

Constraining mixed dark matter models with high-redshift Lyman-alpha forest data

Olga Garcia-Gallego,^{1,2,*} Vid Iršič^{1,3,4,5,6,7,†}, Martin G. Haehnelt^{1,2}, Matteo Viel^{4,5,6,7,8} and James S. Bolton^{1,9}

¹*Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, United Kingdom*

²*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, United Kingdom*

³*Center for Astrophysics Research, Department of Physics, Astronomy and Mathematics, University of Hertfordshire, College Lane, Hatfield AL10 9AB, United Kingdom*

⁴*SISSA—International School for Advanced Studies, Via Bonomea 265, 34136 Trieste, Italy*

⁵*INFN—National Institute for Nuclear Physics, Via Valerio 2, I-34127 Trieste, Italy*

⁶*IFPU, Institute for Fundamental Physics of the Universe, Via Beirut 2, 34151 Trieste, Italy*

⁷*INAF, Osservatorio Astronomico di Trieste, Via G. B. Tiepolo 11, I-34131 Trieste, Italy*

⁸*ICSC—Centro Nazionale di Ricerca in High Performance Computing, Big Data e Quantum Computing, Via Magnanelli 2, Bologna, Italy*

⁹*School of Physics and Astronomy, University of Nottingham, University Park, Nottingham, NG7 2RD, United Kingdom*



(Received 15 April 2025; accepted 23 June 2025; published 1 August 2025)

This study sets new constraints on cold + warm dark matter (CWDM) models by leveraging the small-scale suppression of structure formation imprinted in the Lyman- α forest. Using the Sherwood-Relics suite, we extract high-fidelity flux power spectra from simulated Lyman- α forest data, spanning a broad range of cosmologies and thermal histories. This enables precise constraints on the warm dark matter (WDM) fraction, f_{WDM} , and the mass of the WDM particle, m_{WDM} . A key advancement of our analysis is the integration of a neural network emulator directly at the likelihood level, significantly accelerating Bayesian parameter inference. With new observations of high-redshift ($z = 4.2\text{--}5.0$) quasar spectra from UVES and HIRES, we establish stringent upper limits: for $m_{\text{WDM}} = 1 \text{ keV}$, we find $f_{\text{WDM}} < 0.16$ (2σ), with constraints loosening to 35%, 50%, and 67% for $m_{\text{WDM}} = 2, 3$, and 4 keV , respectively. Our results for pure WDM reaffirm the lower bounds of previous work. Crucially, we account for the fixed resolution of simulations and the impact of patchy reionization, demonstrating their minimal influence on mixed dark matter constraints. This robustness paves the way for tighter bounds with improved statistical samples in the future. Our findings suggest that CWDM models can naturally accommodate mild suppression of matter clustering in the high-redshift Lyman- α forest 1D flux power, potentially offering a resolution to some of the ongoing cosmological tensions at low redshifts, namely the S_8 tension.

DOI: 10.1103/4k29-h991

I. INTRODUCTION

In the standard cosmological model, Λ CDM (cold dark matter), cold nonbaryonic dark matter with negligible primordial velocities constitutes approximately $\Omega_{\text{CDM}} \approx 0.26$ of the total energy density of the Universe. This model provides a pivotal framework for galaxy formation

theory, where the hierarchical merging of nonlinear dark matter structures, or haloes, plays a key role in shaping the large-scale structure of the Universe. In this context, the Λ CDM paradigm accurately predicts the evolution of the matter power spectrum across cosmic time from cosmic microwave background (CMB) observations [1], demonstrating remarkable success in matching observations down to scales comparable to the average spacing between galaxies (\approx a few Mpc).

At smaller scales, however, discrepancies arise between theoretical predictions from N-body simulations and observational data. Some of these inconsistencies stem from the hierarchical nature of the theory. While the apparent mismatch in the number of dwarf galaxies in the Milky Way (missing satellites problem [2,3]) has been

*Contact author: og313@cam.ac.uk

†Contact author: v.irsic@herts.ac.uk

largely resolved through new observations of faint satellites [4] and improved baryonic modeling in high-resolution simulations (e.g., [5,6]), other tensions persist. The most significant among these is related to the amount of dark matter in the inner parts of galaxies (cusp-core problem [7]), particularly the challenge of explaining the diverse rotation curves observed [8]. Moreover, while the dynamics of the most massive Milky Way satellites (too-big-to-fail problem [9,10]) are now better understood through updated cosmological simulations (e.g., [11,12]), the situation remains unclear for satellites outside the vicinity of the Milky Way [13].

Another key challenge of CDM on small scales involves the amplitude of matter density fluctuations, quantified by σ_8 . The parameter combination $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$ has given rise to the so-called S_8 tension, where cosmic shear, galaxy clustering surveys, and their cross-correlation with CMB lensing (e.g., [14–20]) suggest lower values than those inferred from the CMB [1].

A major open question in cosmology is whether both of these problems could be solved by baryonic physics (e.g., [21–27]) or dark matter models beyond the Standard Model (e.g., [28–34]).

Proper modeling of the former, however, is computationally challenging, and while the physics of galaxy formation is understood in broad terms, many of the underlying baryonic processes, such as feedback mechanisms, remain poorly constrained [35–39]. As such, exploring alternative dark matter models is further motivated by two primary considerations: (1) the absence of a viable dark matter candidate within the Standard Model of particle physics, and (2) no direct detection has been achieved thus far, after a decade of experiments aimed at detecting either of the two most popular CDM candidates: WIMPs (weakly interacting massive particles) (LHC) and axions (ADMX [40]).

From a particle physics perspective, the simplest extension of the Standard Model is to accommodate a *heavy* or *sterile* neutrino that can explain the mechanism facilitating neutrino oscillations [41]. Sterile neutrinos can be produced via active neutrinos oscillations in the early Universe, as proposed by [42]. They would be the prototype for warm dark matter (WDM), a hypothetical thermal relic dark matter particle that decouples earlier and is more massive than Standard Model neutrinos, initially proposed in [43–45]. Such a dark matter candidate has a non-negligible primordial velocity distribution at the time of decoupling, and can consequently escape small gravitational perturbations, suppressing structure formation on scales smaller than their free-streaming length, λ_{fs} , which is inversely proportional to the WDM particle mass. For WDM particle masses of the order of a few keV, the free-streaming length roughly corresponds to the size of a dwarf galaxy. WDM follows the same bottom-up structure formation as CDM, making it a promising candidate to address small-scale challenges while preserving the

successes of CDM at larger scales. The most recent constraint from cosmological analysis on this model excludes thermal relics lighter than 5.7 keV (2σ) [46].

For this reason, the interest in a hybrid model that interpolates between CDM-like and WDM cosmology has reemerged. The cold + warm dark matter (CWDM) model was first introduced by [47–50]. The model is parametrized by the fraction of the warm (*lighter*) counterpart, $f_{\text{WDM}} = \Omega_{\text{WDM}}/\Omega_M$, and its particle mass m_{WDM} , which has the same properties as in a pure WDM cosmology.

In some variants of the production mechanism of sterile neutrinos, the resulting velocity dispersion of the hypothetical particle consists of cold and warm components [51,52]. Within this framework, a structured dark matter sector could exist, featuring components in the keV and GeV mass regimes, making sterile neutrinos viable candidates for CWDM [53]. Such a dual-component nature of dark matter has also been proposed for axions [33], reinforcing the motivation to investigate CWDM.

The free-streaming property of CWDM can be constrained observationally from the Lyman- α forest, a powerful cosmological probe of matter density fluctuations in the weakly nonlinear regime through the high-redshift and underdense intergalactic medium (IGM). In the past, the Lyman- α forest has been used, combined with CMB data, to constrain the matter power spectrum (e.g., [54–56]), nature of dark matter [57–62], and to test physics beyond the Standard Model using the 3D correlation function to measure the scale of baryon acoustic oscillations (e.g., [63–68]).

At the scales of interest for the dark matter models, the Lyman- α forest is influenced by two key factors: the thermal evolution of the IGM (e.g., [69]) and the underlying dark matter density field, which drives structure formation through gravitational interactions (e.g., [46,70]). Consequently, accurately interpreting observations of the Lyman- α forest at these scales requires hydrodynamic simulations to model the complex gas dynamics. These simulations have facilitated detailed comparisons between the structures predicted by the CDM paradigm and those observed in the Lyman- α forest down to the Jeans scale [71], while also incorporating the free-streaming properties of dark matter. In the past, constraints on the mass m_{WDM} have been obtained by comparing simulations with quasars (QSOs) datasets of different resolution and signal-to-noise ratios (S/N). This comparison is usually done within a Bayesian framework by interpolating in the parameter space where simulations are available [58,60,72–74], which becomes unfeasible as the dimensionality of the inference problem increases. Instead, surrogate models or emulators can be used based on Gaussian process interpolation [75–80], or a supervised neural network interpolator [81,82], demonstrating significant inference speedup. The data used to carry out this comparison have included, in increasing order of resolution, the Sloan Digital Sky Survey (SDSS) and the Baryon Oscillation Spectroscopic

Survey (BOSS) (e.g., [60,83]), X-Shooter (e.g., [58]) and MIKE spectrograph at the Magellan Telescopes (e.g., [72]), and HIRES spectrograph at the Keck I Telescope (Keck) and UVES spectrograph of the Very Large Telescope (VLT) (e.g., [46]). The higher-resolution data provide stronger bounds on m_{WDM} by probing the small scales where the suppression occurs.

Beyond pure WDM, Ref. [84] first suggested a preference for $f_{\text{WDM}} < 0.7$ mixed models by reinterpreting Lyman- α forest and x-ray data in the presence of sterile neutrinos. Later, Ref. [52] combined the WMAP5 and Lyman- α (SDSS) data to constrain $f_{\text{WDM}} < 0.35$ for the minimum mass of their prior range $m_{\text{WDM}} \sim 1$ keV in the case of thermally decoupled candidates. Reference [83] found that a mass as light as $m_{\text{WDM}} > 0.7$ keV is allowed if $f_{\text{WDM}} < 0.15$. Reference [85] combined Planck, baryon acoustic oscillations, and Milky Way satellite data to constrain $f_{\text{WDM}} < 0.29$ for masses in the range 1–10 keV. Reference [86] used N-body simulations and modeled the effect of baryons on top to predict allowed mixed models to which weak lensing and galaxy clustering power spectra upcoming data will be sensitive to.

In this work, we update previous constraints on CWDM models using high- z Lyman- α forest measurements that are sensitive to matter clustering to two times as small separations as any previous attempts, $k \sim 20h$ Mpc $^{-1}$ [87]; and a set of hydrodynamical simulations that include the effect of a heavy and a light thermal dark matter particle in structure formation, including full gas physics. We compare the synthetic spectra to the data in a Bayesian inference framework. Since the grid of simulations spans a high-dimensional parameter space, we efficiently sample the likelihood by implementing a neural network emulator. We perform a variety of analyses that incorporate numerical correction due to a fixed resolution of simulations (resolution correction), inhomogeneous reionization effects, and a possible misestimation of instrumental noise of the data.

The structure of this paper is as follows: In Sec. II, we present the 1D Lyman- α flux power spectrum measurements. Section III introduces the hydrodynamical simulations, highlighting their key properties and the impact of CWDM on the matter power spectrum. We also describe the synthetic flux power spectrum models derived from these simulations. Section IV details the emulator used for the Monte Carlo Markov chain (MCMC) analysis. In Sec. V, we discuss various analysis choices, comparing their effect on constraints on mixed dark matter models and exploring the degeneracy between the WDM parameters f_{WDM} and m_{WDM} . Finally, Sec. VI summarizes our findings and examines the implications of the allowed CWDM models for the S_8 tension.

II. DATA

We utilize the same dataset as presented in [46], specifically the one-dimensional flux power spectra

introduced in [69]. These spectra were derived from a sample of 15 high-resolution quasar observations conducted with the UVES and HIRES spectrographs per averaged redshift bin, centered at $z = 4.2, 4.6$, and 5.0 . In particular, we analyze the power spectra across 16 k -bins, uniformly spaced by $\Delta_{\log_{10}(k)} = 0.1$, from $\log_{10}(k/\text{km}^{-1} \text{s}) = -2.2$ to $\log_{10}(k/\text{km}^{-1} \text{s}) = -0.7$. The measurements at these k -bins incorporate corrections for finite spectral resolution and pixel size, as detailed in Appendix K of [69]. These corrections were derived by applying the same window function as in [69,88,89]. Furthermore, the nondiagonal elements of the covariance matrix are adopted from [46], which employed the same box size, L_{box} , as in this work to construct the bootstrap cross-correlation coefficient.

The velocity width corresponding to the smallest scale probed in this dataset is approximately 30 km s^{-1} , which is significantly larger than the spectral resolutions of HIRES and UVES, measured as $6 \approx \text{km s}^{-1}$ and $7 \approx \text{km s}^{-1}$, respectively. These characteristics make it an ideal dataset to resolve the small-scale features of the Lyman- α forest. Prior studies have used lower-resolution datasets with a larger number of QSO samples (e.g., [60,73]) or data from the same instruments but with approximately half the number of quasar spectra (e.g., [89,90]). These studies have been limited to scales of up to $k \approx 0.1 \text{ s km}^{-1}$. While such datasets have been effective in constraining both the instantaneous thermal state and the evolution of the IGM, Ref. [69] extends this reach to $k_{\text{max}} = 0.2 \text{ s km}^{-1}$. Crucially, the high k -bins introduced in this work are the ones sensitive to dark matter free-streaming, allowing us to probe power suppression on scales of 10–100 kpc. This extension to smaller scales is made possible by improved modeling of instrumental noise and metal contaminants [69,91,92]. To account for these uncertainties, we marginalize over these effects as part of our analysis.

III. SIMULATIONS

To probe CDM and pure warm dark matter WDM cosmologies, we use the same set of simulations from the Sherwood project as employed in [46], incorporating the fiducial ultraviolet (UV) background from [93]. These simulations feature several advancements over previous models used in earlier studies [58,72], including the implementation of a nonequilibrium thermo-chemistry solver.

In this work, we expand upon the Sherwood simulations framework by conducting additional simulation runs aimed at constraining the CWDM models detailed in Table I. The simulations are performed within a cosmological volume of $20h^{-1}$ Mpc, with 2×1024^3 dark matter and gas particles, which provides sufficient resolution to capture the small-scale structure of the Lyman- α forest [46] (see Sec. III C). Further details regarding the simulations can be found in [46] (Table I) and [94].

TABLE I. Reference thermal history simulations used in this work. A complete list, including all thermal history runs, can be found in Table I of [46]. The columns include the box size ($[h^{-1} \text{ cMpc}]$), the number of particles, the end of reionization redshift, the thermal parameters T_0 and u_0 at redshift $z = 4.6$, the inverse of the WDM mass m_{WDM}^{-1} , and the additional dark matter parameter, the fraction of WDM particle mass f_{WDM} . The first two rows list the main simulations: one with fiducial CAMB initial conditions (as in [46,95]) and the other using the CLASS code [96]. The last row contains the fiducial simulation for mass resolution correction. Box correction runs match those in [46]. Note that the simulations in all three rows extend to $m_{\text{WDM}} = 1 \text{ keV}$, with the first set covering CWDM models. Different CWDM models listed in this table were run for all 12 thermal histories of [46].

Name	L_{box} [$h^{-1} \text{ cMpc}$]	N_{part}	$z_{\text{rei}}^{\text{end}}$	$T_0(z = 4.6)$ [K]	$u_0(z = 4.6)$ [eV m _p ⁻¹]	WDM mass [keV ⁻¹]	f_{WDM}
L20-ref	20.0	2×1024^3	6.00	10,066	7.7	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}, 1]$	$[0, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}, 1]$
L20-ref ^{CLASS}	20.0	2×1024^3	6.00	10,066	7.7	0	0
R10-ref	10.0	$2 \times [1024^3, 512^3]$	6.00	10,066	7.7	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$	$[0, 1]$

In Fig. 1 we show the projected gas and dark matter density fields extracted from simulation runs for CDM and WDM cosmologies. One can see that the overdense regions in the gas distribution, which serve as the seeds for the formation of stars and galaxies, closely follow the distribution of dark matter on large scales. The gas distribution at the small scales appears more diffuse compared to the dark matter distribution in the top-right panel due to the pressure smoothing effect: the gas hydrodynamically reacts to heating from reionization by expanding, suppressing small-scale structures [93]. The neutral hydrogen, which is responsible for the Lyman- α absorption, resides in these filaments or voids. The bottom two panels highlight the primary effect of dark matter free-streaming on structure

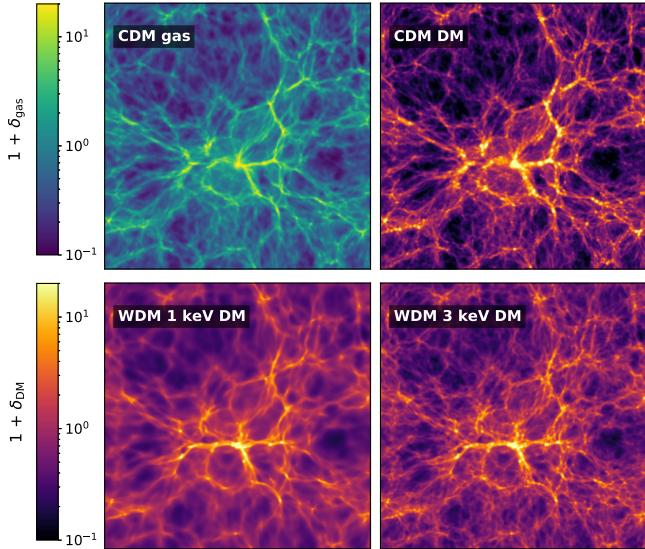


FIG. 1. Slice through a cosmological simulation of size $20h^{-1} \text{ Mpc}$ with 2×1024^3 dark matter (DM) and gas particles at $z = 4.2$ from the Sherwood-Relics simulation suite that was run with the P-Gadget3 code. Top panel: hydrogen gas distribution (left), and dark matter distribution in the standard CDM model (right). Bottom panel: dark matter distribution in a 1 keV WDM model (left) compared to a 3 keV WDM model (right), with brighter colors indicating higher density.

formation. As m_{WDM} decreases in these pure WDM models, the free-streaming length increases. This results in the smoother filamentary structure observed in the figure. Notably, this smoothing effect is more pronounced for the $m_{\text{WDM}} = 1 \text{ keV}$ model compared to the $m_{\text{WDM}} = 3 \text{ keV}$ model.

A. Imprint of mixed dark matter models on the matter power spectrum

The initial conditions for WDM and CWDM simulations differ from those of CDM in the computation of the power spectrum. The free-streaming effect, λ_{fs} , is incorporated through a transfer function, $T(k)$, that modifies the original power spectrum, (e.g., [97]). Specifically, the power spectra for CDM and pure warm thermal relic models are computed using the Boltzmann solver CAMB [95], while for the latter cosmology, we apply the transfer function from [97–99]. The mixed models are generated using the in-built CLASS option, which accounts for the temperature, mass, and density of non-cold dark matter [96,100]. To ensure consistency across different cosmologies simulated with different codes, we correct for minor discrepancies by taking the ratio of the flux power spectrum between CDM models computed with CAMB and CLASS (see Table I).

In the context of CWDM models, the transfer function asymptotically approaches a constant plateau at small scales. The height of this plateau is primarily determined by the WDM fraction, f_{WDM} , due to the presence of a dominant CDM component. The characteristic cutoff scale is governed by the mass of the warm dark matter counterpart, similar to the behavior observed in pure WDM models.

Figure 2 shows on the left the linear matter power spectrum for CDM at $z = 0$, along with the 2σ lower bound on pure WDM and on a mixed model with $m_{\text{WDM}} = 1 \text{ keV}$. The latter is the lightest thermal relic allowed in our analysis. We further show models with progressively heavier warm counterparts (and correspondingly smaller fractions), which bracket the 2σ allowed region found in this work, shaded in blue. This figure illustrates the imprint

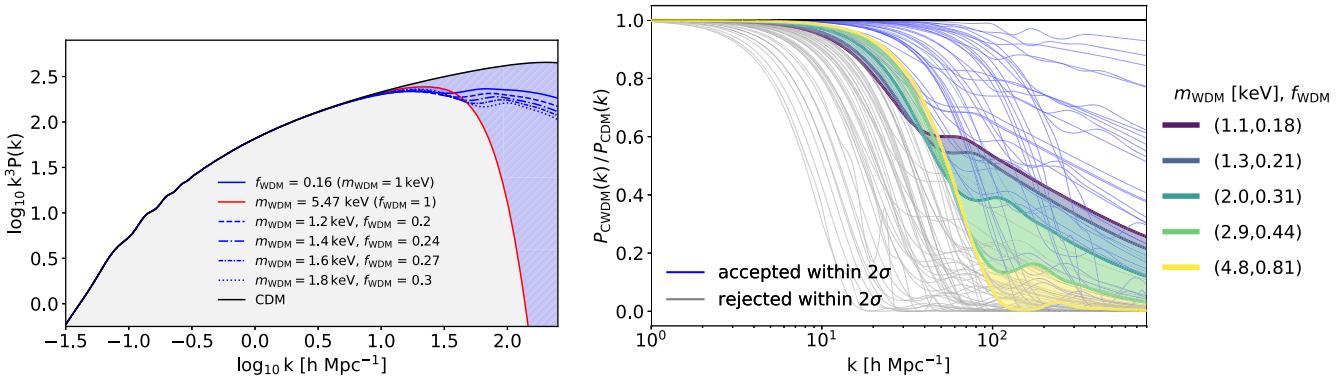


FIG. 2. Left: linear matter power spectrum at $z = 0$ computed using the Boltzmann solver CLASS for CDM (black), a CWDM model (blue), and a pure WDM model (red). The last two correspond to the 2σ bounds found when fixing the highest f_{WDM} ($f_{\text{WDM}} = 1$) and lowest m_{WDM} ($m_{\text{WDM}} = 1 \text{ keV}$) in our analysis. We also present a range of CWDM models, shown in blue with varying line styles, which are drawn from the region of power suppression of allowed models. These models delineate the 2σ allowed region, represented by the blue shaded area. The region of rejected models is shown in gray. Right: the matter transfer function [$T^2(k)$] from random samples of the 2D posterior distribution in the $(m_{\text{WDM}}^{-1} - f_{\text{WDM}})$ plane inside (outside) the 2σ contour in blue (gray). We show five allowed models in color, highlighting that we are sensitive to both the CWDM departure scale from CDM (mainly set by m_{WDM}) and the suppression amplitude (dependent on f_{WDM}). Notably, lower m_{WDM} values can be constrained at the cost of reduced suppression at higher k .

of CWDM in the matter power spectrum: the scale at which it deviates from CDM is set by m_{WDM} , while the suppression depends on the abundance given by f_{WDM} . We show in the right plot of Fig. 2 the distinct suppression patterns of mixed CWDM models relative to CDM. In particular, these hybrid models can remain light as long as f_{WDM} is small, resulting in a weaker suppression of the matter power spectrum at slightly larger scales.

However, CWDM models with stronger suppression at larger scales also show less suppression at small scales, compared to pure WDM models. This is shown for a collection of hand-picked models in the color gradient of Fig. 2. For such models, we find that the Lyman- α forest data in this work constrain both the scale of the small-scale suppression and the shape of the transfer function around the suppression scale.

B. Flux power spectrum models

The synthetic Lyman- α forest spectra, specifically the optical depth field (τ), can be computed by extracting the information on neutral hydrogen number density, temperature, and peculiar velocity contained in each pixel in the box along the different lines of sight. These are parallel to the boundaries of the box and have $N_{\text{bins}} = 2048$ number of bins with width $\delta x = L_{\text{box}}/N_{\text{bins}}$. We extract $N_{\text{LOS}} = 5000$ lines of sight and compute the flux ($F = e^{-\tau}$) field for each. The perturbations in flux, $\delta_F = F/\bar{F} - 1$, are characterized by the 1D flux power spectrum, $P_F(k) = L_{\text{box}}^{-1} \langle \delta_F \delta_F^* \rangle$, where \bar{F} is the mean transmitted flux calculated as an average of the flux field across all the pixels extracted from a simulation at each redshift. We repeat this for the three redshift bins used in our analysis.

Without postprocessing, the number of models that we have established are spanned by the allowed WDM parameters, $m_{\text{WDM}} = [1, 2, 3, 4] \text{ keV}$ and $f_{\text{WDM}} = [0.125, 0.25, 0.5, 1.0]$. These add up to 16 WDM cosmologies +1 from CDM ($f_{\text{WDM}} = 0$). As described above, for each cosmology, we run 12 simulations, each one with different thermal history, characterized by the evolution of the cumulative injected heat into the IGM at mean density [$u_0(z)$]. This amounts to $17 \times 12 = 204$ models per redshift bin.

Under the assumption of photoionization equilibrium, we postprocess the simulations before inferring our cosmology from the flux power spectra models in two steps. First, we exploit the temperature dependence of the recombination coefficient $\alpha_{\text{HII}} \propto T^{-0.72}$ in the Lyman- α forest regime to isolate the impact of the instantaneous temperature of the gas on the power spectrum. We do this by rotating and translating the $T - \rho$ plane of simulation parameters, as has been done in previous work (e.g., [46,69]), recalculating the τ field for a range of T_0 and γ values. In this way, we obtain flux models for a combination of $T_0 \in [5000, 14,000] \text{ K}$ spaced by 1000 K, and $\gamma \in [0.9, 1.8]$ in steps of 0.1.

In the second step, we adopt the common approach of adjusting the mean flux, \bar{F} , derived from the simulated spectra to match the observed measurements of effective optical depth τ_{eff} , in order to account for uncertainties in the UV background. This is done by rescaling the τ field, with \bar{F} defined as $\bar{F} = e^{-\tau_{\text{eff}}}$, since $\tau \propto \Gamma_{\text{HI}}^{-1}$ under the photoionization equilibrium assumption. This method is preferred over direct rescaling of Γ_{HI} because the mean transmitted Lyman- α flux is measurable [101–103], providing valuable prior knowledge of τ_{eff} to calibrate simulations and infer physically motivated posteriors on model parameters.

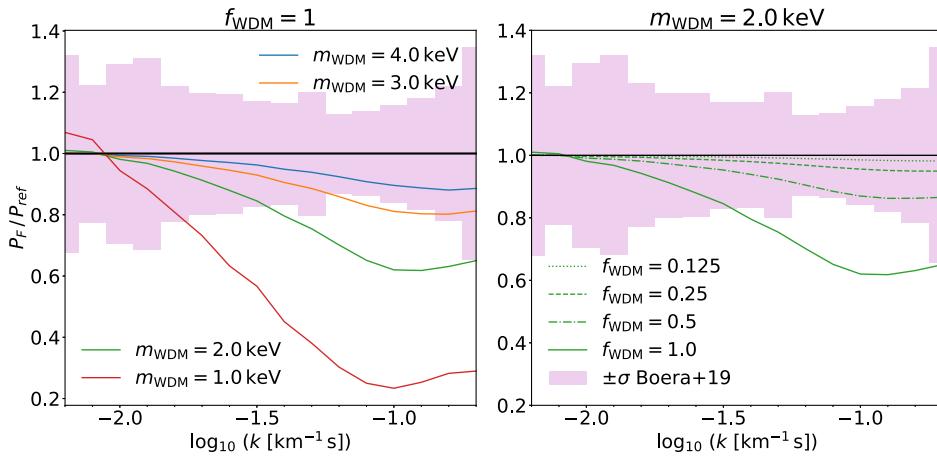


FIG. 3. Flux power spectrum ratio of simulated models to the reference model for CDM cosmology with postprocessed parameters fixed to fiducial values ($T_0 = 10,700$ K, $\gamma = 1.2$) at $z = 4.2$. Left: the WDM particle mass varied with fixed $f_{\text{WDM}} = 1$. Right: CWDW models with varying f_{WDM} for $m_{\text{WDM}} = 2$ keV.

After postprocessing, the total number of models span a grid of $15 \times 10 \times 10 \times 204$ models = 306,000 parameter values.

In Fig. 3, we show the relative ratio of these simulated flux power spectra for WDM and CWDW cosmologies to the reference CDM simulation at $z = 4.2$ and for a fixed thermal history ($u_0 = 8.14$ eV, $T_0 = 10,700$ K, $\gamma = 1.2$), and the mean flux (\bar{F}) is fixed such that the effective optical depth matches that of [69], i.e., $\tau_{\text{eff}} \equiv -\ln \bar{F} = \tau_{\text{eff}}^{\text{B}}$. One can see that the suppression in power at small scales (high k values) is enhanced for WDM models with lower m_{WDM} . This agrees with what one expects for a lighter DM particle that has a larger free-streaming length. The magnitude of the suppression increases, as the larger free-streaming strength suppresses the structure at even larger scales. This can also be seen in the behavior of the power spectrum for different f_{WDM} values, namely, the smaller f_{WDM} , the weaker the effect on the power spectrum, eventually recovering CDM. The 1σ uncertainties from [69] are also shown to illustrate the sensitivity of the data to the WDM models. The sensitivity is higher for “hotter” dark matter models, whereas the “colder” dark matter models differ by less than the error margin, making it difficult for the data to distinguish between them effectively.

C. Mass resolution and box size correction

Prior to performing the inference, we test whether the models extracted from simulations are able to resolve the Lyman- α forest structure by running numerical convergence tests. The two relevant parameters are L_{box} and N_p [104–106]. The first parameter sets the number of k -modes on large scales and influences all scales due to nonlinear coupling between the modes. When L_{box} is kept fixed, the smallest resolvable scale is governed by the particle number, N_p .

Following [46], we splice together the flux power from large and small box size simulations, $L_{\text{box}} = 40$ and $10h^{-1}$ Mpc, varying $N_p = [1024^3, 512^3]$ [107]. We check that [69] is sensitive only to the intermediate splicing regime. Therefore, we correct the flux power spectra obtained with our reference simulation $P_{20,1024}$ by using the ratio of a simulation with better resolution, $P_{10,1024}$, and an equally limited resolution as the reference, $P_{10,512}$, $R_s = P_{10,1024}/P_{10,512}$, where R_s is the mass resolution correction. The correction ranges from 5 to 14% over the scales of interest, and it increases with redshift [108]. These models are shown in Table I in [46].

Furthermore, the finite box size naturally introduces a truncation of the flux power spectrum at large scales but also affects the small scales. This systematic is corrected with the box size correction from the ratio $P_{40,2048}/P_{20,1024}$. As shown in [46], the correction at the low- k modes is at the level of 3–5%. We further check that the effect of mean flux rescaling in postprocessing on the box size correction is $\leq 0.5\%$. Therefore, we correct our models at fixed mean flux corresponding to $\tau_{\text{eff}}^{\text{B}}$.

D. Patchy correction

Cosmic reionization is not homogeneous but occurs in three stages: the formation of H II bubbles, their overlap due to the increased number of ionizing sources, and the disappearance of the last remaining neutral hydrogen regions as uniform UV background forms. As pointed out in previous work that has used the same suite of simulations (e.g., [46,74]), the nonuniform nature of the reionization process leaves an imprint on the Lyman- α forest at both large and small scales. The former consists of an enhancement of power at scales $k < 0.03$ km^{-1} s, more noticeable at higher redshifts, as well as thermal histories with an earlier end of reionization. This can be understood

through the temperature dependence of the recombination coefficient already mentioned in Sec. III B, which means that fluctuations in the gas temperature contribute to the variations in the neutral hydrogen fraction. At small scales, patchy reionization causes suppression in the power spectrum mainly through thermal broadening and differences in peculiar velocity for regions that reionized recently [74,109]. We account for this extra power by computing the correction tabulated in [74] following [46].

IV. STATISTICAL MEASURES OF THE FLUX DISTRIBUTION

The flux power spectrum models are defined on a six-dimensional grid per redshift bin, for each combination of model parameters $(T_0^{z_i}, \gamma^{z_i}, \tau_{\text{eff}}^{z_i}, u_0^{z_i}, m_{\text{WDM}}^{-1}, f_{\text{WDM}}) = \vec{\theta}$. Comparing power spectra derived from observations and simulations enables us to constrain the free parameters of the model, $\vec{\theta}$, a process typically conducted within a Bayesian MCMC framework. In this study, we utilize the Cobaya code [110] to sample the Gaussian multivariate likelihood function from [69] at each redshift. However, our approach differs from previous works in the interpolation scheme used to generate simulated power spectra at each grid point. Specifically, we employ a machine learning model capable of predicting the power spectra within a fraction of a second on a single CPU.

A. Neural network emulator

Given that $\vec{\theta} \in \mathbb{R}^6$, evaluating the likelihood at each grid point is limited by the computation cost of interpolating in such a high-dimensional space. We circumvent this problem by employing a neural network. This neural network maps the input parameters θ to the output parameters $P_F(k_m)$, in the form $P_F(k_m) = f[\vec{\theta}, \phi]$, where ϕ are the hyperparameters that describe the relationship between input and output, and k_m is a finite set of wave numbers matching the ones reported in the measurements of [69].

As an initial framework, we adopt the model architecture proposed by [82] and make two key observations. First, the original work assessed the emulation error using k -fold cross-validation on the training set itself. This approach likely leads to overfitting, reducing the model's ability to generalize to unseen data, which is precisely the role of an independent test set. Second, we can enhance the performance of the emulator by using a combination of baseline models, also known as ensemble learning. In particular, we consider averaging the output from a range of models trained using k -fold cross-validation.

We begin by randomly splitting the dataset into a 90% training set and a 10% test set, with the validation set excluded within the training subset. We then divide the training set into k subsets, and perform the k -fold cross-validation. In this way, each neural network is trained on

$k-1$ folds of the training data, with the remaining one being used as validation. We use $k = 10$, therefore repeating the process iteratively 10 times. This results in 10 neural network emulators, each validated on a different subset of the mock spectra. We build our neural network model using PyTorch in Python, which enables GPU execution. Before training, we apply min-max normalization to rescale the input data, ensuring that $\vec{\theta} \in [0, 1]$, while the output (labels) undergo a logarithmic transformation. These pre-processing steps mitigate data variability, thereby enabling the network to learn more efficiently. We perform cross-validation for several hyperparameter combinations and, after monitoring the convergence of the loss function, we settle on a neural network with [120,120,120] units and a rectified linear unit activation function (see Fig. 4). We further find that the minimum of the validation loss is generally achieved after $\Delta t = 500$ epochs. However, by allowing for early stopping, we note that the lowest redshift bin completes training earlier, whereas higher redshift bins do not meet the early stopping criteria. We further optimize training by allowing the initial learning rate to decrease from an initial value of 10^{-3} to a minimum value of 10^{-6} if the validation loss does not improve for $\Delta t = 30$ epochs. We choose the Adam optimization algorithm and a batch size = 18. Even though this batch size seems small compared to the total data size, we check that larger batches are not able to achieve as small losses for even larger numbers of epochs.

This process of hyperparameter tuning is, therefore, done within cross-validation, where each neural network is validated using a different subset of the original training set, while the test set remains unseen. The building process of the emulator is shown in Fig. 4.

In this way, we end up with 10 neural networks per redshift bin at our disposal. To obtain the final prediction, we average the prediction extracted from each model using the arithmetic mean. Since each model was trained in parallel and on a different training set in cross-validation, the combination of these models will give a better average performance than each model individually. We also use this method to estimate the combined emulator's error, as shown in Fig. 5. This shows that the error from the emulator is significantly smaller than the error on the data shown in Fig. 3 and the one used in [82] without model ensembling and a separate test set.

V. RESULTS

The results of the main analyses for thermal and dark matter parameters are shown in Figs. 6 and 7. We employ a uniform prior on m_{WDM}^{-1} , specifically $\mathcal{U} \sim [0, 1] \text{ keV}^{-1}$, after verifying that the emulator's predictions are unreliable in extrapolated regions of the parameter space. Similarly, we impose a uniform prior on f_{WDM} , $\mathcal{U} \sim [0, 1]$. Figure 6 shows the 1 and 2σ recovered posterior projected in the

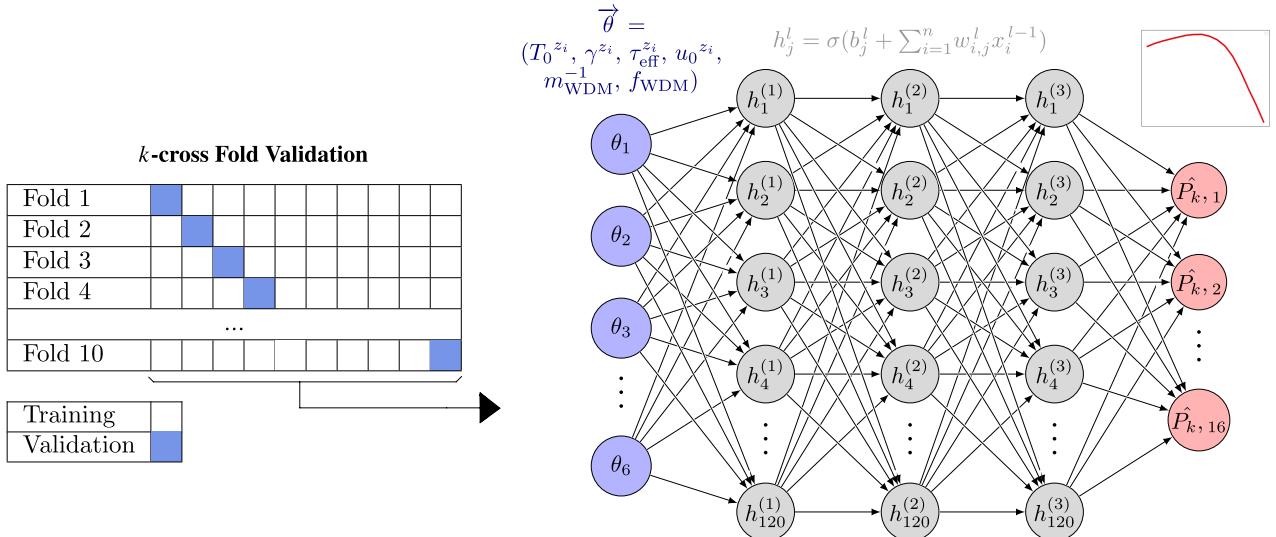


FIG. 4. Illustration of the emulator building process used in this work. Left: k -fold cross-validation. Right: neural network trained on each fold $i \in [1, 10]$. We first extract the training set from the total input and output data to perform a k -fold cross-validation, as shown on the left. Each fold is divided into 10 subsets, nine are used to train the neural network model depicted on the right, while the remaining one is used as a validation set to monitor the convergence. We feed the cosmological model parameters $\vec{\theta}$ in the input blue layer. The hidden layers are shown in gray. For hidden neuron j in layer l with $n = 6$ input parameters, the inputs to each node are combined in the weighted linear combination h_j^l shown on top in gray, where \vec{b} is the bias vector and \mathbf{w} is the weight matrix, such that $\phi = [\vec{b}, \mathbf{w}]$, x^{l-1} are the outputs from the previous layer, and σ is the so-called activation function, which permits nonlinear relations between the input and the output. This step results in 10 models validated on different subsets of the training data. Then, we feed the unseen test data into each model and average the predictions across all the emulators. The outputs are the power spectrum values in 16 k -bins shown in red, corresponding to the output layer. Note that this is repeated for the three redshift bins, $z = 4.2, 4.6, 5.0$, of the data.

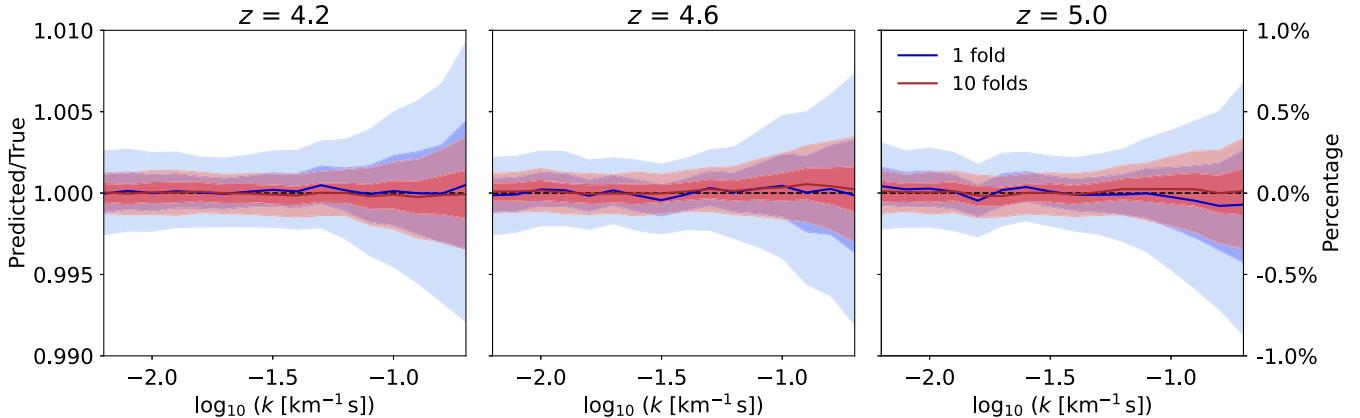


FIG. 5. Characterizing the performance of the neural network emulator with the predicted versus true power spectra ratio distribution on the test set before and after performing k -fold cross-validation. Lighter blue and red regions correspond to the 95% confidence intervals for one fold (before cross-validation) and for ten folds (after cross-validation), respectively. The same applies to the darker shaded regions, which correspond to the 68% confidence interval. The median of the distribution is shown, as well as a solid line. Perfect recovery is indicated by the gray dashed line.

$(m_{\text{WDM}}^{-1} - f_{\text{WDM}})$ plane. The shape of the contours matches the expected degeneracy between the two parameters, where the data allow for a very light thermal relic as long as its abundance is small ($m_{\text{WDM}}^{-1} \rightarrow 1$), or a heavy thermal relic with large abundance ($f_{\text{WDM}} \rightarrow 1$), with

CDM at $(m_{\text{WDM}}^{-1}, f_{\text{WDM}}) = (0, 0)$ in this parametrization. This degeneracy is well fitted by a power law of the form $f_{\text{WDM}} = A \times (1 \text{ keV}/m_{\text{WDM}})^b$. For the 2σ contours of the default analysis, we find $A = 0.14 \pm 0.0007$ for the normalization, and $b = -1.1 \pm 0.0033$ for the index

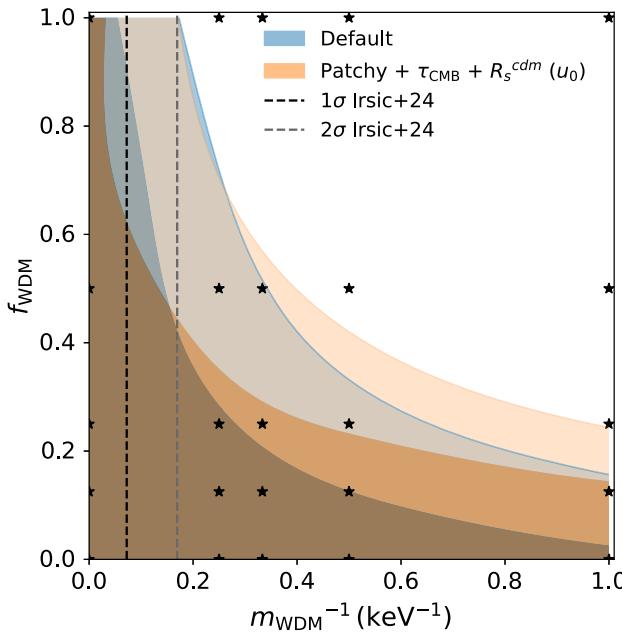


FIG. 6. The 2D posterior distribution in the f_{WDM} and m_{WDM}^{-1} plane for our default analysis (Gaussian priors on T_0) and for patchy + $\tau_{\text{CMB}} + R_s^{\text{cdm}}(u_0)$. Vertical black and gray dashed lines correspond to the 1σ and 2σ constraints from [46]. The grid of simulations is also shown as the starred black points.

parameter. We further identify a combination of f_{WDM} and m_{WDM}^{-1} that is perpendicular to the above degeneracy axis, and is best-constrained by the data, through the parameter $W_{\text{WDM}} \approx f_{\text{WDM}} (1 \text{ keV}/m_{\text{WDM}})^{3.4}$. The 1σ models in the 2D posterior of Fig. 6 result in the upper bound $W_{\text{WDM}} < 0.27$ at 1σ .

The default analysis uses Gaussian priors on $T_0^{z_i}$ centered on fitted $T_0^{z_i}$ measurements from observations at low- z [112] and high- z [113], following [46]. The second analysis (orange) represents the model that incorporates all the physical considerations in addition to the default priors. These include a thermal-dependent resolution correction, patchy reionization correction, and $u_0 - T_0$ informed through the measurement of τ_{CMB} from Planck [114]. The following sections outline each correction and the resulting m_{WDM} and f_{WDM} constraints from the analyses, which are summarized in Fig. 8. For the two main models in Fig. 6, we show their fit to the data in Fig. 9. We further show that the predicted IGM thermal evolution is in agreement with independent numerical and observational predictions in Figs. 10 and 11. A summary of the results from all the analyses, including thermal and dark matter parameters along with their best-fit χ^2 values, is presented in Table II.

A. Thermal priors

To effectively constrain the $(m_{\text{WDM}}^{-1} - f_{\text{WDM}})$ plane, we need to marginalize over the two astrophysical effects

that impact the small-scale power spectrum, apart from dark matter free-streaming.

The thermal histories within our hydrodynamical simulations were constructed to bracket observational constraints on the thermal state of the IGM and effective optical depth evolution. These models can provide a useful prior in the $(u_0 - T_0)$ plane, corresponding to the gray band in the 2D contours in the top row of Fig. 7, as was done in [46].

Alternatively, as proposed by [46], one can impose a prior on the thermal state of the IGM represented by T_0 , based on observational data and, therefore, independent of the specific physical models used in our analysis. We apply these priors in our default analysis shown in Fig. 7.

Finally, one can impose informed thermal priors through measurements of the Thomson scattering optical depth, denoted as τ_{CMB} , derived from CMB data. This parameter offers a means to assess the IGM thermal history by quantifying the integrated scattering rate of photons by electrons over time, a quantity which is inherently dependent on the electron number density [94]. Consequently, for each thermal history model, there exists a mapping between the integrated and instantaneous thermal history, which is parametrized by $u_0^{z_i}$ and $T_0^{z_i}$, and the electron scattering optical depth. We use this mapping to interpolate τ_{CMB} for each sampled combination $(u_0^{z_i}, T_0^{z_i})$. At the likelihood level, we introduce an additional Gaussian term for τ_{CMB} , centered on $\tau_{\text{CMB}} = 0.054 \pm 0.007$, as inferred by Planck [114]. While the derived τ_{CMB} is heavily dependent on our simulated models, the prior effectively constrains the thermal parameters in a manner that ensures the preferred cosmology is consistent with CMB measurements of this parameter.

We find similar results to [46] in that the contour in the $(u_0 - T_0)$ plane for the default analysis expands along the anticorrelation direction, with slightly colder models allowed within 1σ at $z = 4.2$ and 4.6 , as shown by the blue contours in the upper row of Fig. 7. This then results in a slight upward shift in the $(u_0 - m_{\text{WDM}}^{-1})$ plane in the middle row, with colder models showing less pressure smoothing. The effect of different thermal priors does not have a noticeable effect in the $(m_{\text{WDM}}^{-1} - f_{\text{WDM}})$ plane (bottom row); therefore, the 2σ constraints for pure WDM (see Table II) are also very similar. The addition of τ_{CMB} only slightly shrinks the posterior in the thermal parameter space at $z = 4.2$, resulting in smoother 1σ contours, but the inferred parameters are practically the same. For mixed DM models, we find that $f_{\text{WDM}} \lesssim 16\%, 35\%, 50\%$, and 67% for fixed $m_{\text{WDM}} = [1, 2, 3, 4] \text{ keV}$, respectively, where the choice of different thermal priors only changes these constraints by $\leq 5\%$.

B. Effect of small-scale cut

The previous datasets used to infer the WDM constraints extend to $k \approx 0.1 \text{ s/km}$ [46, 57, 116]. We check whether we reproduce their results on pure WDM thermal relics by

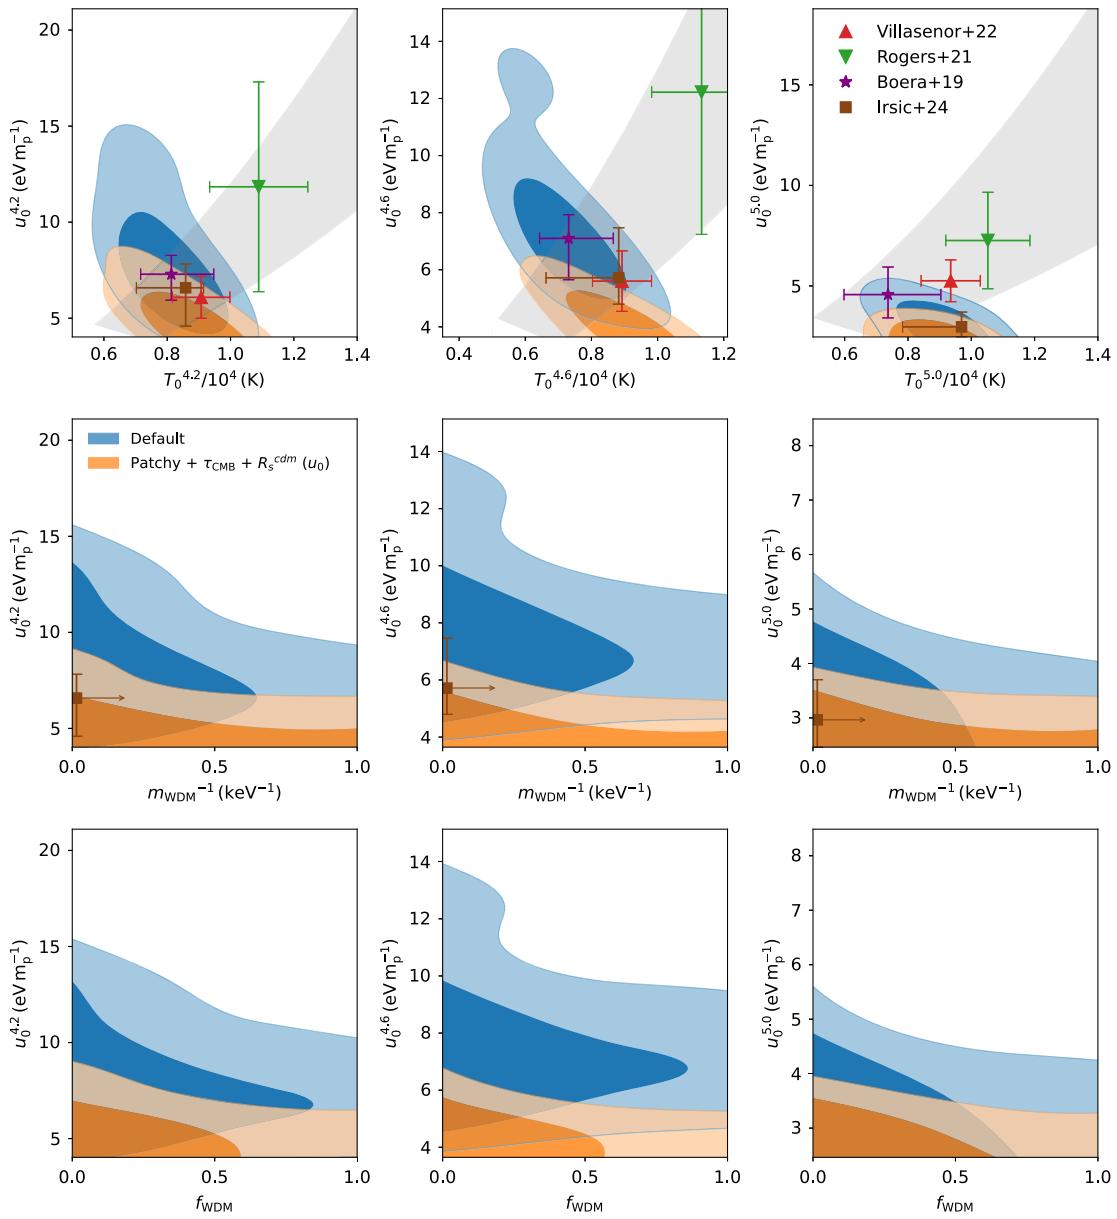


FIG. 7. The 2D posterior distributions for thermal parameters and dark matter parameters for the two analyses shown in Fig. 6. The analyses shown are corrected for resolution, R_s^{cdm} , *a posteriori*. The gray band represents the prior band imposed in the $u_0 - T_0$ plane. Best-fit values for the same datasets are also shown in red [111], green [62], purple [69], and blue [46]. In the middle row, we show the 2σ upper limit on m_{WDM}^{-1} from [46] slightly shifted to the right for illustrative purposes.

performing an analysis with both uninformative and default thermal priors and without the last three k -bins in the data. These scales are key to disentangle thermal effects from dark matter free-streaming, meaning that without them, constraints in this parameter space will be dominated by the priors. The best-fit model for the default analysis with the small-scale cut precisely recovers $T_0^{z_i}$, within the Gaussian prior's peak. The case for uninformed priors expands along the $u_0 - T_0$ degeneracy axes, effectively showing no sensitivity to these parameters at the scales considered. The recovered χ^2 when including or not including thermal priors are 12.3/20 and 11.6/20, respectively. We further

recover weaker constraints for pure $m_{\text{WDM}} < 3.67$ and $m_{\text{WDM}} < 3.82$ keV (2σ) for each analysis, consistent with [46,70]. In terms of mixed dark matter models, the small-scale cut weakens the allowed f_{WDM} by 15%, with a stronger effect for $m_{\text{WDM}} = 3$ keV, where the difference is $\approx 30\%$, as shown in Fig. 8.

C. Effect of noise

The small-scale power in the Lyman- α power spectrum includes contributions from metal lines, instrumental resolution, and potentially underestimated instrumental noise.

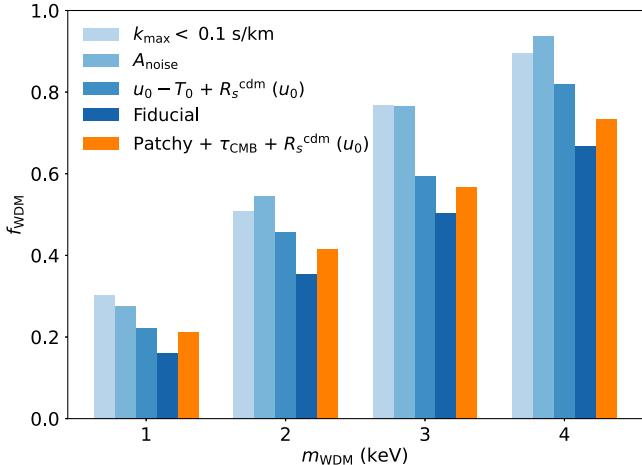


FIG. 8. Comparison of 2σ upper bounds on f_{WDM} at fixed m_{WDM} bins.

As mentioned in Sec. II, we correct for the instrumental resolution and pixel size of the data following [69]. Furthermore, we use the analysis results of [69] that correct for the effect of the intervening metal lines. However, to address instrumental noise, Ref. [69] estimated the ‘‘noiseless’’ power spectrum by subtracting the white noise power corresponding to the mean of the noise distribution, which dominates the measurement at high k values. Reference [46] fitted the noise data for each redshift bin, as estimated by [69], to a log-normal distribution. To assess whether the noise flux power may be underestimated in the data, Ref. [46] added an additional term to the theoretical model, which accounts for deviations from the mean of the noise distribution. We model the noise in the same way,

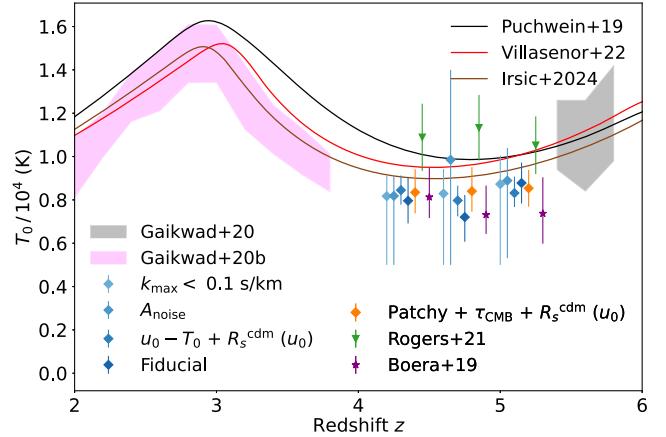


FIG. 10. Marginalized mean of the posteriors and corresponding error bars for $T_0(z)$ with other thermal evolution constraints. We show in different tones of blue the inferred temperature and mean density for a variety of analyses. The gray and pink bands show independent measurements from [112,113]. The black line shows the predicted evolution from our simulated models [115]. Inferred values from other works, such as [46,62,69,105], are also shown. We have shifted the redshift axis for each new measurement for visual clarity.

including an extra parameter with log-normal priors, $A_{\text{noise}}^{z_i}$, alongside the power spectra predicted by the emulator, as given by Eq. (4) in [46]. This analysis excludes pure thermal relics with $m_{\text{WDM}} < 3.82$ keV (2σ), in agreement with [46].

Moreover, we recover the same behavior for $A_{\text{noise}}^{z_i}$, which remains largely insensitive to m_{WDM}^{-1} (that is, the free-streaming effect), while exhibiting a correlation with the thermal parameters $T_0^{z_i}$ and $u_0^{z_i}$. This correlation

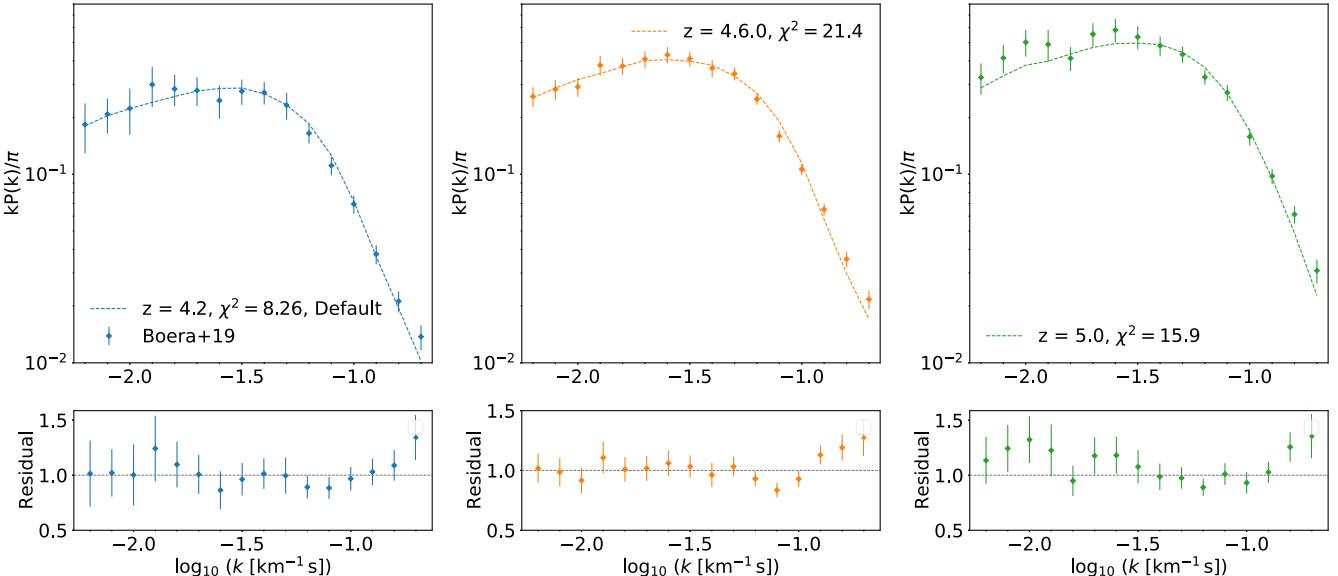


FIG. 9. Best-fit plot with individual χ^2 for each redshift bin for our default analysis (Gaussian T_0 priors + R_s^{cdm}) with four redshift-dependent parameters (T_0 , γ , τ_{eff} , u_0) and fixed dark matter parameters, m_{WDM}^{-1} and f_{WDM} . The combined weighted χ^2 is 45.6/34.

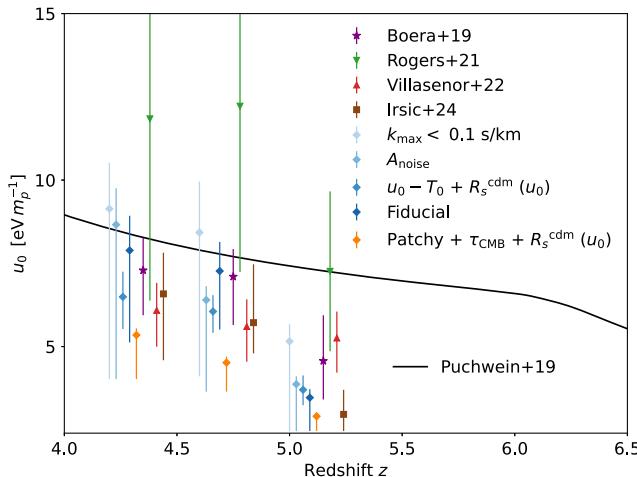


FIG. 11. Marginalized mean of the posteriors and corresponding error bars for $u_0(z)$ with other inferred constraints at the same redshift bins from [46,62,69,105].

with the IGM temperature favors hotter models (see Table II), whereas its anticorrelation with $u_0^{z_i}$ suggests a preference for later reionization, implying smaller heat injection. Importantly, we obtain mean values for the posteriors of $A_{\text{noise}}^{z=4.2} = 0.72^{+0.98}_{-0.34}$, $A_{\text{noise}}^{z=4.6} = 1.19^{+1.61}_{-0.85}$, and $A_{\text{noise}}^{z=5.0} = 1.19^{+1.65}_{-0.78}$, suggesting that the noise in the data might indeed be underestimated. Compared to [46], the 1σ error is broader, primarily due to the large uncertainty in $T_0^{z_i}$, which allows for a very warm IGM temperature, thereby weakening constraints on the WDM mass. The preference toward less heat injected (lower u_0) is particularly noticeable at $z = 4.6$, as shown in Figs. 10 and 11.

Regarding CWDW constraints, we find that the upper bounds on f_{WDM} decrease by up to 25% (see Fig. 8),

yielding very similar results to those obtained in the small-scale cut analysis, as described in Sec. V B.

D. Effect of mass resolution

Our emulator provides an averaged prediction of the absolute binned flux power spectrum, as described in Sec. IV A. Assuming that the box size and mass resolution corrections are independent of the parameters in our flux models, we apply these corrections *a posteriori*, that is, after obtaining the predicted flux power spectrum from the emulator. This corresponds to the default analysis resulting in the fit shown in Fig. 9.

However, this simplified method does not account for the fact that the mass resolution correction depends on thermal parameters. The missing small-scale perturbations in the flux power spectrum are influenced by the initial perturbations in this regime. These perturbations contribute to the formation of small-scale structures, which are smoothed out in models with more cumulative injected heat (higher u_0). As a result, models where the hydrodynamic response naturally suppresses small-scale perturbations tend to be more numerically converged, leading to a smaller resolution correction. This effect was explored in [46], which allowed this work to impose even tighter constraints on dark matter free-streaming. We recompute the models and find that scenarios with earlier reionization require a correction that is about 5% smaller than that required for colder models. This refined resolution correction is denoted as $R_s^{\text{cdm}}(u_0)$. In practice, we also apply this correction *a posteriori*, interpolating in $u_0^{z_i}$ on the fly during sampling.

We find that including the thermal dependence in the mass resolution tightens the posteriors in the thermal parameter space while broadening the contours in the dark matter parameter space, ultimately leading to weaker

TABLE II. Best-fit constraints on thermal and dark matter model parameters for different analyses at $z = 4.6$. Default corresponds to using Gaussian T_0 priors and mass resolution correction R_s^{cdm} , except for the thermal-resolution-dependent analysis, where $R_s^{\text{cdm}}(u_0)$ is applied instead. The remaining analyses are additive on top of the default.

Analysis	T_0 [10 ⁴ K]	u_0 [eV m _p ⁻¹]	m_{WDM} (2 σ) [keV] ($f_{\text{WDM}} = 1$)	A_{noise}	$\chi^2/\text{d.o.f.}$
Default	$0.714^{+0.113}_{-0.099}$	$6.461^{+2.10}_{-0.58}$	>5.47	...	45.6/34
τ_{CMB}	$0.749^{+0.066}_{-0.137}$	$6.362^{+2.149}_{-0.331}$	>5.52	...	47.6/34
$u_0 - T_0$	$0.749^{+0.101}_{-0.033}$	$6.349^{+0.612}_{-0.644}$	>5.42	...	47.1/34
$k_{\text{max}} = 0.1$ s/km	$1.18^{+0.22}_{-0.68}$	$5.03^{+4.93}_{-0.9}$	>3.58	...	11.6/20
A_{noise}	$1.30^{+0.1}_{-0.8}$	$3.87^{+2.93}_{-0.221}$	>3.82	$1.30^{+0.1}_{-0.09}$	15.3/31
$\tau_{\text{CMB}} + R_s^{\text{cdm}}(u_0)$	$0.862^{+0.021}_{-0.201}$	<5.92	>5.22	...	43/34
$u_0 - T_0 + R_s^{\text{cdm}}(u_0)$	$0.893^{+0.107}_{-0.275}$	<5.75	>4.66	...	39.2/34
patchy + $u_0 - T_0$	$0.754^{+0.068}_{-0.07}$	<6.28	>5.86	...	55.7/34
patchy + $u_0 - T_0 + R_s^{\text{cdm}}(u_0)$	$0.788^{+0.043}_{-0.089}$	<4.70	>5.11	...	46/34
patchy + $\tau_{\text{CMB}} + R_s^{\text{cdm}}(u_0)$	$0.885^{+0.055}_{-0.1345}$	$3.69^{+1.01}_{-0.94}$	>5.39	...	45.2/34

constraints on m_{WDM} and f_{WDM} . This conclusion holds for the default Gaussian T_0 priors, and when including $u_0 - T_0$ priors on top of the default prior, with the latter exhibiting only a slightly reduced constraining power in the high m_{WDM}^{-1} regime due to the posteriors in $u_0 - T_0$ being confined to the envelope of the simulations. For a pure WDM cosmology, we obtain a lower bound of $m_{\text{WDM}} > 4.66$ keV for $f_{\text{WDM}} = 1$ when accounting for the thermal dependence of mass resolution, similar to bounds found by [46]. The constraints on f_{WDM} and fixed mass bins become weaker by 5–10% (Fig. 8).

E. Effect of patchy correction

We account for the effect of patchy reionization on the default analysis by repeating the inference using a new set of emulators trained on models that include the patchy correction described in Sec. III D. With default thermal priors, we find that the data tend to prefer cold models with very low u_0 values. This was observed in [46], and occurs due to patchy-corrected models showing less structure at the small scales as a result of the peculiar velocity field structure [117]. These models therefore match a very late end of reionization, eventually becoming unphysical. If we impose a $u_0 - T_0$ prior, the posterior still shifts to lower values but within the envelope of simulations, leaving more room for dark matter free-streaming. The bounds on WDM now become $m_{\text{WDM}} > 5.86$ keV, with little change of $\leq 5\%$ on mixed models.

Adding the thermal-dependent resolution correction shifts the posterior in the other direction, since resolution correction increases flux power at small scales, increasing with decreasing heat injected during reionization. As a result, slightly warmer IGM temperatures are allowed with u_0 , especially at high z , pushing the lower bound. As a result, the constraints slightly weaken to $m_{\text{WDM}} > 5.11$ keV with the $u_0 - T_0$ prior, and to a lesser extent ($m_{\text{WDM}} > 5.39$ keV) with the τ_{CMB} prior. We show this behavior in orange for the analysis with thermal priors using τ_{CMB} , patchy, and a thermal-dependent resolution correction in Fig. 7. We further note that the χ^2 does not change noticeably compared to the default analysis when both patchy and R_s^{cdm} are modeled (see Table II). The CWDM constraints, as shown in Fig. 8, are stronger than the other analyses and comparable to the default one. From the 2σ contour in Fig. 6, the patchy + $\tau_{\text{CMB}} + R_s^{\text{cdm}}$ (u_0) orange model is more inflated than the fiducial model toward high m_{WDM}^{-1} , yielding slightly weaker constraints at the fixed mass bins shown in Fig. 8. In general, the difference with respect to other analyses is more noticeable in the intermediate mass bin, $m_{\text{WDM}} = 3$ keV, corresponding to the region in Fig. 6 where the degeneracy between f_{WDM} and m_{WDM}^{-1} is stronger. While these models incorporate physical effects applied *a posteriori*, the resulting constraints remain unchanged, as the opposing effects of patchy reionization and resolution correction

counterbalance each other. To the best of our knowledge, this approach represents the most comprehensive level of modeling. Investigating a cosmology-dependent resolution correction, alongside other effects, including peculiar velocity [46], will improve our understanding of the recovered bounds for each mass bin in the context of CWDM models. This is particularly crucial in the regime $k > 0.1$ km $^{-1}$ s, where the sensitivity to dark matter free-streaming is the highest.

F. Comparison to previous CWDM analyses

The main results of this work are shown in Fig. 8. A thermal relic as light as $m_{\text{WDM}} = 1$ keV is allowed by the data if it contributes $< 16\%$ at 2σ to the total dark matter energy density, with increasing f_{WDM} for increasing mass. The pure WDM limit in the default analysis results in a $m_{\text{WDM}} > 5.47$ keV (2σ), which is slightly weaker than the mass bound found by [46]. However, within our set of analyses, we also recover the $m_{\text{WDM}} > 5.86$ keV (2σ) constraints (Table II). These lower bounds are highly sensitive to the modeling of $P_F(k)$, including the corrections described in Secs. III C and III D. Moreover, incorporating f_{WDM} into the $\vec{\theta}$ model parameters requires sampling a higher-dimensional parameter space during MCMC compared to [46]. This introduces nontrivial degeneracies, particularly the one illustrated in Fig. 6, which requires a substantial number of samples to achieve convergence. Notably, the quoted limits have been derived by filtering the original samples from CWDM runs, in contrast to fixed mass bin analyses conducted for axions [118,119]. A similar degeneracy between the fraction and the axion mass has been observed also in previous work based on simulations of mixed dark matter [120], with updated constraints by [62] using the same Lyman- α data from [69] as in this work. The degeneracy in the context of CWDM models, however, appears in the keV mass range.

The main two analyses carried out in the context of Lyman- α CWDM constraints are Refs. [52,83]. The former found constraints on thermal relics that are significantly weaker, allowing the smallest thermal-relic mass probed, $m_{\text{WDM}} = 1.1$ keV, for $f_{\text{WDM}} < 40\%$, with the pure WDM limit $m_{\text{WDM}} \geq 1.7$ keV (95% C.L.). These constraints are not competitive with the results presented in Fig. 6 mainly due to the lower resolution of the data used (SDSS and the VISTA Hemisphere Survey with $k_{\text{max}} = 0.03$ km $^{-1}$ s).

Reference [83] analyzed BOSS DR9 data combined with XQ-100, HIRES, and MIKE, finding $m_{\text{WDM}} > 0.7$ keV for $f_{\text{WDM}} < 0.10$. Their m_{WDM} bound falls outside our simulation grid in the $(f_{\text{WDM}}, m_{\text{WDM}}^{-1})$ plane. Their power-law fit to the posterior in this plane, given by $f_{\text{WDM}} = A \times (1 \text{ keV}/m_{\text{WDM}})^b$, yielded the same A value but smaller b compared to the fit described in Sec. V. This yields an inflated posterior for $f_{\text{WDM}} \rightarrow 1$ (larger α), leading to

weaker WDM constraints, which declines more steeply, with $f_{\text{WDM}} \leq 0.17$ at $m_{\text{WDM}} = 1 \text{ keV}$, almost identical to our constraints. Compared to [83], the work presented in this paper uses improved simulations that vary thermal history for all the CWDM models considered. The effects of thermal history are marginalized over a more physical and conservative prior range than in [83]. This would weaken the resulting constraint on CWDM models, a point discussed in previous works [58,83]. However, the improved resolution (higher k_{max}) and treatment of systematics (e.g., noise and metal contamination) in [69] counterbalances this effect, yielding slightly stronger constraints for these mixed models.

The results by [83] show that, as m_{WDM}^{-1} increases, the 1σ contour in the $(f_{\text{WDM}}-m_{\text{WDM}}^{-1})$ plane shrinks when using higher-resolution data. This would, in principle, mean that there exists a light thermal relic for which the contour closes, meaning that Lyman- α forest data would be able to completely exclude CWDM models beyond a m_{WDM} value.

In our analysis, however, we use data that reaches to higher k_{max} than in [83]. While the posterior in Fig. 6 shrinks due to the small-scale information available from [69] data, we expect the high m_{WDM}^{-1} regime to achieve a constant plateau. This arises from the imprint that these light thermal relics ($m_{\text{WDM}} < 1 \text{ keV}$) leave on the matter power spectrum, as discussed in Sec. III. For large m_{WDM}^{-1} , suppression in the matter power spectrum shifts to lower k -modes. While the Lyman- α forest data are sensitive to the mildly nonlinear scales, and therefore not to the free-streaming scale from these light thermal relics, the plateau reached by these mixed dark matter models shifts the overall amplitude of $P_F(k)$. This effect is degenerate with other astrophysical effects, mainly with the mean IGM transmission, τ_{eff} . The difference between these flux models is, therefore, smaller than the error on the data, shown in Fig. 3, meaning that the data will not be sensitive to them. Consequently, even if additional simulations were run for $m_{\text{WDM}} > 1 \text{ keV}$, we expect the data to maintain the constraint on the fraction f_{WDM} at the same level. This is likely true for the models where the half-mode scale is smaller than 0.05 Mpc^{-1} , at which point marginalizing over other cosmological parameters (e.g., A_s , n_s) would become important.

VI. CONCLUSIONS

In this work, we present updated constraints on CWDM models using new 1D flux power spectrum measurements from the high-redshift Lyman- α forest provided by [69]. The statistical power of this dataset is limited by cosmic variance as a result of the small number of QSO sightlines available. However, the number of sightlines has more than doubled compared to previous analyses using the same spectrographs, and the improvement in the systematic modeling provides access to scales a factor of ≈ 2 smaller

than those probed in previous WDM and CWDM constraints. This enables direct measurement of the flux power spectrum at wave numbers $k \geq 10h \text{ Mpc}^{-1}$, precisely where free-streaming effects from the WDM component are expected to suppress structure formation in these hybrid models.

Our analysis indicates that a thermal relic with $m_{\text{WDM}} = 1 \text{ keV}$ is allowed if its contribution to the total dark matter density does not exceed 16% (2σ). For higher mass bins ($m_{\text{WDM}} = 2, 3, 4 \text{ keV}$), increasingly larger abundances are permitted, with $f_{\text{WDM}} = 0.35, 0.50, \text{ and } 0.67$ at 2σ . These findings refine previous constraints from [52], which were based on lower-resolution SDSS and VHS data, as well as a limited number of simulations for low m_{WDM} values. Additionally, our results align closely with those of [83] on CWDM in the low- m_{WDM} regime explored by our simulations, with minor differences arising from variations in thermal history treatment, simulation runs at lower m_{WDM} grid points, and improved data from HIRES.

On the pure WDM side, we also obtain stronger constraints than previous studies [52,70,116], due to the inclusion of higher k -bin information. These constraints remain consistent with those reported by [46], although fiducial analysis results in slightly weaker lower bounds. This is primarily attributed to the extensive sampling required to explore the high-dimensional parameter space, particularly the degeneracy between f_{WDM} and m_{WDM}^{-1} . However, in certain analyses, we derive even more stringent constraints. For instance, the patchy + $u_0 - T_0$ model yields $m_{\text{WDM}} > 5.86 \text{ keV}$, compared to $m_{\text{WDM}} > 5.10 \text{ keV}$ in [46]. This suggests that incorporating additional physical effects may alleviate the strong degeneracies among sampled parameters. Specifically, patchy models consistently favor scenarios with reduced heat injection, which may, in turn, enhance the efficiency of likelihood sampling.

Beyond Lyman- α forest analyses, CWDM has also been constrained by combining BOSS data with Planck and estimates of the number of satellite galaxies [85]. Our constraint at the lowest mass bin probed in this analysis, $m_{\text{WDM}} = 1 \text{ keV}$, is 10% stronger. More recently, Ref. [121] demonstrated that the combination of CMB data and the predicted 21 cm signal power spectrum yields very weak constraints in the $(f_{\text{WDM}} - m_{\text{WDM}}^{-1})$ plane, with a lower bound of $m_{\text{WDM}} > 1.8 \text{ keV}$ ($f_{\text{WDM}} = 1$), in contrast to the significantly stronger constraints from the Lyman- α forest given in this work. In particular, the main findings of this work can be summarized as follows:

- (i) We provide new updated constraints on CWDM cosmology, where a thermal relic with $m_{\text{WDM}} = 1 \text{ keV}$ is allowed for $f_{\text{WDM}} < 0.16$ (2σ). Higher mass bins are viable with increasing abundances, leading to a strong degeneracy between f_{WDM} and m_{WDM}^{-1} parameters. The 2σ contour $(f_{\text{WDM}} - m_{\text{WDM}}^{-1})$ plane is well described by $f_{\text{WDM}} = 0.14 (1 \text{ keV}/m_{\text{WDM}})^{-1.1}$.

- (ii) We update a previous Bayesian inference framework for 1D flux power spectrum Lyman- α forest analysis by integrating a neural network emulator. Model ensembling via k -fold cross-validation reduces the emulator error to $<0.5\%$ at 2σ across all k -bins.
- (iii) We recover a thermal state and evolution of the IGM consistent with previous Lyman- α forest analyses [62,69,70], and external measurements [112,113], with very mild variations in T_0 and u_0 due to the strong degeneracy between these two parameters. The mean IGM transmission is also in agreement with independent measurements from, e.g., [101,103,122].
- (iv) Limiting the amount of information in the high k -modes results in weaker WDM constraints of $m_{\text{WDM}} < 3.6$ keV, consistent with previous work by [70]. We further recover similar findings by [46] that the instrumental noise on the data might be underestimated by $\approx 30\%$.
- (v) Modeling of patchy reionization and thermally dependent resolution correction can vary the 2σ constraints on f_{WDM} by $\approx 5\%$ in most cases. Further work on investigating the effect of peculiar velocity on small scales, as pointed out by [46], as well as on a cosmology-dependent resolution correction, will help mitigate these differences.

Our results build upon the latest constraints on pure WDM from [46], to further highlight the Lyman- α forest as a unique and sensitive probe of matter clustering on subgalactic scales. Extending this analysis to CWDM cosmologies, we find that the updated constraints have broader implications for any dark matter model that suppresses the amount of clustering on small scales. Our findings suggest that the scale of suppression can be moved to larger scales, but only if the level of the suppression on smaller scales is lower. A logical next step is to construct a general framework that quantifies the allowed suppression amplitude and its scale dependence, enabling comparison across a wide class of dark matter models exhibiting small-scale power suppression [62,123,124]. A parametrization of this kind was introduced by [125] for a large set of both thermal and nonthermal relics. We leave the development of such a generalized framework for future work.

The main interest in constraining this broad class of dark matter models lies in their potential to resolve the S_8 tension when the Lyman- α forest is combined with other large-scale structure probes (e.g., [126]). For instance, the

tension has been reported to be alleviated when incorporating CMB data in the presence of axion DM [34] or with weak lensing data [127]. A detailed investigation of the implications for S_8 within the particular CWDM models allowed in this work is also left for future study.

ACKNOWLEDGMENTS

The authors would like to thank David Chemaly, Tom Hehir, Harry Bevins, and Will Handley for helpful conversations. V. I. acknowledges visitor support from the Kavli Institute for Cosmology, Cambridge, where part of this work was completed. M. V. is supported by the INFN PD51 INDARK grant. Support by ERC Advanced Grant 320596, “The Emergence of Structure During the Epoch of Reionization,” is gratefully acknowledged. M. G. H. has been supported by STFC consolidated grants ST/N000927/1 and ST/S000623/1. J. S. B. is supported by STFC consolidated grant ST/X000982/1. The simulations used in this work were performed using the Joliot Curie supercomputer at the Trés Grand Centre de Calcul (TGCC) and the Cambridge Service for Data Driven Discovery (CSD3), part of which is operated by the University of Cambridge Research Computing on behalf of the STFC DiRAC HPC Facility [128]. We acknowledge the Partnership for Advanced Computing in Europe (PRACE) for awarding us time on Joliot Curie in the 16th call. The DiRAC component of CSD3 was funded by BEIS capital funding via STFC capital grants ST/P002307/1 and ST/R002452/1 and STFC operations grant ST/R00689X/1. This work also used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility. The equipment was funded by BEIS capital funding via STFC capital grants ST/P002293/1 and ST/R002371/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

[1] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), *Astron. Astrophys.* **652**, C4 (2021).

[2] A. Klypin, A. V. Kravtsov, O. Valenzuela, and F. Prada, *Astrophys. J.* **522**, 82 (1999).

[3] B. Moore, S. Ghigna, F. Governato, G. Lake, T. Quinn, J. Stadel, and P. Tozzi, *Astrophys. J. Lett.* **524**, L19 (1999).

[4] J. D. Simon, *Annu. Rev. Astron. Astrophys.* **57**, 375 (2019).

[5] C. Engler, A. Pillepich, A. Pasquali, D. Nelson, V. Rodriguez-Gomez, K. T. E. Chua, E. K. Grebel, V. Springel, F. Marinacci, R. Weinberger *et al.*, *Mon. Not. R. Astron. Soc.* **507**, 4211 (2021).

[6] A. S. Font, I. G. McCarthy, and V. Belokurov, *Mon. Not. R. Astron. Soc.* **505**, 783 (2021).

[7] W. J. G. de Blok, *Adv. Astron.* **2010**, 789293 (2010).

[8] I. M. E. Santos-Santos, J. F. Navarro, A. Robertson, A. Benítez-Llambay, K. A. Oman, M. R. Lovell, C. S. Frenk, A. D. Ludlow, A. Fattahi, and A. Ritz, *Mon. Not. R. Astron. Soc.* **495**, 58 (2020).

[9] M. Boylan-Kolchin, J. S. Bullock, and M. Kaplinghat, *Mon. Not. R. Astron. Soc.* **415**, L40 (2011).

[10] M. Boylan-Kolchin, J. S. Bullock, and M. Kaplinghat, *Mon. Not. R. Astron. Soc.* **422**, 1203 (2012).

[11] A. R. Wetzel, P. F. Hopkins, J. Kim, C.-A. Faucher-Giguère, D. Kereš, and E. Quataert, *Astrophys. J. Lett.* **827**, L23 (2016).

[12] J. Samuel, A. Wetzel, E. Tollerud, S. Garrison-Kimmel, S. Loebman, K. El-Badry, P. F. Hopkins, M. Boylan-Kolchin, C.-A. Faucher-Giguère, J. S. Bullock *et al.*, *Mon. Not. R. Astron. Soc.* **491**, 1471 (2020).

[13] L. V. Sales, A. Wetzel, and A. Fattahi, *Nat. Astron.* **6**, 897 (2022).

[14] M. White, R. Zhou, J. DeRose, S. Ferraro, S.-F. Chen, N. Kokron, S. Bailey, D. Brooks, J. García-Bellido, J. Guy *et al.*, *J. Cosmol. Astropart. Phys.* **02** (2022) 007.

[15] S. Ferraro, N. Sailer, A. Slosar, and M. White, [arXiv:2203.07506](https://arxiv.org/abs/2203.07506).

[16] A. Amon, N. C. Robertson, H. Miyatake, C. Heymans, M. White, J. DeRose, S. Yuan, R. H. Wechsler, T. N. Varga, S. Bocquet *et al.*, *Mon. Not. R. Astron. Soc.* **518**, 477 (2023).

[17] M. M. Ivanov, O. H. E. Philcox, G. Cabass, T. Nishimichi, M. Simonović, and M. Zaldarriaga, *Phys. Rev. D* **107**, 083515 (2023).

[18] J. U. Lange, A. P. Hearn, A. Leauthaud, F. C. van den Bosch, E. Xhakaj, H. Guo, R. H. Wechsler, and J. DeRose, *Mon. Not. R. Astron. Soc.* **520**, 5373 (2023).

[19] G. S. Farren, A. Krolewski, N. MacCrann, S. Ferraro, I. Abril-Cabezas, R. An, Z. Atkins, N. Battaglia, J. R. Bond, E. Calabrese *et al.*, *Astrophys. J.* **966**, 157 (2024).

[20] S.-F. Chen, M. M. Ivanov, O. H. E. Philcox, and L. Wenzl, *Phys. Rev. Lett.* **133**, 231001 (2024).

[21] S. Mashchenko, H. M. P. Couchman, and J. Wadsley, *Nature (London)* **442**, 539 (2006).

[22] A. M. Brooks, M. Kuhlen, A. Zolotov, and D. Hooper, *Astrophys. J.* **765**, 22 (2013).

[23] A. Fattahi, J. F. Navarro, T. Sawala, C. S. Frenk, L. V. Sales, K. Oman, M. Schaller, and J. Wang, [arXiv:1607.06479](https://arxiv.org/abs/1607.06479).

[24] A. Pontzen and F. Governato, *Nature (London)* **506**, 171 (2014).

[25] M. P. van Daalen, J. Schaye, C. M. Booth, and C. Dalla Vecchia, *Mon. Not. R. Astron. Soc.* **415**, 3649 (2011).

[26] V. Springel, R. Pakmor, A. Pillepich, R. Weinberger, D. Nelson, L. Hernquist, M. Vogelsberger, S. Genel, P. Torrey, F. Marinacci *et al.*, *Mon. Not. R. Astron. Soc.* **475**, 676 (2018).

[27] N. E. Chisari, A. J. Mead, S. Joudaki, P. G. Ferreira, A. Schneider, J. Mohr, T. Tröster, D. Alonso, I. G. McCarthy, S. Martin-Alvarez *et al.*, *Open J. Astrophys.* **2**, 4 (2019).

[28] K.-Y. Su and P. Chen, *J. Cosmol. Astropart. Phys.* **08** (2011) 016.

[29] L. Wang, V. Gonzalez-Perez, L. Xie, A. P. Cooper, C. S. Frenk, L. Gao, W. A. Hellwing, J. Helly, M. R. Lovell, and L. Jiang, *Mon. Not. R. Astron. Soc.* **468**, 4579 (2017).

[30] M. R. Lovell, V. Gonzalez-Perez, S. Bose, A. Boyarsky, S. Cole, C. S. Frenk, and O. Ruchayskiy, *Mon. Not. R. Astron. Soc.* **468**, 2836 (2017).

[31] J. Liu, S. Bird, J. M. Zorrilla Matilla, J. C. Hill, Z. Haiman, M. S. Madhavacheril, A. Petri, and D. N. Spergel, *J. Cosmol. Astropart. Phys.* **03** (2018) 049.

[32] E. O. Nadler, D. Yang, and H.-B. Yu, *Astrophys. J. Lett.* **958**, L39 (2023).

[33] A. Laguë, B. Schwabe, R. Hložek, D. J. E. Marsh, and K. K. Rogers, *Phys. Rev. D* **109**, 043507 (2024).

[34] K. K. Rogers and V. Poulin, *Phys. Rev. Res.* **7**, L012018 (2025).

[35] M. P. van Daalen, J. Schaye, C. M. Booth, and C. Dalla Vecchia, *Mon. Not. R. Astron. Soc.* **415**, 3649 (2011).

[36] T. Naab and J. P. Ostriker, *Annu. Rev. Astron. Astrophys.* **55**, 59 (2017).

[37] M. Vogelsberger, F. Marinacci, P. Torrey, and E. Puchwein, *Nat. Rev. Phys.* **2**, 42 (2020).

[38] I. Medlock, C. Neufeld, D. Nagai, D. Anglés-Alcázar, S. Genel, B. D. Oppenheimer, X. Sims, P. Singh, and F. Villaescusa-Navarro, *Astrophys. J.* **980**, 61 (2025).

[39] L. Bigwood, M. A. Bourne, V. Irsic, A. Amon, and D. Sijacki, [arXiv:2501.16983](https://arxiv.org/abs/2501.16983).

[40] D. Lyapustin, [arXiv:1112.1167](https://arxiv.org/abs/1112.1167).

[41] A. Boyarsky, M. Drewes, T. Lasserre, S. Mertens, and O. Ruchayskiy, *Prog. Part. Nucl. Phys.* **104**, 1 (2019).

[42] S. Dodelson and L. M. Widrow, *Phys. Rev. Lett.* **72**, 17 (1994).

[43] G. R. Blumenthal, H. Pagels, and J. R. Primack, *Nature (London)* **299**, 37 (1982).

[44] P. J. E. Peebles, *Astrophys. J. Lett.* **263**, L1 (1982).

[45] J. R. Bond, A. S. Szalay, and M. S. Turner, *Phys. Rev. Lett.* **48**, 1636 (1982).

[46] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, M. Molaro, E. Puchwein, E. Boera, G. D. Becker, P. Gaikwad, L. C. Keating *et al.*, *Phys. Rev. D* **109**, 043511 (2024).

[47] M. Davis, F. J. Summers, and D. Schlegel, *Nature (London)* **359**, 393 (1992).

[48] A. N. Taylor and M. Rowan-Robinson, *Nature (London)* **359**, 396 (1992).

[49] A. van Dalen and R. K. Schaefer, *Astrophys. J.* **398**, 33 (1992).

[50] A. Klypin, J. Holtzman, J. Primack, and E. Regos, *Astrophys. J.* **416**, 1 (1993).

[51] X. Shi and G. M. Fuller, *Phys. Rev. Lett.* **82**, 2832 (1999).

[52] A. Boyarsky, J. Lesgourgues, O. Ruchayskiy, and M. Viel, *J. Cosmol. Astropart. Phys.* **05** (2009) 012.

[53] M. Shaposhnikov and I. Tkachev, *Phys. Lett. B* **639**, 414 (2006).

[54] R. A. C. Croft, D. H. Weinberg, M. Bolte, S. Burles, L. Hernquist, N. Katz, D. Kirkman, and D. Tytler, *Astrophys. J.* **581**, 20 (2002).

[55] P. McDonald, U. Seljak, R. Cen, D. Shih, D. H. Weinberg, S. Burles, D. P. Schneider, D. J. Schlegel, N. A. Bahcall, J. W. Briggs *et al.*, *Astrophys. J.* **635**, 761 (2005).

[56] A. Lidz, C.-A. Faucher-Giguère, A. Dall'Aglio, M. McQuinn, C. Fechner, M. Zaldarriaga, L. Hernquist, and S. Dutta, *Astrophys. J.* **718**, 199 (2010).

[57] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, *Phys. Rev. D* **88**, 043502 (2013).

[58] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, S. Cristiani, G. D. Becker, V. D'Odorico, G. Cupani, T.-S. Kim, T. A. M. Berg *et al.*, [arXiv:1702.01764](https://arxiv.org/abs/1702.01764).

[59] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, and G. D. Becker, *Phys. Rev. Lett.* **119**, 031302 (2017).

[60] N. Palanque-Delabrouille, C. Yèche, N. Schöneberg, J. Lesgourgues, M. Walther, S. Chabanier, and E. Armengaud, *J. Cosmol. Astropart. Phys.* **04** (2020) 038.

[61] A. Garzilli, A. Magalich, O. Ruchayskiy, and A. Boyarsky, *Mon. Not. R. Astron. Soc.* **502**, 2356 (2021).

[62] K. K. Rogers and H. V. Peiris, *Phys. Rev. Lett.* **126**, 071302 (2021).

[63] A. Slosar, A. Font-Ribera, M. M. Pieri, J. Rich, J.-M. Le Goff, É. Aubourg, J. Brinkmann, N. Busca, B. Carithers, R. Charlassier *et al.*, *J. Cosmol. Astropart. Phys.* **09** (2011) 001.

[64] A. Slosar, V. Iršič, D. Kirkby, S. Bailey, N. G. Busca, T. Delubac, J. Rich, É. Aubourg, J. E. Bautista, V. Bhardwaj *et al.*, *J. Cosmol. Astropart. Phys.* **04** (2013) 026.

[65] N. G. Busca, T. Delubac, J. Rich, S. Bailey, A. Font-Ribera, D. Kirkby, J. M. Le Goff, M. M. Pieri, A. Slosar, É. Aubourg *et al.*, *Astron. Astrophys.* **552**, A96 (2013).

[66] H. du Mas des Bourboux, J. Rich, A. Font-Ribera, V. de Sainte Agathe, J. Farr, T. Etourneau, J.-M. Le Goff, A. Cuceu, C. Balland, J. E. Bautista *et al.*, *Astrophys. J.* **901**, 153 (2020).

[67] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.*, *J. Cosmol. Astropart. Phys.* **01** (2025) 124.

[68] M. Abdul-Karim, J. Aguilar, S. Ahlen, C. Allende Prieto, O. Alves, A. Anand, U. Andrade, E. Armengaud, A. Aviles *et al.* (DESI Collaboration), [arXiv:2503.14739](https://arxiv.org/abs/2503.14739).

[69] E. Boera, G. D. Becker, J. S. Bolton, and F. Nasir, *Astrophys. J.* **872**, 101 (2019).

[70] B. Villasenor, B. Robertson, P. Madau, and E. Schneider, *Phys. Rev. D* **108**, 023502 (2023).

[71] M. Viel, J. S. Bolton, and M. G. Haehnelt, *Mon. Not. R. Astron. Soc.* **399**, L39 (2009).

[72] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, *Phys. Rev. D* **88**, 043502 (2013).

[73] C. Yèche, N. Palanque-Delabrouille, J. Baur, and H. du Mas des Bourboux, *J. Cosmol. Astropart. Phys.* **06** (2017) 047.

[74] M. Molaro, V. Iršič, J. S. Bolton, L. C. Keating, E. Puchwein, P. Gaikwad, M. G. Haehnelt, G. Kulkarni, and M. Viel, *Mon. Not. R. Astron. Soc.* **509**, 6119 (2022).

[75] S. Bird, K. K. Rogers, H. V. Peiris, L. Verde, A. Font-Ribera, and A. Pontzen, *J. Cosmol. Astropart. Phys.* **02** (2019) 050.

[76] M. Walther, J. Oñorbe, J. F. Hennawi, and Z. Lukić, *Astrophys. J.* **872**, 13 (2019).

[77] K. K. Rogers, H. V. Peiris, A. Pontzen, S. Bird, L. Verde, and A. Font-Ribera, *J. Cosmol. Astropart. Phys.* **02** (2019) 031.

[78] M. Walther, E. Armengaud, C. Ravoux, N. Palanque-Delabrouille, C. Yèche, and Z. Lukić, *J. Cosmol. Astropart. Phys.* **04** (2021) 059.

[79] M. A. Fernandez, M.-F. Ho, and S. Bird, *Mon. Not. R. Astron. Soc.* **517**, 3200 (2022).

[80] S. Bird, M. Fernandez, M.-F. Ho, M. Qezlou, R. Monadi, Y. Ni, N. Chen, R. Croft, and T. Di Matteo, *J. Cosmol. Astropart. Phys.* **10** (2023) 037.

[81] L. Cabayol-Garcia, J. Chaves-Montero, A. Font-Ribera, and C. Pedersen, *Mon. Not. R. Astron. Soc.* **525**, 3499 (2023).

[82] M. Molaro, V. Iršič, J. S. Bolton, M. Lieu, L. C. Keating, E. Puchwein, M. G. Haehnelt, and M. Viel, *Mon. Not. R. Astron. Soc.* **521**, 1489 (2023).

[83] J. Baur, N. Palanque-Delabrouille, C. Yèche, A. Boyarsky, O. Ruchayskiy, É. Armengaud, and J. Lesgourgues, *J. Cosmol. Astropart. Phys.* **12** (2017) 013.

[84] A. Palazzo, D. Cumberbatch, A. Slosar, and J. Silk, *Phys. Rev. D* **76**, 103511 (2007).

[85] R. Diamanti, S. Ando, S. Gariazzo, O. Mena, and C. Weniger, *J. Cosmol. Astropart. Phys.* **06** (2017) 008.

[86] G. Parimbelli, G. Scelfo, S. K. Giri, A. Schneider, M. Archidiacono, S. Camera, and M. Viel, *J. Cosmol. Astropart. Phys.* **12** (2021) 044.

[87] The 1D Lyman- α clustering is sensitive only to line-of-sight Fourier modes k_{\parallel} . For the sake of simplicity we adopt a shorthand notation of k in velocity units to denote wave numbers along the line-of-sight.

[88] N. Palanque-Delabrouille, C. Yèche, J. Baur, C. Magneville, G. Rossi, J. Lesgourgues, A. Bilde, E. Burtin, J.-M. LeGoff, J. Rich *et al.*, *J. Cosmol. Astropart. Phys.* **11** (2015) 011.

[89] V. Iršič, M. Viel, T. A. M. Berg, V. D'Odorico, M. G. Haehnelt, S. Cristiani, G. Cupani, T.-S. Kim, S. López, S. Ellison *et al.*, *Mon. Not. R. Astron. Soc.* **466**, 4332 (2017).

[90] M. Viel, J. Schaye, and C. M. Booth, *Mon. Not. R. Astron. Soc.* **429**, 1734 (2013).

[91] A. Day, D. Tytler, and B. Kambalur, *Mon. Not. R. Astron. Soc.* **489**, 2536 (2019).

[92] N. G. Karaçaylı, N. Padmanabhan, A. Font-Ribera, V. Iršič, M. Walther, D. Brooks, E. Gaztañaga, R. Kehoe, M. Levi, P. Ntelis *et al.*, *Mon. Not. R. Astron. Soc.* **509**, 2842 (2022).

[93] E. Puchwein, F. Haardt, M. G. Haehnelt, and P. Madau, *Mon. Not. R. Astron. Soc.* **485**, 47 (2019).

[94] E. Puchwein, J. S. Bolton, L. C. Keating, M. Molaro, P. Gaikwad, G. Kulkarni, M. G. Haehnelt, V. Iršič, T. Šoltinský, M. Viel *et al.*, *Mon. Not. R. Astron. Soc.* **519**, 6162 (2023).

[95] A. Lewis and A. Challinor, *Astrophys. J.* **538**, 473 (2000).

[96] D. Blas, J. Lesgourgues, and T. Tram, *J. Cosmol. Astropart. Phys.* **07** (2011) 034.

[97] M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese, and A. Riotto, *Phys. Rev. D* **71**, 063534 (2005).

[98] P. Bode, J. P. Ostriker, and N. Turok, *Astrophys. J.* **556**, 93 (2001).

[99] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, *Phys. Rev. D* **88**, 043502 (2013).

[100] J. Lesgourgues and T. Tram, *J. Cosmol. Astropart. Phys.* **09** (2011) 032.

[101] G. D. Becker, P. C. Hewett, G. Worseck, and J. X. Prochaska, *Mon. Not. R. Astron. Soc.* **430**, 2067 (2013).

[102] S. E. I. Bosman, X. Fan, L. Jiang, S. Reed, Y. Matsuoka, G. Becker, and M. Haehnelt, *Mon. Not. R. Astron. Soc.* **479**, 1055 (2018).

[103] S. E. I. Bosman, F. B. Davies, G. D. Becker, L. C. Keating, R. L. Davies, Y. Zhu, A.-C. Eilers, V. D’Odorico, F. Bian, M. Bischetti *et al.*, *Mon. Not. R. Astron. Soc.* **514**, 55 (2022).

[104] A. Borde, N. Palanque-Delabrouille, G. Rossi, M. Viel, J. S. Bolton, C. Yèche, J.-M. LeGoff, and J. Rich, *J. Cosmol. Astropart. Phys.* **07** (2014) 005.

[105] B. Villasenor, B. Robertson, P. Madau, and E. Schneider, *Astrophys. J.* **912**, 138 (2021).

[106] C. C. Doughty, J. F. Hennawi, F. B. Davies, Z. Lukić, and J. Oñorbe, *Mon. Not. R. Astron. Soc.* **525**, 3790 (2023).

[107] P. McDonald, *Astrophys. J.* **585**, 34 (2003).

[108] J. S. Bolton, E. Puchwein, D. Sijacki, M. G. Haehnelt, T.-S. Kim, A. Meiksin, J. A. Regan, and M. Viel, *Mon. Not. R. Astron. Soc.* **464**, 897 (2017).

[109] N. K. Khan, G. Kulkarni, J. S. Bolton, M. G. Haehnelt, V. Iršič, E. Puchwein, and S. Asthana, *Mon. Not. R. Astron. Soc.* **530**, 4920 (2024).

[110] J. Torrado and A. Lewis, Cobaya: Bayesian analysis in cosmology (2019), <https://ui.adsabs.harvard.edu/abs/2019ascl.soft10019T>.

[111] B. Villasenor, B. Robertson, P. Madau, and E. Schneider, *Phys. Rev. D* **108**, 023502 (2023).

[112] P. Gaikwad, M. Rauch, M. G. Haehnelt, E. Puchwein, J. S. Bolton, L. C. Keating, G. Kulkarni, V. Iršič, E. Bañados, G. D. Becker *et al.*, *Mon. Not. R. Astron. Soc.* **494**, 5091 (2020).

[113] P. Gaikwad, R. Srianand, M. G. Haehnelt, and T. R. Choudhury, *Mon. Not. R. Astron. Soc.* **506**, 4389 (2021).

[114] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020).

[115] E. Puchwein, F. Haardt, M. G. Haehnelt, and P. Madau, *Mon. Not. R. Astron. Soc.* **485**, 47 (2019).

[116] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, S. Cristiani, G. D. Becker, V. D’Odorico, G. Cupani, T.-S. Kim, T. A. M. Berg *et al.*, *Phys. Rev. D* **96**, 023522 (2017).

[117] M. Molaro, V. Iršič, J. S. Bolton, L. C. Keating, E. Puchwein, P. Gaikwad, M. G. Haehnelt, G. Kulkarni, and M. Viel, *Mon. Not. R. Astron. Soc.* **509**, 6119 (2022).

[118] S. Zelmer, E. Artis, E. Bulbul, S. Grandis, V. Ghirardini, A. von der Linden, Y. E. Bahar, F. Balzer, M. Brüggen, I. Chiu *et al.*, [arXiv:2502.03353](https://arxiv.org/abs/2502.03353).

[119] E. Calabrese, J. C. Hill, H. T. Jense, A. La Posta, I. Abril-Cabezas, G. E. Addison, P. A. R. Ade, S. Aiola, T. Alford, D. Alonso *et al.*, [arXiv:2503.14454](https://arxiv.org/abs/2503.14454).

[120] T. Kobayashi, R. Murgia, A. De Simone, V. Iršič, and M. Viel, *Phys. Rev. D* **96**, 123514 (2017).

[121] G. Facchinetto, [arXiv:2503.11261](https://arxiv.org/abs/2503.11261).

[122] A.-C. Eilers, F. B. Davies, and J. F. Hennawi, *Astrophys. J.* **864**, 53 (2018).

[123] M. Archidiacono, D. C. Hooper, R. Murgia, S. Bohr, J. Lesgourgues, and M. Viel, *J. Cosmol. Astropart. Phys.* **10** (2019) 055.

[124] D. C. Hooper, N. Schöneberg, R. Murgia, M. Archidiacono, J. Lesgourgues, and M. Viel, *J. Cosmol. Astropart. Phys.* **10** (2022) 032.

[125] R. Murgia, V. Iršič, and M. Viel, *Phys. Rev. D* **98**, 083540 (2018).

[126] M. Esposito, V. Iršič, M. Costanzi, S. Borgani, A. Saro, and M. Viel, *Mon. Not. R. Astron. Soc.* **515**, 857 (2022).

[127] F. Hervas Peters, A. Schneider, J. Bucko, S. K. Giri, and G. Parimbelli, *Astron. Astrophys.* **687**, A161 (2024).

[128] <https://www.dirac.ac.uk>