and R. Scoccimarro, *Transients from Initial Conditions in Cosmological Simulations*, *Mon. Not. Roy. Astron. Soc.* **373** (2006) 369 [[astro-ph/0606505](#)] [[INSPIRE](#)].
- [29] M. Crocce, S. Pueblas and R. Scoccimarro, *2LPTIC: 2nd-order Lagrangian Perturbation Theory Initial Conditions*, Astrophysics Source Code Library, record ascl:1201.005 (2012).
- [30] N. Katz, L. Hernquist and D.H. Weinberg, *Galaxies and gas in a cold dark matter universe*, *Astrophys. J. Lett.* **399** (1992) L109 [[INSPIRE](#)].
- [31] 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](#)].
- [32] G. Kulkarni et al., *Characterizing the Pressure Smoothing Scale of the Intergalactic Medium*, *Astrophys. J.* **812** (2015) 30 [[arXiv:1504.00366](#)] [[INSPIRE](#)].
- [33] B. Friesen et al., *In situ and in-transit analysis of cosmological simulations*, *Comput. Astrophys. Cosmol.* **3** (2016) 4 [[INSPIRE](#)].
- [34] 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](#)].
- [35] J. Oñorbe et al., *Inhomogeneous Reionization Models in Cosmological Hydrodynamical Simulations*, *Mon. Not. Roy. Astron. Soc.* **486** (2019) 4075 [[arXiv:1810.11683](#)] [[INSPIRE](#)].
- [36] P. Upton Sanderbeck and S. Bird, *Inhomogeneous He ii reionization in hydrodynamic simulations*, *Mon. Not. Roy. Astron. Soc.* **496** (2020) 4372 [[arXiv:2002.05733](#)] [[INSPIRE](#)].

- [37] SDSS collaboration, *The Linear theory power spectrum from the Lyman-alpha forest in the Sloan Digital Sky Survey*, *Astrophys. J.* **635** (2005) 761 [[astro-ph/0407377](#)] [[INSPIRE](#)].
- [38] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [39] A. Rorai et al., *Measurement of the small-scale structure of the intergalactic medium using close quasar pairs*, *Science* **356** (2017) 418 [[arXiv:1704.08366](#)] [[INSPIRE](#)].
- [40] P. Gaikwad, R. Srianand, M.G. Haehnelt and T.R. Choudhury, *A consistent and robust measurement of the thermal state of the IGM at $2 \leq z \leq 4$ from a large sample of Ly α forest spectra: evidence for late and rapid He II reionization*, *Mon. Not. Roy. Astron. Soc.* **506** (2021) 4389 [[arXiv:2009.00016](#)] [[INSPIRE](#)].
- [41] G.D. Becker, P.C. Hewett, G. Worseck and J.X. Prochaska, *A Refined Measurement of the Mean Transmitted Flux in the Ly-alpha Forest over $2 < z < 5$ Using Composite Quasar Spectra*, *Mon. Not. Roy. Astron. Soc.* **430** (2013) 2067 [[arXiv:1208.2584](#)] [[INSPIRE](#)].
- [42] S. Ambikasaran et al., *Fast Direct Methods for Gaussian Processes*, *IEEE Trans. Pattern Anal. Machine Intell.* **38** (2016) 252 [[arXiv:1403.6015](#)].
- [43] T. Brinckmann and J. Lesgourgues, *MontePython 3: boosted MCMC sampler and other features*, *Phys. Dark Univ.* **24** (2019) 100260 [[arXiv:1804.07261](#)] [[INSPIRE](#)].
- [44] F. Schmidt, *Effect of relative velocity and density perturbations between baryons and dark matter on the clustering of galaxies*, *Phys. Rev. D* **94** (2016) 063508 [[arXiv:1602.09059](#)] [[INSPIRE](#)].
- [45] O. Hahn, C. Rampf and C. Uhlemann, *Higher order initial conditions for mixed baryon-CDM simulations*, *Mon. Not. Roy. Astron. Soc.* **503** (2021) 426 [[arXiv:2008.09124](#)] [[INSPIRE](#)].
- [46] O. Hahn et al., *MUSIC2-monofonIC: 3LPT initial condition generator*, Astrophysics Source Code Library, record ascl:2008.024 (2020).
- [47] M.A. Fernandez, S. Bird and P. Upton Sanderbeck, *Effect of separate initial conditions on the lyman- α forest in simulations*, *Mon. Not. Roy. Astron. Soc.* **503** (2021) 1668 [[arXiv:2009.09119](#)] [[INSPIRE](#)].
- [48] P. McDonald et al., *Physical effects on the Ly-alpha forest flux power spectrum: Damping wings, ionizing radiation fluctuations, and galactic winds*, *Mon. Not. Roy. Astron. Soc.* **360** (2005) 1471 [[astro-ph/0407378](#)] [[INSPIRE](#)].
- [49] M. Viel, M.G. Haehnelt, R.F. Carswell and T.S. Kim, *The effect of (strong) discrete absorption systems on the lyman alpha forest flux power spectrum*, *Mon. Not. Roy. Astron. Soc.* **349** (2004) L33 [[astro-ph/0308078](#)] [[INSPIRE](#)].
- [50] M. Viel, J. Schaye and C.M. Booth, *The impact of feedback from galaxy formation on the Lyman-alpha transmitted flux*, *Mon. Not. Roy. Astron. Soc.* **429** (2013) 1734 [[arXiv:1207.6567](#)] [[INSPIRE](#)].
- [51] S. Chabanier et al., *The impact of AGN feedback on the 1D power spectra from the Ly α forest using the Horizon-AGN suite of simulations*, *Mon. Not. Roy. Astron. Soc.* **495** (2020) 1825 [[arXiv:2002.02822](#)] [[INSPIRE](#)].
- [52] Y. Dubois et al., *Dancing in the dark: galactic properties trace spin swings along the cosmic web*, *Mon. Not. Roy. Astron. Soc.* **444** (2014) 1453 [[arXiv:1402.1165](#)] [[INSPIRE](#)].
- [53] SDSS collaboration, *The Lyman-alpha forest power spectrum from the Sloan Digital Sky Survey*, *Astrophys. J. Suppl.* **163** (2006) 80 [[astro-ph/0405013](#)] [[INSPIRE](#)].

- [54] BOSS collaboration, *The one-dimensional Ly-alpha forest power spectrum from BOSS*, *Astron. Astrophys.* **559** (2013) A85 [[arXiv:1306.5896](#)] [[INSPIRE](#)].
- [55] C.-A. Faucher-Giguere et al., *A Direct Precision Measurement of the Intergalactic Lyman- α Opacity at $2 < z < 4.2$* , *Astrophys. J.* **681** (2008) 831 [[arXiv:0709.2382](#)] [[INSPIRE](#)].
- [56] U. Seljak, P. McDonald and A. Makarov, *Cosmological constraints from the CMB and Ly- α forest revisited*, *Mon. Not. Roy. Astron. Soc.* **342** (2003) L79 [[astro-ph/0302571](#)] [[INSPIRE](#)].
- [57] W. Turner et al., *New measurements of the Lyman- α forest continuum and effective optical depth with LyCAN and DESI Y1 data*, *Astrophys. J.* **976** (2024) 143 [[arXiv:2405.06743](#)] [[INSPIRE](#)].