

Unveiling dark matter free streaming at the smallest scales with the high redshift Lyman-alpha forest

Vid Iršič^{1,2,*}, Matteo Viel^{3,4,5,6,7}, Martin G. Haehnelt^{1,8}, James S. Bolton⁹, Margherita Molaro⁹, Ewald Puchwein¹⁰, Elisa Boera^{5,6}, George D. Becker¹¹, Prakash Gaikwad¹², Laura C. Keating¹³, and Girish Kulkarni¹⁴

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

²Cavendish Laboratory, University of Cambridge, 19 J. J. Thomson Avenue, Cambridge CB3 0HE, United Kingdom

³SISSA—International School for Advanced Studies, Via Bonomea 265, 34136 Trieste, Italy

⁴INFN—National Institute for Nuclear Physics, Via Valerio 2, I-34127 Trieste, Italy

⁵IFPU, Institute for Fundamental Physics of the Universe, via Beirut 2, 34151 Trieste, Italy

⁶INAF, Osservatorio Astronomico di Trieste, Via G. B. Tiepolo 11, I-34131 Trieste, Italy

⁷ICSC—Centro Nazionale di Ricerca in High Performance Computing, Big Data e Quantum Computing, Via Magnanelli 2, Bologna, Italy

⁸Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, United Kingdom

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

¹⁰Leibniz-Institut für Astrophysik Potsdam, An der Sternwarte 16, 14482 Potsdam, Germany

¹¹Department of Physics and Astronomy, University of California, Riverside, California 92521, USA

¹²Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

¹³Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, United Kingdom

¹⁴Tata Institute of Fundamental Research, Homi Bhabha Road, Mumbai 400005, India



(Received 15 September 2023; accepted 9 January 2024; published 8 February 2024)

This study introduces novel constraints on the free streaming of thermal relic warm dark matter (WDM) from Lyman- α forest flux power spectra. Our analysis utilizes a high resolution, high redshift sample of quasar spectra observed using the HIRES and UVES spectrographs ($z = 4.2\text{--}5.0$). We employ a Bayesian inference framework and a simulation-based likelihood that encompasses various parameters including the free streaming of dark matter, cosmological parameters, the thermal history of the intergalactic medium, and inhomogeneous reionization to establish lower limits on the mass of a thermal relic WDM particle of 5.7 keV (at 95% CL). This result surpasses previous limits from the Lyman- α forest through reduction of the measured uncertainties due to a larger statistical sample and by measuring clustering to smaller scales ($k_{\text{max}} = 0.2 \text{ km}^{-1} \text{ s}$). The approximately two-fold improvement due to the expanded statistical sample suggests that the effectiveness of Lyman- α forest constraints on WDM models at high redshifts are limited by the availability of high quality quasar spectra. Restricting the analysis to comparable scales and thermal history priors as in prior studies ($k_{\text{max}} < 0.1 \text{ km}^{-1} \text{ s}$) lowers the bound on the WDM mass to 4.1 keV. As the precision of the measurements increases, it becomes crucial to examine the instrumental and modeling systematics. On the modeling front, we argue that the impact of the thermal history uncertainty on the WDM particle mass constraint has diminished due to improved independent observations. At the smallest scales, the primary source of modeling systematic arises from the structure in the peculiar velocity of the intergalactic medium and inhomogeneous reionization.

DOI: [10.1103/PhysRevD.109.043511](https://doi.org/10.1103/PhysRevD.109.043511)

*vi223@cam.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

I. INTRODUCTION

The Lyman- α forest is the main manifestation of the high redshift intergalactic cosmic web. It is visible in the spectra of quasars (QSOs) and produced by the scattering of the background photons with the neutral hydrogen atoms along the line of sight [1,2]. The Lyman- α forest is a unique probe of geometry and the dynamical state of the Universe,

probing diffuse matter around galaxies and in the intergalactic medium (IGM) in regimes which are not covered by other observables, both in terms of redshifts and scales.

In the last decade, we have witnessed tremendous progress in the cosmological investigation of the Lyman- α forest, mainly along two different directions which are connected to fundamental physics. For example, the discovery of baryonic acoustic oscillations in the 3D correlation function of the transmitted flux has offered the possibility to constrain new physics beyond the standard cosmological model, in the context of allowing curvature or an evolution of the equation of state for dark energy [3]. Another important research line, following the work of [4,5], has focussed on the 1D flux power spectrum used to probe the growth of structure down to the smallest scales to see to which extent dark matter free streaming could be constrained.

In this work, we investigate this second aspect and present new results based on a new set of simulations which incorporate the most important physical ingredients [6,7], and a new comprehensive analysis of high resolution high redshift data down to the smallest scales. A key goal is to disentangle the different roles of the physical processes able to affect the 1D flux power: the thermal broadening, which is a 1D effect acting along the line of sight and is sensitive to the instantaneous gas temperature, and two 3D effects, the gas pressure smoothing that depends on the whole thermal history of the IGM and the dark matter (DM) free streaming.

The possibility of constraining the nature of DM by using the Lyman- α forest has motivated a series of works which were able to constrain the models further, explore different particle physics dark matter candidates, and combine likelihoods with other experiments able to constrain the nature of dark matter with strong lensing or flux ratio anomalies [8,9]. One of the main reasons to explore warm dark matter (WDM) models was to solve or ease putative problems of cold dark matter at small scales [10,11]. However most of these tensions must be discussed also in the context of baryonic physics [12], with processes like galactic feedback playing a major role. Moreover, it appears that minimal extensions of the standard model of particle physics could also accommodate particles like sterile neutrinos or a scalar field [13–16], which could suppress or erase power at small scales, effectively acting as WDM.

For thermal WDM masses in the keV range, the power suppression happens at the small nonlinear scales sampled by the Lyman- α forest. In particular, QSO data sets with different resolution and signal-to-noise properties have been used in order to tighten the constraints. The low resolution SDSS and BOSS data sets [5,17,18], the medium resolution X-Shooter sample [19], and the high resolution and high signal-to-noise Keck/HIRES and UVES/VLT QSO spectra [20–24] have all played a major role in the

advancement of the field. For example, while the low and medium resolution data are not particularly effective in sampling the scales of the cutoff fully, they nevertheless are sensitive to the thermal history and can return very tight constraints especially when combined with data at smaller scales. The goal of this paper is to give a comprehensive state-of-the-art analysis focusing on high resolution data [25]. In Sec. II, we describe the data set, while in Sec. III we present the suite of hydrodynamical simulations used. Section IV contains our new results which are extensively discussed in terms of the thermal history of the IGM, the dependence on mass resolution, patchy reionization, instrumental effects (including modeling of the noise), and consistency with results in the literature. We conclude in Sec. V.

II. DATA

We apply our analysis to the measurements presented in [25]. Their 1D flux power spectrum is estimated using 15 high signal-to-noise spectra observed by VLT/UVES [26] and Keck/HIRES [27]. The measurements span the high redshift range of $z = 4.2\text{--}5.0$ in bins of $\Delta z = 0.4$. In each of the redshift bins, the flux power spectrum is measured in 15 k -bins equidistantly spaced in $\log_{10} k$ in the range of $\log_{10}(k/[\text{km}^{-1} \text{ s}]) = -2.2$ to $\log_{10}(k/[\text{km}^{-1} \text{ s}]) = -0.7$, with logarithmic spacing of $\Delta \log_{10}(k/[\text{km}^{-1} \text{ s}]) = 0.1$. Unless specified otherwise, we use the full extent of the data, resulting in 45 data points across three redshift bins.

The spectrograph resolution in these observations is very high, with $R \sim 50,000$ (FWHM of $\sim 6 \text{ km s}^{-1}$ for HIRES and $\sim 7 \text{ km s}^{-1}$ for UVES). As already pointed out in the study of [25], the effects of resolution uncertainty are very small, even for the highest wave number power spectrum bin measured. A conservative estimate of the 10% uncertainty on the resolution leads only to 1% (5%) uncertainty on the 1D flux power spectrum at scales of $k = 0.1(0.2) \text{ km}^{-1} \text{ s}$. The power spectrum measurements of [25] were reported both with and without instrumental resolution correction. In this analysis, we use the measurements with instrumental resolution corrected and propagate this correction through the covariance matrix. The reported measurements are also corrected for power spectrum due to metal contaminants.

The typical flux noise estimated in these measurements is white noise, with its power spectrum amplitude of $0.1\text{--}0.2 \text{ km s}^{-1}$. This is comparable to the estimated level of the models at the highest wave numbers. Characterizing and accounting for the noise levels is of key importance and has been one of the factors restricting previous analyses to smaller wave numbers.

The k range of [25] covers the smallest scales measured with the 1D flux power spectrum, extending to $k \sim 0.2 \text{ km}^{-1} \text{ s}$, a factor of two higher wave number than in previous studies [21–23]. These studies have shown that the constraining power on WDM models from the Ly α forest

is dominated by high redshifts and the smallest scales, making this an ideal data set to exploit.

III. SIMULATIONS

The absorption features of the Ly α forest contain a wealth of information regarding cosmology and the nature of dark matter, as well as the thermal state of the intergalactic gas. Due to the high sensitivity of the spectrographic instruments, it provides a unique window into clustering at the smallest scales. Accessing that information, however, is a nontrivial task. The standard approaches of clustering analysis that invoke biasing schemes typically rely on perturbation theory [28] or build an approximate clustering scheme [29]. While very informative in a qualitative sense, these schemes cannot capture the complexity of the data that are highly sensitive to nonlinear structure evolution and gas physics, such as Doppler broadening and thermal pressure smoothing [30].

Such a task requires simulating the expected Ly α forest in different thermal and cosmological models, spanning a wide, multidimensional parameter space, and comparing it to the data. In this work, we carry out the comparison within the framework of Bayesian inference analysis, which describes—according to Bayes’ theorem—the posterior probability $p(\theta|D)$ having parameters θ given observed data D as

$$p(\theta|D) \propto \mathcal{L}(D|\theta) \times \pi(\theta), \quad (1)$$

where $\mathcal{L}(D|\theta)$ is the likelihood, and $\pi(\theta)$ is the prior on each parameter.

In this work, we expand upon the Bayesian inference setup adopted in [7,31] to evaluate the likelihood and prior at each parameter combination in the sampler. The likelihood is evaluated jointly at all the observed data points. This is based on the Monte Carlo Markov Chain (MCMC) sampler, combined with the Metropolis-Hastings algorithm by dynamically learning the proposal matrix from the covariance that was introduced in [22]. The precision of the thermal parameter recovery with this simulation based emulator was shown to be in good agreement with more advanced machine-learning augmented emulator models [31].

The priors $\pi(\theta)$ we adopt in our analysis are described in Sec. IV. The likelihood is modeled as a Gaussian likelihood, determined by the data and its covariance, and a theoretical prediction for the flux power spectrum. The latter is estimated using hydrodynamical numerical simulations.

We use simulations from the Sherwood-Relics project [6]. These are a series of high resolution cosmological hydrodynamical simulations that use a customized version of P-Gadget3 (see [32] for the original Gadget-2 reference). We use cosmological boxes of size $20 h^{-1}$ Mpc with 2×1024^3 dark matter and gas particles. The box size and resolution have been chosen to adequately resolve the small-scale structure that contributes to the flux power

spectrum of the Ly α forest, while still retaining a cosmologically relevant volume [33–36]. We further correct the numerical convergence with both box size and resolution with a series of additional simulations summarized in Table I. In all models, we use a simple, computationally efficient star-formation scheme—often called `Quick_lya`—where gas particles are converted into collisionless star particles if they reach overdensities $\Delta = 1 + \delta > 10^3$ and temperatures $T < 10^5$ K [37]. We assume a flat Λ CDM cosmology with $\Omega_\Lambda = 0.692$, $\Omega_m = 0.308$, $\Omega_b = 0.0482$, $\sigma_8 = 0.829$, $n_s = 0.961$, $h = 0.678$, and a primordial helium mass abundance of $Y_p = 0.24$ [38]. The initial conditions for the CDM simulations are identical to those used in the earlier Sherwood simulation project [35]. We use the WDM transfer function approximation of [21].

A set of simulations is constructed using modifications to the spatially uniform UV background synthesis model introduced by [39]. These simulations are similar to models used in earlier works [16,22,40,41], with the main improvements being the larger dynamic range of the simulations, the use of a nonequilibrium thermochemistry solver [42], and improved treatment of the IGM opacity that consistently captures the transition between neutral and ionized IGM.

In addition to running a model with the fiducial UV background, we also vary the photoheating rates to achieve models with different gas temperatures and ends of reionization, following the approach described in [6,43]. This approach results in 12 models with varying thermal histories (see Table I). For *each* of the thermal history models with fiducial Λ CDM cosmology, we also run models varying the WDM particle mass ($m_{\text{WDM}} = [2, 3, 4]$ keV), amplitude of Λ CDM matter clustering ($\sigma_8 = [0.754, 0.804, 0.854, 0.904]$), and spectral index of inflation ($n_s = [0.921, 0.941, 0.981, 1.001]$). This results in a total of $12 \times (3 + 2 \times 4 + 1) = 144$ simulations.

In order to construct a sufficiently well sampled grid of models spanning the entire multidimensional parameter range, we postprocess the 144 simulations (12 simulations for each cosmology) to obtain different parameter combinations. We follow the method of [25,44] in order to interpolate in the temperature-density plane. Briefly, we rotate and translate the line-of-sight particles in the temperature-density plane to obtain models with different temperature at mean density T_0 and temperature-density power-law indices γ [the values of T_0 and γ are inferred from the line-of-sight gas properties; a power-law relation is fitted to points in the temperature density plane in the range of gas overdensity ($0.1 < \Delta_g < 1.0$) and neutral fraction weighted gas temperature ($T < 10^5$ K)]. This preserves the temperature-density cross-correlation coefficient, allowing for an inexpensive construction of models with different thermal parameters on a finely spaced grid.

In postprocessing, we also vary the redshift evolution of the mean transmission $\langle F \rangle$, by rescaling the optical depth of

TABLE I. List of simulations used in this work (see also [6]). From left to right, the columns list the simulation name, the box size in h^{-1} cMpc, the number of particles, the redshift of reionization (defined as the redshift when the volume averaged ionized fraction $1 - x_{\text{HI}} \leq 10^{-3}$), the gas temperature at the mean density, T_0 , the cumulative energy input per proton mass at the mean density, u_0 , for $4.6 \leq z \leq 13$ [cf. [25]], and the cosmological model described by Λ CDM parameters (σ_8, n_s) and a WDM parameter for the inverse of the WDM particle mass of a thermal relic (m_{WDM}^{-1}). The upper section of the table lists the models in the first set of simulations that we use for our MCMC analysis (see text for details). The lower section of the table lists our second set of simulations, which includes mass resolution (R10) and box size (B40) corrections to the predicted flux power spectrum. The dark matter and gas particle mass are $5.37 \times 10^5 h^{-1} M_\odot$ and $9.97 \times 10^4 h^{-1} M_\odot$, respectively, for L20, B40, and a subset of R10 runs (2×512^3). The cosmology parameter ranges for σ_8 include five runs [0.754, 0.804, 0.829, 0.854, 0.904] and similarly 5x runs for n_s [0.921, 0.941, 0.961, 0.981, 1.001]. The WDM mass in keV^{-1} of 0 indicates a CDM run. The other WDM runs are for 2, 3, and 4 keV WDM particle mass.

Name	L_{box} [h^{-1} cMpc]	N_{part}	$z_{\text{rei}}^{\text{end}}$	$T_0(z=4.6)$ [K]	$u_0(z=4.6)$ [eV m_p^{-1}]	σ_8	n_s	WDM mass [keV^{-1}]
L20-ref	20.0	2×1024^3	6.00	10,066	7.7	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-late	20.0	2×1024^3	5.37	10,069	6.6	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-early	20.0	2×1024^3	6.70	10,050	9.6	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-very early	20.0	2×1024^3	7.40	10,003	11.4	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-ref- cold	20.0	2×1024^3	5.98	6598	4.3	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-late- cold	20.0	2×1024^3	5.35	6409	3.6	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-early- cold	20.0	2×1024^3	6.69	6803	5.4	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-very early-cold	20.0	2×1024^3	7.39	6806	6.4	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-ref-hot	20.0	2×1024^3	6.01	13,957	14.4	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-late-hot	20.0	2×1024^3	5.38	13,451	12.5	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-early-hot	20.0	2×1024^3	6.71	14,369	17.8	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
L20-very early-hot	20.0	2×1024^3	7.41	14,624	21.1	[0.754–0.904]	[0.921–1.001]	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
B40-ref	40.0	2×2048^3	6.00	10,063	7.7	0.829	0.961	0
R-set	[5.0, 10.0, 20.0]	$2 \times [1024^3, 768^3, 512^3]$	6.00	10,066	7.7	0.829	0.961	0
R10-ref	10.0	$2 \times [1024^3, 512^3]$	6.00	10,066	7.7	0.829	0.961	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
R10-late	10.0	$2 \times [1024^3, 512^3]$	5.37	10,069	6.6	0.829	0.961	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
R10-early	10.0	$2 \times [1024^3, 512^3]$	6.70	10,050	9.6	0.829	0.961	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
R10-ref-cold	10.0	$2 \times [1024^3, 512^3]$	5.98	6598	4.3	0.829	0.961	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
R10-late-cold	10.0	$2 \times [1024^3, 512^3]$	5.35	6409	3.6	0.829	0.961	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$
R10-ref-hot	10.0	$2 \times [1024^3, 512^3]$	6.01	13,957	14.4	0.829	0.961	$[0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}]$

$\text{Ly}\alpha$ absorption ($\tau_{\text{Ly}\alpha}$) obtained from simulations to match observed values of the effective optical depth $\tau_{\text{eff}} = -\ln\langle F \rangle$. Uncertainties in the background photoionization rate mean a rescaling is commonly used to match the simulations to observations [34,45]. Note that this step is only a good approximation after reionization, as it implicitly assumes that the gas in the low density IGM is in photoionization equilibrium, such that $\tau_{\text{Ly}\alpha} \propto x_{\text{HI}} \propto \Gamma_{\text{HI}}^{-1}$. The redshift evolution that we adopt for τ_{eff} is

$$\tau_{\text{eff}} = 1.56 \times \left(\frac{1+z}{5.75} \right)^4, \quad (2)$$

taken from [25] and similar to the evolution reported in [40,46].

Using the methods described above, we construct a $15 \times 10 \times 10$ grid of parameter values on top of *each* of the 144 simulations (upper section of Table I). This grid of models consists of 10 values of T_0 spanning the range from 5000 to 15,000 K in steps of 1000 K, 10 values of γ spanning the range from 0.9 to 1.8 in steps of 0.1, and 15 values of τ_{eff} in the range from 0.3 to 1.8 times the value in Eq. (2), in multiplicative steps of 0.1. This gives a total of $15 \times 10 \times 10 \times 12 \times (1 + 3 + 2 \times 4) = 216,000$ models. Since we do not extrapolate outside of this grid of models, we have implicit priors on T_0 between 5000 and 15,000 K, γ between 0.9 and 1.8, and for u_0 between (4.03, 21.12),

(3.65, 21.08), and (2.46, 18.73) eV/ m_p for redshifts 4.2, 4.6, and 5.0, respectively.

A. Flux power spectrum models

From the grid of models, we extract 5000 lines of sight in different orientations through the box. The flux field along each skewer is Fourier transformed, and the resulting power spectrum is averaged over all the lines of sight, resulting in the predicted 1D flux power spectrum for a given model. In order to compare the simulated models to the data, we construct an emulator that interpolates the 1D flux power spectrum between the models, allowing us to explore the parameter space spanned by the simulated models. The emulator is based on linear interpolation [22]. Since the grid of models fills the parameter space in a uniform fashion, the interpolation error is small as demonstrated on the subset of the models in [7]. Neglecting quadratic terms in the interpolation leads to at most 1.2% correction at high (k, z) in the flux power spectrum, well below the statistical uncertainty on the data.

Figure 1 shows the 1D flux power spectra when varying the parameters that govern the three main scales of suppression of the flux power. In the left panel, increasing the temperature of the gas at mean density increases the suppression on small scales (high k), while inducing a small increase in power at large scales (low k). The latter is due to keeping τ_{eff} fixed, while the former can be understood in the context of thermal broadening of the lines—the transmission profile of the Lyman- α scattering is determined by the random motion of the gas at a finite temperature. The higher the temperature is, the larger the velocity dispersion of the thermal motion is, leading to more extended profiles that erase small-scale structure.

A related effect, shown in the central panel of Fig. 1, is the effect of pressure smoothing. As the gas is heated during reionization, it hydrodynamically responds to the resulting increase in its temperature and pressure by expanding [6,47]. The more heat injected, the more the gas expands, erasing more small-scale structure. In our models, we parametrized this effect with the cumulative heat injected per proton by a given redshift (u_0) [41]. The exact redshift range of u_0 parameters is the same as in [25].

The small-scale structure in the gas could further be affected by the free streaming of nonstandard dark matter models such as WDM. The lighter the mass of a thermal relic WDM particle is, the longer the particles will free stream, from when they decouple from the thermal bath until they become nonrelativistic. The longer this time is, the larger are the scales affected, and the stronger is the suppression in the small-scale power. This is shown in the right-hand side panel of Fig. 1, where the proxy for the free-streaming scale used is the inverse of the particle mass, m_{WDM}^{-1} .

B. Mass resolution and box size

Since our models are built from the results of hydrodynamical simulations, it is important to understand whether the results of these simulations are numerically converged. Two main factors limit this convergence [34–36]—the size of the simulated box limits the number of large-scale modes and affects the convergence on large scales, and the mass or particle resolution of the simulation limits the smallest resolved scale.

We have supplemented our simulation suite with additional calibration runs varying the size of the simulated box at fixed mass resolution. Our fiducial grid of simulations




FIG. 1. The relative ratios of the 1D flux power spectra of the simulated models relative to a reference simulation run when varying one parameter at a time: T_0 (left), u_0 (center), and m_{WDM} (right). In each panel, all the other parameters are kept fixed. The scale dependence of the flux power spectrum changes in response to changes in the input parameters. The left panel shows the effect of thermal broadening on the absorption features of the forest. The center and right panels show the emergence of a small-scale enhancement of the relative flux power spectrum in simulations with varying reionization history and WDM free streaming. The cumulative heat injection values (center panel) correspond to reionization ending at 5.25, 6.0, 6.75, and 7.5 (top to bottom) for the ionizing UV background model of [39].




FIG. 2. The effect of mass resolution in the simulations, shown as a flux power spectrum decrement as a function of wave number for simulations of varying particle numbers. The mass resolution decrement of the flux power spectrum is largest for the lowest resolution simulations (blue solid) and smallest for the highest resolution simulations (green dashed). The decrement as a function of mass resolution decreases, indicating convergence. The fiducial grid of simulations (20, 1024) used in this work is converged at the 5%–10% level at $k = 0.2 \text{ km}^{-1} \text{ s}$. The default mass resolution correction uses models with higher mass resolution (10, 1024) that are converged at 2%–5% at $k = 0.2 \text{ km}^{-1} \text{ s}$. The shaded regions show the observational 1σ uncertainty on the flux power spectrum from [21] (pink) and [25] (violet). The vertical dashed lines indicate the k_{max} of different data sets.

uses a box size of $20 h^{-1} \text{ Mpc}$. We have applied the splicing correction, [48], using the $40 h^{-1} \text{ Mpc}$ box with the same resolution as L20-ref, which results in a correction of the level of $\leq 3\%$ on the 1D flux power spectrum in the low- k regime. We have further verified that at the scales of interest for the analysis of [25] data; further corrections using 80 and $160 h^{-1} \text{ Mpc}$ box sizes were negligible. This was not an unexpected result and has been observed in several previous studies [21,22,25].

Of more importance for the studies of the small-scale 1D flux power spectrum is the mass resolution of the simulations (R_s). The grid of simulations was run with the fiducial gas mass resolution of $9.97 \times 10^4 h^{-1} M_\odot$, corresponding to 2×10^{24} baryon and dark matter particles. These models are converged at 5%–10% at the smallest scales used in the analysis. We have complemented these models with additional simulations varying the number of simulated particles at different fixed box sizes.

Figure 2 shows the 1D flux power spectrum decrements between different models. The poorer the mass resolution of the simulation is, the larger the suppression of the small-scale flux power spectrum is relative to a higher resolution simulation. The mass resolution correction is larger at higher redshifts, and at smaller scales, in agreement with

previous results in the literature (e.g., [35,36]). The mass resolution correction (R_s) and the 1D flux power decrements shown in Fig. 2 are connected as $R_s^{-1} = 1 + \Delta P/P$. The grid of our simulations at the resolution of (20, 1024) was corrected for the residual mass resolution with (10, 1024) model (R-set; see Table I), corresponding to gas mass resolution of $1.25 \times 10^4 h^{-1} M_\odot$. Additional correction due to higher-resolution simulations [e.g. (5, 1024)] adds less than a few percent to the total mass resolution correction.

IV. RESULTS

The new results on the free streaming of warm dark matter are summarized in Fig. 3. The six panels show the 2D posteriors for three redshift bins of the data [25], with the redshift label referring to the label of the thermal parameters that are independent in each redshift bin. The bottom row shows the constraints in the thermal parameter space of gas temperature and pressure smoothing (through the proxy of cumulative injected heat), whereas the top row shows the constraints spanning the parameter space of pressure smoothing and free streaming.

The fiducial analysis choice assumes priors on the thermal history in the $u_0 - T_0$ plane as an envelope around our fiducial grid of simulations (see below). We also




FIG. 3. The 2D posterior distributions of the best-fit analysis for the 1D flux power spectrum measurements of [25] using UVES/HIRES quasar spectra. The blue contours show the default analysis, and the orange contours show the analysis that captures our best knowledge of the thermal history (Sec. IV D), inhomogeneous reionization (Sec. IV G), and mass resolution corrections (Sec. IV F). The three columns correspond to the three different redshifts of $z = 4.2, 4.6$, and 5.0 (from left to right). The bottom row shows the contours in the thermal parameter space, with the violet band showing the envelope around the physically motivated simulations, shown as gray points (squares, circles, and triangles). This band serves as a prior in the thermal parameter space in the default model. The colored points correspond to the measurements in the literature from the same data set: from [25] (purple), [15] (in green), and [24] (in red). The top panels show the 1 and 2σ contours in the parameter space of free streaming and pressure smoothing (heat injection). The vertical dotted line and surrounding gray band indicate the best-fit measurements of [25]. The intersecting gray dashed and dot-dashed lines show typical degeneracy axes between the parameters. The cutoff at small u_0 and small T_0 values comes from the implicit prior imposed by the extent of the grid of models (see text for details).

assume Planck [49] priors on CDM cosmology parameters (σ_8, n_s). For the default analysis, we use mass resolution correction using a fiducial thermal history with CDM cosmology (R-set; see Table I). We also do not include any correction due to inhomogeneous reionization. These assumptions were chosen as our reference analysis in order to facilitate better comparison with previous analysis. The additional work presented in this paper which includes patchy correction, thermal dependence of the mass resolution correction [$R_s(u_0)$], and observationally informed thermal priors (T_0 prior) is also shown in Fig. 3 (orange contours) and discussed in more detail in subsections below.

Our measurements of the thermal state of the gas largely agree with independent measurements in the literature [15, 24, 25] within $1 - 2\sigma$. The data prefer a slightly colder temperature at mean density of $T_0 = 8,000$ (7500; 7800) K at redshift $z = 4.2$ (4.6; 5.0) as a best fit (see Table II). At the same time, the cumulative heat injected is constrained to

be $u_0 = 7.2$ (6.8; 5.2) eV/ m_p between redshifts 4.2 and 12.0 (4.6 and 12.0; 6.0 and 13.0). The result is consistent with the analysis of [25]. However, models with slightly hotter temperature consistent with [44, 50] and less pressure smoothing [51] are within the 2σ contours.

The measurements of effective optical depth, τ_{eff} , from the flux power spectrum are also consistent with direct observations of the transmitted flux [46, 52]. The derived measurement of the mean transmitted flux at $z = 5.0$ is $\langle F_{\text{Ly}\alpha} \rangle = 0.1764^{+0.0177}_{-0.0171}$. This is consistent at $1 - 2\sigma$ with the measurement of [52] of $\langle F_{\text{Ly}\alpha} \rangle = 0.1581^{+0.0082}_{-0.0089}$ which used almost four times the number of sightlines compared to [25].

This analysis also varies the power law of the temperature-density relation (γ) as a free parameter in each redshift bin. The data, however, are not constraining this parameter well, and its posterior is dominated by the prior. This result

TABLE II. List of different models used in the analysis with their corresponding best-fit warm dark matter constraints. The table shows the name of the model and the resulting 2σ lower bound on the WDM particle mass (m_{WDM}), along with best-fit values of the thermal parameters at $z = 4.6$ for the effective optical depth (τ_{eff}), gas temperature at mean density (T_0), the slope of the temperature-density relation (γ), and the cumulative injected heat (u_0). For the model where extra instrumental noise in the data was modeled with a free parameter, the best-fit value is shown as well (A_{noise}). The last column displays the best-fit χ^2 value and the degrees of freedom.

Name	m_{WDM} [keV] (2σ)	$\tau_{\text{eff}}(z=4.6)$	$T_0(z=4.6)$ [10^4 K]	$\gamma(z=4.6)$	$u_0(z=4.6)$ [eV/ m_p]	$A_{\text{noise}}(z=4.6)$	χ^2/dof
Default	>5.72	$1.502^{+0.061}_{-0.061}$	$0.743^{+0.041}_{-0.075}$	$1.35^{+0.24}_{-0.19}$	$6.19^{+0.68}_{-0.68}$	...	40.7/34
$k_{\text{max}} < 0.1 \text{ km}^{-1} \text{ s}$	>4.10	$1.501^{+0.060}_{-0.074}$	$0.840^{+0.095}_{-0.340}$	$1.28^{+0.09}_{-0.28}$	$8.91^{+1.57}_{-5.26}$	...	10.2/20
A_{noise}	>3.91	$1.458^{+0.053}_{-0.074}$	$0.966^{+0.156}_{-0.466}$	$1.23^{+0.06}_{-0.23}$	$5.93^{+0.38}_{-2.28}$	$1.12^{+0.49}_{-0.29}$	18.4/31
T_0 prior	>5.85	$1.494^{+0.062}_{-0.077}$	$0.770^{+0.110}_{-0.120}$	$1.31^{+0.10}_{-0.31}$	$6.50^{+1.00}_{-1.60}$	...	47.6/34
$R_s(u_0)$ mass resolution	>4.44	$1.531^{+0.073}_{-0.064}$	$0.617^{+0.007}_{-0.118}$	$1.38^{+0.28}_{-0.13}$	$7.90^{+1.70}_{-2.30}$	...	30.7/34
Patchy reionization	>5.10	$1.486^{+0.058}_{-0.068}$	$0.686^{+0.046}_{-0.080}$	$1.33^{+0.17}_{-0.26}$	$5.32^{+0.58}_{-0.52}$	...	41.0/34
$R_s(u_0) + T_0$ prior	>4.24	$1.473^{+0.056}_{-0.076}$	$0.83^{+0.11}_{-0.11}$	$1.28^{+0.09}_{-0.28}$	$5.53^{+0.73}_{-1.2}$	...	39.4/34
Patchy + $R_s(u_0) + T_0$ prior	>5.90	$1.450^{+0.051}_{-0.070}$	$0.828^{+0.098}_{-0.098}$	$1.26^{+0.08}_{-0.26}$	$4.87^{+0.52}_{-0.71}$	...	40.8/34

was also found in previous studies of high redshift Lyman- α forest data (e.g., [22,25]).

The panels at the bottom of Fig. 3 also show the $u_0 - T_0$ combinations of hydrodynamical simulations as gray markers (L20; see Table I). The thermal and reionization histories were chosen to bracket the observed flux distribution of high redshift quasar spectra [53], as well as the electron optical depth inferred from the cosmic microwave background as reported by Planck [49,54]. Through the postprocessing technique described in Sec. III, the likelihood is able to sample the full span of the $u_0 - T_0$ parameter space on a (nonuniform) grid; however, in order to avoid unphysical parts of the $u_0 - T_0$ parameter space, we consider a prior defined as an envelope around the simulations' results (indicated in Fig. 3 by the gray band).

A. Degeneracy axes

The simulated models exhibit a tight correlation between the IGM temperature at a given time and the integrated injected heat up until that time. The positive correlation between the thermal parameters (dot-dashed lines in bottom panel of Fig. 3) can be well described by $u_0 \propto T_0^{1.7}$, and the parameter anticorrelation (dashed lines in bottom panel of Fig. 3) is well described by $u_0 \propto -T_0$. The anticorrelation also indicates the degeneracy axis we would expect from the measurement of the 1D flux power spectrum. At a given observed redshift, the flux power suppression can be explained by either higher injected heat and, therefore, a larger pressure smoothing scale, or it can be explained by a higher temperature and therefore larger thermal broadening. The Lyman- α forest provides constraints in the direction perpendicular to that degeneracy axis, along the direction of the positive correlation between u_0 and T_0 .

Similarly to the degeneracy between the thermal broadening and pressure scales, we observe a correlation between the pressure smoothing and the free-streaming scales, as shown in the top panels of Fig. 3. The vertical black dashed line and gray shaded region indicate measurements of the cumulative injected heat, u_0 , in a CDM analysis of [25]. A negative correlation (dot-dashed lines in top panel of Fig. 3) between the two smoothing scales can be understood as a consequence of both physical mechanisms reducing the small-scale power of the 3D density field. The pressure smoothing scale is typically described as an exponential suppression of the power, $P_g \sim P_m \exp(-k^2 \lambda_F^2)$ [6,55], at a typical filtering scale λ_F . The larger the heat injected into the gas is, the more the gas expands due to the pressure, resulting in a positive correlation between the filtering scale and the injected heat u_0 . Such a relation was explored in the simulations of [25], where it was found that $\lambda_F \sim 20 \text{ ckpc} \times \sqrt{1 + 2u_0/(1 \text{ eV}/m_p)}$.

Equivalently, the warm dark matter transfer function can be approximated by $T_{\text{WDM}} \sim [1 + (\alpha k)^{2\mu}]^{-5/\mu}$, with $\mu = 1.12$ and the typical free-streaming scale, $\alpha = 70 \text{ ckpc} \times (m_{\text{WDM}}/(1 \text{ keV}))^{-1.11}$, given by [21].

The total power suppression in the 3D field on small scales is a product of both the pressure smoothing and free-streaming transfer functions. Expanding the product in powers of k , the lowest scale dependent coefficient scales as $\propto k^2$, with the amplitude of $c_2^2 = \lambda_F^2 + 10\alpha^2$, where we have approximated $\mu \sim 1$. The anticorrelation between pressure smoothing and the free streaming that we observe in the data is driven by being sensitive to the total shape of the suppression; thus, $c_2^2 = \text{constant}$. This can be interpreted as the smoothing being driven by either higher

pressure smoothing or larger free-streaming length and is shown as dot-dashed gray lines in Fig. 3.

Whereas the shape of the power spectrum suppression is poorly constrained by the current data, the data are able to constrain the scale where the suppression occurs—shown as dashed lines in Fig. 3. We estimate this positive correlation between the parameters (dashed lines in top panel of Fig. 3) by matching the scale where the pressure smoothing and free-streaming transfer functions equal one half (e.g., $T_{\text{WDM}}(k_{1/2}) = 1/2$, or $P_{\text{WDM}}(k_{1/2}) = 1/4$). The two scales are given by $k_{1/2}^g = \sqrt{2 \log 2} / \lambda_F$ and $k_{1/2}^{\text{WDM}} = (-1 + 2^{\mu/5})^{\mu/2} / \alpha$. Equating the two leads to a relation $m_{\text{WDM}}^{-2.22} \propto 1 + 2u_0/(1 \text{ eV}/m_p)$ that defines the directional axis along which the Lyman- α forest data give the tightest constraints.

B. Best-fit model

Figure 4 shows 1D flux power spectrum corresponding to the best-fit model overplotted on the data. The model fits the data reasonably well, with a total χ^2 of 40.7 and 34 degrees of freedom (see Table II). Furthermore, the model is in excellent agreement with the data up to $k \sim 0.1 \text{ km}^{-1} \text{ s}$ and describes the position and shape of the flux power spectrum suppression on small scales. To illustrate this, we can compare the model that is fit to all the data points and re-evaluate the χ^2 for the points up to $k < 0.1 \text{ km}^{-1} \text{ s}$. In this case, the fit gives χ^2 of 20.4 with 20 degrees of freedom. All three redshift bins show an increase in the measured power relative to the model at $k > 0.1 \text{ km}^{-1} \text{ s}$. This indicates a possible shortcoming of the model on the smallest scales, or else a signal in the data that is not part of the model.

The best-fit model excludes $m_{\text{WDM}} < 5.7 \text{ keV}$ (95% CL) and provides the tightest constraints on the thermal relics WDM particle mass to date (see Table II). The model

constraints exclude masses of 3.73 keV and 3.18 keV at 3σ and 5σ , effectively excluding the much discussed 3 keV WDM model (e.g., [56,57]) at more than a 5σ confidence level.

C. Improvement on WDM constraints

In Fig. 5, we compare the results of this work to the existing constraints on the WDM mass from the literature. The main result of this work results in a WDM mass bound that excludes WDM masses below $m_{\text{WDM}} < 5.7 \text{ keV}$ at 2σ confidence level. It provides improved constraints on WDM mass coming from the matter power spectrum suppression in the Lyman- α forest analyses [18,22] as well as non-Lyman- α constraints such as the flux ratios of strong lensed systems [58] and stellar streams in the Milky Way [59].

The new constraint is stronger than the studies using low- z [60] or a combination of low- z and high- z [21,61] Lyman- α data, especially when comparing to similar choices in the thermal history priors. The new data are in fact producing a strong enough constraint that, even when relaxing the prior on the astrophysical parameters, the WDM mass bound remains stronger or competitive with past studies that used strong priors on the, e.g., temperature evolution with redshift [18,22].

In the regime of the high redshift Lyman- α forest analysis, the current analysis tightens the constraint on the WDM particle mass compared to previous analyses. In comparison to older analyses using HIRES/MIKE data [21,22], we see an improvement in the number of the observed quasar spectra by almost a factor of 2 [25]. For a factor of 2 improvement in the number of sightlines, we would expect the uncertainty on the flux power spectrum to improve by $\sim 1/\sqrt{2}$, at least in the limit that statistical uncertainty dominates the error budget. From Fig. 2, we see that this is indeed the case in the high- k regime of the data




FIG. 4. The best-fit model compared to the data [25]. The three panels correspond to three redshift bins, with the bottom panels showing the residuals of the data over the model. The total χ^2 is 40.7 with 34 degrees of freedom. The data were compared to a simulation based model that varies three thermal parameters and mean transmission independently in each redshift bin ($\tau_{\text{eff}}, T_0, \gamma, u_0$) and three cosmology parameters (σ_8, n_s) and (m_{WDM}) (see text for details).




FIG. 5. The 2σ constraints on the thermal relic warm dark matter mass. The arrows indicate the exclusion limits on the WDM particle mass in keV. The bottom panel shows a compilation of constraints from high redshift Lyman- α forest 1D flux power spectrum. The black arrows at the very bottom indicate the results of this study, for three different analysis choices pertaining the measured flux power at highest wave numbers [default, (P, R_s, T_0) that represents corrections due to patchy reionization, thermal dependence on the mass resolution and independent T_0 prior, and $k_{\max} < 0.1 \text{ km}^{-1} \text{ s}$ data scale cut analysis]. The resulting lower bounds on the WDM particle mass are stronger or comparable to those previously published in the literature, including studies that combined low- and high- z Lyman- α forest data to increase the redshift lever arm (middle panels). The top panel shows a compilation of results from non-Lyman- α studies.

that is most sensitive to the free-streaming effect of WDM. In fact, in linear theory the sensitivity to the WDM mass scales as $P_{\text{L,wdm}}/P_{\text{L,CDM}} \sim m_{\text{wdm}}^{20} k^{-20}$ in the limit of $k \gg 14(m_{\text{wdm}}/1 \text{ keV}) \text{ Mpc}^{-1}$.

However, the nonlinear mapping between the linear density field and the nonlinear flux field is complex. For a range of redshifts ($4.2 < z < 5.0$) and scales ($0.01 < k/[\text{km}^{-1} \text{ s}] < 0.2$) considered, the flux power spectrum suppression in our simulations (L20-ref) approximately scales as

$$\frac{P_{\text{F,wdm}}}{P_{\text{F,cdm}}} \sim \begin{cases} 1 - 0.1 \left(\frac{1+z}{5} \right)^4 \left(\frac{k}{0.1} \right)^{\frac{3}{4}} \left(\frac{m_{\text{wdm}}}{4} \right)^{-1}, & m_{\text{wdm}} > 3 \text{ keV} \\ 1 - 0.1 \left(\frac{1+z}{5} \right)^3 \left(\frac{k}{0.1} \right)^{\frac{1}{2}} \left(\frac{m_{\text{wdm}}}{4} \right)^{-\frac{3}{2}}, & m_{\text{wdm}} < 3 \text{ keV}, \end{cases} \quad (3)$$

with line-of-sight wave number k in units of $[\text{km}^{-1} \text{ s}]$ and m_{wdm} in units of [keV]. For higher WDM masses, the flux power suppression due to WDM increases rapidly with redshift, but only linearly with the WDM particle mass. The scaling changes at around the WDM mass of 3 keV, when the scaling with mass becomes stronger, and the redshift dependence slightly weaker. The wave number dependence is roughly the same, and not dominant in this range of scales. The scaling is only approximately valid at a fixed thermal history (L20-ref), and the transition between the two scales, as well as the power-law dependencies, can vary across thermal histories. However, at the higher WDM mass limit, the sensitivity to the particle mass increases with redshift relatively quickly in the redshift range of the data, improving the linear sensitivity to the mass. As a result, the constraining power on WDM mass improves by more than a factor of $\sim 1/\sqrt{2}$.

D. Thermal history

The signal at high k in the 1D Lyman- α forest flux power spectrum depends on the thermal parameters. The fiducial priors on the thermal history limit the possible combinations in the $u_0 - T_0$ plane to the volume of physically motivated simulation results [6].

A different approach would be to instead apply a prior based on independent measurements of the thermal history, for example, using the measurements of $T_0(z)$ from different datasets and different statistical methods. To achieve that, we use improved and precise measurements of $T_0(z)$ that span the redshift range $z < 3.8$ [44] and $z > 5.2$ [50]. In order to predict viable models in the redshift range covered by our data ($4.2 < z < 5.0$), we rescale and shift the photoheating and photoionization rates of our fiducial thermal history model [39]. A similar methodology was employed in [62] in order to fit the flux power spectra measurements over a range of redshifts. Figure 6 shows the two models from the literature, as well as our new fit calibrated directly against $T_0(z)$ measurements. The best fit prefers slightly higher temperatures in the redshift range $4.2 < z < 5.0$ than the measurements of [25] using the flux power spectra data used in this work. In order to construct informative priors on $T_0(z)$ parameters in our model, we use the best-fit values of the new thermal history as central points of a Gaussian distribution at each redshift, with a fixed standard deviation of 1000 K. While the standard deviation is somewhat arbitrarily chosen, it roughly matches the typical uncertainty found in more recent works [44,50]. Even if a realistic uncertainty is slightly lower at lower redshift, and slightly higher at higher redshift, due to the decreasing numbers of quasar spectra available, this should not impact the main conclusion of this exercise, which is to highlight the effect of independent, observationally informed thermal priors.

The results of the analyses using different thermal prior choices are shown in Fig. 7. The fiducial model (green




FIG. 6. The thermal evolution of $T_0(z)$ for various models, compared to the independent measurements of [44] (high z), [50] (low z), and temperature measurements of [25]. The models shown are that of [39] (in black; our reference simulation run), [62] (in blue), and a new fit to [44,50] data (in red). The new fit was obtained by rescaling the model of [39] and serves as a prior in the analysis shown in Fig. 7, with prior values shown as red circles and error bars at $z = 4.2, 4.6, 5.0$ with values of 9155.5, 8986.5, and 9286.5 K, respectively. The uncertainty propagated in the prior is 1000 K at each of the redshifts.

contours) uses thermal priors in the form of an envelope around the simulations (gray band). Replacing these priors by simpler priors on T_0 at each redshift results in slightly more elongated constraints on the $u_0 - T_0$ plane (orange

contours), with the posterior expanding along the degeneracy direction. The mean and best fit of the posterior, however, change only marginally compared to the standard analysis. Even though the posterior in the u_0 direction expands slightly, the posterior on m_{WDM}^{-1} remains roughly the same at the 2σ level, resulting in a very similar constraint on the WDM mass of $m_{\text{WDM}} > 5.85$ keV (2σ) compared to the default analysis choice of thermal priors. The main difference is that the thermal priors have now been informed by the measured T_0 evolution with redshift from other observational studies, rather than by our suite of hydrodynamical simulations. The model with the T_0 prior (see Table II) excludes low WDM masses of 3.75 keV and 3.21 keV at 3σ and 5σ , respectively.

Figure 7 also illustrates the effect of not imposing any thermal priors on the analysis. This resulting posterior (blue contours) is shifted to lower IGM temperatures and relatively higher values of the cumulative injected heat. This part of the thermal parameter space is unphysical as we expect u_0 and T_0 to be correlated for physically reasonable IGM heating scenarios. Rather counterintuitively, the constraints on the WDM mass become much stronger if we impose no thermal prior. This result can be understood by considering degeneracy axis along which the posterior distributions move. Colder temperatures and enhanced amount of pressure smoothing leave much less room for a WDM model to accommodate the amount of




FIG. 7. Effect of thermal priors on the posterior. The two panels show 2D posterior distributions for redshift $z = 4.2$ in the plane of temperature and heat injection (left) and warm dark matter mass and heat injection (right). In the thermal parameter space (left), the default analysis (blue contours) uses thermal priors that envelop the simulations (the envelope is shown as a violet shaded area). A similar result can be obtained by instead imposing a $T_0(z)$ prior (orange contours) using independent temperature measurements [44,50]. The warm dark matter particle mass constraints get slightly stronger if a temperature prior is used instead of the envelope prior in the $u_0 - T_0$ plane. As a reference, we also show an analysis without imposing any thermal priors (green contours). The vertical shaded band on the left panel indicates measurement of u_0 from [25].

flux power spectrum suppression in the data. Therefore, models with strong WDM suppression are excluded more strongly.

E. Effect of small-scale peculiar velocity

The enhancement of the small-scale power in the models is associated with the enhancement of the small-scale structure in the peculiar velocities. Figure 8 (left) shows the relative effect of the peculiar velocity field on the flux power spectrum ratios. The models of early ($z_{\text{rei}} = 7.5$) and reference ($z_{\text{rei}} = 6.0$) reionization (blue solid line) show a relative enhancement of power compared to the ratio of the two flux power spectra when the effects of peculiar velocities are not included in the calculation of the optical depth. This effect of setting $v_{\text{pec}} = 0$ has also been seen in [24]. Figure 8 further illustrates that the amplitude of this feature at $k > 0.1 \text{ km}^{-1} \text{ s}$ is sensitive to the amplitude of the peculiar velocity field changing with the amount of pressure smoothing. Furthermore, the feature in the flux power can be associated with the emerging feature in the 1D power spectra of the peculiar velocities (Fig. 8; right), with the strength of the feature exhibiting a positive correlation between its amplitude and the cumulative injected heat. That the feature is stronger for later ending

reionization, and weaker for earlier reionization, suggests that this behavior is due to the hydrodynamic response of the gas to the photoheating.

In terms of the constraints on the WDM particle mass, this suggests that a certain caution has to be exercised when pushing the models to $k > 0.1 \text{ km}^{-1} \text{ s}$. While the peculiar velocity feature might be related and correlated with the existing thermal parameters, it is not *a priori* obvious that this new scale in the model is properly covered within the range of existing simulations, and therefore not properly marginalized over.

F. Thermal dependence of the mass resolution

The results of Sec. IV E show that a small-scale peculiar velocity structure can modify the amount of small-scale flux power. However, this is also the regime where the mass resolution (R_s) of our simulations has the biggest effect.

The mass resolution correction of the simulations should depend on the thermal history in this high- k regime of the model, i.e., $R_s = R_s(u_0)$. This is perhaps not surprising—the mass resolution corrections essentially describe how much small-scale structure is missing in the (power spectrum) statistics as a result of not resolving the structure at very small scales and its nonlinear coupling to larger scales. If the field in configuration space is smoothed out




FIG. 8. Effect of peculiar velocities on the small-scale ($k > 0.1 \text{ km}^{-1} \text{ s}$) suppression of power. Left: the ratio of 1D flux power spectra of very early ($z_{\text{rei, end}} = 7.5$) and reference ($z_{\text{rei, end}} = 6.0$) reionization models (solid blue). The dashed lines show the effect of replacing the peculiar velocity fields of both simulations simultaneously. The dashed blue line shows the effect where no peculiar velocities are included. The colored lines show the effect of using peculiar velocity fields corresponding to different thermal histories—in particular, different heat injection values. All models show a relative increase of small-scale power ratio compared to the model with no peculiar velocities. The strength of this relative increase correlates with heat injected during reionization, with v_{pec} coming from a late reionization run (low heat injection; red dashed line) giving the strongest signal. Right: the relative increase in the small-scale structure of the 1D flux power spectrum is related to small-scale structure in the peculiar velocity field. A feature is present in the power spectrum of the peculiar velocity gradient ($\eta = \nabla v_{\text{pec}}$), where the peak shifts from $k \sim 0.15 \text{ km}^{-1} \text{ s}$ for early reionization models (with higher heat injection) to $k \sim 0.30 \text{ km}^{-1} \text{ s}$ for late reionization models (with lower heat injection). Models with cumulative injected heat of $u_0 = 8.14, 11.6, 7.11, 21.1 \text{ eV}/m_p$ correspond to the L20-ref, L20-very early, L20-late, and L20-very early-hot models, respectively. Similarly, the models from [6] with $u_0 = 6.04 \text{ eV}/m_p$ correspond to their late reionization model ($z_{\text{rei, end}} = 5.3$).




FIG. 9. The effect of mass resolution in the simulations, shown as a flux power spectrum decrement as a function of wave number for different models with varying the thermal history and WDM particle mass. The more heat is injected into the IGM, the more the gas is smoothed due to pressure effects. Larger pressure smoothing results in less structure in the flux power spectrum at small scales and depends less on the mass resolution. The same happens if the matter density is smoother due to the presence of free streaming, which results in a smaller required mass resolution correction.

due to physical effects—such as higher pressure smoothing or larger free-streaming scale—the amount of missing small-scale flux power will also be smaller. To estimate this effect, we repeated the resolution correction exercise for different models in our suite of simulations (R10-; see Table I). The results, in Fig. 9 show that indeed the mass resolution correction exhibits a strong dependence on the thermal history at $k > 0.1 \text{ km}^{-1} \text{ s}$. In particular, late reionization models with less pressure smoothing (or lower cumulative injected heat) can show up to 5% larger mass resolution corrections compared to the fiducial correction used in the analysis. This trend is more prominent at higher redshifts and less important at $z \leq 4.2$. On the other hand, models with larger pressure smoothing scales require consistently smaller resolution corrections at small scales, by up to 2%.

Similarly to the effect of the thermal history, the smoothing of the density field due to free streaming also decreases the required mass resolution correction. As shown in Fig. 9, a 2 keV WDM model on average requires a 5% lower mass resolution correction at $k \sim 0.1 \text{ km}^{-1} \text{ s}$ at $z = 5.0$. This effect is reduced at lower redshifts.

These results imply that applying a mass resolution corrections that depends on the thermal history widens the range of P_{ID} at $k > 0.1 \text{ km}^{-1} \text{ s}$ for models within a given section of the parameter space, ultimately resulting in

higher sensitivity to thermal parameters and lower sensitivity to free streaming. On the other hand, the mass resolution correction that depends on m_{WDM} shows stronger sensitivity of the P_{ID} at $k > 0.1 \text{ km}^{-1} \text{ s}$, which leads to stronger constraints on the lower bound WDM mass. Current bounds on the WDM mass lie in the range of $\sim 4\text{--}6 \text{ keV}$, however, and the effect of mass resolution dependence on WDM free streaming is severely reduced. Thus, most of the effect of the mass resolution that depends on thermal history and WDM mass comes from the thermal history dependence.

The results of applying free streaming and thermal history dependent mass resolution correction are shown in Fig. 10. Compared to the analysis without any thermal priors shown in Fig. 7, the new mass resolution correction does not shift the posterior in the thermal parameters, suggesting that the effect of peculiar velocities is not completely explained by accounting for the thermal dependence in the mass resolution. On the other hand, the WDM constraints are weakened, roughly to the same level as when a sensible thermal prior is applied to the model, resulting in $m_{\text{WDM}} > 4.44 \text{ keV}$ at 2σ . Low WDM masses of 3.19 and 2.78 keV are excluded at 3σ and 5σ , respectively. Applying both the physical thermal prior (T_0 prior) and the new mass resolution correction leads to $m_{\text{WDM}} > 4.24 \text{ keV}$ at 2σ .




FIG. 10. Effect of using a mass resolution correction that depends on the IGM thermal history1 ($R_s(u_0)$). The two panels show 2D posterior distributions for redshift $z = 4.2$ in the plane of temperature and heat injection (left) and warm dark matter mass and heat injection (right). In the thermal parameter space (left), the mass resolution correction adds more power for models with less pressure smoothing, which makes such models easier to fit the data. As a result, the posterior still lies along the $u_0 - T_0$ degeneracy axis, but the amplitude of this axis moves to lower u_0 values. Similarly, the degeneracy between free streaming and pressure smoothing (right) opens up along the degeneracy axis, because the models with low u_0 require larger mass resolution corrections that increase the power, even for lower warm dark matter masses.

G. Patchy reionization

The original Sherwood suite of simulations used in this study evolves the reionization homogeneously throughout the simulated volume. In reality, the Universe reionizes in a more complex, inhomogeneous manner, where local ionized bubbles first appear around the sources of ionizing photons [6]. Observations of the Lyman- α forest at higher redshifts can thus still be affected by relic fluctuations of the reionization persisting for a time after most of the Universe has been reionized. This topic has been a focus of several studies over the years [7,31,63–66]. The main effect of the patchy nature of reionization on the 1D flux power spectrum of the Lyman- α forest has been found to be an enhancement of power on large scales that traces the fluctuations in the temperature and ionized fraction of hydrogen gas. The conclusions of recent works [7,31] suggest that the enhancement of power appears at larger scales ($k < 5 \times 10^{-3} \text{ km}^{-1} \text{ s}$) than those observed in [25], i.e., the flux power spectrum measurements used in this study.

While the large-scale effect of ionization fluctuations and their effect on the Lyman- α forest have seen a certain agreement between different methods and simulations, the same is not true for the effect of inhomogeneous reionization on small scales. Spatial fluctuations of the photoionization rate result in spatial fluctuations of the temperature density relation. Regions of the IGM that are ionized later heat up later as well, while regions that reionized and

heated up earlier had time to cool down, due primarily to the expansion of the Universe and inverse Compton scattering [63,64,67]. As a result, regions ionizing later would exhibit stronger suppression of the flux power spectrum due to thermal Doppler broadening than IGM regions that have ionized long ago. As pointed out by [6,65], a competing effect to the thermal fluctuations is that the IGM regions that ionized earlier had more time to hydrodynamically respond to the injected heat, resulting in a larger pressure smoothing scale. More pressure smoothing also reduces the small-scale power. It has been suggested that these two effects might largely cancel each other out, leaving small-scale power unchanged compared to homogeneous reionization models.

In this study, we make use of the tabulated correction to the 1D Lyman- α flux power spectrum from [7] that is based on the Sherwood Relics simulation suite [6]. The effect of a patchy reionization correction in that study results in a $\sim 10\%$ suppression of Lyman- α flux power at $k > 0.1 \text{ km}^{-1} \text{ s}$. The amplitude and shape of the suppression are largely independent of the thermal history models used in that study. Aside from the temperature fluctuations, [7] found that the dominant effect of the small-scale suppression was due to the effect of the peculiar velocity field.

The results of this analysis are presented in Fig. 11. Since patchy reionization suppresses the small-scale power, one could expect that lower values of WDM masses might be even further excluded by the data. However the small-scale




FIG. 11. Effect of including a correction due to inhomogeneous reionization. The two panels show 2D posterior distributions for redshift $z = 4.2$ in the plane of temperature and heat injection (left) and warm dark matter mass and heat injection (right). The effect on the analysis mostly comes from the small-scale suppression of power due to the peculiar velocity structure of the gas, rather than the large-scale enhancement of power due to photoionization and temperature fluctuations. Different thermal history models have an almost identical suppression of power due to inhomogeneous reionization compared to the equivalent homogeneous reionization model. The net effect is a systematic shift along the (positive) $u_0 - T_0$ degeneracy axis in the thermal parameter space (left), in the direction of lower temperature and lower pressure smoothing. Similarly, the degeneracy between free streaming and pressure smoothing (right) opens up along the degeneracy axis, because the models with lower thermal and pressure smoothing leave more freedom for WDM suppression to accommodate the data.

suppression induced by inhomogeneous reionization affects *all models equally*, including the models with different thermal and pressure broadening. The main effect on the Lyman- α data analysis is to move the peak of the T_0 posterior to lower values. The reason for this is as follows: the higher the T_0 value is, the stronger the suppression due to thermal broadening is. The flux power spectrum models for low T_0 values that on their own do not exhibit enough suppression to explain the data now achieve enough suppression through patchy reionization correction. Therefore, the first conclusion is that lower T_0 models that were excluded before now fit the data, and the posterior of the T_0 parameter expands towards lower T_0 values. On the other hand, the models with high T_0 values would now show too strong of a suppression, and models that fit the data without the correction due to patchy reionization are now in tension with the data. The posterior of T_0 therefore shrinks for high T_0 values. Due to the prior in the $u_0 - T_0$ plane, shifting the T_0 posterior to lower values also shifts the u_0 posterior to lower values at each redshift, resulting in data preferring less pressure smoothing. Since both the thermal and pressure smoothing effects are reduced, the posterior of the WDM mass expands to compensate for the fact that somewhat lower WDM mass models are now no longer in tension with the data.

While the specific result shown in Fig. 11 depends on the choice of thermal priors, the main conclusion would remain

the same even in the light of less stringent priors. As the T_0 posterior systematically shifts to lower values, the $u_0 - T_0$ anticorrelation direction is preserved as it depends on the fact that both parameters increase the small-scale suppression. The resulting posterior in a scenario with wider thermal priors would therefore only extend further along the $u_0 - T_0$ anticorrelation direction, but still resulting in reduced sensitivity to m_{WDM} .

From Fig. 11, we also observe that a $\sim 10\%$ suppression of power in all the models results in only $\sim 0.5\sigma$ shift of the posterior in the $u_0 - T_0$ plane, along the positive degeneracy axis (shift between the blue and orange contours). The constraints on the WDM mass are thus slightly weaker, with $m_{\text{WDM}} > 5.10$ keV at (2σ) . However, as was highlighted in [7], the 10% small-scale suppression is mainly driven by the peculiar velocity field differences between the inhomogeneous and homogeneous reionization models. As we have shown in previous sections, the exact nature of the peculiar velocity structure on small scales has implications for the WDM mass inference and can affect both the mass resolution correction of the simulations as well as parametrization of the thermal history on the smallest scales probed by the Lyman- α forest (~ 50 – 100 ckpc/h). While the nature of the peculiar velocity field structure requires further study, it is reassuring that the effect on the WDM constraints is small ($\sim 10\%$).

Combining the corrections due to inhomogeneous reionization and the thermal history dependence of the mass resolution $[R_s(u_0)]$ with the thermal priors coming from independent $T_0(z)$ observations (T_0 prior), we get a combined constraint on the WDM particle mass of > 5.9 keV (95% CL). Even though individually both the patchy reionization and the $R_s(u_0)$ correction reduce the WDM constraining power, together with the T_0 prior, they are pushed in the parameter space of higher T_0 values and a lower pressure smoothing scale (or late reionization), which leaves little room for additional suppression due to WDM free streaming. While these constraints are the strongest presented in this paper, they rely on our first attempt at both a patchy reionization and $R_s(u_0)$ corrections. With their impact on the WDM particle mass, these results provide additional incentive to improve on the modeling of the small-scale thermal history in the inhomogeneous reionization models.

H. Instrumental effects

Several instrumental and observational effects can potentially systematically alter the small-scale flux power spectrum: mis-estimation of the observed flux noise, contamination due to metal lines, or the instrument resolution. Of the three, the instrument resolution has been one of the more studied effects, as it has a large impact on large Lyman- α surveys that observe spectra at lower spectral resolution [19,61,68,69]. For a typical line-spread function shape, the correction of the instrument resolution on the flux power spectrum is well described by a Gaussian kernel $P_F \rightarrow P_F/W_k^2 = P_F \exp k^2 \sigma_R^2$, with the Gaussian width of the resolution $\sigma_R = \text{FWHM}_R / (2\sqrt{2 \ln 2})$ given as function of the FWHM resolution element ($\text{FWHM}_R = c/R$) or resolving power (R). For the scales of $k \gtrsim \sigma_R^{-1}$, the correction due to resolution becomes ~ 1 , dominating the total signal in the instrument. While of significant concern for lower resolution instruments such as X-Shooter ($R \sim 8,000$), the resolving power of Keck/HIRES and VLT/UVES is high enough ($R \sim 50,000$ – $80,000$) to not play a major role in flux power spectrum measurements for scale cuts $k < 0.1 \text{ km}^{-1} \text{ s}$ [21,25,70]. Indeed assuming the fiducial value of the resolution $\text{FWHM}_R = 6 \text{ km/s}$ ($R \sim 50,000$) of the data [25], this translates into $\sigma_R = 2.55 \text{ km/s}$. In order to improve the fit to the data at small scales, the resolution width would have to be overestimated by 30%–40%. Typically the resolution is estimated to $\sim 10\%$, and a factor three to four seems unlikely to be an explanation for excess small-scale power.

Similarly, the contribution from contaminating metal absorption in the Lyman- α forest has been studied in both low-resolution [61,71] and high resolution data [19,25,69]. The contamination can be split into two main groups: (a) metals that have a rest-frame wavelength transition close to the Lyman- α line (e.g., SiIII) [48,61] and (b) metals situated at a lower redshift and associated with either IGM

or circumgalactic medium contributions [48,72]. The first group (a), imprints an oscillatory feature on the flux power spectrum. The frequency of this feature increases with scale and is typically averaged over many periods in measurements of the high- k flux power spectra, leaving distinct features observable only at low k , $k < 0.01 \text{ km}^{-1} \text{ s}$. The second group (b) is important at all redshifts and scales and due to large differences in redshift can be subtracted statistically by measuring the flux power spectrum on the red side of the Lyman- α emission line. The metal flux power spectra are typically dominated by CIV and SiIV doublets [73] and are smaller than the Lyman- α flux power spectrum by one or two orders of magnitude. The amplitude of the metal power spectrum would need to be larger by a factor of 5–10, in order to have an impact on Lyman- α flux power spectrum parameter estimation. While some studies suggest that the red side metal power spectrum captures only about half of the contaminated metal content in the Lyman- α forest [74], it is difficult to argue on observational grounds that the small-scale enhancement of the small-scale power spectrum due to metals could significantly affect our analysis.

The flux noise estimation has received somewhat less attention as a source of systematic uncertainty in high resolution and high signal-to-noise quasar spectra. It plays a crucial role in the low signal-to-noise spectra of large surveys (e.g., [61]). The flux power spectrum of the noise is 2–3 orders of magnitude lower than the Lyman- α forest flux power spectrum signal in medium ($S/N > 20$) and high ($S/N > 40$) quality data (e.g., [19]). The signal decays exponentially towards higher wave numbers, suggesting that the noise flux power quickly becomes a bigger contribution to the signal as the analysis is pushed towards higher k .

The analysis of [25] estimated the noise power on a per quasar sightline basis in $20 h^{-1} \text{ Mpc}$ sections. This was achieved by measuring the raw or total flux power spectrum in each section of the Lyman- α forest and estimating the asymptote level at high k . This method relies on the assumption that the noise power is white—an assumption that is largely validated in other studies (e.g., [19,60,61])—and that it dominates at high wave numbers. The method contends with several challenges, from a noisy estimation of the measured noise power spectrum in individual 20 cMpc/h sections, to the fact that the asymptote levels at high k will also include the metal contamination, as well as the very signal that one wishes to measure.

A careful analysis of uncertainty propagation is warranted, especially for a signal dominated by the highest wave numbers such as is the case in this WDM study. Figure 12 shows the probability distribution of the noise power estimates from the individual 20 cMpc/h sightline sections, in each of the three redshift bins. The vertical black lines indicate the effective average P_{noise} assumed in the analysis of [25]. This is simply a result of averaging the




FIG. 12. The probability distribution of the measured noise power spectra from [25] in each of the redshift bins. The measured distribution is reasonably well approximated by a log-normal distribution at each redshift, shown as a solid black line. For comparison, we also show a normal distribution with mean and variance computed from the first two moments of the measured distribution. The distributions are fairly broad; however, the inclusion or removal of tails beyond the range of measured P_{noise} has a negligible effect.

difference of raw and noise power per section over all the sightlines in a given redshift bin. As the average was not weighted by the signal to noise, the estimated average P_{noise} is simply the mean of the distribution. One immediate conclusion of Fig. 12 is that the distribution of the noise power is not Gaussian around the mean, with the bulk of the distribution typically peaking at lower than average P_{noise} values. The distributions at each redshift are also relatively broad. The mean of the distributions, $\langle P_{\text{noise}} \rangle$ are 0.08, 0.1, and 0.12 for $z = 4.2$, 4.6, and 5.0, respectively. This corresponds to roughly 5% of the total power at the highest wave number in the data. We approximate the P_{noise} distributions with a log-normal model with the min/max range of the measured values (solid black lines in each of the redshift bins in Fig. 12).

The noise power spectrum distribution in Fig. 12 is dominated primarily by the distribution of signal to noise in the data, as well as the mean transmission variations among the $20 h^{-1} \text{Mpc}$ segments of the Lyman- α forest. This has been verified in mock data with pathlength and redshift ranges of observed quasars reported in [25]. The methodology of estimating the noise power asymptote within each $20 h^{-1} \text{Mpc}$ segment is noisy; it is therefore unlikely that the uncertainty on the noise power estimation exceeds the width of the distributions in Fig. 12. As such, we use the distribution of the noise power as a conservative prior.

In order to assess the potential impact of noise misestimation in the data, we add a constant term $A_{\text{noise}}(z)$ to the model of the 1D flux power spectrum. This term is scale independent, but is modeled separately for each redshift bin. Since the mean of the noise flux power spectra distributions were already subtracted from the data, this constant term measures the deviation of the noise flux power from this mean value. In the data analysis step, the noise is subtracted from the raw power before the resolution correction of the instrument is deconvolved. The theoretical 1D flux power spectrum model is modified as follows:

$$P_{F,1D}^{\text{tot}}(k, z) = P_{F,1D}^{\text{Ly}\alpha}(k, z) + A_{\text{noise}}(z) \frac{\langle P_{\text{noise}} \rangle(z)}{W^2(k)}, \quad (4)$$

where $P_{F,1D}^{\text{Ly}\alpha}$ is the Lyman- α flux power spectrum as given by the emulator, $\langle P_{\text{noise}} \rangle(z)$ are the means of the P_{noise} distribution in each of the redshift bins, and $W^2(k)$ is the instrumental correction due to resolution and pixel size (following [25], we use a pixel size of 2.5 km/s and a FWHM resolution of 6.0 km/s using top-hat and Gaussian kernels for the two corrections, respectively).

Figure 13 shows the results of the analysis where three A_{noise} parameters were added to the theoretical model (one for each redshift bin), and the parameters' priors were assumed to be given by the approximate log-normal model of the P_{noise} distribution. The resulting WDM mass constraint is slightly weakened, and the lower WDM mass bound is $m_{\text{WDM}} > 3.91 \text{ keV}$. The thermal constraints are significantly degraded along the $u_0 - T_0$ degeneracy axis. This is because at every individual redshift, the noise parameter A_{noise} strongly correlates (anticorrelates) with the IGM temperature (cumulative heat injection), whereas the correlation with the WDM mass parameter is weaker. The marginalized mean of the posteriors (and their best-fit) values of A_{noise} are $0.74^{+0.49}_{-0.49}$ (0.24), $1.12^{+0.49}_{-0.29}$ (1.73), $0.87^{+0.56}_{-0.15}$ (1.38). The data prefer values of $A_{\text{noise}} > 0$, implying noise was underestimated. The best-fit values also show a slight increase with redshift, suggesting that the effect was larger for higher redshift quasar spectra. The typical values of A_{noise} are of the order of unity, suggesting that the noise subtraction of the data performed by [25] may be incomplete. The sensitivity of thermal parameters to this relatively small noise contribution is quite large, possibly implying that measurements of the IGM temperature and reionization are very sensitive to noise subtraction in the data. There is also sensitivity of the WDM mass constraints




FIG. 13. Effect of marginalizing over the noise uncertainty distribution. The two panels show 2D posterior distributions for redshift $z = 4.2$ in the plane of temperature and heat injection (left) and warm dark matter mass and heat injection (right). As the noise affects the amount of small-scale power, it effectively removes the information from those scales, which leads to poorer constraints on thermal parameter as well as warm dark matter mass. The noise distribution is marginalized over the measured distribution from [25]. However, the results remain largely unchanged if the shape of the distribution was changed to a normal distribution with a standard deviation of 10% of the measured power. This can also be included at the level of the covariance matrix.

to this effect, although somewhat reduced compared to the thermal parameters.

The sampling of the P_{noise} distribution is relatively sparse, measured in $20 h^{-1} \text{Mpc}$ sections in only 15 quasar sightlines in each redshift bin. This marks a significant improvement on previous measurements, but is nonetheless sensitive to sample variance. To understand the sensitivity of the conclusions of this analysis step, we modify the prior choice to be a Gaussian distribution, with the mean and the standard deviation estimated from the average and variance of the samples in each redshift bin. Since the posteriors of A_{noise} for the highest two redshifts are dominated by the upper limit on the prior range, we further allow this Gaussian prior to have no min/max limits other than the physical requirement that noise is larger or equal to zero $P_{\text{noise}} \geq 0$. This allows for a tail of the A_{noise} distribution to arbitrarily large values. The mean of the posterior distributions of A_{noise} parameters however do not move significantly. The main difference is the tail of the posteriors towards higher A_{values} . Note that the WDM constraint changes less than 1%.

If the signal at the high- k end of [25] data is indeed due to undersubtracted noise power, then the situation should improve with better and more data. If the flux uncertainties are dominated by the read-out noise component, then increasing the signal-to-noise (S/N) ratio of individual quasar sightlines quickly reduces the level of noise power ($P_{\text{noise}} \propto (S/N)^{-2}$) [75]. Future studies should thus suffer less from the impact of instrumental effects on the flux

power spectrum measurements, allowing for exploration of data to high k_{max} .

I. Small-scale data cuts

In order to facilitate a more direct comparison between the new analysis using [25] data and previous analyses using high redshift HIRES/MIKE data [21,22], a consistency check can be performed by limiting the new analysis to the same scale cuts ($k < 0.1 \text{ km}^{-1} \text{ s}$). We further compare such an analysis to [21,22] that used $k_{\text{max}} = 0.088 \text{ km}^{-1} \text{ s}$, and a similar redshift range. The previous analysis extended to $z = 5.4$. However, the flux power spectrum uncertainty at this highest redshift was considerably larger, and most of the constraining power came from the $z = 4.2, 4.6, 5.0$ redshift bins, which are also the ones used in this study. Furthermore, we limit the comparison to the thermal history priors where T_0 is varied independently in each redshift bin. In [22], the resulting lower bound on the WDM mass was $\sim 2.1 \text{ keV}$ (2σ) [MIKE/HIRES Iršič+17 + wide thermal prior (Fig. 5); [22]].

A similar test was performed in [24], where the reported value on the WDM mass bound sits at 3.6 keV . The same scale cuts were used, using the quasar spectra data of [25]. However, somewhat different thermal history priors were applied.

Figure 14 shows the result of small-scale data cuts in this analysis. Using the same scale cuts and treatment of the thermal history with the new data improves the constraint to




FIG. 14. Effect of varying $k_{\max}$ of the analysis. The two panels show 2D posterior distributions for redshift $z = 4.2$ in the plane of temperature and heat injection (left) and warm dark matter mass and heat injection (right). In the thermal parameter space (left), limiting the analysis to the value of $k_{\max} < 0.1 \text{ km}^{-1} \text{ s}$ —similar to previous analyses using Ly α forest data—has a similar effect to marginalizing over noise or thermal dependence of the resolution correction. The posterior stretches in the degeneracy direction of $u_0 - T_0$. Similarly, the warm dark matter mass constraints relax as more thermal support can accommodate the data.

$m_{\text{WDM}} > 4.09 \text{ keV}$ (2σ) [$k_{\max} < 0.1 \text{ km}^{-1} \text{ s}$; this work (Fig. 5)]. The right-hand side panel of Fig. 14 illustrates that imposing conservative scale cuts reduces the sensitivity to m_{WDM}^{-1} and pressure smoothing scale as probed by the injected heat u_0 . This reduced sensitivity to the pressure smoothing can be understood in the thermal parameter space (left panel of Fig. 14) as expanding of the posterior along the $u_0 - T_0$ degeneracy axis. The posterior in this parameter space also shifts by $\sim 0.2\sigma$ along the positive $u_0 - T_0$ relation that exists in hydrodynamical simulations. The shift, however, is small and can at least in part be attributed to reaching the corner of the priors at low u_0 and low T_0 values.

In the case of conservative scale cuts ($k_{\max} < 0.1 \text{ km}^{-1} \text{ s}$), the best fit improves over the fit to all the data, with the $\chi^2/\text{d.o.f.} = 10.2/20$ (see Table II). The fit prefers slightly warmer temperatures of the IGM ($T_0(z = 4.6) \sim 8400 \text{ K}$) and slightly higher values of cumulative injected heat. The posterior of the u_0 parameter is very wide, however, suggesting that with conservative scale cuts the data are not sensitive to this parameter anymore. This has been observed in previous analysis using older data sets that did not extend beyond $k \sim 0.1 \text{ km}^{-1} \text{ s}$.

V. DISCUSSION AND CONCLUSIONS

This study presents new constraints on the free streaming of WDM using a simulation based likelihood and Bayesian analysis of the VLT/UVES and Keck/HIRES Lyman- α

forest flux power spectrum measurements of [25]. The new constraints of our fiducial analysis on the mass of a thermal relic WDM particle mass, $m_{\text{WDM}} > 5.7 \text{ keV}$, are the strongest to date. For the fixed shape of the WDM transfer function used in this study, the bound on the WDM particle mass translates into a wave number scale below which the matter power spectrum cannot drop by more than 5%, $k_{0.05} = 14.35 h^{-1} \text{ Mpc}$.

Comparing to the previous high redshift Lyman- α forest data from HIRES/MIKE [21], the new data comprise a larger number of quasar spectra in the range $4.2 < z < 5.0$ and are probing small scales up to a wave number of $k_{\max} = 0.19 \text{ km}^{-1} \text{ s}$ —almost a factor of two improvement. Limiting the analysis to the same $k_{\max}$ cuts, and thermal state priors, as in previous HIRES/MIKE analyses, we find the constraint to be $m_{\text{WDM}} > 4.1 \text{ keV}$ (this work), compared to $m_{\text{WDM}} > 2.0 \text{ keV}$ (e.g., [22]). This factor two improvement on the bound on the WDM particle mass is consistent with the expected improvement of the statistical power of the high redshift Lyman- α forest data and WDM mass sensitivity at $0.1 \text{ km}^{-1} \text{ s}$ dominated by statistical uncertainty. This result is qualitatively similar to the recent analysis of [24] using the same scale cuts and a different thermal state parametrization.

Recent studies of [24,76] have found a preference for nonzero m_{WDM}^{-1} in their default analysis, indicating a preference for a WDM cosmology. This warrants further study and rigorous tests on both the data and theory side. Our findings here lead us to suggest that one possibility is

that the nonzero preference in the WDM parameter space is a result of the restricted variations in the thermal parameters. For example, the analysis of [76] assumes a thermal history with very low cumulative injected heat and therefore a small amount of pressure smoothing. Limiting the amount of pressure smoothing can be of interest in specific applications, but in terms of a WDM particle mass constraint, a thorough marginalization over the parameter space should be more robust. The analysis of [24] follows a similar simulation setup and parameter space variation as in this work, except for two main differences: (a) the variations in the redshift of hydrogen reionization were much narrower than explored in this work, and (b) parametrization in the $z_{\text{rei}} - T_0$ as opposed to $u_0 - T_0$ plane only allowed for more restricted thermal histories.

The HIRES/UVES data of [25] have also recently been used to provide constraints on a slightly different class of dark matter models—ultralight axion dark matter [15]. While the transfer functions of the two dark matter models are different enough that a direct comparison is nontrivial, the results of [15] suggest a strong bound on the thermal WDM mass, while at the same time recovering a hotter thermal history compared to both results in the literature [24,25] and the results of this study. A more thorough investigation is required, but a major difference in the simulation setup of [15] is the initial condition generation with MP-Gadget [77] which uses glass initial conditions for the gas component. This introduces spurious small-scale power [78].

Additionally, [15] uses a Lyman- α spectral extraction code `fake_spectra` [79] that uses a different optical depth assignment scheme that leads to additional enhancement of small-scale power in the 1D Lyman- α forest flux power spectrum [80]. These differences in the simulated flux power spectrum suggest further investigation is required for the comparison with the ultralight axion constraints of [15].

Further improvement in the WDM constraint comes from the smallest scales, $k > 0.1 \text{ km}^{-1} \text{ s}$. The sensitivity to the WDM mass is increased at smaller scales, resulting in potentially stronger constraining power. However, the regime of $k > 0.1 \text{ km}^{-1} \text{ s}$ is also more sensitive to observational and modeling systematics. In this study, we have reviewed several aspects of the observational and instrumental systematics, of which the observational flux noise subtraction in the Lyman- α forest flux power spectrum is potentially the most likely to affect the results. An average of a factor of two increase in the noise power (or 40% increase in the level of noise) would on its own explain the small-scale signal observed, with no additional cosmological information beyond $k_{\text{max}} > 0.1 \text{ km}^{-1} \text{ s}$. While this appears not very likely, it illustrates the need for improved treatment of the noise power at smallest scales in future Lyman- α forest data analyses.

The thermal history priors have previously been identified as the dominant source of modeling systematics in the Lyman- α forest flux power spectrum. In this study, we revisited this, by exploring thermal priors motivated by different assumptions: a prior in the plane of IGM temperature and cumulative injected heat that envelopes physically motivated simulations consistent with the still rather weak constraints on the evolution of the neutral hydrogen fraction during the epoch of reionization, or a simple prior on the IGM temperature as interpolated from the measurements of the IGM temperature at $z < 4.2$ and $z > 5.0$. This was possible due to the improved range of simulations as well as postprocessing techniques to expand on the number of models. While the posterior distributions are indeed sensitive to the choice of these priors, the WDM mass only changes by 2%. This suggests that reasonable thermal priors lead to a stable and robust constraint on the WDM particle mass. We further point out that not imposing any thermal priors leads to a stronger and not weaker bound on the WDM particle mass.

With the data extending to $k > 0.1 \text{ km}^{-1} \text{ s}$, we have also identified a new source of modeling systematics—the gas peculiar velocity field as modified by inhomogeneous reionization. The effect of peculiar velocities is present, although different in amplitude, in both homogeneous and inhomogeneous models of reionization. The peculiar velocity field induces a knee in the flux power spectrum, that appears sensitive to the cumulative injected heat, indicating that the timing and process of reionization are an important factor in the peculiar velocity structure. This high- k regime of the models is also sensitive to the numerical mass resolution of the simulations (at a level of up to 20%). The peculiar velocity field structure is sensitive to the mass resolution at the level of as much as 50%, resulting in thermal history dependent mass resolution corrections. All of these statements result in a similar effect on the WDM mass inference, weakening the constraint to $m_{\text{WDM}} > 4.44(5.10) \text{ keV}$ at 95% CL for including the mass resolution and inhomogeneous reionization, respectively. Combining inhomogeneous reionization and mass resolution corrections together with observationally motivated thermal prior results in a WDM constraint that is not very different from our fiducial analysis, while at the same time preferring a slightly hotter IGM and reionization histories that end later. These statements are, however, somewhat model and simulation dependent, and indicate that further work into the origin and impact of the small-scale peculiar velocity structure is required. However, an important result for the particle astrophysics modeling is that WDM particle masses of 2.5 keV are ruled out at more than 5σ , and 3 keV at more than 3σ , for any of the analysis choices presented in this paper. In fact, the 3 keV is ruled out at 5σ for any of the reasonable choices of thermal priors (e.g., a T_0 prior).

We summarize the main conclusion points as follows:

- (i) Our fiducial analysis leads to improved WDM mass constraints from high redshift quasar spectra of $m_{\text{WDM}} > 5.7$ keV at 95% CL.
- (ii) Using small-scale data cuts, limiting the analysis to $k_{\text{max}} < 0.1 \text{ km}^{-1} \text{ s}$ results in a WDM constrain of $m_{\text{WDM}} > 4.1$ keV at 95% CL, a factor of two stronger constraint than previously published for the same choice of thermal priors and redshift range of the data [21,22].
- (iii) The 50% higher WDM constraint coming from small-scale $k > 0.1 \text{ km}^{-1} \text{ s}$ has been explored with a variety of checks for instrumental systematics. We find that the flux noise may be underestimated by 40% in the data, reducing the constraining power at $k > 0.1 \text{ km}^{-1} \text{ s}$. It should be possible to mitigate this in future surveys by a careful study of the instrumental noise, as well as by obtaining higher signal-to-noise spectra.
- (iv) The modeling uncertainties on the small-scale peculiar velocity structure can weaken the constraining power on the WDM mass by as much as 25%. The effect of thermal history and inhomogeneous nature of reionization on the peculiar velocity fields of the baryonic gas is still poorly explored and needs further study.

The Lyman- α forest data continue to push the frontier on astrophysical constraints on the CDM paradigm. The new constraints on the free streaming of dark matter in this study both improve on existing constraints and demonstrate that a larger number of quasar sightlines should translate into strong improvements on the WDM particle mass bound. As has been done in previous studies [18,21,22], the high redshift Lyman- α forest data can be further combined with the low redshift ($z < 4.0$) flux power spectrum measurements in order to increase the redshift leverage arm and further push the constraints on the free streaming of dark matter. We leave such a study for future work.

ACKNOWLEDGMENTS

The authors would like to thank Matthew McQuinn, Simeon Bird, Steven Gratton, and Anson D’Aloisio for helpful conversations. V.I. acknowledges support by the Kavli Foundation. 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 No. ST/N000927/1 and No. ST/S000623/1. J. S. B. is supported by STFC consolidated Grants No. ST/T000171/1 and No. ST/X000982/1. G. D. B. is supported by the National Science Foundation through Grant No. AST-1751404. G. K. is partly supported by the Department of Atomic Energy (Government of India) research project with Project Identification Number RTI 4002 and by the Max Planck Society through a Max Planck Partner Group. The simulations used in this work were performed using the Joliot Curie supercomputer at the Tré 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 ([81]). 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 No. ST/P002307/1 and No. ST/R002452/1 and STFC operations Grant No. 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 No. ST/P002293/1 and No. ST/R002371/1, Durham University, and STFC operations Grant No. ST/R000832/1. DiRAC is part of the National e-Infrastructure.

-
- [1] A. A. Meiksin, *Rev. Mod. Phys.* **81**, 1405 (2009).
 - [2] M. McQuinn, *Annu. Rev. Astron. Astrophys.* **54**, 313 (2016).
 - [3] É. Aubourg, S. Bailey, J. E. Bautista, F. Beutler, V. Bhardwaj, D. Bizyaev, M. Blanton, M. Blomqvist, A. S. Bolton, J. Bovy *et al.*, *Phys. Rev. D* **92**, 123516 (2015).
 - [4] V. K. Narayanan, D. N. Spergel, R. Davé, and C.-P. Ma, *Astrophys. J. Lett.* **543**, L103 (2000).
 - [5] U. Seljak, A. Slosar, and P. McDonald, *J. Cosmol. Astropart. Phys.* **10** (2006) 014.
 - [6] 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).
 - [7] 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).
 - [8] W. Enzi, R. Murgia, O. Newton, S. Vegetti, C. Frenk, M. Viel, M. Cautun, C. D. Fassnacht, M. Auger, G. Despali *et al.*, *Mon. Not. R. Astron. Soc.* **506**, 5848 (2021).
 - [9] J. W. Hsueh, W. Enzi, S. Vegetti, M. W. Auger, C. D. Fassnacht, G. Despali, L. V. E. Koopmans, and J. P. McKean, *Mon. Not. R. Astron. Soc.* **492**, 3047 (2020).
 - [10] P. Colín, V. Avila-Reese, and O. Valenzuela, *Astrophys. J.* **542**, 622 (2000).
 - [11] P. Bode, J. P. Ostriker, and N. Turok, *Astrophys. J.* **556**, 93 (2001).

- [12] D. H. Weinberg, J. S. Bullock, F. Governato, R. Kuzio de Naray, and A. H. G. Peter, *Proc. Natl. Acad. Sci. U.S.A.* **112**, 12249 (2015).
- [13] R. Adhikari, M. Agostini, N. A. Ky, T. Araki, M. Archidiacono, M. Bahr, J. Baur, J. Behrens, F. Bezrukov, P. S. Bhupal Dev *et al.*, *J. Cosmol. Astropart. Phys.* **01** (2017) 025.
- [14] K. K. Boddy, M. Lisanti, S. D. McDermott, N. L. Rodd, C. Weniger, Y. Ali-Haïmoud, M. Buschmann, I. Cholis, D. Croon, A. L. Erickcek *et al.*, *J. High Energy Astrophys.* **35**, 112 (2022).
- [15] K. K. Rogers and H. V. Peiris, *Phys. Rev. Lett.* **126**, 071302 (2021).
- [16] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, and G. D. Becker, *Phys. Rev. Lett.* **119**, 031302 (2017).
- [17] J. Baur, N. Palanque-Delabrouille, C. Yèche, A. Boyarsky, O. Ruchayskiy, É. Armengaud, and J. Lesgourgues, *J. Cosmol. Astropart. Phys.* **12** (2017) 013.
- [18] N. Palanque-Delabrouille, C. Yèche, N. Schöneberg, J. Lesgourgues, M. Walther, S. Chabanier, and E. Armengaud, *J. Cosmol. Astropart. Phys.* **04** (2020) 038.
- [19] 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).
- [20] M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese, and A. Riotto, *Phys. Rev. D* **71**, 063534 (2005).
- [21] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, *Phys. Rev. D* **88**, 043502 (2013).
- [22] 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).
- [23] A. Garzilli, A. Magalich, O. Ruchayskiy, and A. Boyarsky, *Mon. Not. R. Astron. Soc.* **502**, 2356 (2021).
- [24] B. Villaseñor, B. Robertson, P. Madau, and E. Schneider, *Phys. Rev. D* **108**, 023502 (2023).
- [25] E. Boera, G. D. Becker, J. S. Bolton, and F. Nasir, *Astrophys. J.* **872**, 101 (2019).
- [26] H. Dekker, S. D'Odorico, A. Kaufer, B. Delabre, and H. Kotzlowski, in *Optical and IR Telescope Instrumentation and Detectors*, edited by M. Iye and A. F. Moorwood, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 4008 (SPIE, Bellingham, 2000), pp. 534–545.
- [27] S. S. Vogt, S. L. Allen, B. C. Bigelow, L. Bresee, B. Brown, T. Cantrall, A. Conrad, M. Couture, C. Delaney, H. W. Epps *et al.*, in *Instrumentation in Astronomy VIII*, edited by D. L. Crawford and E. R. Craine, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 2198 (SPIE, Bellingham, 1994), p. 362.
- [28] M. Garny, T. Konstandin, L. Sagunski, and M. Viel, *J. Cosmol. Astropart. Phys.* **03** (2021) 049.
- [29] V. Iršič and M. McQuinn, *J. Cosmol. Astropart. Phys.* **04** (2018) 026.
- [30] N. Y. Gnedin and L. Hui, *Mon. Not. R. Astron. Soc.* **296**, 44 (1998).
- [31] 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).
- [32] V. Springel, *Mon. Not. R. Astron. Soc.* **364**, 1105 (2005).
- [33] J. S. Bolton and G. D. Becker, *Mon. Not. R. Astron. Soc.* **398**, L26 (2009).
- [34] Z. Lukić, C. W. Stark, P. Nugent, M. White, A. A. Meiksin, and A. Almgren, *Mon. Not. R. Astron. Soc.* **446**, 3697 (2015).
- [35] 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).
- [36] C. C. Doughty, J. F. Hennawi, F. B. Davies, Z. Lukić, and J. Oñorbe, *Mon. Not. R. Astron. Soc.* **525**, 3790 (2023).
- [37] M. Viel, M. G. Haehnelt, and V. Springel, *Mon. Not. R. Astron. Soc.* **354**, 684 (2004).
- [38] 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).
- [39] E. Puchwein, F. Haardt, M. G. Haehnelt, and P. Madau, *Mon. Not. R. Astron. Soc.* **485**, 47 (2019).
- [40] M. Viel, J. Schaye, and C. M. Booth, *Mon. Not. R. Astron. Soc.* **429**, 1734 (2013).
- [41] F. Nasir, J. S. Bolton, and G. D. Becker, *Mon. Not. R. Astron. Soc.* **463**, 2335 (2016).
- [42] E. Puchwein, J. S. Bolton, M. G. Haehnelt, P. Madau, G. D. Becker, and F. Haardt, *Mon. Not. R. Astron. Soc.* **450**, 4081 (2015).
- [43] G. D. Becker, J. S. Bolton, M. G. Haehnelt, and W. L. W. Sargent, *Mon. Not. R. Astron. Soc.* **410**, 1096 (2011).
- [44] 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).
- [45] J. S. Bolton, M. G. Haehnelt, M. Viel, and V. Springel, *Mon. Not. R. Astron. Soc.* **357**, 1178 (2005).
- [46] G. D. Becker, P. C. Hewett, G. Worseck, and J. X. Prochaska, *Mon. Not. R. Astron. Soc.* **430**, 2067 (2013).
- [47] L. Hui and N. Y. Gnedin, *Mon. Not. R. Astron. Soc.* **292**, 27 (1997).
- [48] 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).
- [49] 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).
- [50] P. Gaikwad, R. Srianand, M. G. Haehnelt, and T. R. Choudhury, *Mon. Not. R. Astron. Soc.* **506**, 4389 (2021).
- [51] G. Kulkarni, L. C. Keating, M. G. Haehnelt, S. E. I. Bosman, E. Puchwein, J. Chardin, and D. Aubert, *Mon. Not. R. Astron. Soc.* **485**, L24 (2019).
- [52] 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).
- [53] 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).
- [54] R. de Belsunce, S. Gratton, W. Coulton, and G. Efstathiou, *Mon. Not. R. Astron. Soc.* **507**, 1072 (2021).
- [55] L. Hui, *Astrophys. J.* **516**, 519 (1999).

- [56] R. Adhikari, M. Agostini, N. A. Ky, T. Araki, M. Archidiacono, M. Bahr, J. Baur, J. Behrens, F. Bezrukov, P. S. Bhupal Dev *et al.*, *J. Cosmol. Astropart. Phys.* **01** (2017) 025.
- [57] W. Enzi, R. Murgia, O. Newton, S. Vegetti, C. Frenk, M. Viel, M. Cautun, C. D. Fassnacht, M. Auger, G. Despali *et al.*, *Mon. Not. R. Astron. Soc.* **506**, 5848 (2021).
- [58] D. Gilman, S. Birrer, A. Nierenberg, T. Treu, X. Du, and A. Benson, *Mon. Not. R. Astron. Soc.* **491**, 6077 (2020).
- [59] N. Banik, G. Bertone, J. Bovy, and N. Bozorgnia, *J. Cosmol. Astropart. Phys.* **07** (2018) 061.
- [60] C. Yèche, N. Palanque-Delabrouille, J. Baur, and H. du Mas des Bourboux, *J. Cosmol. Astropart. Phys.* **06** (2017) 047.
- [61] N. Palanque-Delabrouille, C. Yèche, A. Borde, J.-M. Le Goff, G. Rossi, M. Viel, É. Aubourg, S. Bailey, J. Bautista, M. Blomqvist *et al.*, *Astron. Astrophys.* **559**, A85 (2013).
- [62] B. Villasenor, B. Robertson, P. Madau, and E. Schneider, *Astrophys. J.* **933**, 59 (2022).
- [63] L. C. Keating, L. H. Weinberger, G. Kulkarni, M. G. Haehnelt, J. Chardin, and D. Aubert, *Mon. Not. R. Astron. Soc.* **491**, 1736 (2020).
- [64] J. Oñorbe, F. B. Davies, Z. Lukić, J. F. Hennawi, and D. Sorini, *Mon. Not. R. Astron. Soc.* **486**, 4075 (2019).
- [65] X. Wu, M. McQuinn, and D. Eisenstein, *J. Cosmol. Astropart. Phys.* **02** (2021) 042.
- [66] C. Cain, A. D'Aloisio, N. Gangolli, and G. D. Becker, *Astrophys. J. Lett.* **917**, L37 (2021).
- [67] M. McQuinn, *Annu. Rev. Astron. Astrophys.* **54**, 313 (2016).
- [68] 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).
- [69] B. Wilson, V. Iršič, and M. McQuinn, *Mon. Not. R. Astron. Soc.* **509**, 2423 (2022).
- [70] M. Walther, J. F. Hennawi, H. Hiss, J. Oñorbe, K.-G. Lee, A. Rorai, and J. O'Meara, *Astrophys. J.* **852**, 22 (2018).
- [71] V. Iršič and A. Slosar, *Phys. Rev. D* **89**, 107301 (2014).
- [72] S. S. Tie, J. F. Hennawi, K. Kakiichi, and S. E. I. Bosman, *Mon. Not. R. Astron. Soc.* **515**, 3656 (2022).
- [73] N. G. Karaçaylı, P. Martini, D. H. Weinberg, V. Iršič, J. Aguilar, S. Ahlen, D. Brooks, A. de la Macorra, A. Font-Ribera, S. Gontcho A Gontcho *et al.*, *Mon. Not. R. Astron. Soc.* **522**, 5980 (2023).
- [74] A. Day, D. Tytler, and B. Kambalur, *Mon. Not. R. Astron. Soc.* **489**, 2536 (2019).
- [75] M. McQuinn and M. White, *Mon. Not. R. Astron. Soc.* **415**, 2257 (2011).
- [76] A. Garzilli, A. Magalich, T. Theuns, C. S. Frenk, C. Weniger, O. Ruchayskiy, and A. Boyarsky, *Mon. Not. R. Astron. Soc.* **489**, 3456 (2019).
- [77] Y. Feng, S. Bird, L. Anderson, A. Font-Ribera, and C. Pedersen, MP-Gadget/MP-Gadget: A tag for getting a DOI, [10.5281/zenodo.1451799](https://doi.org/10.5281/zenodo.1451799) (2018).
- [78] N. K. Khan, G. Kulkarni, J. S. Bolton, M. G. Haehnelt, V. Iršič, and S. Asthana, [arXiv:2310.07767](https://arxiv.org/abs/2310.07767).
- [79] S. Bird, FSFE: Fake Spectra Flux Extractor, Astrophysics Source Code Library, record ascl:1710.012 (2017).
- [80] S. Bird (private communication).
- [81] www.dirac.ac.uk.