

New constraints on warm dark matter from the Lyman- α forest power spectrum

Bruno Villasenor¹,* Brant Robertson², and Piero Madau¹

*Department of Astronomy and Astrophysics, University of California,
Santa Cruz, 1156 High Street, Santa Cruz, California 95064, USA*

Evan Schneider³

*Department of Physics and Astronomy and Pittsburgh Particle Physics, Astrophysics,
and Cosmology Center (PITT PACC), University of Pittsburgh,
Pittsburgh, Pennsylvania 15260, USA*



(Received 28 September 2022; accepted 9 June 2023; published 5 July 2023)

The forest of Lyman- α absorption lines detected in the spectra of distant quasars encodes information on the nature and properties of dark matter and the thermodynamics of diffuse baryonic material. Its main observable—the 1D flux power spectrum (FPS)—should exhibit a suppression on small scales and an enhancement on large scales in warm dark matter (WDM) cosmologies compared to standard Λ CDM. Here, we present an unprecedented suite of 1080 high-resolution cosmological hydrodynamical simulations run with the graphics processing unit-accelerated code CHOLLA to study the evolution of the Lyman- α forest under a wide range of physically motivated gas thermal histories along with different free-streaming lengths of WDM thermal relics in the early Universe. A statistical comparison of synthetic data with the forest FPS measured down to the smallest velocity scales ever probed at redshifts $4.0 \lesssim z \lesssim 5.2$ [E. Boera *et al.*, Revealing reionization with the thermal history of the intergalactic medium: New constraints from the Ly α flux power spectrum, *Astrophys. J.* **872**, 101 (2019)] yields a lower-limit $m_{\text{WDM}} > 3.1$ keV (95% C.L.) for the WDM particle mass and constrains the amplitude and spectrum of the photoheating and photoionizing background produced by star-forming galaxies and active galactic nuclei at these redshifts. Interestingly, our Bayesian inference analysis appears to weakly favor WDM models with a peak likelihood value at the thermal relic mass of $m_{\text{WDM}} = 4.5$ keV. We find that the suppression of the FPS from free-streaming saturates at $k \gtrsim 0.1$ s km⁻¹ because of peculiar velocity smearing, and this saturated suppression combined with a slightly lower gas temperature provides a moderately better fit to the observed small-scale FPS for WDM cosmologies.

DOI: [10.1103/PhysRevD.108.023502](https://doi.org/10.1103/PhysRevD.108.023502)

I. INTRODUCTION

The Λ -cold dark matter (Λ CDM) cosmological paradigm has been immensely successful at matching across cosmic time observations spanning physical scales from the horizon length [1] all the way down to galaxy scales [e.g., [2–4]], and a vast menagerie of hypothetical non-baryonic elementary particles has been proposed to explain the astrophysical data [5]. Cold dark matter particles have negligible thermal velocities in the matter-dominated era and therefore clump gravitationally even on the smallest subgalactic scales, a property that has caused persistent challenges with observations of the abundances and density profiles of dwarf galaxies in the local Universe [e.g., [6]]. Warm dark matter (WDM) is a simple modification of CDM that has been proposed to suppress small-scale power

and alleviate some of these problems [7]. WDM particles of a few keV have significant intrinsic velocities from having decoupled as thermal relics or been produced by non-equilibrium processes, and one of the effects of their Mpc-range *free-streaming length* is to limit the gravitational collapse of structures and produce a cutoff in the matter power spectrum.

Intergalactic hydrogen at redshift $2 \lesssim z \lesssim 5$ scatters Lyman- α radiation and produces absorption features in the spectra of distant quasars. This “Lyman- α forest” is a powerful cosmological probe as it traces density fluctuations, the underlying dark matter weblike distribution (the “cosmic web”), and the ionization state and temperature of the intergalactic medium (IGM) at scales and redshifts that cannot be probed by any other observable [8–11]. The primary statistic derived from spectroscopic data is the 1D power spectrum (FPS) of the flux distribution in the forest—the Fourier transform of the fractional flux autocorrelation

*brvillas@ucsc.edu

function in velocity space—which arguably provides the best tool for distinguishing between CDM and WDM models [12–16]. The FPS is observationally accessible over a wide range of redshifts, involves the fundamental, well-known physics of the hydrogen atom, and is largely free from the uncertain baryonic physics (star formation, feedback, and metal cooling) that affects, e.g., the abundance of Milky Way satellites and the central density cores of dwarf galaxies [17]. Interpreting such observations requires expensive hydrodynamical simulations of the IGM that cover an extensive range of uncertain IGM photoionization and photoheating histories and consistently maintain high resolution throughout a statistically representative subvolume of the Universe, a traditional limiting factor in previous analyses [e.g., [13,18–21]]. At present, the tightest lower limit to the mass of a thermal WDM relic, $m_{\text{WDM}} > 1.9$ keV (95% C.L.), is obtained for highly conservative thermal histories of intergalactic gas and without a marginalization over the properties of the cosmic ionizing background [16].

In this paper, we revisit these constraints in light of recent measurements with the *Keck* and *VLT* telescopes of the Lyman- α FPS down to the smallest scales ever probed at redshifts $4 \lesssim z \lesssim 5$ [22], using a massive suite of 1080 high-resolution cosmological hydrodynamics simulations that are part of the CHOLLA IGM photoheating simulations (CHIPS) suite [23,24]. High-redshift observations provide better limits on the mass of a WDM particle, as free streaming becomes more prominent in velocity space and the effect of the nonlinear evolution of the matter density field—resulting in a transfer of power from large scales to small scales—is weaker at these epochs. There is also observational evidence for a local minimum of the temperature of the IGM—and therefore a corresponding minimum in the thermal cutoff length scale—in the redshift range $4 < z < 5$ [24,25], making this era an optimal epoch for deriving bounds on dark matter properties.

II. BASIC EQUATIONS AND CHARACTERISTIC SCALES

The hydrogen and helium photoheating \mathcal{H}_i and photoionization Γ_i rates per atom i depend on the intensity of a uniform UV background radiation field $J(\nu, z)$ as

$$\Gamma_i(z) = 4\pi \int_{\nu_i}^{\infty} \frac{J(\nu, z)}{h\nu} \sigma_i(\nu) d\nu, \quad (1)$$

$$\mathcal{H}_i(z) = 4\pi \int_{\nu_i}^{\infty} \frac{J(\nu, z)}{h\nu} (h\nu - h\nu_i) \sigma_i(\nu) d\nu, \quad (2)$$

where ν_i and $\sigma_i(\nu)$ are the threshold frequency and photoionization cross section, respectively. The total photoheating rate \mathcal{H} is summed over the species $i = \text{HI}, \text{He I},$ and He II each of proper number density n_i , $\mathcal{H} = \sum_i n_i \mathcal{H}_i$. These rates, together with radiative recombinations, adiabatic compression, Compton and expansion cooling, and

the gas density and peculiar velocity fields shaped by gravity and by pressure forces, fully determine the hydrogen Lyman- α absorption optical depth in velocity space v [26],

$$\tau_v = \sigma_\alpha c \int \frac{n_{\text{HI}}}{\sqrt{\pi} b H(z)} e^{-(v-u-v_p)^2/b^2} du. \quad (3)$$

Here, $\sigma_\alpha = \pi e^2 f_{12} / (m_e c \nu_\alpha)$, ν_α and f_{12} are the frequency and upward oscillator strength of the Lyman- α transition, $b = (2k_B T / m_{\text{H}})^{1/2}$ is the Doppler parameter, T the gas temperature, $H(z)$ the Hubble parameter, u and v_p are, respectively, the real-space coordinate (in km s^{-1}) and the peculiar velocity along the line of sight, and we have assumed a Gaussian profile. Denoting with \bar{z} some mean redshift of interest, e.g. the redshift of a simulation output or the average redshift of any given data subset, one can write

$$u(x) = \frac{H(\bar{z})}{(1 + \bar{z})} (x - \bar{x}), \quad (4)$$

where \bar{x} is the comoving position at which the redshift from Hubble expansion is exactly \bar{z} , and $H(\bar{z})$ is the Hubble expansion rate. Line center in the rest frame of intervening hydrogen occurs at velocity $v = u + v_p$, which is related to the observed frequency ν_0 by

$$\nu_\alpha = \nu_0 (1 + \bar{z}) \left(1 + \frac{v}{c} \right). \quad (5)$$

Because of peculiar velocities, a photon observed at ν_0 can have the same rest-frame frequency ν_α at more than one place in its trajectory from the quasar to the observer.

In practice, only a limited range of u values contributes to τ_v , and one can replace in Eq. (3) $H(z)$ with $H(\bar{z})$. We can then define the flux contrast $\delta_F(v) = F(v) / \langle F \rangle - 1$, where $F(v) = \exp(-\tau_v)$ is the flux at velocity v and $\langle F \rangle$ is the mean transmitted flux at a given redshift, and decompose each absorption spectrum into Fourier modes $\tilde{\delta}_F(k)$. Their variance as a function of the Fourier wave number $k = 2\pi/v$ is the FPS over some velocity interval Δv ,

$$P_F(k) = \Delta v \langle \tilde{\delta}_F(k)^2 \rangle, \quad (6)$$

which is commonly expressed in terms of the dimensionless quantity $\Delta_F^2(k) = k P_F(k) / \pi$.

Four physical effects (cf. [15]) act to erase small-scale power in the FPS:

- (1) *Doppler broadening caused by gas thermal velocities along the line of sight.* Assuming a Gaussian smoothing kernel in the $\tilde{\delta}_F(k)$ field of the form $\exp(-k^2 \sigma_{\text{th}}^2 / 2)$, where $\sigma_{\text{th}} = (k_B T / m_{\text{H}})^{1/2}$ is the broadening velocity scale, a cutoff in the FPS arises then at the proper wave number

$$k_{\text{th}} = \frac{1}{\sigma_{\text{th}}} = 0.11T_4^{-1/2} \text{ km}^{-1} \text{ s}, \quad (7)$$

where $T_4 \equiv T/10^4 \text{ K}$.

- (2) *Pressure support.* Pressure smooths gas density fluctuations in the spatial Fourier domain with a Gaussian kernel [27,28]

$$W(\tilde{k}, \tilde{k}_F) = \exp(-\tilde{k}^2/\tilde{k}_F^2), \quad (8)$$

where \tilde{k} is a comoving wave number,

$$\tilde{k}_F = g\tilde{k}_J = g\frac{a}{c_s} \sqrt{4\pi G\bar{\rho}} \quad (9)$$

is the ‘filtering scale’, and \tilde{k}_J is the wave number corresponding to the Jeans scale. Here, $c_s = [5k_B T/(3\mu m_H)]^{1/2}$ is the sound speed, μ is the mean molecular weight ($\mu = 0.61$ for an admixture of ionized hydrogen and singly ionized helium), $\bar{\rho}$ is the mean total (dark matter plus baryons) mass density of the Universe, and a is the scale factor. The g factor accounts for the fact that the filtering scale depends on the prior thermal history of the IGM, and not just the instantaneous Jeans scale. Typically, $g > 1$ after reionization [28].

Mapping the 3D filtering scale to a broadening velocity scale $\sigma_F = \sqrt{2}Ha/\tilde{k}_F$ along the line of sight results in a cutoff of the FPS at proper wave number

$$k_F = \frac{1}{\sigma_F} \simeq \frac{\sqrt{3}g}{2c_s} = 0.06gT_4^{-1/2} \text{ km}^{-1} \text{ s}, \quad (10)$$

where the second equality is valid at high enough redshifts. Under the physically-motivated assumptions that the gas temperature decays as $1/a$ after reionization completes around redshift $z \sim 6$, the factor $g \sim 4\text{--}10$ at the redshifts $4 \lesssim z \lesssim 5.2$ of interest here [27]. Pressure filtering, always lagging behind the growth of the Jeans length, is therefore subdominant compared to the thermal broadening of the absorption spectrum [22,29–31].

- (3) *Nonlinear peculiar velocities in the gas.* Peculiar motions smear power along line-of-sight (the ‘fingers-of-God’ effect) over a range of scales comparable with the one-dimensional velocity dispersion. A fit to the results of numerical simulations with a Gaussian in Fourier space yields [32]

$$\sigma_k = 0.25 \text{ km}^{-1} \text{ s} \left(\frac{k}{\text{km}^{-1} \text{ s}} \right)^{1/2}. \quad (11)$$

Setting $k = \sqrt{2}\sigma_k$, a cutoff in the FPS arises at proper wave number

$$k_v = 0.13 \text{ km}^{-1} \text{ s}. \quad (12)$$

- (4) *Dark matter free streaming.* This effect can be described by the wave number, $\tilde{k}_{1/2}$, at which the WDM matter power spectrum is suppressed relative to the CDM case by a factor of 2 [33],

$$\tilde{k}_{1/2} = 4.25 \left(\frac{m_{\text{WDM}}}{1 \text{ keV}} \right)^{1.11} \text{ Mpc}^{-1} \quad (13)$$

for a standard thermal relic. The conversion to a proper wave number in inverse km s^{-1} , $k_{\text{FS}} = \tilde{k}_{1/2}/(aH)$, yields a cutoff in the FPS at

$$k_{\text{FS}} = 0.05 \left(\frac{m_{\text{WDM}}}{1 \text{ keV}} \right)^{1.11} \left(\frac{6}{1+z} \right)^{1/2} \text{ km}^{-1} \text{ s}. \quad (14)$$

Note how the free-streaming cutoff moves toward larger scales (smaller k) at higher redshift.

The numerical values of k_{th} , k_F , k_v , and k_{FS} above suggest that free streaming, thermal and velocity broadening effects all set in at similar scales. Thus, comparing the observations of the FPS to a large set of models that accurately depict the suppression on the FPS in the non-linear regime associated to each of these mechanisms is critical to constrain the WDM particle mass. In the following section, we present our grid of simulations that evolve a total of 1080 models, which simultaneously vary the impact of free streaming, Doppler broadening, and pressure smoothing. A comparison between our grid of models and the observed FPS will provide new constraints on the mass of the dark matter particle that, for the first time, marginalize over all the important confounding physical effects.

III. SIMULATIONS

The simulations used in this work were performed using the GPU-native MPI-parallelized code CHOLLA [34], and evolve the equations of hydrodynamics on a uniform Cartesian grid while simultaneously tracking the nonequilibrium ionization states of hydrogen and helium using the GRACKLE library [35]. A spatially uniform, time-dependent UV radiation background was assumed in the form of redshift-dependent photoionization and photoheating rates per ion, Γ_i and H_i , for the species H I, He I, and He II. The initial conditions at $z = 100$ were generated using the MUSIC code [36] for a flat cosmology with parameters $h = 0.6766$, $\Omega_M = 0.3111$, $\Omega_\Lambda = 0.6889$, $\Omega_b = 0.0497$, $\sigma_8 = 0.8102$, and $n_s = 0.9665$, consistent with constraints from Planck data [1]. The initial conditions for all runs were generated from identical random number




FIG. 1. Impact of particle free streaming on baryonic structures at $0 < z < 5.2$. The evolution of the gas density from a slice through the IGM was obtained from a subset of 8 CHIPS simulations where the mass of the warm dark matter particle m_{WDM} was increased from 0.3 keV to ∞ . All simulations assume the same gas thermal history from the best-fit model presented in [24]. Due to thermal pressure, the gas distribution is smoothed relative to the dark matter.

seeds to preserve the same amplitude and phase for all initial Fourier modes across the simulation suite. The volume and numerical size of the simulations correspond to $L = 25 h^{-1} \text{ Mpc}$ and $N = 2 \times 10^4$ cells and particles. All simulations were run on the *Summit* computing system at the Oak Ridge National Laboratory. The effect of a thermally produced WDM particle was introduced by modifying the input transfer function following the fitting formula of [33], and no attempt was made to explicitly incorporate particle thermal velocities in the initial conditions [37]. WDM thermal velocities will be only a few percent of the velocities due to gravitational acceleration at the cosmic time of our initial conditions, and should not be important on the scales resolved by our simulations.

To compare the properties of the IGM in our simulations to observations, we extracted synthetic Lyman- α forest spectra measuring HI densities, temperatures, and peculiar velocities of the gas using 4096 skewers randomly distributed for each simulation along each axis of the box. The optical depth τ_ν along each discretized line of sight was estimated as described in [23]. From the set of lines of sight, we computed the transmitted flux and the mean

one-dimensional power spectrum of the fractional transmission (see Sec. 5.4 in [23] for a detailed description).

A. WDM effects on the IGM

The impact of particle free streaming on the gas density structure of the IGM at fixed thermal history is shown in Fig. 1. The image displays a slice through the gaseous cosmic web at $0 \leq z \leq 5.2$ generated from a set of eight simulations that vary the mass of the dark matter particle m_{WDM} increasing from 0.3 keV to ∞ . The simulations follow the thermal history produced by the baseline photoheating and photoionization history presented in [24].

The main effect of the free streaming of keV particles on the matter density distribution is to suppress the formation of small-scale structure, reducing the clumpiness of the cosmic web. As free streaming washes out small-scale inhomogeneities and decreases the concentration of overdense regions, baryon redistribution leaves an excess (relative to CDM) of close-to-mean density gas to permeate the IGM. Figure 2 (left) shows the change in the volume-weighted probability distribution function of the gas

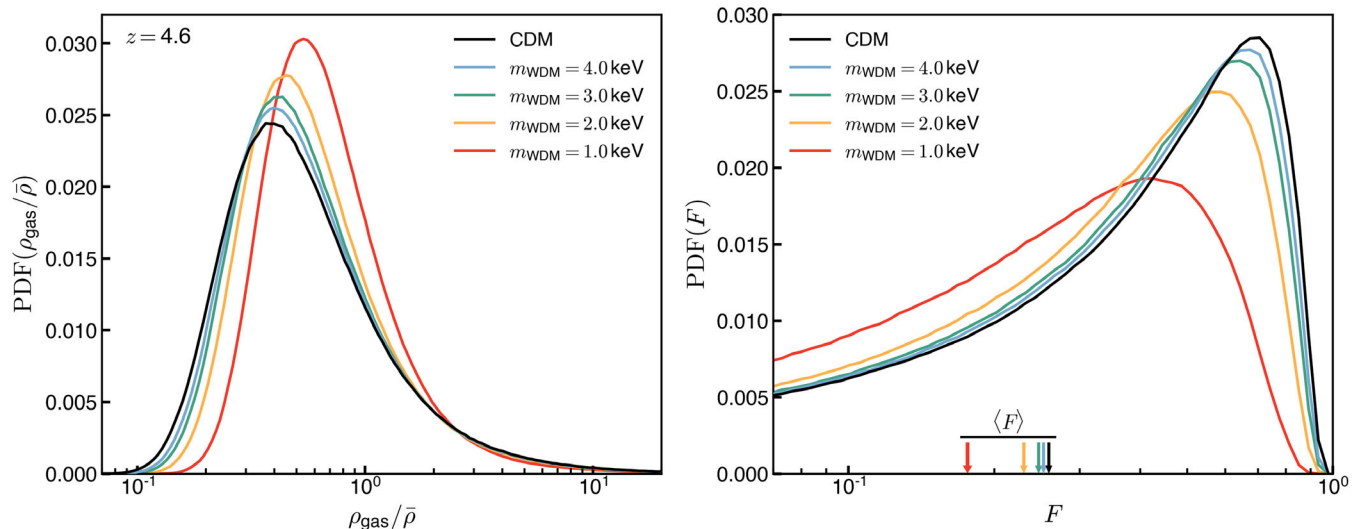


FIG. 2. *Left*: Volume-weighted probability distribution function of the gas density at $z = 4.6$ in WDM cosmologies compared to Λ CDM. Due to free streaming, gas that otherwise would collapse into small-scale structures redistributes on larger scales. This suppressed collapse shifts the peak of the distribution closer to the mean density $\bar{\rho}$. *Right*: Gas redistribution in WDM cosmologies results in an increase of the Lyman- α scattering optical depth of the IGM and a decrease of the transmitted flux F . Arrows display the decrease of the mean transmission $\langle F \rangle$ as m_{WDM} decreases.

density at $z = 4.6$ for different m_{WDM} cosmologies compared to a CDM simulation with the same thermal history. The enhanced fraction of close-to-mean density as a consequence of free streaming increases the opacity of the IGM to Lyman- α scattering. Figure 2 (right) shows the probability distribution function of the Lyman- α transmitted flux F along the skewer set for the WDM and CDM runs at $z = 4.6$. The distribution of transmitted flux shifts to lower values with increasing free-streaming lengths. The arrows at the bottom of the figure display the decrease of the mean transmission $\langle F \rangle$ of the forest as m_{WDM} decreases.

While free streaming (in the regime $m_{\text{WDM}} \geq 1$ keV) changes the IGM opacity it has a negligible effect on its density-temperature relation, often approximated with a power-law model $T = T_0(\rho/\bar{\rho})^{\gamma-1}$, and the parameters T_0 and γ are virtually unchanged—for fixed photoionization and photoheating histories—from the CDM case.

B. WDM effects on the flux power spectrum

The suppression of small-scale density fluctuations from free streaming translates into a decrease of the small-scale FPS compared to CDM. Figure 3 shows the dimensionless FPS of the Lyman- α transmitted flux, $\pi^{-1}kP(k)$, measured at $z = 4.6$ in simulations with different WDM particle mass. The ratio of $P(k)$ with respect to the standard CDM case is depicted in the bottom panel. Free streaming affects the FPS differently on small and large scales. On small scales, $k \gtrsim 0.02$ s km $^{-1}$, the suppression of density fluctuations reduces power. On large scales, $k \lesssim 0.02$ s km $^{-1}$, there is an increase of the overall normalization of $P(k)$ associated with the reduction of the mean transmitted flux $\langle F \rangle$.

C. Simulation grid

To constrain WDM cosmologies, we have performed an unprecedented grid of 1080 high-resolution simulations (with $L = 25 h^{-1}$ Mpc and $N = 2 \times 1024^3$) for a variety of particle masses m_{WDM} and thermal histories of the IGM that alter the impact of free streaming, Doppler broadening, and pressure smoothing on the forest FPS. Here, we use the fiducial model for the hydrogen and helium photoionization and photoheating rates of [24] (hereafter V22) as a template to generate different gas reionization and thermal histories. The V22 photorates are a modification of the rates by [38] and were determined by running hundreds of cosmological simulations to produce a reheating history in agreement with observations of the hydrogen Lyman- α FPS as well as the effective opacity of the He II Lyman- α forest (see Ref. [24] for details).¹

Our grid of models is based on three different transformations of the V22 fiducial rates:

- (1) We rescale the hydrogen (H I) and helium (He I) photoionization rates Γ_i by a constant factor β . This rescaling mainly impacts the neutral hydrogen

¹In our modeling, the ionization and thermal evolution of the IGM is primarily determined by the radiation emitted by star-forming galaxies and active galactic nuclei (AGNs) over cosmic history. The photoionization and photoheating rates of [38] were computed from the intensity of the UV background (UVB) radiation field, which was in turn determined by the emissivity of the radiating sources. An improved treatment of the IGM opacity to ionizing radiation that consistently captures the transition from a neutral to an ionized IGM was also adopted in [38]. This UVB model, when applied to cosmological simulations, results in a hydrogen reionization era that completes by redshift ~ 6 (V22).




FIG. 3. Transmitted FPS at $z = 4.6$ from simulations that vary the WDM particle mass at fixed thermal history. The bottom panel displays the fractional difference relative to the CDM case. The suppression of small-scale structure due to free-streaming results in reduced power compared to CDM for $k \gtrsim 0.02 \text{ s km}^{-1}$. On large scales, the increase of the overall normalization compared to CDM is associated with the excess of close-to-mean density gas and resulting reduction of the mean transmitted flux $\langle F \rangle$. The gray bands denote the regions that fall outside the observational measurements of [22].

(and helium) fractions, changing the mean transmitted flux $\langle F \rangle$ and therefore the overall normalization of the FPS.

- (2) We modify the H I and He I photoheating rates \mathcal{H}_i by rescaling the mean energy of the ejected photoelectrons $\mathcal{H}_i(z)/\Gamma_i(z)$ by a factor α_E . This modification changes the temperature of the IGM and the impact of thermal Doppler broadening on the small-scale FPS.
- (3) We change the timing of hydrogen reionization by offsetting the H I and He I photoionization and photoheating rates by an amount Δz —a value $\Delta z > 0$ shifts reionization to earlier times. This shift affects the time available for the IGM to cool by adiabatic expansion and respond to thermally induced pressure gradients, changing the impact of pressure smoothing on the small-scale FPS.

The transformations applied to the H I and He I photoheating rates described above can be expressed as

$$\begin{aligned} \Gamma_i(z) &\rightarrow \beta \Gamma_i^{Y22}(z - \Delta z), \\ \mathcal{H}_i(z) &\rightarrow \beta \alpha_E \mathcal{H}_i^{Y22}(z - \Delta z). \end{aligned} \quad (15)$$

In this work we shall compare the results of our simulations to observational determinations of the FPS in the redshift

range $4.2 \leq z \leq 5.0$. In the V22 model the IGM cools roughly adiabatically after reionization at $z \sim 6$ until AGNs begin to photoionize He II and heat intergalactic gas at $z \sim 4.5$, resulting in a local minimum of the IGM temperature at this epoch. The effect of He II reheating becomes dominant, however, only at $z < 4$ and does not significantly impact the FPS at $z \geq 4.2$. Accordingly, we do not marginalize over the reionization history of helium, and fix the He II photoionization rates to those of V22. We note that since our simulations assume a spatially uniform UV background, there are no couplings between, e.g., the UV radiation field, the dark matter density, and peculiar motions that would affect the interpretation of our power spectrum results.

Each model in our simulation suite is therefore characterized by the four-dimensional parameter vector $\theta = \{m_{\text{WDM}}, \beta, \alpha_E, \Delta z\}$, and our parameter grid is shown in Table I. The range of dark matter masses were chosen to span from very light WDM ($m_{\text{WDM}} = 1 \text{ keV}$) to CDM. After running some test calculations the ranges of thermal parameters (α, β) were chosen to encompass the parameter space with large likelihood. The range of Δz shifts span about 200 million years of cosmic time around the end of reionization. Full cosmological hydrodynamical simulations were performed for all possible combinations of these four parameters, requiring the suite of 1080 simulations used for our analysis.

D. Effects of model variations on the flux power spectrum

Suppression of small-scale fluctuations in the Lyman- α forest can be attributed to decreased small-scale inhomogeneities in the intergalactic gas due to the free streaming of WDM particles, but also to Doppler broadening of absorption lines and pressure smoothing of gas overdensities. Providing accurate constraints on the WDM particle mass from observations of the FPS requires differentiating the impact that these processes have on $P(k)$, and sampling over a dense set of models that simultaneously vary the effect of each mechanism.

Figure 4 illustrates the impact on the FPS of free streaming (left), thermal broadening (center), and pressure smoothing (right) at $z = 4.6$. Shown is the ratio of the $P(k)$ measured in simulations that independently vary the model

TABLE I. WDM-CHIPS simulation grid.

Parameter	Parameter values
m_{WDM} [keV]	1, 2, 3, 4, 5, 6, 8, 12, 20, 40, 80, CDM
β	0.6, 0.8, 1.0, 1.2, 1.4, 1.8
α_E	0.1, 0.5, 0.9, 1.3, 1.7
Δz	-0.5, 0.0, 0.5

NOTE: The grid of models consist of the 1080 possible combinations of the parameters $\theta = \{m_{\text{WDM}}, \beta, \alpha_E, \Delta z\}$. Each simulation evolves an $L = 25h^{-1} \text{ Mpc}$ and $N = 2 \times 1024^3$ box.




FIG. 4. Effects of free streaming (left), thermal broadening (center) and pressure smoothing (right) on $P(k)$ at $z = 4.6$, compared with the $P_0(k)$ measured from a simulation that adopts the V22 fiducial model (Λ CDM, $\beta = 1$, $\alpha_E = 1$, and $\Delta z = 0$). Gray bands show the regions outside the observational measurements from [22]. Free streaming has the largest effect on the shape of $P(k)$ relative to the fiducial CDM model, with suppression of 10%–60% at $k \sim 0.1$ s km $^{-1}$ for $m_{\text{WDM}} \sim 5$ –2 keV. On smaller scales the free-streaming suppression saturates due to peculiar velocities (see text and Fig. 5). Thermal broadening of lines induces a 10%–20% effect at $k \sim 0.1$ s km $^{-1}$ for mean IGM temperature changes of $\Delta T_0 \sim 1,000$ –2,000 K. Shifting the redshift of reionization by $\Delta z \sim 0.5$ leads to differences in the impact of pressure smoothing of 1%–2% at $k \sim 0.1$ s km $^{-1}$.

parameters m_{WDM} (left), α_E (center) and Δz (right) to the $P_0(k)$ measured from a simulation that adopts the fiducial V22 model corresponding to a Λ CDM cosmology and parameters $\beta = 1$, $\alpha_E = 1$, and $\Delta z = 0$. To isolate small-scale effects, for this comparison the Lyman- α optical depths of the skewers from each simulation were rescaled to produce the same $\langle F \rangle$ of the fiducial model as this equalizes $P(k)$ on large scales. Additionally, for the simulations that vary Δz (right), we separate the impact of pressure smoothing by also rescaling the instantaneous gas temperature along the skewers to have the same value at mean density T_0 as the fiducial model. Variations of the model parameter β mainly impact the ionization fraction of intergalactic hydrogen, changing the mean transmitted flux and thereby rescaling the overall normalization of the FPS. For a detailed analysis on this effect we refer the reader to Appendix B of [24].

As shown in Fig. 4 (center), the effect from increasing/decreasing the IGM temperature is to suppress/boost the FPS for $k \gtrsim 0.01$ s km $^{-1}$ relative to the fiducial model. On these scales $P(k)$ is also altered by pressure smoothing (right). An earlier (later) reionization epoch suppresses (boosts) power as overdensities have more (less) time to respond to thermally-induced pressure gradients. Figure 4 also makes clear the small-scale $P(k)$ behavior of WDM models (left) relative to the fiducial CDM model. On (quasi)linear scales (e.g., $k < 0.1$ s km $^{-1}$), the effect on the power spectrum of increasing the WDM particle mass is well-captured by the exponential cutoff of Eq. (14). The suppression of $P(k)$ due to free-streaming increases with decreasing m_{WDM} , and extends to $k \gtrsim 0.1$ s km $^{-1}$ where it saturates owing to peculiar motions. We argue that this important feature could be key for constraining the nature of dark matter.

Interestingly, we find that if $P(k)$ is measured in configuration (‘real’) space instead of redshift space (ignoring gas peculiar velocities) then the suppression of $P(k)$ continues to $k \sim 0.2$ s km $^{-1}$. Figure 5 displays the ratio $P(k)/P_{\text{CDM}}(k)$ for simulations with WDM particle mass $m_{\text{WDM}} = 5$ keV (blue) and $m_{\text{WDM}} = 4$ keV (green). Dashed lines show the same ratio $P(k)/P_{\text{CDM}}(k)$ but in this case the FPS is computed from synthetic spectra in real space. Here, the effective opacity of the skewers of the




FIG. 5. Suppression of $P(k)$ due to free streaming relative to CDM at $z = 4.6$. The full lines show the ratio $P(k)/P_{\text{CDM}}(k)$ where $P(k)$ has been measured from the Lyman- α transmitted flux computed in redshift space. For the dashed lines the transmitted flux from the CDM and WDM case is computed in real space (ignoring peculiar velocities). In real space the suppression of $P(k)$ due to free streaming continues to $k \gtrsim 0.2$ s km $^{-1}$, while in redshift space the suppression saturates at $k \sim 0.1$ s km $^{-1}$.

WDM simulations has been rescaled to match CDM on large scales.

The comparison presented in Fig. 5 suggests that peculiar velocities play an influential role shaping the structure of the Lyman- α forest on small scales. We conclude that for $k \gtrsim 0.1 \text{ s km}^{-1}$ peculiar velocities are more important for shaping $P(k)$ than free streaming. Since peculiar velocities and thermal broadening on the smallest scales affect WDM and CDM models similarly, the decrease in $P(k)$ from WDM relative to CDM saturates at $k \gtrsim 0.1 \text{ s km}^{-1}$.

IV. STATISTICAL COMPARISON

In this section we present the observations of $P(k)$ used for this work and we describe the methodology used for the Bayesian approach used in this work to constrain the models used for our simulation suite.

A. Observational flux power spectrum

For this work we employ observations of the Lyman- α forest power spectrum from the Keck observatory and the Very Large Telescope, presented in [22]. These measurements represent the highest resolution determination of $P(k)$ to date and probe the structure of the forest in the redshift range $4.2 \leq z \leq 5.0$. Since the impact of free streaming on $P(k)$ is greater at high redshift the data set from [22] is arguably the most constraining determination of $P(k)$ to infer the nature of dark matter currently available. The measurements of $P(k)$ from [22] along our best-fit determination obtained from our analysis are shown in Fig. 6.

B. MCMC inference

To constrain WDM cosmologies from the observations of $P(k)$, we use an MCMC approach to sample over our suite of models and obtain best-fit distributions for our parameters. The likelihood function for the model given by the parameters $\theta = \{m_{\text{WDM}}^{-1}, \beta, \alpha_E, \Delta z\}$ is evaluated as

$$\ln \mathcal{L}(\theta) = -\frac{1}{2} \sum_z [\Delta^T \mathbf{C}^{-1} \Delta + \ln \det(\mathbf{C}) + N \ln 2\pi], \quad (16)$$

where Δ denotes the difference vector between the observational $P(k)$ and the model $\Delta = P_{\text{obs}}(z, k) - P(z, k|\theta)$, and \mathbf{C} corresponds to the covariance matrix associated with the observation of $P(k)$ as reported by [22]. These authors have considered how systematic uncertainties may affect their covariance matrix, and find that the level of potential systematics is $\lesssim 1\%$. To compute $P(z, k|\theta)$ for arbitrary values of the parameters θ not directly simulated by our grid, we perform a four-dimensional linear interpolation of the FPS measured from the sixteen neighboring simulations in parameter space. The methodology for our inference approach differs from previous works in a number of key respects:

- (1) Our WDM-CHIPS simulation grid captures a wide range of free-streaming lengths and UVB models to sample over a variety of cosmic structure formation and gas thermal histories, and thereby produce different statistical properties for the Lyman- α forest. Due to the computational cost of the simulations, evolving such a large grid of high-resolution cosmological hydrodynamical simulations had not been achieved before. With the advent of efficient




FIG. 6. Observational determination of $P(k)$ in the redshift range $4.2 \leq z \leq 5.0$ used in this work to constrain WDM cosmologies. The best-fit from the full WDM-CHIPS grid (purple) and the subset CDM grid (orange) are shown as lines and shaded regions which correspond to 95% confidence range of $P(k)$ marginalized over the posterior distribution. From our analysis, the $P(k)$ from WDM cosmologies with $m_{\text{WDM}} = 4.5 \text{ keV}$ is preferred over the best-fit CDM model. We note that the WDM best-fit results in lower IGM temperatures compared to CDM (see Fig. 10). The bottom panels show the ratios between the observations and the best-fit WDM FPS (blue points) and the best-fit CDM and WDM FPS (orange line), along with the WDM and CDM model 95% confidence intervals (shaded regions).

numerical codes like CHOLLA and capable systems like Summit, it is now possible to simulate thousands of models.

- (2) We evaluate the performance of a given model in matching the observations of $P(k)$ by marginalizing over the self-consistently evolved hydrogen reionization history in the redshift interval $4.2 \leq z \leq 5.0$. The IGM thermal properties at one redshift cannot be disentangled from its properties at previous epochs, thus the marginalization over the parameter posterior distributions should not be performed independently at each redshift [cf. [22,31,39–42]].
- (3) We do not modify, in postprocessing, the mean transmitted flux $\langle F \rangle$ in the forest by rescaling the Lyman- α optical depth, nor do we assume an instantaneous power-law temperature-density relation (cf. [13,19,21,22]). The latter does not accurately reproduce the thermal state of IGM gas in the range $-1 \leq \log_{10} \Delta \leq 1$. Instead, in our approach simulations self-consistently evolve the ionization and thermal structure of the IGM determined by the wide range of photoionization and photoheating histories applied in our model grid.

V. RESULTS

By comparing the high-resolution observation of the flux power spectrum from [22] to our suite of simulations that simultaneously vary the impact of free streaming from WDM cosmologies on the matter distribution and different reionization and thermal histories of the IGM, we infer new constraints on the WDM particle mass m_{WDM} . In this section we present the posterior distribution from our MCMC analysis and the marginalized $P(k)$ and thermal evolution of the IGM for our best-fit models. To complement our results, we repeat our analysis by modifying $P(k)$ from the simulations to account for a nonuniform UV background, and finally, we show hypothetical constraints on m_{WDM} from an artificially increased number of quasars.

A. Distribution of the model parameters

The posterior distribution of our four model parameters $\theta = \{m_{\text{WDM}}^{-1}, \beta, \alpha_E, \Delta z\}$ resulting from the Bayesian inference procedure is shown in Fig. 7. A clear global maximum of the 1D marginalized distributions is shown for m_{WDM} and the parameters β and α_E responsible for rescaling the photoionization and photoheating rates. The distribution of the parameter Δz is not fully contained by our grid of models, but we argue that this issue does not represent a significant challenge to our conclusions. Values of $\Delta z < -0.5$ would reduce the suppression of the small-scale $P(k)$ from pressure smoothing which could be compensated by free streaming or thermal broadening, possibly shifting our constraint of m_{WDM} to slightly lower values.

The marginalized distribution for m_{WDM}^{-1} (top left panel, Fig. 7) is well-constrained and it peaks at $m_{\text{WDM}}^{-1} = 1/(4.5 \text{ keV})$. Arguably, the preference for a nonzero free-streaming length is weakly significant as the CDM cosmology is contained within the 3σ interval of the marginalized distribution. The principal result from our analysis is the lower limit $m_{\text{WDM}} > 3.0 \text{ keV}$ (at the 2σ level) obtained for the WDM particle mass. We have also performed our MCMC analysis sampling in m_{WDM} space (instead of m_{WDM}^{-1} space) applying a flat prior in mass the mass range of $1 < m_{\text{WDM}} < 80 \text{ keV}$, and we find that our conclusions are unchanged. We also note that the sensitivity of the best-fit WDM mass on the UVB parameters (and hence on the reionization history) is largely a thermal Doppler broadening (and to a less extent, a pressure) effect.

To find the best-fit model to $P(k)$ for a CDM cosmology, we repeat our MCMC approach but restricted to the models that evolve the Λ CDM cosmology varying only the photoionization and photoheating history. Figure 8 shows the one- and two-dimensional projections of the posterior distribution for the UV background (UVB) parameters $\{\beta, \alpha_E, \Delta z\}$ obtained from the comparison of $P(k)$ from [22] to the grid of 90 CDM simulations. The posterior distribution of the UVB parameters from our CDM fit differs slightly from the distribution of parameters obtained from the inference that included WDM cosmologies (see Fig. 7). The main difference between the two distributions is that the parameters β and α_E peak at slightly shifted values. In particular, the product $\beta\alpha_E$, which rescales the photoheating rates, is higher for the CDM case. Higher photoheating rates directly result in increased IGM temperatures (see Sec. V C).

B. Best-fit power spectrum

The marginalized flux power spectrum $P(k)$ over the posterior distribution obtained from our MCMC analysis along the observational data used to constrain the model is shown in Fig. 6. Lines and shaded regions show the best-fit $P(k)$ and the 95% confidence range. The result from the fit performed by sampling over the entire WDM model grid (purple) and restricting to the CDM grid (orange) are shown separately. As shown, both results provide a good match to the observed $P(k)$ nevertheless, the relative difference between the model and the data quantified as $\chi^2 = \sum_z \Delta^T \mathbf{C}^{-1} \Delta$ is slightly lower for the best-fit WDM model ($\chi^2 = 38.3$), compared to $\chi^2 = 40.9$ from the CDM best-fit model. Here \mathbf{C} denotes the covariance matrix taken from [22] and Δ is the difference vector between the observed and model $P(k)$.

We find that the preference for a nonzero free-streaming scale derived from our analysis is driven by the smallest scales probed by the high-resolution measurement of $P(k)$ from [22]. We repeated our inference methodology but

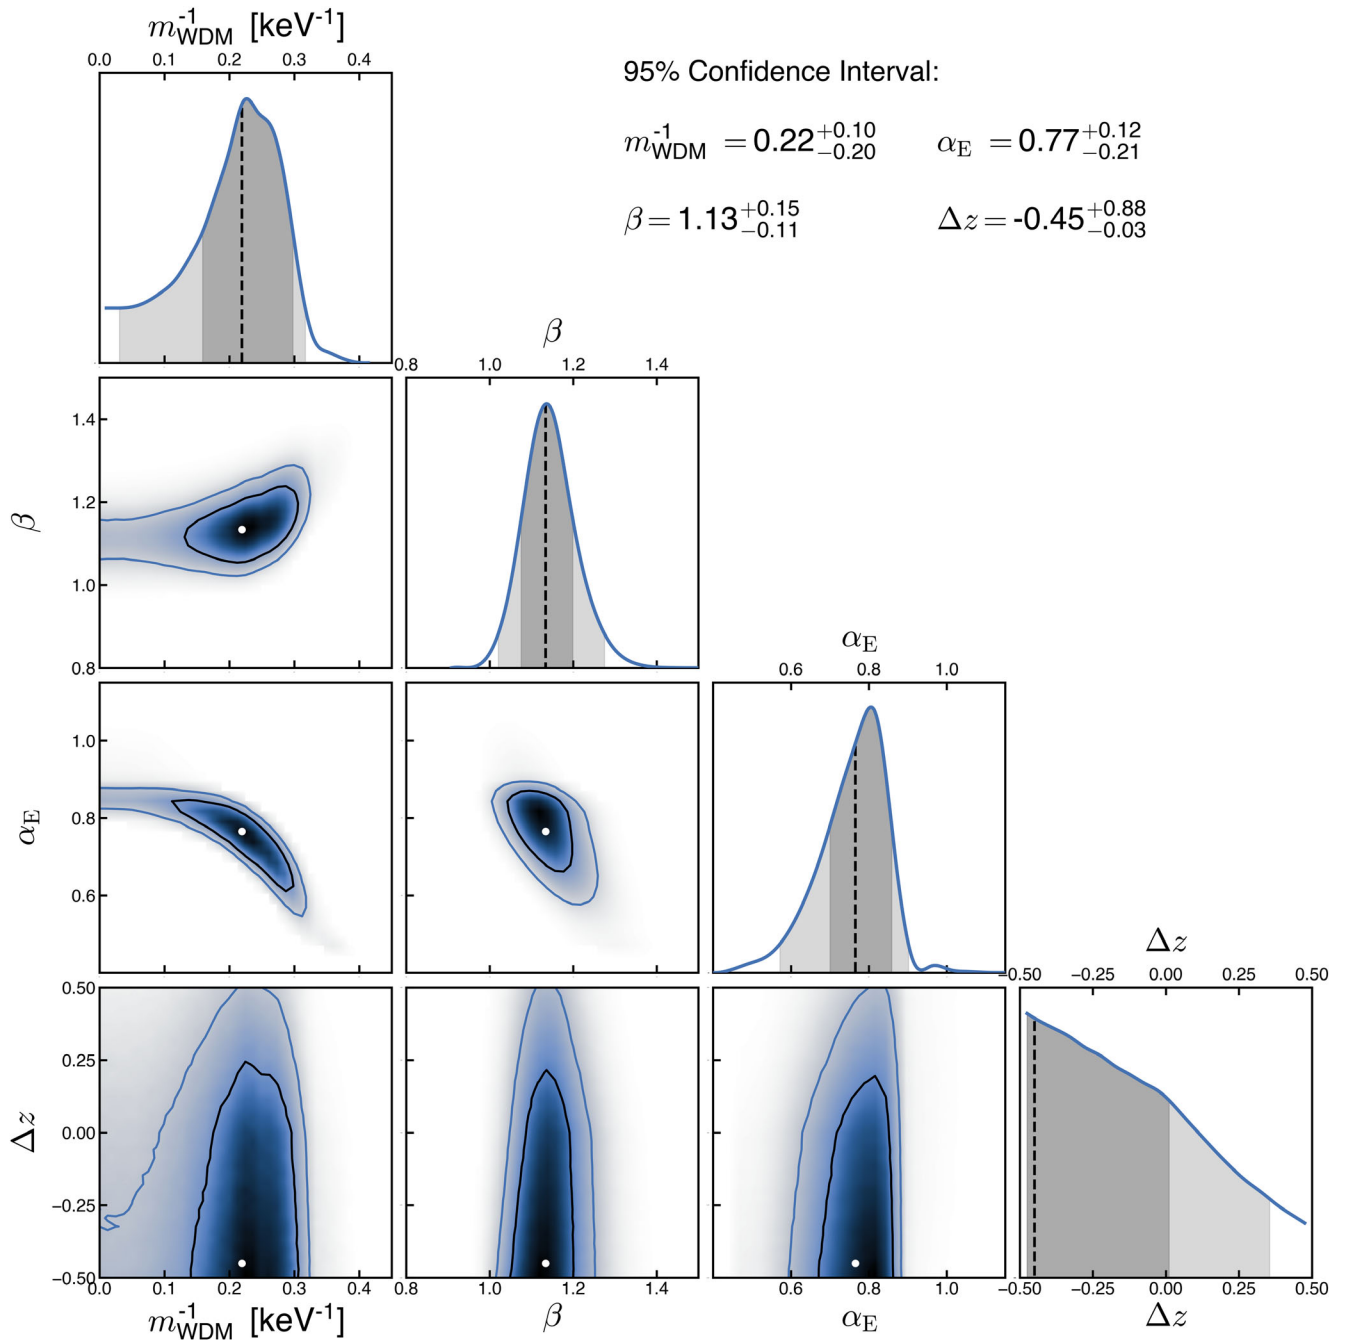


FIG. 7. Results from the Bayesian inference procedure, showing one- and two-dimensional projections of the posterior distributions for the parameters $\theta = \{m_{\text{WDM}}^{-1}, \beta, \alpha_E, \Delta z\}$ recovered by fitting synthetic flux power spectra from our grid of WDM-CHIPS simulations to observations of the Lyman- α forest from [22]. Gray bands in the 1D distributions show the 1σ and 2σ intervals. The marginalized likelihood for $1/m_{\text{WDM}}$ peaks at $1/(4.5 \text{ keV})$. The preference for a nonvanishing free-streaming length is only weakly statistically significant, as the CDM case is contained within the 3σ interval of the distribution. Our main result is a lower bound $m_{\text{WDM}} > 3.1 \text{ keV}$ at the 2σ confidence level.

excluding the three last data points ($k > 0.1 \text{ s km}^{-1}$) from the likelihood calculation. Here we found that excluding the high- k measurements places ΛCDM as the preferred cosmology. The marginalized posterior distributions for m_{WDM} from the fits limited to $k \leq 0.1 \text{ s km}^{-1}$ and to full data set are displayed in Fig. 9. As shown, the likelihood

peaks at CDM and the lower limit constraint at the 2σ level shifts to $m_{\text{WDM}} > 3.6 \text{ keV}$.

In Sec. III D we showed that free streaming and Doppler broadening impact $P(k)$ differently on small scales ($k > 0.1 \text{ s km}^{-1}$). As the temperature of the gas is increased/decreased, the impact of thermal broadening is to

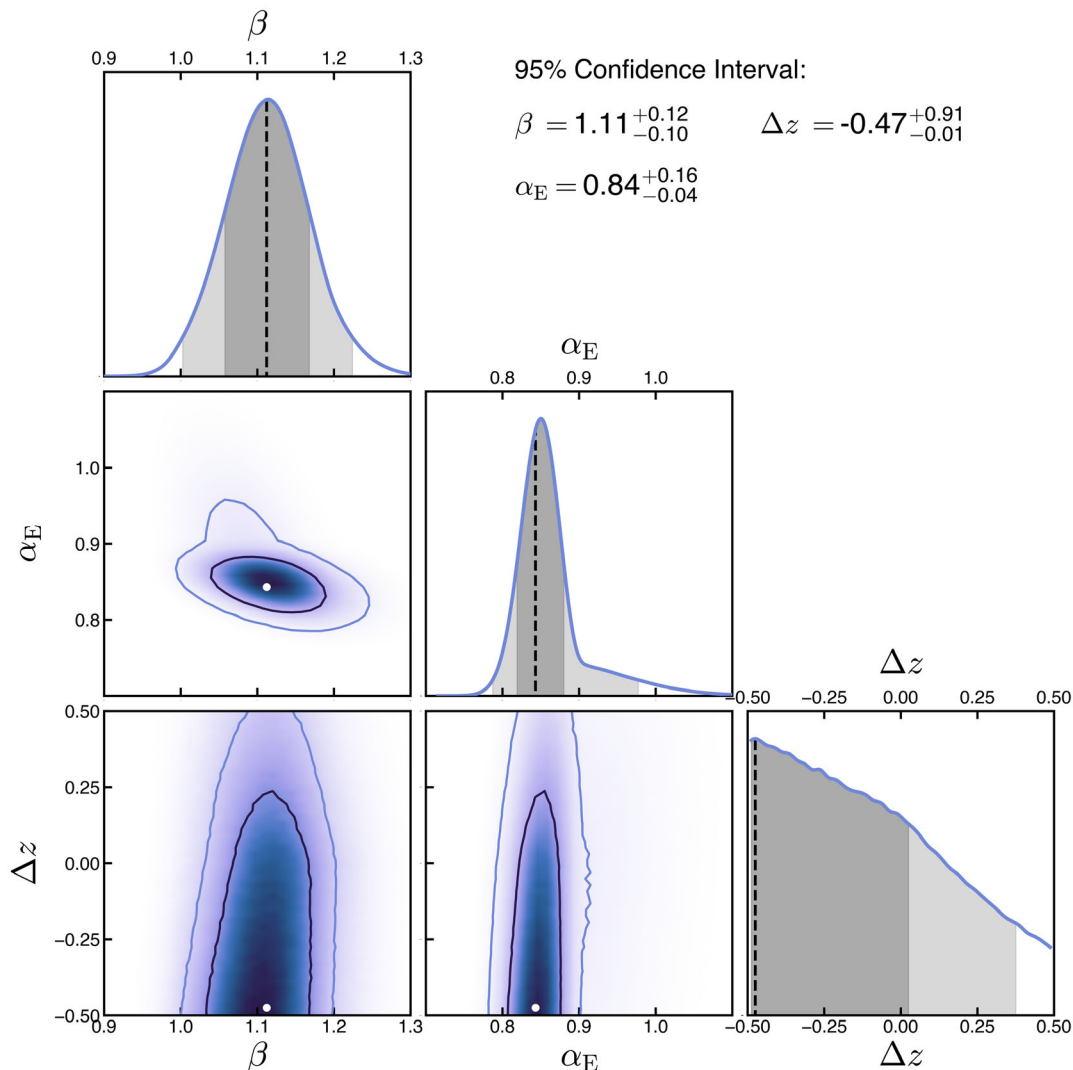


FIG. 8. Results from our MCMC procedure, showing one- and two-dimensional projections of the posterior distributions for the UVB parameters $\theta = \{\beta, \alpha_E, \Delta z\}$. Here the sampling of models is restricted to the subset of 90 simulations that evolve a CDM cosmology. Gray bands in the 1D distributions show the 1σ and 2σ intervals.

decrease/increase the small-scale $P(k)$. While free streaming decreases small-scale density fluctuations, its effect on the flux power spectrum saturates at $k \gtrsim 0.1 \text{ s km}^{-1}$ due to the impact from peculiar velocities (see Figs. 4 and 5).

We have shown that the preference for a WDM cosmology over Λ CDM originates from the $k > 0.1 \text{ s km}^{-1}$ measurement of $P(k)$. Also, we showed that on these scales ($k > 0.1 \text{ s km}^{-1}$) the suppression of $P(k)$ due to free-streaming saturates while the suppression of $P(k)$ from thermal broadening increases monotonically. Therefore, we conclude that the saturated suppression of $P(k)$ due to a $m_{\text{WDM}} \sim 4.5 \text{ keV}$ combined with a moderately lower IGM temperature provides a slightly better fit to the $k > 0.1 \text{ s km}^{-1}$ observation of $P(k)$ compared to the Λ CDM best-fit model with higher IGM temperatures at $4 \leq z \leq 5$ (see Fig. 10).

C. Thermal evolution of the IGM

One of the advantages of our approach compared to previous works that aim to constrain the WDM free streaming from observations of the Lyman- α forest power spectrum is that our simulations self-consistently evolve the IGM thermal and ionization history during and after reionization by sampling over a large grid models that vary the UVB photoionization and photoheating rates. Instead, the methodology adopted by previous works [13,15,16,19] was to change the temperature T_0 and the ionization fraction of the IGM in post-processing by rescaling T_0 and the effective optical depth τ_{eff} of the simulated skewers.

Figure 10 shows the redshift evolution of the temperature of the gas at mean density T_0 marginalized over the posterior distribution from our MCMC analysis; shaded bars show the 1σ interval. The best-fit model for WDM cosmologies with $m_{\text{WDM}} = 4.5 \text{ keV}$ results in moderately




FIG. 9. One-dimensional posterior likelihood for the WDM particle mass obtained by limiting the observational measurements of $P(k)$ [22] to $k \lesssim 0.1 \text{ s km}^{-1}$ compared to the likelihood distribution from fitting to the complete data set (dashed blue). The limited data set for the observational $P(k)$ is constructed by excluding the last three points of each redshift snapshot. Gray bands show the 1σ and 2σ interval from the fit to the limited data. In this case ΛCDM is the preferred cosmology and the lower limit at 95% confidence shifts to $m_{\text{WDM}} > 3.6 \text{ keV}$ compared to $m_{\text{WDM}} > 3.0 \text{ keV}$ from the fit to the full data set.

lower (5%–10%) IGM temperatures (purple) due to the slightly reduced photoheating rates compared to the ΛCDM best-fit model (orange). Lower temperatures are expected from the WDM models as they decrease the impact of thermal broadening in suppressing $P(k)$ which compensates for the suppression due to free streaming.

In Fig. 10, we compare the IGM temperature T_0 inferred by this work to other measurements at $z > 4$ [22,25]. The high-redshift inference presented by [25] was obtained by characterizing the transmission spikes observed in $z > 5$ spectra. The thermal histories obtained from our analysis are consistent with the results from [25], suggesting a peak in T_0 due to hydrogen reionization at $z \sim 6$.

The measurement of T_0 presented in [22] for $4.2 \leq z \leq 5.0$ was made by fitting the observed $P(k)$ to simulated spectra according to a ΛCDM cosmology where the instantaneous density-temperature distribution of the gas is modified in post-processing changing parameters T_0 and γ from the power-law relation $T = T_0(\rho/\bar{\rho})^\gamma$. Despite using the same observational determination of $P(k)$ for our inference, the IGM temperatures obtained in this work during $4.2 \leq z \leq 5.0$ are slightly higher than those inferred by [22]. Restricting to a ΛCDM cosmology, we find a best-fit T_0 is 15%–20% higher at redshift $z = 5.0$ and $z = 4.6$ compared to the result from [22]. At redshift $z = 4.2$ the difference lowers to 10%–15% and our result agrees with their inference within 1σ . The evolution of T_0 from our best-fit WDM model is also higher (5%–10%) than the results from [22], but in this case our results also agree to within 1σ .




FIG. 10. Redshift evolution of the temperature T_0 of IGM gas at the mean density from the best-fit ΛCDM and WDM models (solid lines) and 1σ interval (colored bands) obtained from our MCMC analysis. The data points show the values of T_0 inferred from observations of the Lyman- α forest by [22,25]. The best-fit from WDM cosmologies have moderately lower temperatures compared to CDM, which compensates for the suppression of small-scale $P(k)$ due to free streaming.

D. Thomson scattering optical depth

Our WDM and CDM-Only best-fit models result in very similar reionization histories with Thomson scattering optical depths that are consistent with constraints from the Planck satellite [43,44]. Figure 11 shows such constraints together with the marginalized Thomson scattering optical depth, τ_e , for our WDM best-fit model (black line, with the shaded regions displaying the 1σ interval), and for our CDM-Only best-fit model (blue line). We measure $\tau_e = 0.0586^{+0.0043}_{-0.0012}$ and $\tau_e = 0.0584^{+0.0041}_{-0.0010}$ from our WDM and CDM-Only best-fit simulations, respectively.

E. Constraining WDM with increased quasar sightlines

Upcoming surveys of the Lyman- α forest (e.g., DESI [45], WEAVE [46], EUCLID [47], LSST) will drastically increase the available observations of quasar sightlines, which will significantly improve the statistics of measurements derived from the forest. The improved statistics will allow tighter constraints on cosmological parameters as well as on WDM and the sum of the neutrino masses from their role in suppressing small-scale structure in the forest.

To assess how better statistics would impact the constraining power of high-redshift and high-resolution observations of the Lyman- α forest for WDM cosmologies, we repeat our analysis but decreasing the uncertainty on the observed $P(k)$ by a factor of one half, which would correspond to increasing the number of observed quasar spectra from fifteen used by [22] for their measurement of $P(k)$, to about sixty (a factor of four increase).

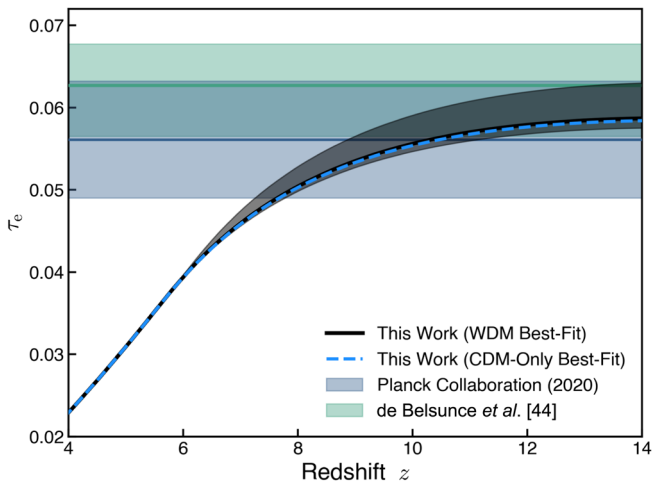


FIG. 11. Electron scattering optical depth, τ_e , to reionization. Black line and shaded bar: WDM best-fit model and 1σ interval. Dashed blue line: CDM-Only best-fit model. Also shown are the observational constraints from the Planck satellite [43,44].

In our approach, we approximate a more constraining dataset by rescaling the covariance matrix \mathbf{C} of $P(k)$ by a factor of one-fourth in Eq. (16). Note that the observational measurements of the FPS are not altered, and only the covariance matrix is reduced. Figure 12 shows the marginalized likelihood for m_{WDM}^{-1} obtained from our analysis using the reduced uncertainty. In this hypothetical case, the improved statistics of the $P(k)$ measurement provide a tighter constraint for the WDM particle mass measured as




FIG. 12. One-dimensional posterior likelihood for the WDM particle mass obtained by fitting to the measurements of the FPS from [22] but with an artificially reduced uncertainty motivated by the improved statistics of Lyman- α spectra from upcoming surveys. Here we rescale the covariance matrix \mathbf{C} of $P(k)$ by a factor of one-fourth. In this hypothetical case, the constraint on m_{WDM} is tighter, measured as $m_{\text{WDM}} = 4.5_{-1.0}^{+1.9}$ keV at 95% confidence level. The dashed line shows our result from fitting to $P(k)$ with the reported uncertainty from [22].

$m_{\text{WDM}} = 4.5_{-1.0}^{+1.9}$ keV at the 95% confidence level. This exercise demonstrates that increasing the sample of high- z and high-resolution observations of the Lyman- α forest should place tight constraints on WDM cosmologies.

F. Modified $P(k)$ for inhomogeneous UVB

One of the main limitations of our analysis is that the WDM-CHIPS simulations evolve under a homogeneous UVB in the form of uniform photoionization and photoheating rates. We evaluate the effect of this assumption on our conclusions by repeating our analysis but using the modification to $P(k)$ prescribed by [48] to account for the impact of a nonuniform UVB on the simulated $P(k)$.

In [48], a set of simulations that apply a uniform UVB were compared to a set that follows a hybrid-RT method. For the latter, spatially varying maps for the H I photoionization and photoheating rates were computed in post-processing and used as input for a rerun of the base simulation, incorporating the response of the gas to the non-uniform photoionization and photoheating. For the comparison, pairs of uniform and nonuniform UVB simulations were calibrated to have the same average ionization and thermal history. The authors concluded that simulations with a patchy reionization show a suppression (10%–15%) of the FPS on small scales ($k \sim 0.1 \text{ s km}^{-1}$) with respect to the uniform UVB case. This effect is mainly driven by the Doppler broadening associated with the high temperatures of recently ionized regions and the divergent peculiar velocities of thermally pressurized gas [49]. On large scales ($k > 0.03 \text{ s km}^{-1}$) the variation of the IGM neutral fraction due to large-scale fluctuations of the gas temperature leads to an increase on $P(k)$.

The likelihood distribution from our MCMC inference where we modify the $P(k)$ from each one of the simulations in our WDM-CHIPS grid to account for an inhomogeneous reionization according to the transformation presented in [48] is shown in Fig. 13. Dashed blue lines correspond to the one-dimensional marginalized distributions obtained from fitting the base $P(k)$ from the simulations with a uniform reionization.

Figure 13 shows that modifying the models $P(k)$ to account for patchy reionization mainly affects the likelihood of m_{WDM} . While WDM cosmologies are still preferred over Λ CDM when fitting the nonuniform UVB models, the likelihood shifts to higher values of m_{WDM} compared to the result from uniform UVB models. An inclination for models with reduced free-streaming compensates for the reduction of small-scale $P(k)$ due to patchy reionization. In this case the maximum likelihood occurs at $m_{\text{WDM}} = 7.1$ keV and the lower limit at 95% level is at $m_{\text{WDM}} = 3.8$ keV. Notably, for the patchy UVB models, the Λ CDM cosmology is contained within the 95% interval of the likelihood distribution.

We note that the difference between the observed and model $P(k)$ defined as $\chi^2 = \sum_z \Delta^T \mathbf{C}^{-1} \Delta$ (see Sec. IV B is

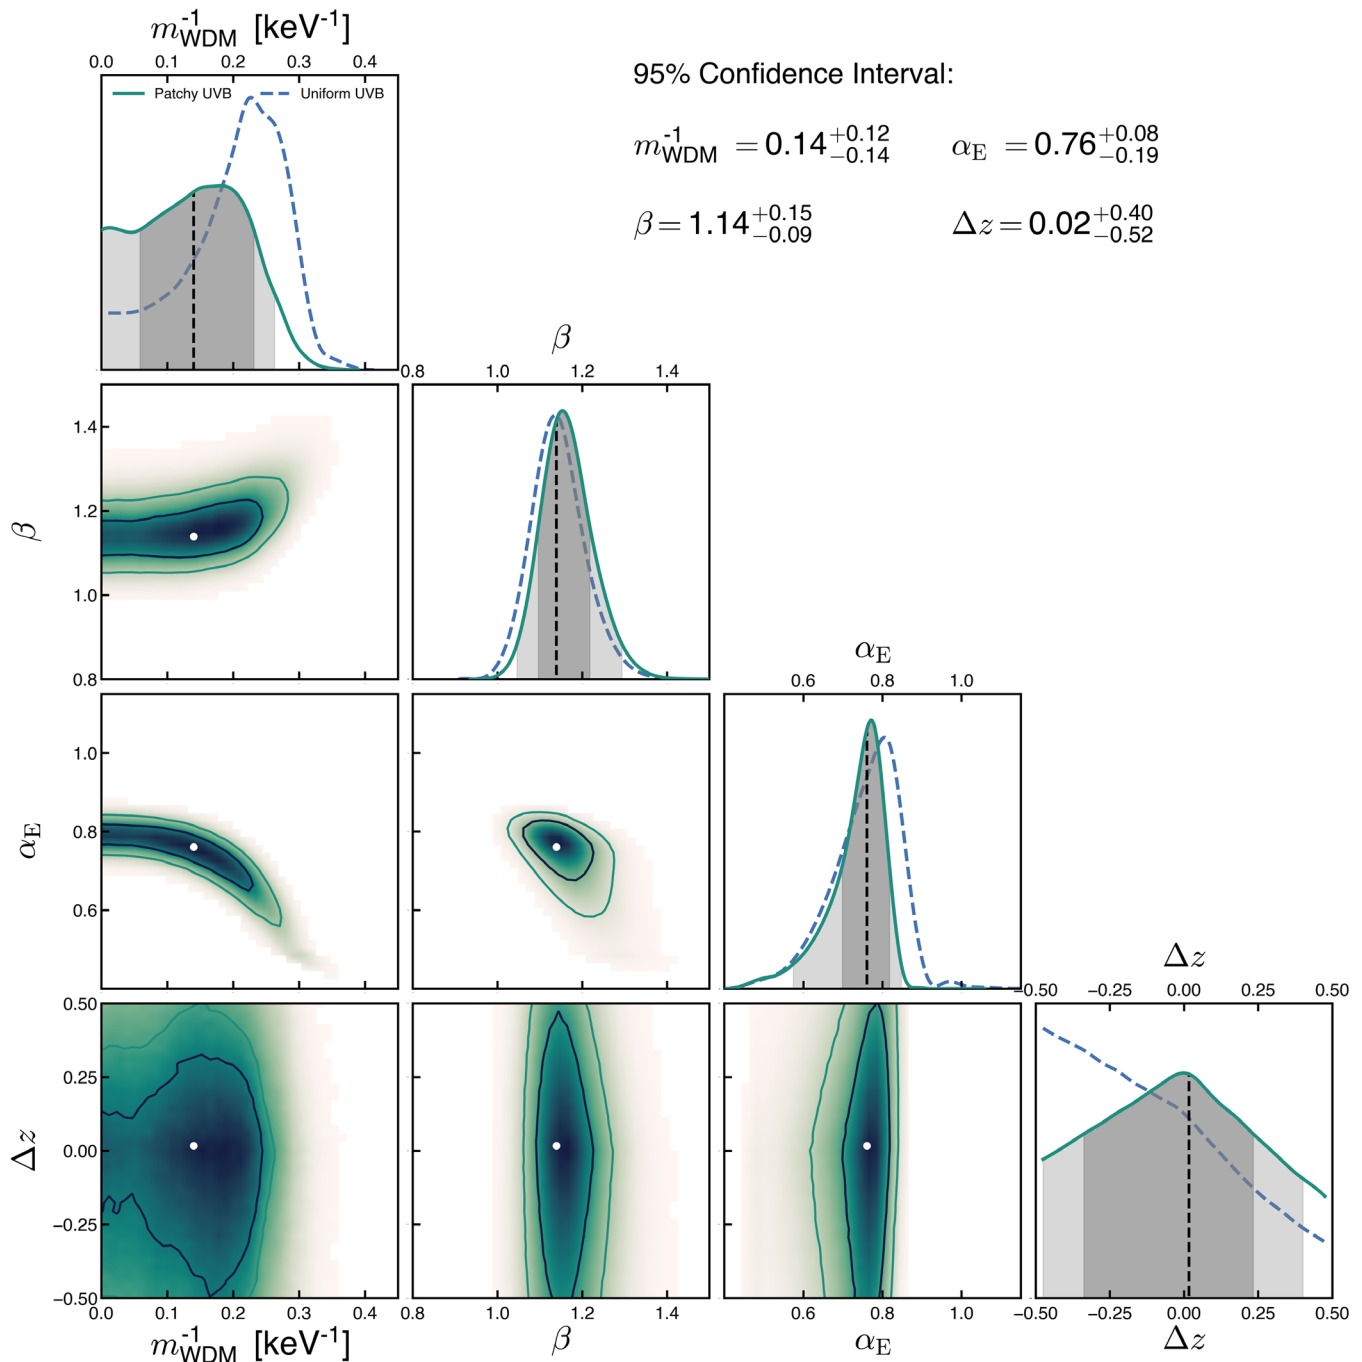


FIG. 13. Posterior distribution of the parameters $\theta = \{m_{\text{WDM}}^{-1}, \beta, \alpha_E, \Delta z\}$ from fitting the observed $P(k)$ from [22] with models that account for a patchy reionization according to the modification to $P(k)$ presented in [48]. In this case, the preference for a WDM cosmology persists but the distribution shifts to higher mass. The likelihood peaks at $m_{\text{WDM}} = 7.1$ keV and the lower limit at the 95% levels is $m_{\text{WDM}} = 3.8$ keV. In this case the Λ CDM cosmology is contained within the 95% interval of the distribution. The dashed blue lines show the one-dimensional likelihood distributions obtained from fitting $P(k)$ measured directly from the uniform UVB simulations.

larger for the nonuniform UVB modified best-fit models compared to the best-fit models sampling from the uniform UVB simulated $P(k)$. When fitting with $P(k)$ measured directly from the uniform UVB simulations, we measure $\chi^2 = 38.3$ for the best-fit model with $m_{\text{WDM}} = 4.5$ keV and a slightly higher $\chi^2 = 40.9$ for the Λ CDM best-fit model.

On the other hand, when sampling over the modified $P(k)$ to account for patchy reionization, we measure larger values $\chi^2 = 46.4$ and $\chi^2 = 46.6$ for the WDM and CDM best-fit models, respectively. The higher values of χ^2 obtained for the patchy reionization modified models show that, in the context of this work, $P(k)$ measured directly

from the uniform UVB simulations provide a better match to the observation from [22], therefore we use the likelihood distribution obtained from the sampling over the uniform UVB models to construct the main results from our analysis.

VI. CONCLUSIONS

To constrain cosmological models where dark matter free streaming smooths the matter distribution in the Universe, we have used the GPU-native code `CHOLLA` to perform a massive suite of high-resolution hydrodynamical cosmological simulations that simultaneously vary the effect of free streaming from WDM particles and the IGM reionization history. We compare the power spectrum of the synthetic Lyman- α forest from our simulations to the high-resolution observational measurement presented by [22] to determine via a likelihood analysis the optimal model for cosmological free streaming that best matches the observation. A summary of the efforts and conclusions from this work follows.

- (i) We present the WDM-CHIPS suite consisting of a grid of 1080 high-resolution simulations ($L = 25h^{-1}$ Mpc, $N = 2 \times 1024^3$) that vary the free streaming from WDM cosmologies and the photoionization and photoheating rates from the metagalactic UVB. The UVB rates applied for our grid use the model from [24] as a template, and use three parameters that control a rescaling amplitude and redshift-timing of the hydrogen photoionization and photoheating rates. Combined with the WDM particle mass m_{WDM} , our four-dimensional grid of models densely sample a wide range of self-consistently evolved reionization and thermal histories of the IGM. This represents a significant improvement over previous studies that aimed to constrain WDM cosmologies from observations of the Lyman- α forest [13,15,16,19,33].
- (ii) The large range of thermal histories produced by the different UVB models in our grid of simulations results in synthetic measurements of Lyman- α spectra where the impact from Doppler broadening and pressure smoothing on suppressing the small-scale $P(k)$ varies widely. This flexibility is important as these mechanisms have similar effects as free streaming decreasing small-scale fluctuations in the Lyman- α forest. Additionally, our approach does not require an assumption of a power-law relation for the density-temperature distribution of the gas. Self-consistent evolution of the IGM phase structure proves to be important as we find that a single power law does not accurately describe the $\rho_{\text{gas}} - T$ distribution in the density range relevant to generating the signal of the Lyman- α forest (see Appendix E of [24]).

- (iii) We compare our grid of models to the high-redshift ($4.2 \leq z \leq 5.0$) observational measurement of the Lyman- α forest power spectrum from [22]. The observations presented in [22] provide the highest resolution measurement of $P(k)$ at the time that work was executed. We perform a Bayesian MCMC sampling to determine the best-fit model for free streaming due to WDM cosmologies and the IGM photoionization and photoheating history.
- (iv) From our MCMC analysis, we find that a cosmological model with nonvanishing free streaming is preferred over Λ CDM. The likelihood function peaks at the thermal relic mass of $m_{\text{WDM}} = 4.5$ keV and we measure a lower limit of $m_{\text{WDM}} = 3.1$ keV at the 95% C.L. We find a weak (3σ) preference for WDM over Λ CDM, but both are statistically consistent with the currently available data.
- (v) We repeat our MCMC analysis, restricting to Λ CDM only but with a variable UVB. We find that the best-fit Λ CDM model mainly differs from the WDM optimal model in that the IGM temperatures are slightly lower (5%–10%) for the WDM case. We find that both thermal histories are in good agreement with the inference from [25] at $5.4 \lesssim z \lesssim 5.8$ and moderately higher than the temperatures from [22] at $4.2 \leq z \leq 5.0$.
- (vi) We introduce the effect of a patchy reionization in our models by modifying $P(k)$ from our uniform UVB simulations according to the prescription presented in [48]. The likelihood distribution from our MCMC approach using the modified $P(k)$ shows that the preference for a WDM cosmology is maintained but the distribution shifts to higher values of m_{WDM} . The maximum likelihood and the 95% lower limit constraint are $m_{\text{WDM}} = 7.1$ keV and $m_{\text{WDM}} = 3.8$ keV, respectively. Additionally, the Λ CDM case is 2σ consistent with the best-fit WDM model when accounting for an inhomogeneous UVB. Notably, the different χ^2 between the observed and best-fit $P(k)$ model is higher for the modified model to account for a nonuniform UVB, and we therefore use the likelihood distribution obtained from the uniform UVB models as the basis for our results.

ACKNOWLEDGMENTS

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725, using Summit Allocations No. AST169 and No. AST175. An award of computer time was provided by the INCITE program, via Project

No. AST175. We acknowledge use of the *lux* supercomputer at UC Santa Cruz, funded by NSF Major research instrumentation (MRI) Grant No. AST1828315, and support from NASA Theoretical and computational astrophysics networks (TCAN) Grant No. 80NSSC21K0271. B. V. is supported in part by the UC MEXUS-CONACyT doctoral fellowship.

B. E. R. acknowledges support from NASA Contract No. NNG16PJ25C and Grants No. 80NSSC18K0563 and No. 80NSSC22K0814. E. E. S. acknowledges support from NASA Grant No. 80NSSC22K0720 and STScI Grant No. HST-AR-16633.001-A. We wish to thank Nick Gnedin and Avery Meiksin for many useful inputs and valuable discussions.

-
- [1] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [2] S. Chabanier, M. Millea, and N. Palanque-Delabrouille, Matter power spectrum: From Ly α forest to CMB scales, *Mon. Not. R. Astron. Soc.* **489**, 2247 (2019).
- [3] P. Colín, V. Avila-Reese, and O. Valenzuela, Substructure and halo density profiles in a warm dark matter cosmology, *Astrophys. J.* **542**, 622 (2000).
- [4] V. Avila-Reese, P. Colín, O. Valenzuela, E. D’Onghia, and C. Firmani, Formation and structure of halos in a warm dark matter cosmology, *Astrophys. J.* **559**, 516 (2001).
- [5] J. L. Feng, Dark matter candidates from particle physics and methods of detection, *Annu. Rev. Astron. Astrophys.* **48**, 495 (2010).
- [6] J. S. Bullock and M. Boylan-Kolchin, Small-scale challenges to the Λ CDM paradigm, *Annu. Rev. Astron. Astrophys.* **55**, 343 (2017).
- [7] P. Bode, J. P. Ostriker, and N. Turok, Halo formation in warm dark matter models, *Astrophys. J.* **556**, 93 (2001).
- [8] L. Hernquist, N. Katz, D. H. Weinberg, and J. Miralda-Escudé, The Lyman-Alpha forest in the cold dark matter model, *Astrophys. J. Lett.* **457**, L51 (1996).
- [9] R. A. C. Croft, D. H. Weinberg, N. Katz, and L. Hernquist, Recovery of the power spectrum of mass fluctuations from observations of the Ly α forest, *Astrophys. J.* **495**, 44 (1998).
- [10] A. A. Meiksin, The physics of the intergalactic medium, *Rev. Mod. Phys.* **81**, 1405 (2009).
- [11] M. McQuinn, The evolution of the intergalactic medium, *Annu. Rev. Astron. Astrophys.* **54**, 313 (2016).
- [12] U. Seljak, A. Makarov, P. McDonald, and H. Trac, Can Sterile Neutrinos be the Dark Matter?, *Phys. Rev. Lett.* **97**, 191303 (2006).
- [13] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt, Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data, *Phys. Rev. D* **88**, 043502 (2013).
- [14] J. Baur, N. Palanque-Delabrouille, C. Yèche, C. Magneville, and M. Viel, Lyman-alpha forests cool warm dark matter, *J. Cosmol. Astropart. Phys.* **08** (2016) 012.
- [15] A. Garzilli, A. Magalich, T. Theuns, C. S. Frenk, C. Weniger, O. Ruchayskiy, and A. Boyarsky, The Lyman- α forest as a diagnostic of the nature of the dark matter, *Mon. Not. R. Astron. Soc.* **489**, 3456 (2019).
- [16] A. Garzilli, A. Magalich, O. Ruchayskiy, and A. Boyarsky, How to constrain warm dark matter with the Lyman- α forest, *Mon. Not. R. Astron. Soc.* **502**, 2356 (2021).
- [17] J. S. Bullock and M. Boylan-Kolchin, Small-scale challenges to the Λ CDM paradigm, *Annu. Rev. Astron. Astrophys.* **55**, 343 (2017).
- [18] Z. Lukić, C. W. Stark, P. Nugent, M. White, A. A. Meiksin, and A. Almgren, The Lyman α forest in optically thin hydrodynamical simulations, *Mon. Not. R. Astron. Soc.* **446**, 3697 (2015).
- [19] 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, S. López, S. Ellison, L. Christensen, K. D. Denney, and G. Worseck, New constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman- α forest data, *Phys. Rev. D* **96**, 023522 (2017).
- [20] J. S. Bolton, E. Puchwein, D. Sijacki, M. G. Haehnelt, T.-S. Kim, A. Meiksin, J. A. Regan, and M. Viel, The Sherwood simulation suite: Overview and data comparisons with the Lyman α forest at redshifts $2 \leq z \leq 5$, *Mon. Not. R. Astron. Soc.* **464**, 897 (2017).
- [21] M. Walther, E. Armengaud, C. Ravoux, N. Palanque-Delabrouille, C. Yèche, and Z. Lukić, Simulating intergalactic gas for DESI-like small scale Lyman α forest observations, *J. Cosmol. Astropart. Phys.* **04** (2021) 059.
- [22] E. Boera, G. D. Becker, J. S. Bolton, and F. Nasir, Revealing reionization with the thermal history of the intergalactic medium: New constraints from the Ly α flux power spectrum, *Astrophys. J.* **872**, 101 (2019).
- [23] B. Villaseñor, B. Robertson, P. Madau, and E. Schneider, Effects of photoionization and photoheating on Ly α forest properties from CHOLLA cosmological simulations, *Astrophys. J.* **912**, 138 (2021).
- [24] B. Villaseñor, B. Robertson, P. Madau, and E. Schneider, Inferring the thermal history of the intergalactic medium from the properties of the hydrogen and helium Lyman-alpha forest, *Astrophys. J.* **933**, 59 (2022).
- [25] P. Gaikwad, M. Rauch, M. G. Haehnelt, E. Puchwein, J. S. Bolton, L. C. Keating, G. Kulkarni, V. Iršič, E. Bañados, G. D. Becker, E. Boera, F. S. Zahedy, H.-W. Chen, R. F. Carswell, J. Chardin, and A. Rorai, Probing the thermal state of the intergalactic medium at $z > 5$ with the transmission

- spikes in high-resolution Ly α forest spectra, *Mon. Not. R. Astron. Soc.* **494**, 5091 (2020).
- [26] L. Hui, N. Y. Gnedin, and Y. Zhang, The statistics of density peaks and the column density distribution of the Ly α forest, *Astrophys. J.* **486**, 599 (1997).
- [27] N. Y. Gnedin and L. Hui, Probing the Universe with the Ly α forest—I. Hydrodynamics of the low-density intergalactic medium, *Mon. Not. R. Astron. Soc.* **296**, 44 (1998).
- [28] N. Y. Gnedin, E. J. Baker, T. J. Bethell, M. M. Drosback, A. G. Harford, A. K. Hicks, A. G. Jensen, B. A. Keeney, C. M. Kelso, M. C. Neyrinck, S. E. Pollack, and T. P. van Vliet, Linear gas dynamics in the expanding universe, *Astrophys. J.* **583**, 525 (2003).
- [29] M. S. Peeples, D. H. Weinberg, R. Davé, M. A. Fardal, and N. Katz, Pressure support versus thermal broadening in the Lyman α forest—I. Effects of the equation of state on longitudinal structure, *Mon. Not. R. Astron. Soc.* **404**, 1281 (2010).
- [30] G. Kulkarni, J. F. Hennawi, J. Oñorbe, A. Rorai, and V. Springel, Characterizing the pressure smoothing scale of the intergalactic medium, *Astrophys. J.* **812**, 30 (2015).
- [31] F. Nasir, J. S. Bolton, and G. D. Becker, Inferring the IGM thermal history during reionization with the Lyman α forest power spectrum at redshift $z = 5$, *Mon. Not. R. Astron. Soc.* **463**, 2335 (2016).
- [32] N. Y. Gnedin and A. J. S. Hamilton, Matter power spectrum from the Lyman-alpha forest: Myth or reality?, *Mon. Not. R. Astron. Soc.* **334**, 107 (2002).
- [33] M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese, and A. Riotto, Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman- α forest, *Phys. Rev. D* **71**, 063534 (2005).
- [34] E. E. Schneider and B. E. Robertson, CHOLLA: A new massively parallel hydrodynamics code for astrophysical simulation, *Astrophys. J. Suppl. Ser.* **217**, 24 (2015).
- [35] B. D. Smith, G. L. Bryan, S. C. O. Glover, N. J. Goldbaum, M. J. Turk, J. Regan, J. H. Wise, H.-Y. Schive, T. Abel, A. Emerick, B. W. O’Shea, P. Anninos, C. B. Hummels, and S. Khochfar, GRACKLE: A chemistry and cooling library for astrophysics, *Mon. Not. R. Astron. Soc.* **466**, 2217 (2017).
- [36] O. Hahn and T. Abel, Multi-scale initial conditions for cosmological simulations, *Mon. Not. R. Astron. Soc.* **415**, 2101 (2011).
- [37] A. V. Macciò, O. Ruchayskiy, A. Boyarsky, and J. C. Muñoz-Cuartas, The inner structure of haloes in cold + warm dark matter models, *Mon. Not. R. Astron. Soc.* **428**, 882 (2013).
- [38] E. Puchwein, F. Haardt, M. G. Haehnelt, and P. Madau, Consistent modelling of the meta-galactic UV background and the thermal/ionization history of the intergalactic medium, *Mon. Not. R. Astron. Soc.* **485**, 47 (2019).
- [39] J. S. Bolton, G. D. Becker, M. G. Haehnelt, and M. Viel, A consistent determination of the temperature of the intergalactic medium at redshift $z = 2.4$, *Mon. Not. R. Astron. Soc.* **438**, 2499 (2014).
- [40] H. Hiss, M. Walther, J. F. Hennawi, J. Oñorbe, J. M. O’Meara, A. Rorai, and Z. Lukić, A new measurement of the temperature–density relation of the IGM from Voigt profile fitting, *Astrophys. J.* **865**, 42 (2018).
- [41] M. Walther, J. Oñorbe, J. F. Hennawi, and Z. Lukić, New constraints on IGM thermal evolution from the Ly α forest power spectrum, *Astrophys. J.* **872**, 13 (2019).
- [42] P. Gaikwad, R. Srianand, M. G. Haehnelt, and T. R. Choudhury, A consistent and robust measurement of the thermal state of the IGM at $2 \leq z \leq 4$ from a large sample of Ly α forest spectra: Evidence for late and rapid He II reionization, *Mon. Not. R. Astron. Soc.* **506**, 4389 (2021).
- [43] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [44] R. de Belsunce, S. Gratton, W. Coulton, and G. Efstathiou, Inference of the optical depth to reionization from low multipole temperature and polarization Planck data, *Mon. Not. R. Astron. Soc.* **507**, 1072 (2021).
- [45] A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, and C. Balland (DESI Collaboration), The DESI experiment part I: Science, targeting, and survey design, [arXiv:1611.00036](https://arxiv.org/abs/1611.00036).
- [46] M. M. Pieri, S. Bonoli, J. Chaves-Montero, I. Pâris, M. Fumagalli, J. S. Bolton, M. Viel, P. Noterdaeme, J. Miralda-Escudé, N. G. Busca, H. Rahmani, C. Peroux, A. Font-Ribera, and S. C. Trager, WEAVE-QSO: A massive intergalactic medium survey for the William Herschel Telescope, in *SF2A-2016: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, edited by C. Reylé, J. Richard, L. Cambrésy, M. Deleuil, E. Pécontal, L. Tresse, and I. Vauglin (2016), pp. 259–266, [arXiv:1611.09388](https://arxiv.org/abs/1611.09388).
- [47] J. Amiaux *et al.*, Euclid mission: Building of a reference survey, in *Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8442, edited by M. C. Clampin, G. G. Fazio, H. A. MacEwen, J. Oschmann, and M. Jacobus (2012), p. 84420Z, [arXiv:1209.2228](https://arxiv.org/abs/1209.2228).
- [48] M. Molaro, V. Iršič, J. S. Bolton, L. C. Keating, E. Puchwein, P. Gaikwad, M. G. Haehnelt, G. Kulkarni, and M. Viel, The effect of inhomogeneous reionization on the Lyman α forest power spectrum at redshift $z > 4$: Implications for thermal parameter recovery, *Mon. Not. R. Astron. Soc.* **509**, 6119 (2022).
- [49] X. Wu, M. McQuinn, R. Kannan, A. D’Aloisio, S. Bird, F. Marinacci, R. Davé, and L. Hernquist, Imprints of temperature fluctuations on the $z \sim 5$ Lyman- α forest: A view from radiation-hydrodynamic simulations of reionization, *Mon. Not. R. Astron. Soc.* **490**, 3177 (2019).