PHYSICAL REVIEW D **88,** 043502 (2013)

Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman-α forest data

Matteo Viel,[1,2] George D. Becker,[3] James S. Bolton,[4] and Martin G. Haehnelt[3]
[1]*INAF - Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, I-34131 Trieste, Italy*
[2]*INFN/National Institute for Nuclear Physics, Via Valerio 2, I-34127 Trieste, Italy*
[3]*Kavli Institute for Cosmology and Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, United Kingdom*
[4]*School of Physics and Astronomy, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom*
(Received 11 June 2013; published 5 August 2013)

We present updated constraints on the free-streaming of warm dark matter (WDM) particles derived from an analysis of the Lyman-α flux power spectrum measured from high-resolution spectra of 25 $z > 4$ quasars obtained with the Keck High Resolution Echelle Spectrometer and the Magellan Inamori Kyocera Echelle spectrograph. We utilize a new suite of high-resolution hydrodynamical simulations that explore WDM masses of 1, 2 and 4 keV (assuming the WDM consists of thermal relics), along with different physically motivated thermal histories. We carefully address different sources of systematic error that may affect our final results and perform an analysis of the Lyman-α flux power with conservative error estimates. By using a method that samples the multidimensional astrophysical and cosmological parameter space, we obtain a lower limit $m_{\mathrm{WDM}} \gtrsim 3.3$ keV (2σ) for warm dark matter particles in the form of early decoupled thermal relics. Adding the Sloan Digital Sky Survey Lyman-α flux power spectrum does not improve this limit. Thermal relics of masses 1, 2 and 2.5 keV are disfavored by the data at about the 9σ, 4σ and 3σ C.L., respectively. Our analysis disfavors WDM models where there is a suppression in the linear matter power spectrum at (nonlinear) scales corresponding to $k = 10h/\mathrm{Mpc}$ which deviates more than 10% from a Lambda cold dark matter model. Given this limit, the corresponding "free-streaming mass" below which the mass function may be suppressed is $\sim 2 \times 10^8 h^{-1} \mathrm{M}_{\odot}$. There is thus very little room for a contribution of the free-streaming of WDM to the solution of what has been termed the small scale crisis of cold dark matter.

DOI: 10.1103/PhysRevD.88.043502 PACS numbers: 98.80.Cq, 95.35.+d, 98.62.Ra

I. INTRODUCTION

The Lambda cold dark matter (ΛCDM) paradigm, in many respects, has proven to been an immensely successful cosmological model. ΛCDM is based on a cosmological constant plus "cold" dark matter, i.e. dark matter particles whose streaming velocities are negligible for most astrophysical considerations. On large scales, the *Planck* mission has just delivered another ringing endorsement of this model with its first-year cosmology results [1]. On scales below a few (comoving) Mpc, however, the matter power spectrum is still difficult to probe, and it has been repeatedly suggested that dark matter is perhaps "warm," with a free-streaming length that affects the properties of low-mass (dwarf) galaxies. Warm dark matter (WDM) could alleviate the apparent difficulties of ΛCDM models in reproducing some observations related to the matter power spectrum on scales of a few Mpc and below. The most notable possible tensions under ΛCDM are the excess of the number of galactic satellites, the cuspiness and high (phase-space) density of galactic cores, the luminosities of the Milky Way's satellites and the properties of galaxies filling voids (e.g. [2–5]).

The main effect of the larger velocities of WDM particles, and the resulting significant free-streaming length, would be to suppress structures on Mpc scales and below. The last few years have seen a reintensified discussion of this possibility, particularly in light of improvements in numerical models and observations of the mass and internal structure of Local Group satellites [6]. It has been suggested that a free-streaming length corresponding to that of a thermal relic WDM particle with a mass of 1–2 keV (and in some cases as low as 0.5 keV) provides better agreement between the most recent data and numerical simulations [7,8]. The difficulties associated with the cold dark matter paradigm, however, arise on scales where the matter spectrum is highly nonlinear at $z \sim 0$, and where very uncertain baryonic physics is known to play an important role [9–11].

The Lyman-α absorption produced by intergalactic neutral hydrogen in the spectra of distant quasars (QSOs)—the so called "Lyman-α forest"—provides a powerful alternative tool for constraining dark matter properties, particularly the free-streaming of dark matter particles on the scales in question. The Lyman-α forest probes the matter power spectrum in the mildly nonlinear regime over a large range of redshifts ($z = 2$–6 in ground-based data) down to the small scales of interest (1–80h^{-1} Mpc) [12,13]. On large scales the SDSS-III BOSS Collaboration has recently measured the baryonic acoustic oscillations (BAOs) scale in the 3D correlation function of the Lyman-α forest in $\sim$50 000 QSOs at $z \sim 2.2$ [14]. These findings further emphasize the value of the Lyman-α forest as a tracer of cosmological large-scale structure. Constraints on the matter power spectrum from Lyman-α forest data on small scales are only limited by the thermal cutoff in the flux power spectrum introduced by pressure and thermal motions of baryons in the photoionized intergalactic medium (IGM). The IGM has a characteristic temperature of $\sim 10^4$ K. While not trivial, modeling the relevant physics with numerical simulations is reasonably straightforward; the power spectrum at the relevant redshifts ($z \sim 2$–5) and scales is only mildly nonlinear and stellar feedback effects are much less important than at lower redshifts [15,16].

The basic property of WDM, which impacts on both large-scale structure formation and the internal structure of dark matter haloes and the galaxies they are hosting, is the significant "thermal" velocities of the WDM particles (see [2]). The resulting "free-streaming" eliminates density fluctuations on scales below a characteristic comoving wave number:

$$k_{\rm FS} \sim 15.6 \frac{h}{\rm Mpc}\left(\frac{m_{\rm WDM}}{1\ {\rm keV}}\right)^{4/3}\left(\frac{0.12}{\Omega_{\rm DM}h^2}\right)^{1/3}, \qquad (1)$$

and leads to a very distinctive cutoff in the matter power spectrum at a corresponding scale. For example, the wave number at which the linear WDM suppression reaches 50% in terms of matter power, $k_{1/2}$, with respect to the ΛCDM case can be approximated as

$$k_{1/2} \sim 6.5 \frac{h}{\rm Mpc}\left(\frac{m_{\rm WDM}}{1\,{\rm keV}}\right)^{1.11}\left(\frac{\Omega_{\rm DM}}{0.25}\right)^{-0.11}\left(\frac{h}{0.7}\right)^{1.22}, \qquad (2)$$

where this equation uses the numerical results of Ref. [17]. For standard thermal relics, the shape of the cutoff is therefore well characterized in the linear regime and there is an unambiguous relation between the mass of the thermal relic WDM particle and a well-defined free-streaming length (e.g. [17]). Note that we will also quote a free-streaming mass, which is the mass at the mean density enclosed in a half-wavelength mode corresponding to $k_{1/2}$.

For the sake of simplicity, the analysis here is presented in terms of the mass of a thermal relic dark matter particle, for which there is a one-to-one correspondence between the free-streaming length and the particle mass. We should point out, however, that in recent years sterile neutrinos and other nonthermal particles have become popular WDM candidates. Some of these models are actually more similar to mixed dark matter models with cold and warm dark matter components. The shape of the free-streaming "cutoff" can then be quite different from that of a thermal relic, and may instead correspond to a downward step in the power spectrum rather than a cutoff (see [8,18]). There is also no universal relation between free-streaming length and mass of the WDM particles in these models, and the normalization and functional form of this relation varies greatly between different nonthermal WDM candidates. Unfortunately, this has led to considerable confusion in the literature when WDM models, characterized by their model-dependent WDM particle masses, are compared between each other and/or thermal relic models and in particular with Lyman-α forest data. For example Ref. [8] quotes a sterile neutrino mass of 2 keV for their thermal relic WDM model which corresponds, however, to a thermal relic mass of 1.4 keV. For convenience and ease of comparison with the literature, in this work we therefore consider only matter power spectra with cutoff shapes expected for thermal relic WDM particles and quote the unambiguous thermal relic masses (and corresponding cutoff scale) to characterize our WDM models.

The Lyman-α forest, due to its spectral nature, probes the matter power spectrum in velocity space. With increasing redshift the ratio of a given (comoving) free-streaming length in velocity space to the thermal cutoff length scale at a given temperature increases as $\propto (1+z)^{1/2}$. There is furthermore strong observational evidence that the temperature of the IGM decreases toward higher redshift over the range $3 < z < 5$ [19], and therefore a corresponding decrease in the thermal cutoff length scale. Despite observational difficulties, pushing to high redshift allows the models to probe smaller free-streaming lengths and thus to improve the limits on WDM masses. The high redshift regime has also the advantage of probing structures that are more linear and the WDM cutoff is more prominent at high redshift compared to low redshift [10].

Lyman-α forest data to constrain WDM were first used in Ref. [20] where a limit of 750 eV was obtained by using N-body simulations only. In previous work, Ref. [17], we used instead two samples of high-resolution QSO Lyman-α forest spectra at $z \sim 2.5$ to set a lower limit of 550 eV for the mass of a thermal WDM candidate. Following this, Ref. [21,22], using higher-redshift QSO spectra from the Sloan Digital Sky Survey (SDSS) and applying a different analysis method, significantly improved this limit by a factor ~ 4. As already noted, however, care has to be taken in the correct modeling of the free-streaming properties of "nonthermal" candidate WDM particles of a given model-dependent mass, such as the popular sterile neutrino. In Ref. [18] the authors have focused on constraints on a range of such models. Because of a nonzero mixing angle between active and sterile flavor states, x-ray flux observations can also constrain the abundance and decay rate of such WDM particles (e.g. [23]). The joint constraints from Lyman-α forest data and those from the x-ray fluxes of astrophysical objects now put considerable tension on the parameter space allowed for a sterile neutrino particle with the phase-space distribution proposed by Dodelson and Widrow [24,25], although other, possibly more physical scenarios should be explored [18,26].

In Ref. [27] we presented the most stringent Lyman-α forest limits up to that date on the free-streaming of dark matter, $m_{\rm WDM} > 4$ kev (2σ). That analysis was based on an (at the time) unrivaled sample of high-quality, high-resolution QSO absorption spectra extending to $z \sim 5.5$. The limit is in obvious conflict with many of the recent suggestions for alleviating the difficulties encountered by numerical models in reproducing the observed properties of Local Group satellite galaxies within the cold dark matter paradigm. These models often assume dark matter to be made up by thermal relic WDM with masses in the range 0.5–2 keV (e.g. [8]).

Since our study in [27], the size of our high-quality, high redshift QSO absorption spectra sample, the quality and size (in particular the dynamic range and resolution) of our numerical simulations and our knowledge of the thermal and ionization state of the IGM at the relevant redshifts have all significantly improved. Motivated by these improvements, and in light of the lively debate of dark matter possibly being warm with masses in the range 0.5–2 keV, we present here a new and much more extensive study of the high redshift Lyman-α forest constraints on the free-streaming properties of dark matter. The new study is based on an improved data set, further refined modeling of the flux power spectrum and a large suite of new numerical hydrodynamical simulations. We also perform a comprehensive investigation of the systematic uncertainties related to this measurement.

Finally, it is worth highlighting that WDM would have profound implications in many astrophysical and cosmological contexts. In this respect, IGM constraints are highly complementary to other probes based, for example, on the properties of dark matter haloes [28], the number of satellites and their luminosities [3,29], strong lensing, the velocity function in the local environment [30], phase-space density constraints [31], the formation of the first stars [32], the high redshift quasar luminosity function [33], decays of WDM particles in the high redshift universe [34], reionization [35], gamma ray bursts [36], galaxy formation aspects [37] using N-body/hydrodynamical simulations [38,39] or analytical/semianalytical methods [40–43].

The paper is organized as follows. In Sec. II we present our new data set. The simulations are described in Sec. III. The mock quasar sample, which will be important for estimating error amplitude and covariance, is introduced in Sec. IV. Section V discusses the effect of the most important physical parameters on the flux power spectrum, while most of the remaining nuisance parameters and the impact they have in terms of flux power are discussed in an Appendix. Our main results are reported in Sec. VI, together with a description of the Monte Carlo sampling of the likelihood space. We summarize our findings and conclude in Sec. VIII.

II. DATA

Our analysis is based on high-resolution spectra of 25 quasars with emission redshifts $4.48 \leq z_{\rm em} \leq 6.42$. Compared to our previous analysis in Ref. [27] the number of QSO spectra, at these redshifts, has improved by nearly a factor 2. Spectra for fourteen of the objects were taken with the Keck High Resolution Echelle Spectrometer (HIRES) [44], and the remaining eleven were taken with the Magellan Inamori Kyocera Echelle (MIKE) spectrograph on the Magellan Clay telescope [45]. Most of the data have been presented elsewhere [19,46–48]. Here we briefly review the relevant features of the spectra, and describe how the flux power spectra were calculated.

The majority of spectra were reduced using a custom set of interactive data language routines based on optimal sky subtraction [49] and optimal extraction [50] techniques, while a small subset of the HIRES spectra (PSS 0248 + 1802 and BR 1202 − 0725) were reduced using the MAKEE software package. The HIRES and MIKE spectra have spectral resolutions of 6.7 and 13.6 $\mathrm{km\,s^{-1}}$ (FWHM), and the spectra were extracted using 2.1 and 5.0 $\mathrm{km\,s^{-1}}$ spectral bins, respectively. The one-dimensional relative flux-calibrated spectra were then continuum normalized using spline fits based on power-law extrapolations of the continuum redward of the Lyman-α emission line. Median continuum signal-to-noise (S/N) ratios within the Lyman-α forest of each object are typically in the range of 10–20 per pixel. The continuum estimates are necessarily crude due to the high levels of absorption in the Lyman-α forest at these redshifts. We estimate that typical uncertainties in the continuum are of the order ~10%–20%, a point we return to in the power spectrum analysis.

To compute the flux power spectra, we first divided the Lyman-α forest in each quasar spectrum into two regions of equal redshift length. We then computed the power spectrum of the fractional transmission, $\delta_F(z)$, in each region separately, where

$$\delta_F(z) = \frac{F(z) - \langle F(\bar{z})\rangle}{\langle F(\bar{z})\rangle}. \tag{3}$$

Here, $\langle F(\bar{z})\rangle$ is the mean transmitted flux calculated at the mean redshift of each region. We used fixed relations for the mean flux given by $\langle F(z)\rangle = \exp[-\tau_{\rm eff}(z)]$, where

$$\tau_{\rm eff}(z) = \begin{cases} 0.751\left(\frac{1+z}{4.5}\right)^{2.90} - 0.132, & z \leq 4.5 \\ 2.26\left(\frac{1+z}{6.2}\right)^{4.91}, & z > 4.5. \end{cases} \tag{4}$$

The fit to $\tau_{\rm eff}$ at $z \leq 4.5$ is from [51], while the evolution at $z > 4.5$ is based on a fit to the mean flux measured from the data presented here. The latter is similar to the trend in $\tau_{\rm eff}(z)$ presented by [52]. We note that our analysis is not sensitive to our choice of using a fixed relation for $\langle F(\bar{z})\rangle$. In tests where we instead divided each region by the mean flux in that region alone we obtained very similar power spectrum estimates on average. When calculating the flux power spectrum we do not attempt to mask metal lines (see discussion below). We do, however, mask regions of strong telluric absorption (6275–6315, 6865–6939, 7594–7700, 7163–7313, and 8126–8328 Å). The longest-wavelength mask effectively means that we probe up to a maximum redshift of 5.684.

The power spectra for individual regions were averaged over ten logarithmic wave number bins in the range $\log_{10}k(\mathrm{s/km}) = [-2.9, -1.1]$ with 0.2 dex spacing. The power spectra from all regions were then further averaged according to instrument and the mean redshift in each region. We used median redshifts $z = 4.2$, 4.6, 5, 5.4 for a nominal total of eight combined power spectra and 80 data points. In order to be conservative, however, we decided to use a subset of this sample and do not consider the highest redshift bin for the MIKE data set (which has very large error bars) or the flux power measurements at $\log_{10}k(\mathrm{s/km}) < -2.3$, which might be affected by continuum fitting uncertainties. The final data set used in the present analysis thus consists of 49 data points.

Preliminary estimates of the error in the power spectra were calculated using a bootstrap approach. It is known, however, that bootstrapping typical underestimates the true errors (e.g. [53]). To be conservative, we therefore decided to add an additional 30% uncertainty to our estimates of the errors of the observed flux power spectrum for our standard analysis. We will also quote the tighter limits that would be obtained without this increase of the error estimate. As a further check, we used the set of mock QSO spectra described in Sec. IV to determine what the expected covariance in the power spectra should be (within the limits of our finite simulation box) at each redshift for a sample of similar size and quality to the one used here. These estimate were used to correct a few error estimates in the real data that appeared to be too small. With these corrections, the final flux power spectra used here have error bars that are larger than $\sigma(P_\mathrm{F})/P_\mathrm{F} > 0.075$.

III. COSMOLOGICAL HYDRODYNAMICAL SIMULATIONS

We model the flux power spectrum based on a set of hydrodynamical simulations performed with a modification of the publicly available GADGET-II code. This code implements a simplified star formation criterion [54] that turns all gas particles that have an overdensity above 1000 and a temperature below 10^5 K into star particles. This has been first used and extensively tested in Ref. [55].

The reference model, hereafter referred to as (20,512), is a box of length $20h^{-1}$ comoving Mpc with 2×512^3 gas and cold DM particles (with a gravitational softening length of $1.3h^{-1}$ kpc) in a flat ΛCDM universe with cosmological parameters $\Omega_\mathrm{m} = 0.274$, $\Omega_\mathrm{b} = 0.0457$, $n_\mathrm{s} = 0.968$, $H_0 = 70.2\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$ and $\sigma_8 = 0.816$, in agreement both with WMAP-9 yr and Planck data [1,56]. We further explore three different WDM models with masses $m_\mathrm{WDM} = 1$, 2, 4 keV; these models correspond to 50% suppression of power in the linear (redshift independent) matter power spectrum at scales $k_{1/2} \sim 6.9$, 14.7, $32h/\mathrm{Mpc}$, respectively. The initial condition power spectra are generated with CAMB [57] and the suppression and velocity for the WDM particles are implemented using the approach outlined in Ref. [17]. In order to assess convergence and evaluate resolution corrections (which are model dependent), we also perform four additional (20,768) models, one each for the reference and WDM cases, and a single (20,896) model for the 2 keV simulation only. We also have performed three (60,512) simulations for ΛCDM, WDM 1 keV and WDM 2 keV in order to check the flux power spectrum convergence at the largest scales of our smaller boxes.

We explore the impact of different thermal histories on the Lyman-α forest by modifying the ultraviolet (UV) background photoheating rate in the simulations (e.g. [58]). A power-law temperature-density relation, $T = T_0(1+\delta)^{\gamma-1}$, arises in the low density IGM ($1 + \delta < 10$) as a natural consequence of the interplay between photoheating and adiabatic cooling [59]. We consider a range of values for the temperature at mean density, T_0, and the power-law index of the temperature-density relation, γ, based on the observational measurements presented recently by Ref. [19]. These consist of a set of 3 different indices for the temperature-density relation, $\gamma(z = 4.6) \sim 1.0$, 1.3, 1.6, that are kept approximately constant over the redshift range $z = [4.2\text{–}5.6]$ and 3 different temperatures at mean density, $T_0(z = 4.6) \sim 5400$, 8300, 11200 K, which evolve with redshift, yielding a total of 9 different thermal histories. The reference thermal history assumes $(T_0(z = 4.6), \gamma(z = 4.6)) = (8300\ \mathrm{K}, 1.3)$. These 9 thermal histories have been performed for all the 3 WDM models and for the reference ΛCDM case, resulting in a total of 36 simulations.

In addition to these parameters we also consider and vary several other physical parameters for the reference model only, given that these are poorly constrained by the data. These are the redshift of reionization z_re (i.e. the redshift at which the optically thin UV background is switched on in the simulations) which is chosen to be $z_\mathrm{re} = 12$ for the reference case and $z_\mathrm{re} = 8$, 16 for two additional models; the Hubble constant, with two extra simulations with $H_0 = 66.2$, $74.2\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$; the scalar spectral index, with $n_\mathrm{s} = 0.968$, 0.998; the matter content, with $\Omega_\mathrm{m} = 0.24$, 0.30 and the rms amplitude of the matter power spectrum, with $\sigma_8 = 0.77$, 0.87. We note here that varying the redshift of reionization in particular enables us to assess the impact of different integrated thermal histories (i.e. the effect of Jeans smoothing, see [60] for a recent discussion) on our analysis. The effect of different integrated thermal histories on Lyman-α forest constraints was also considered in Ref. [61] using this parametrization. Overall, a total of 54 hydrodynamical simulations have been performed. Approximately 4000 core hours were required for each (20,512) run to reach $z = 2$, with the higher resolution simulations requiring around 5 times longer.

During the simulation runs, we extract the nonlinear matter power spectra in order to compare with [10]. For the reference case only, we additionally extract the position of the haloes with a friends-of-friends halo finding algorithm for our model of the impact of spatial fluctuations in the UV background on the flux power (see the Appendix for further details).

Lastly, we note that the physical properties of the Lyman-α forest obtained from the TREEPM/SPH code GADGET-II are in very good agreement at the percent level with those inferred from the moving-mesh code AREPO [62] and with the Eulerian code ENZO [63].

IV. THE MOCK QSO SAMPLE

The simulated Lyman-α forest spectra are extracted along 5000 random line-of-sights (LOSs) after interpolation of the relevant physical quantities along the LOSs using the smoothed particle hydrodynamics (SPH) formalism. Box-size effects on the flux power are estimated with (60,512) simulations. Note, however, that the necessary box-size correction is below the percent level at the largest scales used. Resolution corrections are, however, important. The flux power spectra are corrected for resolution effects using the (20,768) simulations (see the Appendix for further details).

The mean flux is varied *a posteriori*, after having extracted the spectra, by reproducing 0.8, 1, 1.2 times the observed $\tau_{\rm eff}$ (see Appendix). At the end of the procedure the four-dimensional parameter space in ($m_{\rm WDM}$, $\tau_{\rm eff}$, T_0, γ) is explored fully by means of quadrilinear interpolation performed over the set of 36 hydrodynamical simulations and 108 (36×3 mean flux values) flux models.

In order to get a better understanding of the expected (co)variance properties of the observed data we generate samples of mock QSO absorption spectra which resemble the observational data as closely as possible. The procedure used to create the mock spectra can be summarized as follows: (i) we consider the total redshift path in each redshift bin and combine the short simulated spectra (20 Mpc/h in length) to match the total length of an observed QSO spectrum (approximately 40 spectra are used); (ii) we allow for an optical depth evolution along the LOS (which is absent since our simulated spectra are from snapshots at fixed redshifts) following the scaling expected from the fluctuating Gunn-Peterson approximation, $\tau \propto (1+z)^{4.5}$ (see e.g. [64]); (iii) for each short simulated spectrum we consider a $\pm 20\%$ error on the quasar continuum placement (the continuum is drawn randomly from a Gaussian distribution around the value 1 with a $\sigma = 0.2$); (iv) we smooth the flux with a Gaussian at a given FWHM corresponding to the spectrograph resolution and rebin the spectra with the observed pixel-size; (v) we add Gaussian-distributed noise on top of the flux, matching the signal-to-noise of the observational data. We demonstrate in the Appendix that the (instrumental) effects of noise and finite resolution, which are scale and redshift dependent, are below 20% (6%) at the smallest scales for MIKE (HIRES).

The (co)variance properties of this mock sample are in reasonable agreement with those of our observed sample, both as a function of redshift and wave number. There are only 5 data points that appear to have error bars that are smaller than those obtained from the mock sample: 4 data points from the MIKE sample [$\log k({\rm s/km}) = -1.5$ at $z = 4.2$, $\log k({\rm s/km}) = -2.1$ at $z = 4.6$, $\log k({\rm s/km}) = -1.3, -1.1$ at $z = 5$] and one data point from the HIRES sample ($\log k = -1.1$ at $z = 4.2$). As the observed sample is still small and it is thus expected that the bootstrap errors estimated from the data could be unrealistically small, for these data points we increase the error bars to match those obtained from a mock sample of 30 QSOs in the same redshift bin. These mock data points, with the extra 30% error added, otherwise agree well with those of the observational data, giving us confidence that this is reasonable.

Finally, we note that the hydrodynamical simulations used to construct our mock Lyman-α forest spectra do not incorporate chemical elements other than hydrogen and helium. We have therefore estimated how unidentified, lower redshift metal lines in the Lyman-α forest may bias our result, and in particular how these narrow absorption lines may alter the flux power spectrum at small scales. Furthermore, the simulations also assume a spatially uniform UV background. We therefore also estimate the impact of spatial fluctuations in the UV background on the Lyman-α forest at $z = 4.2$–5.4. We follow a modified version of the approach described in Ref. [65] for exploring fluctuations in the HeIIionizing background at lower redshift, to which we refer the reader for further details. These two systematics effects will also be discussed more extensively in the Appendix.

V. THE FLUX POWER SPECTRUM

A. The WDM and thermal cutoffs

In this section we demonstrate the distinctively different effects that the thermal (i.e. due to the temperature of the photoionized IGM) and WDM cutoff have on the flux power spectrum. We will also check for possible effects due to the limited numerical resolution of our simulations. Firstly, however, it is instructive to consider Fig. 1, where we show the ratio of the nonlinear matter power spectra in the WDM and ΛCDM simulations. The results are shown at three different redshifts, $z = (3, 4.2, 5.4)$. The redshift range $z = 4.2$–5.2 brackets that of the high-resolution data set used in this work. We additionally present the matter power spectrum at $z = 3$ to show the evolution of the nonlinear power at lower redshift. The three WDM models are reported as orange, blue and black curves for masses of 4,2 and 1 keV, respectively, while the green curve shows the *linear* suppression for the 2 keV case taken from Ref. [17]. The power spectra are already clearly somewhat nonlinear at high redshift; the blue and green curves start to

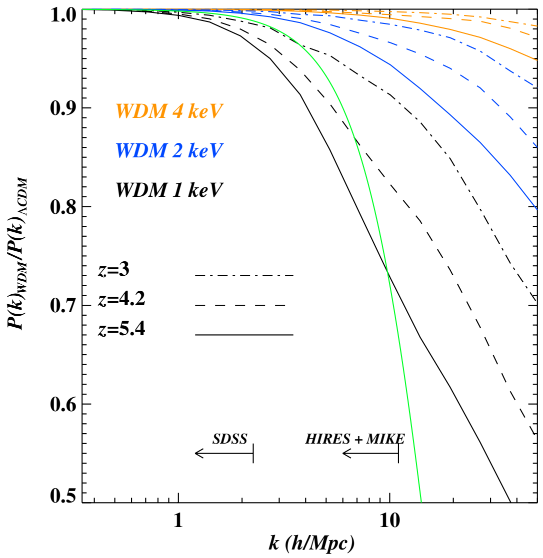

FIG. 1 (color online). Ratio between the 3D nonlinear matter power spectrum of 3 different WDM models (1, 2 and 4 keV, black, blue and orange curves) at 3 different redshifts ($z = 3$, 4.2, 5.4, represented by the dot-dashed, dashed and continuous curves) and the corresponding ΛCDM model. The green curve represents the linear redshift independent suppression in terms of matter power for a $m_{\rm WDM} = 2$ keV model obtained using Eq. (6) of Ref. [17]. The arrows in the bottom part of the figure indicate the maximum value of the wave numbers probed by the SDSS data and by the data set used in the present analysis. This figure refers to the reference (20,512) simulations.

differ significantly at small scales at $k > 3h/\mathrm{Mpc}$. In the bottom part of the panel we show the approximate wave number ranges that are probed by SDSS and the HIRES + MIKE data set used in our analysis. Note that the nonlinear matter suppression is in good agreement with the fitting formula presented in Ref. [10].

In Fig. 2 we qualitatively compare a set of noiseless Lyman-α forest spectra extracted from the ΛCDM, WDM 1 keV and WDM 2 keV models, represented by the green, black and blue curves respectively. It is clear that the amount of small scale substructure in the transmitted flux in the ΛCDM is more prominent with respect to the WDM cases. In the rest of this section we will quantify these differences in terms of the 1D flux power spectrum.

We now turn to Fig. 3, which shows the ratio between the 1D flux power of the WDM and ΛCDM models for the four different redshift bins used in the present analysis (note that we compute the power spectrum of the quantity $\delta_{\rm F} = F/\langle F\rangle - 1$, and we refer to this as the flux power). The suppression of the flux power is larger than that seen in the matter power spectrum. This is due to the fact that the 1D matter power spectrum is an integral of the 3D power spectrum and therefore very sensitive to the small scale cutoff. As expected, the largest differences exist between the 1 keV (black curves) and the ΛCDM model. Note that the flux power also changes at large scales; the requirement of reproducing the same observed mean flux value [given by Eq. (4)] results in an increase of the power at those scales (the power spectrum of the WDM *flux* F, not $\delta_{\rm F}$, does show suppression over all scales when compared to ΛCDM). Furthermore, we also note that there is a substantial redshift evolution of the flux power between $z = 5.4$ and $z = 4.2$.

Numerical convergence for WDM simulations can be particularly difficult to achieve (see Ref. [38]). In Fig. 4 we demonstrate that at the resolution and WDM masses considered in this work, this should, however, not be an issue. Figure 4 compares the flux power extracted from the (20,512) and (20,768) 1 and 2 keV simulations to the corresponding ΛCDM simulations at the same resolution. The agreement between the different resolution simulations is very good, typically at the percent level. The differences are largest for the 1 keV case at the smallest scales probed by our data in the current analysis, where they reach the 10% level. The simulated flux power spectra for both ΛCDM and WDM models have therefore been corrected for resolution effects by multiplying the raw power spectra by the ratio of the results from the (20,768) and (20,512) simulations. In general, we find the requirements for reaching numerical convergence in terms of flux power are more demanding for absolute values of the flux power rather than ratios of different

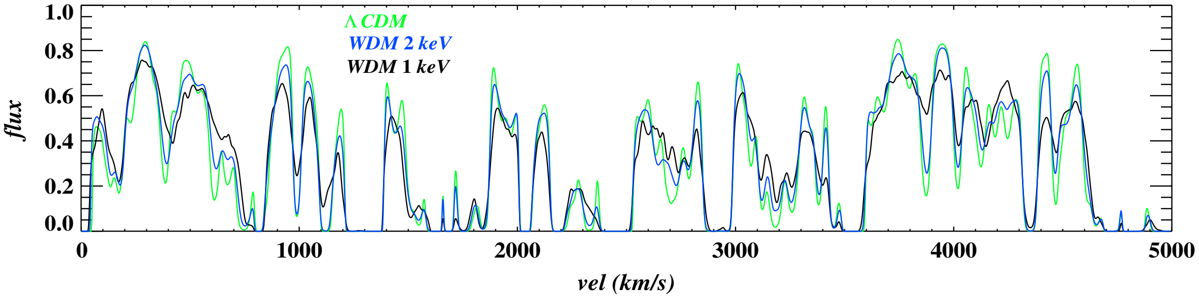

FIG. 2 (color online). Transmitted flux along a set of random LOSs for the ΛCDM (green curve) and WDM 1 keV (black curve) and WDM 2 keV (blue curve) models at $z = 4.6$. This figure refers to the reference (20,512) simulation without adding instrumental noise. The ΛCDM flux is clearly showing more substructure as compared to the WDM models.

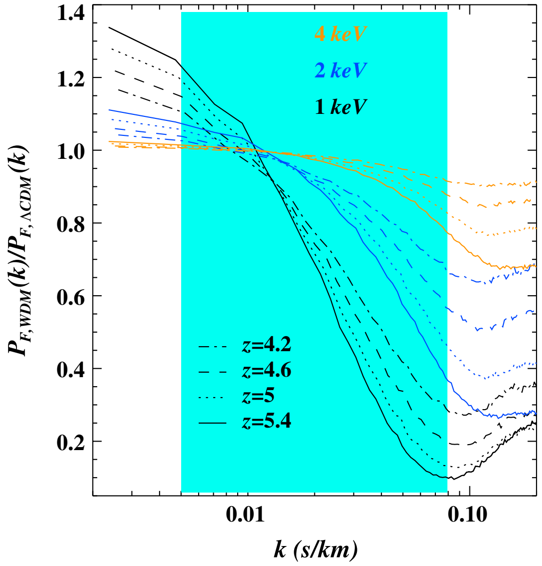

FIG. 3 (color online). The ratio of the 1D flux power spectrum for 3 different WDM models (1, 2 and 4 keV represented in black, blue and orange) at 4 different redshifts ($z = 4.2$, 4.6, 5, 5.4 represented by the continuous, dotted, dashed and dot-dashed curves, respectively) to the corresponding ΛCDM flux power spectra. This figure displays the results from the (20,512) simulations. The mean flux is the same in all models and the shaded area shows the range of wave numbers used in the present analysis.

models with respect to the ΛCDM case. This will be discussed further in the Appendix.

In Figs. 5 and 6 we explore the effects of the two thermal parameters, T_0 and γ, on the flux power spectrum. As discussed earlier, the T-ρ relation is usually parametrized as a power law, $T(z) = T_0(1 + \delta)^{\gamma-1}$. In both of these figures we also plot the WDM 2 keV model in order to emphasize the very distinct differences between the thermal and WDM cutoffs, both in the dependence on wave number and redshift. A hotter (colder) T_0 value produces a suppression (enhancement) in the flux power spectrum with a redshift dependent cutoff. The WDM cutoff is instead more pronounced and steeper than the cutoff induced by a hotter IGM. The dependence of the thermal cutoff on the slope of the temperature density γ in Fig. 6 is also very different from the wave number and redshift dependence of the WDM cutoffs, and is much flatter over the wave number range considered here.

The effects due to changing the mean flux level are discussed in detail in the Appendix (Fig. 17). We conservatively assume a range of ±20% for the observed effective optical depth. We further note that the dependence of changing the mean flux level on wave number is even flatter than that obtained for variations of γ, but shows a weak scale dependence in the highest redshift bin.

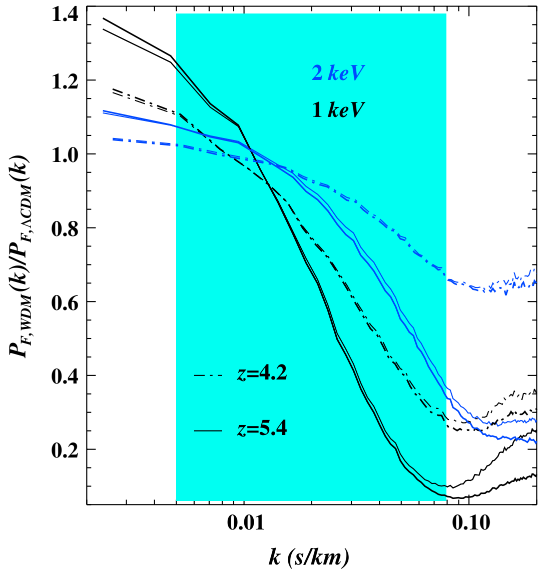

FIG. 4 (color online). The ratio of the 1D flux power spectrum for 2 different WDM models (1 and 2 keV, represented in black and blue) at 2 different redshifts ($z = 4.2$, 5.4 represented by the continuous and dot-dashed curves, respectively) to the corresponding ΛCDM simulations. The thin curves refer to the (20,512) simulations, while the thick curves refer to the high-resolution (20,768) models. The mean flux is the same for all models and the shaded area shows the range of wave numbers used in the present analysis.

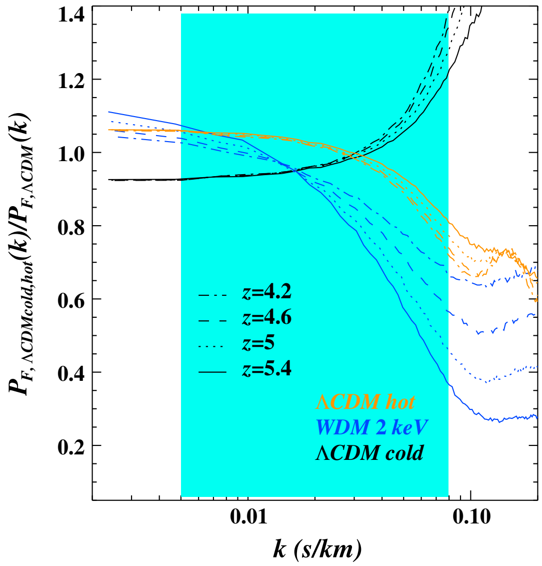

FIG. 5 (color online). The ratio of the 1D flux power spectrum for two ΛCDM models with different temperatures (hot, roughly hotter by 3000 K with respect to the reference simulation, in orange and cold, roughly colder by 3000 K with respect to the reference simulation, in black) and at four different redshifts ($z = 4.2$, 4.6, 5, 5.4 represented by the dot-dashed, dashed, dotted and continuous curves, respectively) to the corresponding ΛCDM simulations. The WDM 2 keV model is also shown in blue. The mean flux is the same for all models, and the shaded area shows the range of wave numbers used in the present analysis.

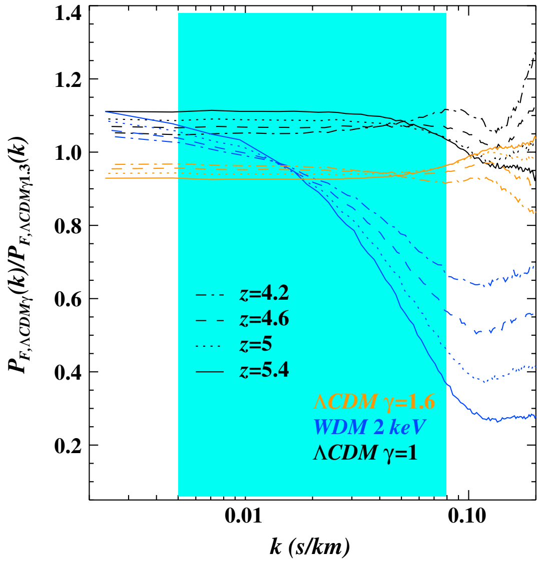

FIG. 6 (color online). The ratio of the 1D flux power spectrum for two ΛCDM models with different slopes of the temperature-density relation ($\gamma = 1.6$ in orange and $\gamma = 1.0$ in black) and at four different redshifts ($z = 4.2$, 4.6, 5, 5.4, represented by the dot-dashed, dashed, dotted and continuous curves, respectively) to the corresponding ΛCDM simulations. The WDM 2 keV model is also shown in blue. The mean flux is the same for all models, and the shaded area shows the range of wave numbers used in the present analysis.

B. Systematic uncertainties

In this section we now briefly discuss the following systematic effects: instrumental resolution; noise; spatial fluctuations in the UV background and metal line contamination. In the Appendix there is a more detailed description of these nuisance effects and how they are modeled. We only summarize the main quantitative results here. Instrumental resolution, which is different for the two sub-data-sets, suppresses the flux power spectrum by at most 20% and 5% for MIKE and HIRES, respectively, at the smallest scales probed, with a negligible redshift dependence. The signal-to-noise ratio impacts at about the 2%–3% level at the smallest scales for $z \leq 5$ while it is at the 7% level for the highest redshift bin.

The UV background fluctuations have been implemented with a deliberately extreme model based on ionizing emission from quasars only. The impact of this extreme model of UV fluctuations on the flux power spectrum is quite scale dependent, and rises considerably at large scales (see also [66]). At the scales of interest here the effect on the flux power spectrum is below the 10% level (see Fig. 15). This should be considered as a generous upper limit. The metal contamination has a much smaller effect on the flux power spectrum, below the 1% level for the whole range of scales considered (see Fig. 16). Apart from the apparently negligible metal contamination, the other nuisance effects have been fully implemented in our analysis.

TABLE I. Summary of the estimates of the relative errors in the flux power spectrum due to a range of nuisance effects: the resolution of the observational data (the MIKE and HIRES data sets have different resolutions), the signal-to-noise ratio of the observational data, the numerical resolution of the simulations, contamination by metal absorbers at lower redshift, the mean flux level, the thermal history of the IGM and the fluctuations in the UV background. The table reports estimates over the wave number range considered and the last three effects are properly marginalized over in the likelihood procedure.

Systematic effects	$\sigma(P_F)/P_F$	Notes
Data resolution	<5–15%	corrected
Data S/N	<3%	corrected
Number resolution	<5%	corrected
Metals	<1%	neglected
Mean flux	~30%	marginalized
Thermal history	~30%	marginalized
UV	<10%	marginalized

In Table I, we summarize the main nuisance parameters and present rough estimates of the relative errors induced in the flux power spectrum. Where possible, effects with a known amplitude such as resolution and noise characteristics are simply incorporated into the mock QSO spectra. The remaining parameters are fully marginalized over in our likelihood analysis.

VI. METHOD AND RESULTS

We now turn to our analysis of the data and discuss our results. For all 108 flux models considered in our analysis, we compute the ratio of these models with respect to the reference model. We then bin this ratio at the same wave numbers as the data. In practice, this means we have a 4D parameter matrix with $(m_{\rm WDM}, \langle F\rangle, T_0, \gamma)$ that summarizes all the results obtained from the hydrodynamical simulations, plus some further parameters for which we have established the effect on the flux power for the reference model only. We parametrize the effect of UV background fluctuations on the flux power with a factor $f_{\rm UV}$ that multiplies the flux power spectrum corrections shown in Fig. 16, constrained to be in the range [0, 1] and applied in addition to the corrections discussed in the previous section ($f_{\rm UV} = 1$ means that the power spectrum is corrected exactly by the amount shown in the lower panel of Fig. 16). We decide to neglect the effect of metal contamination since it is, as we discuss further in the Appendix, very small. We then perform second order Taylor interpolations for the following remaining parameters: $z_{\rm reio}$, $\Omega_{\rm m}$, σ_8, H_0, $n_{\rm s}$, as in Refs. [27,67].

In Fig. 7 we compare our best-guess model (the reference simulation), represented by the orange curves, with the observational data. This best-guess model will be the reference point of our likelihood code that will be described below. Note that there is already rough "visual" agreement with the data, albeit with a poor χ^2 value.

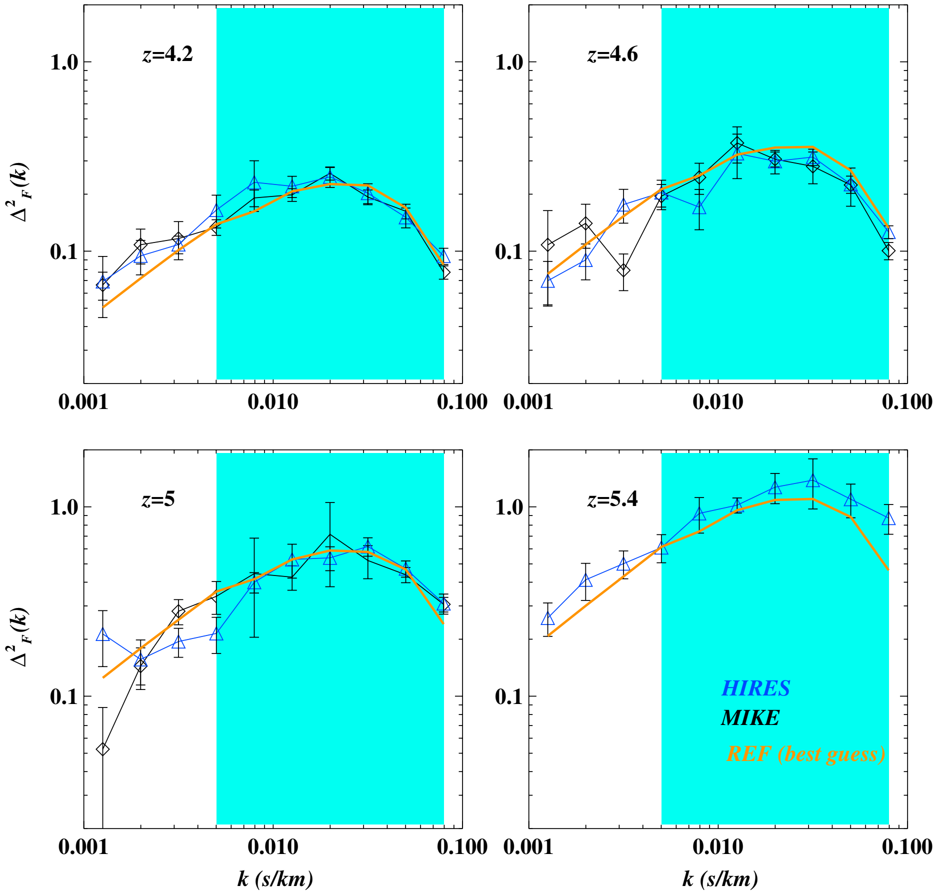

FIG. 7 (color online). The flux power spectrum in dimensionless units, $P_F(k) \times k/!Mp!3$, used in the analysis performed. There are a total of 70 data points at 4 different redshifts. The reference ΛCDM model, which is our best-guess starting point for the Monte Carlo Markov chains, is also shown as orange curves. Only data points in the range $\log_{10} k(\mathrm{s/km}) = [-2.3, -1.1]$ are used in the analysis (shaded area).

We use a modified version of the code COSMOMC [68] to derive parameter likelihoods from the Lyman-α forest data. For the HIRES + MIKE data, we have a set of 15 parameters: 6 cosmological parameters (σ_8, Ω_m, n_s, H_0, z_{reio}, m_{WDM}); 4 parameters describing the thermal state of the IGM, using a power-law parametrization of the temperature-density relation, $T = T_0(z)(1+\delta)^{\gamma(z)-1}$, with parameters $T_0^A(z) = T_0^A[(1+z)/5.5)]^{T_0^S}$ and $\gamma^A(z) = \gamma^A[(1+z)/5.5)]^{\gamma_A^S}$; 4 parameters describing the evolution of the effective optical depth with redshift, since a single power-law has been shown to be a poor approximation over this wide redshift range (see [46]) and one parameter describing the spatial fluctuations in the UV background f_{UV}. We apply strong Gaussian priors to σ_8, Ω_m, n_s in order to mimic Planck constraints: $\Omega_m = 0.315 \pm 0.017$, $\sigma_8 = 0.829 \pm 0.013$, $n_s = 0.9603 \pm 0.0073$. We have checked that these priors do not affect any of our constraints on the free-streaming length/WDM mass, but they are helpful in obtaining faster convergence of the Monte Carlo chains. We vary T_0^A in the range [1000, 20000] K and γ^A in the range [0.7–1.7], and thereby heavily penalize the χ^2 if $\gamma(z = 4.2, 4.6, 5, 5.4)$ is outside the physical range [0.7–1.7]. The values of H_0 and z_{reio} are not constrained by the data and they are prior dependent: the ranges chosen are $H_0 = [50, 100]$ km/s/Mpc and $z_{\mathrm{reio}} = [5, 20]$, respectively.

The covariance matrix calculated from our data set is noisy (especially at high redshift), preventing a reliable inversion. We have therefore regularized the observed covariance matrix with the correlation coefficients as estimated from the simulated spectra as in Ref. [69], $\mathrm{cov_d}(i,j) = r_s(i,j)\sqrt{\mathrm{cov_d}(i,i)\mathrm{cov_d}(j,j)}$ with $r_s(i,j) = \mathrm{cov_s}(i,j)/\sqrt{\mathrm{cov_s}(i,i)\mathrm{cov_s}(j,j)}$, where $\mathrm{cov_d}$ and $\mathrm{cov_s}$ are the covariance matrices of the observed and simulated spectra, respectively.

Our results are summarized in Table II. We obtain a 2σ upper limit on the parameter 1 keV/m_{WDM} of 0.3, which translates into the following constraints: $m_{\mathrm{WDM}} > 3.3$ keV at the 2σ C.L. and $m_{\mathrm{WDM}} > 8.33$ keV at the 1σ C.L., with a best-fit value of $m_{\mathrm{WDM}} = 33$ keV. For a $\sim$3 keV WDM particle the 50% suppression in the 3D linear matter power compared to the ΛCDM case matter power spectrum is at a scale of $k_{1/2} = 22h$/Mpc, while the suppression at $k = 10h$/Mpc is about 10%. If we drop

TABLE II. Marginalized estimates (1σ and 2σ C.L.) and best-fit values for a fit to MIKE + HIRES data using power-law fits for the evolution $\gamma(z)$ and $T_0(z)$. Planck priors on σ_8, n_s and Ω_m have been applied. The best fit χ^2 is 34 for 37 d.o.f. (49 data points − 12 free parameters) which has a probability of 39% of being larger than this value.

Parameter	(1σ)	(2σ)	Best fit
n_s	[0.942, 0.97]	[0.928, 0.984]	0.957
σ_8	[0.806, 0.856]	[0.781, 0.881]	0.822
Ω_m	[0.265, 0.331]	[0.234, 0.362]	0.298
$\tau^A_{\rm eff}(z = 4.2)$	[1.04, 1.16]	[0.98, 1.22]	1.16
$\tau^A_{\rm eff}(z = 4.6)$	[1.19, 1.33]	[1.12, 1.4]	1.32
$\tau^A_{\rm eff}(z = 5)$	[1.76, 1.96]	[1.66, 2.05]	1.91
$\tau^A_{\rm eff}(z = 5.4)$	[2.72, 3.06]	[2.55, 3.21]	3.09
$\gamma^A(z = 4.5)$	[1.38, 1.54]	[1.09, 1.65]	1.64
$\gamma^S(z = 4.5)$	[−0.76, 1.1]	[−2, 2.3]	−0.15
$T^A_0(z = 4.5)(10^3)$ K	[9.1, 10.4]	[7.8, 11.6]	9.2
$T^S_0(z = 4.5)(10^3)$ K	[−3, −2.05]	[−3, −1.1]	−2.5
$f_{\rm UV}$	[0–1]	[0–1]	0.18
$z_{\rm reio}$	[5–11]	[5–16.4]	11.2
1 keV/$m_{\rm WDM}$	[0–0.12]	[0–0.3]	0.03

the 30% additional error applied to the observed flux power spectrum (see Sec. II) we get a tighter lower limit of $m_{\rm WDM} > 4.5$ keV (2σ). The χ^2 gets worse by $\Delta\chi^2 = 14$, but still has a probability of 11% of being this large for the present number of degrees of freedom.

We also took a "frequentist" approach and fixed the values of $m_{\rm WDM}$ to 2.5 and 3.3 keV and found the results in terms of the other parameters: in this case the χ^2 is of course higher than in the ΛCDM model (with $\Delta\chi^2 = 5.6$, 3.8, respectively) but nevertheless compatible with the results obtained in our standard analysis. This is similar to the approach used in Ref. [70], in which an analysis of mixed cold and warm models was performed in both a Bayesian and in a frequentist approach.

The degeneracies between the parameter 1 keV/$m_{\rm WDM}$ and the other parameters are very weak. In Fig. 8 we show the 2D contour plots for the mean likelihood (in color) and the marginalized likelihood (black curves) for $T_0^{\rm A}$ and σ_8 versus 1 keV/$m_{\rm WDM}$. In Fig. 9 we report the 1D mean and marginalized likelihoods for 1 keV/$m_{\rm WDM}$ (continuous and dotted curves, respectively).

We obtain the following evolution for the temperature-density relation $T(z) = 9200[(1 + z)/5.5]^{-2.5}$ K and $\gamma(z) = 1.64[(1 + z)/5.5]^{-0.15}$. The inferred temperature is decreasing with increasing redshift, while the redshift evolution of γ is weak. We stress that the IGM thermal state is just one of several nuisance parameter in our likelihood analysis over which we marginalize. We discuss it here in the context of a consistency check rather than as a measurement. With this in mind, in Fig. 10 we show the recovered redshift evolution for T_0 compared to three input thermal histories used in the simulations and

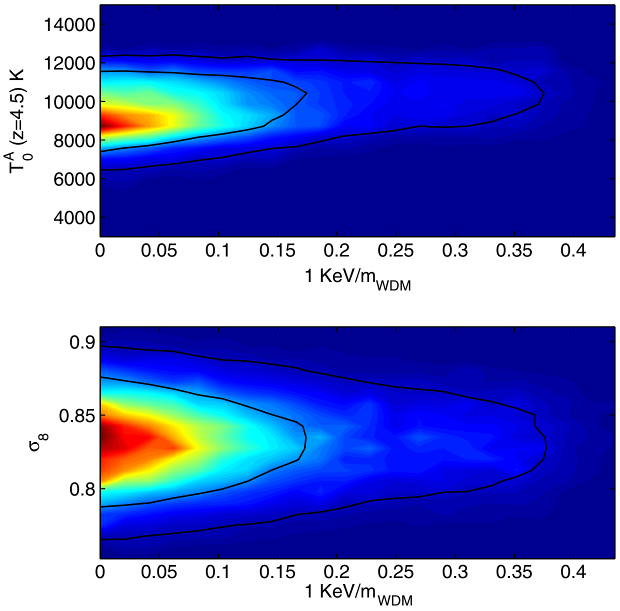

FIG. 8 (color online). The two-dimensional 1σ and 2σ contours for mean (in color) and marginalized (solid black curves) likelihoods for the parameters $1/m_{\rm WDM}$ against $T^A_0(z = 4.5)$ and $1/m_{\rm WDM}$ against σ_8 obtained from the MIKE + HIRES data sets. These results assume a power-law evolution for T_0 and γ, but with $\gamma(z)$ constrained to be in the [0.7, 1.7] range, and refer to a run for which some Planck-like priors on σ_8, n_s and Ω_m have been applied. Note, however, that our results are not sensitive to this choice of prior.

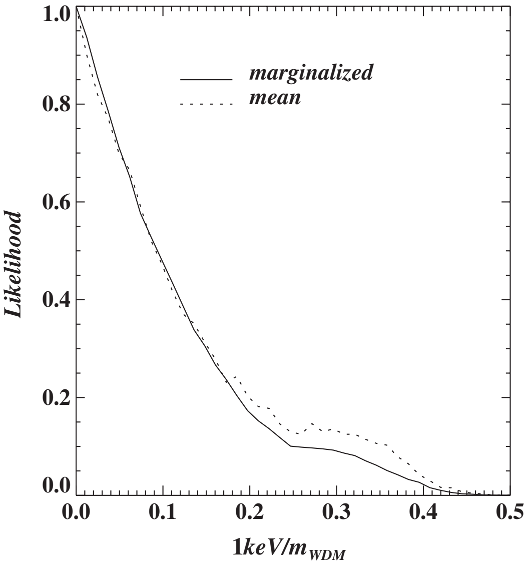

FIG. 9. The one-dimensional mean (dotted curve) and marginalized (continuous curve) likelihoods for the parameter 1 keV/$m_{\rm WDM}$. These results refer to a run for which some Planck-like priors on σ_8, n_s and Ω_m have been applied. Note, however, that our results are not sensitive to this.

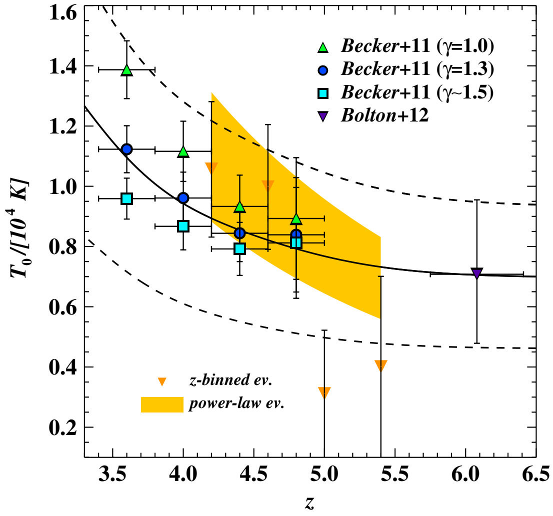

FIG. 10 (color online). The redshift evolution of the temperature at mean density, T_0, used in our reference model is shown as continuous black curve, while the the two dashed line display our cold and hot models. Recent measurements of the IGM temperature at mean density obtained by Ref. [19] are also shown for different values of γ. The measurement at $z \sim 6$ is taken from [71]. The results of our likelihood analysis are shown with the shaded orange area (for the power-law evolution case, $\pm 2\sigma$ ranges), and with orange triangles for model where we left the temperature free in the four redshift bins (1σ error bars). In both cases the temperature values reported are the marginalized results. These results refer to a run for which some Planck-like priors on σ_8, n_s and Ω_m have been applied.

measurements obtained from high-resolution Lyman-α forest data from Refs. [19,71] (note that the power-law index of the temperature-density relation, γ, has not yet been measured directly at $z > 4.2$). The shaded orange area brackets the $\pm 2\sigma$ of the temperatures obtained by our standard likelihood analysis after marginalization. The inferred temperature evolution (parametrized as a power law in redshift) is in good agreement with the measurements from Ref. [19]. We also tested a model where the IGM temperature is left to vary freely in the four redshift bins, shown by the orange data points with error bars in Fig. 10. In the two highest redshift bins this analysis returns temperatures that are rather cold and are disfavored by the data with an unreasonably large temperature jump between $z = 5$ and $z = 4.5$ (see Fig. 10). This suggests that in this case we have introduced too many free parameters and are most likely "overfitting" the flux power spectrum. For completeness we mention that with this apparently unphysical temperature evolution our analysis gives a constraint on $m_{\rm WDM}$ (2σ C.L.) which is about 1 keV lower than for our standard analysis and also returns an unreasonably low reduced χ^2 value. Finally we have also performed a likelihood analysis where the IGM temperature is fixed to be unrealistically cold throughout (3000 K, independent of redshift) to allow for a maximum contribution of the free-streaming of WDM to the observed cutoff in the flux power. Again just for completeness, for this model the constraint on $m_{\rm WDM}$ (2σ C.L.) is lower by about 0.5 keV compared to our standard analysis.

The recovered effective optical depth values at each redshift bin are usually within 20% of the measured optical depth evolution used as the input into the likelihood calculation. The inferred values for the amplitude and slope of the matter power spectrum and for the matter content do not show biases with respect to the Planck-like priors we used. Overall the χ^2 for the best fit model is 34 for 37 d.o.f. which has a reasonably high probability of about 60% of being larger than this value.

Lastly, in Figs. 11 and 12 we show our final best-fit model compared to the data obtained with MIKE and HIRES, respectively. The best fit model is shown as the green curves. We also overplot, for comparison purposes only, three other models that are excluded with very high significance by the present analysis: a model which has a WDM mass of $m_{\rm WDM} = 2.5$ keV (red curves) and a hot model with a temperature value which has been increased by 3000 K with respect to the best fit case (orange curves). When calculating the predicted flux power spectrum for these three models we change only one parameter each time and leave all other parameters fixed at their overall best fit values. In addition, we show the case in which we fix $m_{\rm WDM} = 2.5$ keV and allow all other parameters to assume their best-fitting values under this assumption (blue curves).

Compared to our previous findings obtained in Ref. [27], it is worth stressing the main differences. First of all, from the data side, the sample used here extends to high redshift and doubles the amount of spectra contributing to the signal at $z > 4$. Secondly, both the simulations and the analysis have been refined by: increasing the number of hydrodynamical simulations and their resolution; improving the method in a way that allows a full sampling of the most relevant parameter space (thermal parameters, WDM cutoff and mean flux) compared to a poorer sampling of the parameter space made in Ref. [27]. When considering only the high-resolution data set, we improve the limits by nearly a factor 3 from 1.2 to 3.3 keV at the 2σ C.L.; this is due to both the data and the modeling of the flux power.

VII. JOINT ANALYSIS WITH SDSS DATA

In this section we present the joint analysis with the Sloan Digital Sky Survey (SDSS) 1D flux power spectrum data of Ref. [72] where the authors have presented the flux power spectrum of a sample of 3035 QSO absorption in the redshift range $2 < z < 4$ drawn from the DR1 and DR2 data releases of SDSS. These data have a spectral resolution of $R \sim 2000$, and so typical Lyman-α absorption features, which have a velocity width of ~ 30 km s^{-1}, are not resolved. The wide redshift range, however, makes this data set very constraining in terms of cosmological parameters. As a final result of their analysis they present

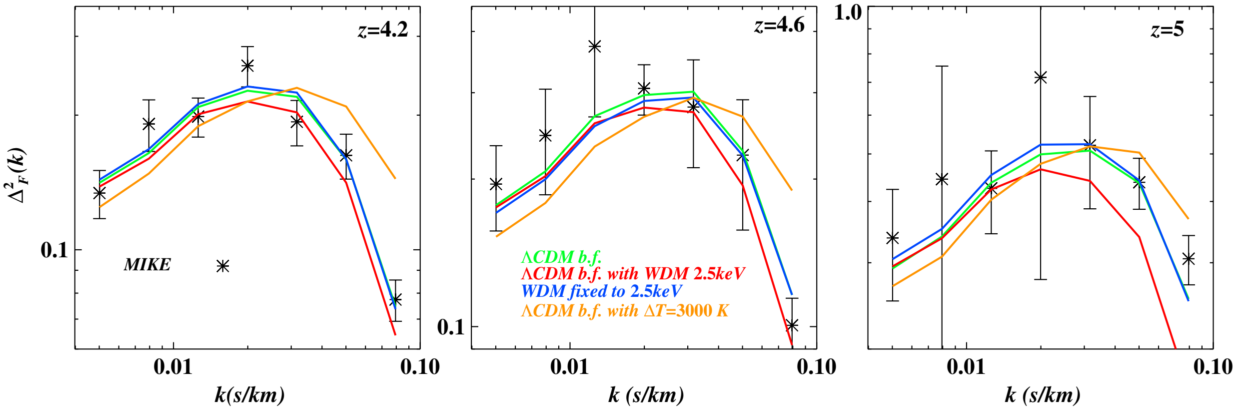

FIG. 11 (color online). The best fit model for the MIKE data set (black crosses) used in the present analysis, shown as the green curves and labeled as "ΛCDM b.f." This model is very close to ΛCDM. We also show for qualitative purposes a few other models: a WDM model that has the same parameters as the best fit model except for the WDM mass (red curves) which is chosen to be 2.5 keV, a model that has a hotter temperature (orange curves) and a model for which the mass of the WDM is fixed to $m_{\mathrm{WDM}} = 2.5$ keV, but for which all other parameters are set to their best-fitting values for this choice (blue curves). Note that for the MIKE data we do not use the $z = 5.4$ redshift bin.

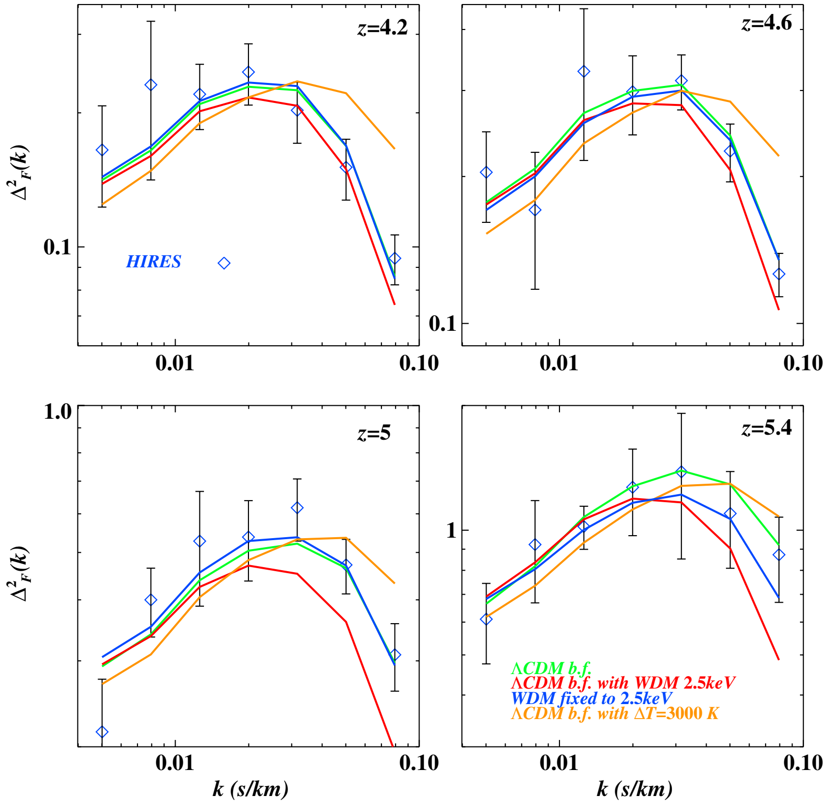

FIG. 12 (color online). The best fit model for the HIRES data set (blue diamonds) used in the present analysis, shown as green curves and labeled as "ΛCDM b.f." As in Fig. 11 we also show a few other models: a WDM model that has the same parameters as the best fit model except for the WDM mass (red curves) which is chosen to be 2.5 keV, a model that has a hotter temperature (orange curves) and a model for which the mass of the WDM is fixed to $m_{\mathrm{WDM}} = 2.5$ keV, but for which all other parameters are set to their best-fitting values for this choice (blue curves). The last model, although visually similar to the ΛCDM model at $z \leq 5$, is excluded at more than 2σ confidence level.

an estimate of flux power spectrum $P_{\mathrm{F}}(k, z)$ at 12 wave numbers in the range $0.00141 < k$ (s/km) < 0.01778, equally spaced in $\Delta \log k = 0.1$ for $z = 2.2, 2.4, 2.6, 2.8, 3, 3.2, 3.4, 3.6, 3.8, 4, 4.2$ for a total of 132 data points. This measurement is likely to improve soon with the new analysis made by the SDSS-III team of a sample which is about 50 times larger than the one we utilize here [73].

For the joint SDSS + MIKE + HIRES analysis we have used a total of 28 parameters: 15 parameters as used for the HIRES/MIKE spectra (without f_{UV} and the two parameters describing the effective optical depth evolution at $z = 5$) plus 13 noise-related parameters: 1 parameter which accounts for the contribution of damped-Lyman-α systems and 12 parameters modeling the resolution and the noise properties of the SDSS data set (see [72]). We do not consider the possible effect of different reionization scenarios on the SDSS flux power. The covariance matrix of the SDSS flux power is provided by the authors of [74]. The 2σ lower limit on m_{WDM} is unchanged at 3.3 keV but now with a $\chi^2 = 183.3$ for the best fit model for a total of 181 data points and 170 degrees of freedom, which has a 23% probability of being this large. Given the fact that the SDSS and MIKE + HIRES data sets do not have redshift overlap there is no significant bias in the other recovered marginalized parameters. Note that for the joint analysis we used the SDSS likelihood based on second order Taylor expansion of the flux power as described in Ref. [15] and used in Ref. [27]. The results are reported in Fig. 13.

Unlike our previous findings obtained in Ref. [27], where the wide redshift range of SDSS data was helpful in breaking the degeneracies between thermal parameters and WDM cutoff, we notice that in this case the SDSS data do not improve the overall constraints. This means that the constraining power of the new high-resolution data set is higher than the low-resolution SDSS data. The joint analysis gives now a lower limit of 3.3 keV, compared to the previous 4 keV value is thereby slightly less stringent: this is due to the different interplay between the data sets and to the relative role of the degeneracies present between IGM thermal state and WDM cutoff. It is also important to stress that the 30% extra error budget on the high-resolution data impacts on the final results also for the joint analysis making the results less stringent than in Ref. [27], where this error was not present.

VIII. DISCUSSION AND CONCLUSIONS

We have presented a comprehensive analysis of the transmitted Lyman-α flux power spectrum extracted from a set of 25 high-resolution QSO spectra taken with the HIRES and MIKE spectrographs. This represents an improved and extended version of the sample originally analysed in [27]. The Lyman-α forest is an excellent probe of the matter distribution at intermediate and high redshift in the mildly nonlinear regime, from sub-Mpc up to BAO scales. In this work we have focused on constraining any possible suppression of the total matter power spectrum which could be induced by the free-streaming of WDM particles in the form of a thermal relic. Due to the nonlinear nature of the relationship between the observed Lyman-α flux and underlying matter density, departures from the standard ΛCDM case are expected over a range of scales that span at least one decade in wave number space and can be constrained by the data used in the present analysis. We model this suppression by using a set of high-resolution hydrodynamical simulations and by marginalizing over a large range of physically motivated thermal histories.

The WDM cutoff exhibits a distinctive behavior which we demonstrate is not degenerate with other physical effects due to its different redshift and scale dependence. We consider possible sources of systematic errors including metal line contamination, spatial fluctuations in the UV background intensity and uncertainties in the mean flux level estimation. Galactic feedback either in the form of supernova driven galactic winds or active galactic nuclei feedback should not impact the flux power spectra at the high redshift considered in this analysis [16].

Our final results are obtained by means of a Monte Carlo Markov chain likelihood analysis around a best-guess reference model. The constraints quoted for m_{WDM} have been calculated after marginalization over the other astrophysical and cosmological parameters. Our analysis is conservative in the following sense: we have dropped the estimates of the power spectrum at the largest scales probed by our sample in order not to be sensitive to continuum fitting uncertainties, we add an additional error of about 30% to our error estimates obtained by bootstrapping to account for the expected underestimation of the real error, and we allow for large fluctuations in the UV background fluctuations, which appear to be the most important nuisance factor. Furthermore, we create a mock QSO sample which resembles as closely as possible the real data including noise and resolution and use the covariance matrix of this mock sample as an estimate of the error properties of the real data. Our final result of this analysis is $m_{\mathrm{WDM}} > 3.3$ keV at the 2σ C.L., where the mass refers to that of a thermal relic. This mass implies that WDM models for which there is a suppression in terms of the 3D linear matter power at scales $k = 10h/\mathrm{Mpc}$ ($k = 22h/\mathrm{Mpc}$) larger than 10% (50%) when compared to the ΛCDM case, are disfavored by the present data sets. The corresponding value of the free-streaming mass is $\sim 2 \times 10^8 M_\odot/h$. A model with a 2.5 keV thermal relic mass is disfavored by the data at about 3σ C.L., a 2 keV mass at about 4σ C.L., and a $m_{\mathrm{WDM}} = 1$ keV model at about 9σ C.L. Our final marginalized estimates and best fit values for m_{WDM} are summarized in Table III.

Overall, the final results presented are similar to those we have obtained in our previous analysis Ref. [27] [3.3 (4.5) keV vs 4 keV previously if we include (do not include) an additional 30% error to account for a possible underestimate of the statistical error from a bootstrapping analysis]. We emphasize, however, that the present analysis is considerably more robust. It uses a larger data set, a much improved analysis based on a broader suite of significantly improved simulations and as well as an extensive analysis of the systematic uncertainties.

Further improvement of the constraints on the free-streaming of dark matter particles from Lyman-α forest data could come mainly from an enlarged set of high-quality, high-resolution spectra, especially at the highest redshifts where the flux power spectrum is most

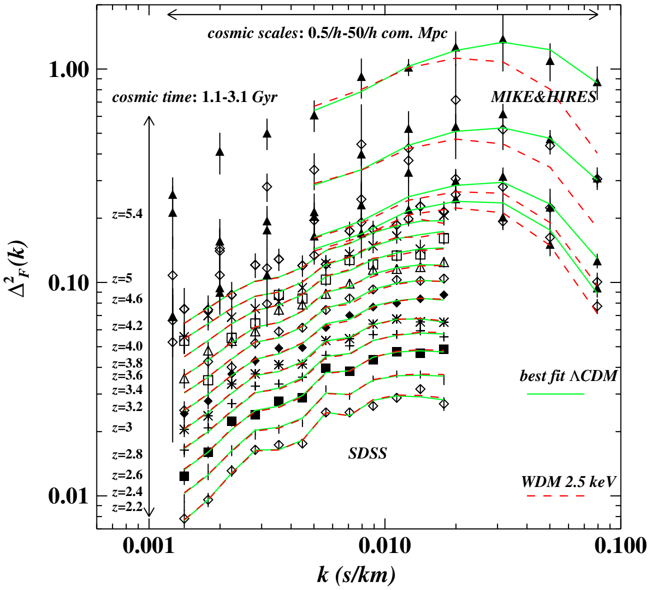

FIG. 13 (color online). Best fit model for the data sets used in the present analysis (SDSS + HIRES + MIKE) shown as green curves. We also show a WDM model that has the best fit values of the green model except for the WDM mass (red dashed curves). These data span about 3 orders of magnitude in scale and the period 1.1–3.1 Gyrs after the big bang.

sensitive to the free-streaming of dark matter. An increase of the dynamical range of the simulations and improved independent constraints on the thermal state and thermal history of the IGM are next on the list as requirements for further corroborating and perhaps pushing the constraints to even larger thermal relic masses. In the future, considerably stronger constraints on WDM may be derived using a baryonic tracer which is colder than the photoionized IGM, thus moving the thermal cutoff to smaller scales in the flux power spectrum. Studies of 21 cm absorption/emission by neutral hydrogen gas before reionization, for example, could eventually fulfill this requirement.

However, with a lower limit of $10^8 M_\odot h^{-1}$ for the mass of dark matter haloes whose abundance could still be significantly affected, the Lyman-α forest data appears to leave already very little room for a contribution of the free-streaming of warm dark matter to the solution of what has been termed the small scale crisis of cold dark matter. In particular, recent suggestions for models with relic masses of 0.5–2 keV are significantly disfavored by our analysis. We finally note that our analysis also suggests that it is unlikely that sterile neutrinos could act in that role.

ACKNOWLEDGMENTS

A part of the observations were made at the W. M. Keck Observatory which is operated as a scientific partnership between the California Institute of Technology and the University of California; it was made possible by the generous support of the W. M. Keck Foundation. This paper also includes data gathered with the 6.5 meter Magellan Telescopes located at Las Campanas Observatory, Chile. The hydrodynamical simulations in this work were performed using the COSMOS Supercomputer in Cambridge (U.K.), which is sponsored by SGI, Intel, HEFCE and the Darwin Supercomputer of the University of Cambridge High Performance Computing Service (http://www.hpc.cam.ac.uk/), provided by Dell Inc. using Strategic Research Infrastructure Funding from the Higher Education Funding Council for England. COSMOS and DARWIN are part of the DIRAC high performance computing facility funded by STFC. M. V. is supported by the FP7 ERC grant "cosmoIGM" GA-257670, PRIN-MIUR and INFN/PD51 grants. G. D. B. acknowledges support from the Kavli Foundation. J. S. B. acknowledges the support of a Royal Society University Research Fellowship. M. H. acknowledges support by the FP7 ERC Grant Emergence-320596.

APPENDIX: SYSTEMATICS

1. Numerical convergence

In Fig. 14 we compare the 1D flux power spectrum extracted from the WDM 2 keV (20,768) and (20,896) models to our reference resolution of (20,512). We focus on this WDM mass since it is excluded at high significance by our analysis, yet it is used by the large-scale structure community in order to solve apparent problems of ΛCDM at small scales. There is at most a 10%–12% correction at

TABLE III. The final summary of the marginalized estimates (1σ and 2σ C.L.) and best fit values for $m_{\rm WDM}$. Planck priors on σ_8, $n_{\rm s}$ and $\Omega_{\rm m}$ have been applied. The REF model refers to our reference conservative analysis; REF. w/o 30% refers to the case in which we do not add an extra 30% uncertainty on the data to account for underestimated bootstrap error bars; REF w/o covmat refers to the case in which we use only the diagonal terms of the covariance matrix; REF + SDSS is the joint analysis of our reference model and SDSS flux power.

Model	(1σ)	(2σ)	Best fit	χ^2/d.o.f.
REF	>8.3 keV	>3.3 keV	33 keV	34/37
REF w/o 30%	>11.1 keV	>4.5 keV	100 keV	48/37
REF w/o covmat	>7.7 keV	>3.1 keV	14.3 keV	33.2/37
REF + SDSS	>7.2 keV	>3.3 keV	42 keV	183.3/170

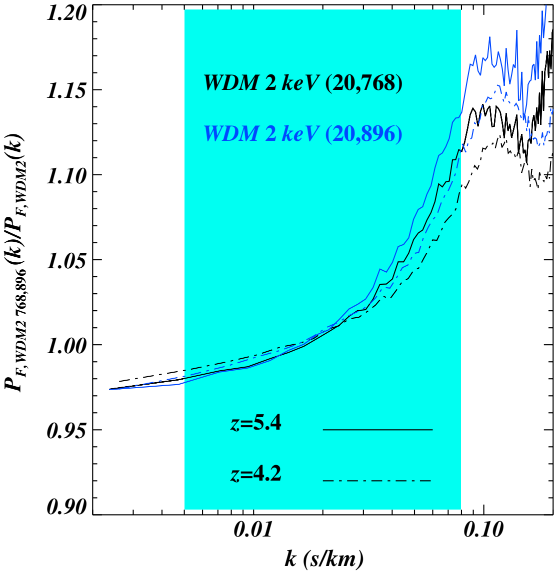

FIG. 14 (color online). The ratio of the 1D flux power spectrum for two WDM 2 keV simulations at different resolutions, (20,768) and (20,896) represented by the black and blue curves respectively, to the reference WDM 2 keV run (20,512) at two redshifts ($z = 4.2$, 5.4 represented by the continuous and dot-dashed curves, respectively). The mean flux is the same in all the models shown, and the shaded area indicates the range of wave numbers used in the present analysis.

the smallest scales probed in our analysis ($k \sim 0.08$ s/km) when the flux power spectra from the (20,768) model are considered, and there is an extra 2% correction when the (20,896) is considered. After correcting for this, we therefore believe that we have reached 5% agreement in terms of the flux power for the smallest scales considered in this work. This is below the 7.75% (1σ) statistical error of the data at the same wave number.

2. Instrumental effects on the flux power

In Fig. 15 we show the effect that instrumental resolution and a given S/N ratio has on the Lyman-α flux power spectrum in the four different redshift bins of our data sample. These results have been obtained from our mock QSO spectra sample. Note that both of these effects have been incorporated into the Lyman-α forest spectra used in our analysis. The S/N ratio results in an increase in the flux power of less than 5% over the range of wave numbers considered in this work, while instrumental resolution effects are particularly important for the MIKE sample. There is a 20% correction at the smallest scales for the MIKE data (for the HIRES data this value is 5%).

3. Systematic effects on the flux power induced by metals and UV background fluctuations

We also consider two important astrophysical nuisance effects in this work: unwanted contamination from metal lines in the Lyman-α forest and the effect of spatial fluctuations in the UV background intensity on the observed Lyman-α forest transmission.

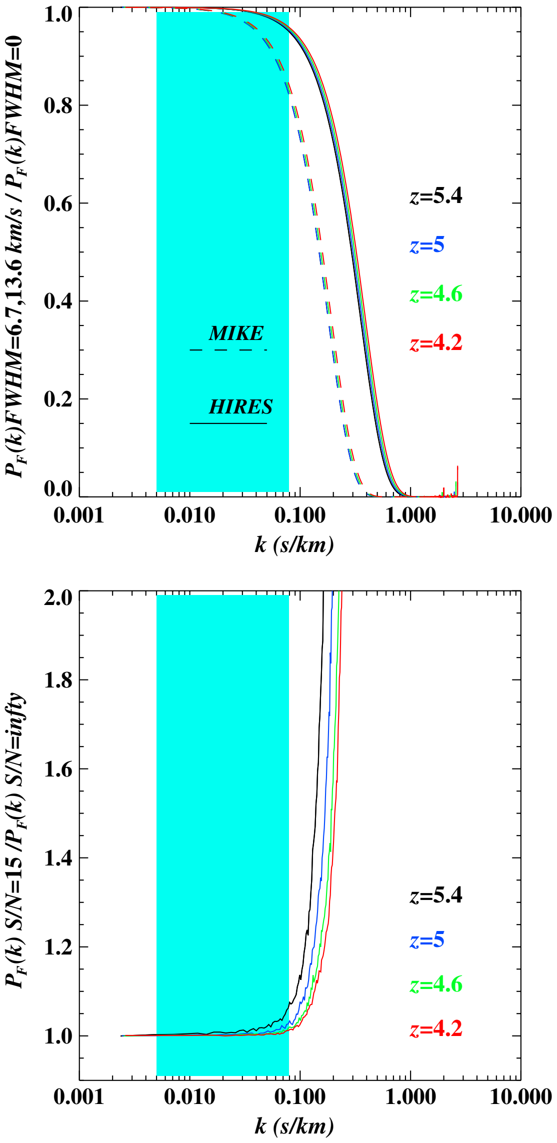

FIG. 15 (color online). Upper panel: The ratio of the 1D flux power spectrum for instrument resolutions corresponding to HIRES (6.7 $\mathrm{km\,s^{-1}}$) and MIKE (13.6 $\mathrm{km\,s^{-1}}$) spectrographs at 4 different redshifts ($z = 4.2$, 4.6, 5, 5.4) to the data with the native resolution of the simulation. Lower panel: The ratio of the 1D flux power spectrum for a S/N value per pixel equal to 15 to a model $\mathrm{S/N} = \infty$. The mock data have been corrected to incorporate both of these effects. The mean flux is the same in all the models shown, and the shaded area indicates the range of wavenumbers used in the present analysis.

In the redshift range we consider in this work, $z = 4.2$–5.4, the most common metal lines in the Lyman-α forest will arise from absorbers at lower redshifts. We therefore consider the effect of absorption from three prominent absorption line doublets; $\mathrm{C_{IV}}(\lambda\lambda 1548, 1551)$, $\mathrm{Si_{IV}}(\lambda\lambda 1394, 1403)$ and $\mathrm{Mg_{II}}(\lambda\lambda 2796, 2804)$, arising over the redshift intervals $z_{\mathrm{CIV}} = 3.08$–4.02, $z_{\mathrm{SiIV}} = 3.54$–4.58 and $z_{\mathrm{MgII}} = 1.26$–1.78, respectively.

We add these metal lines to Lyman-α forest spectra drawn from our reference ΛCDM hydrodynamical model using the following procedure. We firstly integrate fits to the column density distribution functions (CDDFs) presented by [75] from a set of 19 high-resolution VLT/UVES quasar spectra at $z_{\mathrm{qso}} = 2.1$–3.3 over the column density range $\log(N/\mathrm{cm}^{-2}) = 12$–$15$, for all three species. We then multiply the results by the redshift path length of our Lyman-α forest data set to provide an estimate of the number of metal line absorbers in the Lyman-α forest. Note that this approach will likely overestimate the number of $\mathrm{C_{IV}}$ and $\mathrm{Si_{IV}}$ absorbers due to the somewhat lower redshift coverage of the [75] data set relative to this work. Our metal contamination estimates are therefore likely to be conservative in this regard. Next, we Monte Carlo sample column densities and line widths for the appropriate number of lines from the [75] CDDF fits and the Doppler parameter distribution given by Ref. [76], with a chosen value $b_\sigma = 10\ \mathrm{km\,s^{-1}}$ for all three species. Finally, we randomly insert these absorption features into the sight lines in our mock Lyman-α forest data set.

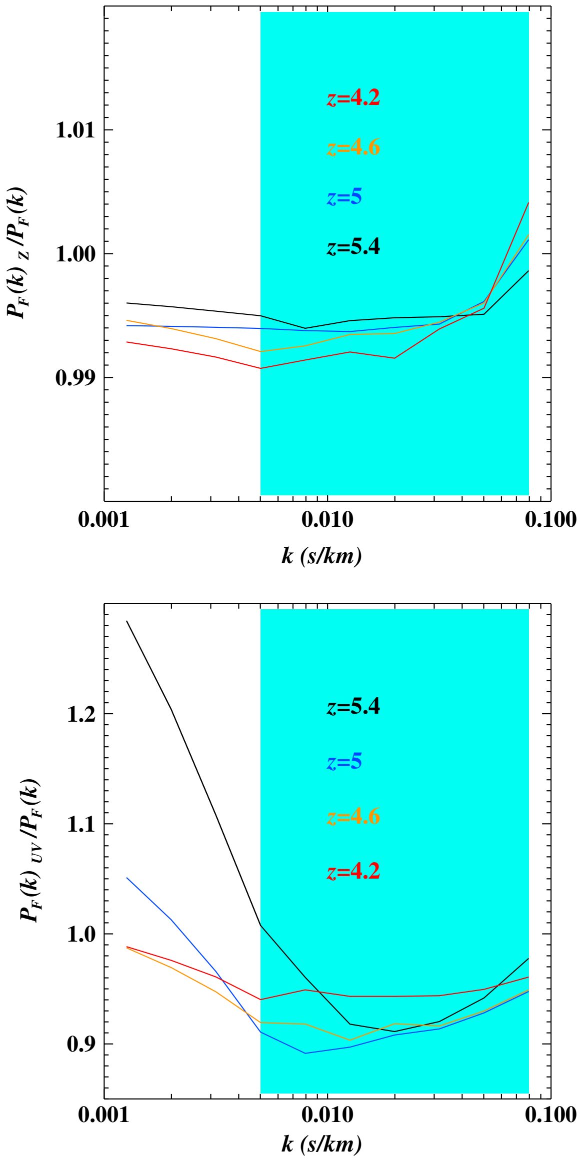

FIG. 16 (color online). Upper panel: The ratio of the 1D flux power spectrum for a model including metal line contamination to the reference model at four different redshifts. Lower panel: The ratio of the 1D flux power spectrum for a model including the effect of spatial fluctuations of the UV background at four different redshifts. Note the different y axis in the two figures. The mean flux is the same in all the models shown, and the shaded area indicates the range of wave numbers used in the present analysis.

In order to estimate the impact of spatial fluctuations in the UV background on the Lyman-α forest at $z = 4.2$–5.4 we use a modified version of the approach described in [65], which was used to examine fluctuations in the HeIIionizing background at lower redshift. We refer the reader to Ref. [65] for further details. The key difference in this work is that we compute the spatially varying $\mathrm{H_I}$ photoionization rate along our simulated sight lines. We achieve this by computing the specific intensity of the ionizing background between 1–4Ry [replacing Eq. (3) in [65]] by solving

$$J(\mathbf{r},\nu) = \frac{1}{4\pi}\sum_{i=1}^{N}\frac{L_i(\mathbf{r}_{\rm i},\nu)}{4\pi|\mathbf{r}_{\rm i}-\mathbf{r}|^2}\,e^{-\frac{|\mathbf{r}_{\rm i}-\mathbf{r}|}{\lambda_{\rm HI}}(\frac{\nu}{\nu_{\rm HI}})^{-3(\beta-1)}}, \quad \text{(A1)}$$

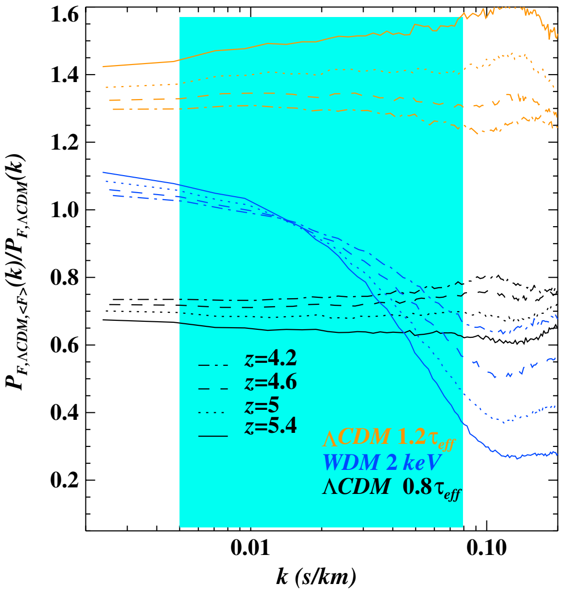

FIG. 17 (color online). The ratio of the 1D flux power spectrum for 2 different ΛCDM models ($\tau_{\mathrm{eff}} = 1.2\tau_{\mathrm{eff,obs}}$ in orange and $\tau_{\mathrm{eff}} = 0.8\tau_{\mathrm{eff,obs}}$ in black) at 4 different redshifts ($z = 4.2$, 4.6, 5, 5.4 represented by the dot-dashed, dashed, dotted and continuous curves, respectively) to the corresponding reference ΛCDM simulation. The WDM 2 keV model is also shown in blue for comparison. The shaded area indicates the range of wave numbers used in the present analysis.

where ν_{HI} is the frequency of the $\mathrm{H_I}$ ionization edge, λ_{HI} is the mean free path for ionising photons presented by [77] and $\beta = 1.5$ is the power-law slope of the $\mathrm{H_I}$ CDDF [78]. The summation in Eq. (A1) is over all quasars with luminosities, L, drawn from the [79] B-band quasar luminosity function. In order to be conservative, in this work we adopt an extreme model which maximizes the effect of the UV fluctuations on the Lyman-α forest by assuming all ionizing photons in the IGM at $z = 4.2$–5.4 are produced by quasars with $M_B < -22$. We therefore ignore the significant contribution to the UV background from the more numerous, fainter star-forming galaxies at these redshifts, which effectively smooth out the large-scale fluctuations produced by the rarer quasars. The spatially fluctuating photoionization rates are then obtained by integrating the specific intensity with respect to frequency, weighted by the photoionization cross section.

In Fig. 16 we show the effect that UV fluctuations and metal contamination have on the flux power spectrum. In the wave number range considered here, the UV background fluctuations have an effect at around the 10% level at the largest scales, dropping to 5% at smallest scales considered in this work. The effect is larger at high redshift ($z = 5$, 5.4) than in the two other redshift bins. Metal contamination affects the flux power very little (at the ±1% level) and is at a level below the statistical error bars of our data.

4. Mean flux level uncertainties

The mean flux level $\langle F\rangle$, or alternatively the effective optical depth $\tau_{\mathrm{eff}} = -\ln\langle F\rangle$, is a key ingredient in our flux modeling procedure and a quantity that needs to be marginalized over in the Monte Carlo Markov chain likelihood estimation. In Fig. 17 we demonstrate the effect that a different mean flux level has on the flux power spectrum. We choose two different mean flux levels with τ_{eff} 20% higher and lower than the reference value (which corresponds to the observed mean flux). A higher (lower) value for τ_{eff} will result in more (less) power relative to the reference value. The trends are similar to those found for the evolution of γ, i.e. rather flat in wave number space, with some weak scale dependence which is only present in the highest and lowest redshift bins. Our final results are not sensitive to the actual choice of the effective optical depth values, since these are marginalized over in the likelihood estimation.

---

[1] P. A. R. Ade *et al.* (Planck Collaboration), arXiv:1303.5076.

[2] P. Bode, J. P. Ostriker, and N. Turok, Astrophys. J. **556**, 93 (2001); B. Moore, T. Quinn, F. Governato, J. Stadel, and G. Lake, Mon. Not. R. Astron. Soc. **310**, 1147 (1999); V. Avila-Reese, P. Colin, O. Valenzuela, E. D'Onghia, and C. Firmani, Astrophys. J. **559**, 516 (2001); M. Miranda and A. V. Macciò, Mon. Not. R. Astron. Soc. **382**, 1225 (2007).

[3] M. Boylan-Kolchin, J. S. Bullock, and M. Kaplinghat, Mon. Not. R. Astron. Soc. **422**, 1203 (2012); J. Pearrubia, A. Pontzen, M. G. Walker, and S. E. Koposov, Astrophys. J. **759**, L42 (2012).

[4] I. Ferrero, M. G. Abadi, J. F. Navarro, L. V. Sales, and S. Gurovich, Mon. Not. R. Astron. Soc. **425**, 2817 (2012).

[5] D. H. Weinberg, J. S. Bullock, F. Governato, R. K. de Naray, and A. H. G. Peter, arXiv:1306.0913.

[6] N. I. Libeskind, A. Di Cintio, A. Knebe, G. Yepes, S. Gottloeber, M. Steinmetz, Y. Hoffman, L. A. Martinez-Vaquero, arXiv:1305.5557.

[7] H. J. de Vega and N. G. Sanchez, arXiv:1304.0759; H. J. de Vega, P. Salucci, and N. G. Sanchez, New Astron. **17**, 653 (2012).

[8] M. R. Lovell, V. Eke, C. S. Frenk, L. Gao, A. Jenkins, T. Theuns, J. Wang, S. D. M. White, A. Boyarsky, and O. Ruchayskiy, Mon. Not. R. Astron. Soc. **420**, 2318 (2012).

[9] K. Heitmann, M. White, C. Wagner, S. Habib, and D. Higdon, Astrophys. J. **715**, 104 (2010); E. Semboloni, H. Hoekstra, J. Schaye, M. P. van Daalen, and I. G. McCarthy, Mon. Not. R. Astron. Soc. **417**, 2020 (2011); M. P. van Daalen, J. Schaye, C. M. Booth, and C. D. Vecchia, Mon. Not. R. Astron. Soc. **415**, 3649 (2011).

[10] M. Viel, K. Markovič, M. Baldi, and J. Weller, Mon. Not. R. Astron. Soc. **421**, 50 (2012).

[11] S. Garrison-Kimmel, M. Rocha, M. Boylan-Kolchin, J. Bullock, and J. Lally, arXiv:1301.3137.

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

[13] U. Seljak, A. Slosar, and P. McDonald, J. Cosmol. Astropart. Phys. 10 (2006) 014; M. Viel, M. G. Haehnelt, and A. Lewis, Mon. Not. R. Astron. Soc. **370**, L51 (2006)

[14] N. Busca *et al.*, Astron. Astrophys. **552**, 18 (2013); A. Slosar *et al.*, J. Cosmol. Astropart. Phys. 04 (2013) 026; K. Dawson *et al.*, Astron. J. **145**, 10 (2013); C. P. Ahn *et al.*, Astrophys. J. Suppl. Ser. **203**, 21 (2012).

[15] M. Viel and M. G. Haehnelt, Mon. Not. R. Astron. Soc. **365**, 231 (2006).

[16] M. Viel, J. Schaye, and C. M. Booth, Mon. Not. R. Astron. Soc. **429**, 1734 (2012).

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

[18] A. Boyarsky, J. Lesgourgues, O. Ruchayskiy, and M. Viel, Phys. Rev. Lett. **102**, 201304 (2009).

[19] G. D. Becker, J. S. Bolton, M. G. Haehnelt, and W. L. W. Sargent, Mon. Not. R. Astron. Soc. **410**, 1096 (2011).

[20] V. K. Narayanan, D. N. Spergel, R. Davé, and C.-P. Ma, Astrophys. J. **543**, L103 (2000).

[21] U. Seljak, A. Makarov, P. McDonald, and H. Trac, Phys. Rev. Lett. **97**, 191303 (2006).

[22] M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese, and A. Riotto, Phys. Rev. Lett. **97**, 071301 (2006).

[23] A. Boyarsky, A. Neronov, O. Ruchayskiy, and M. Shaposhnikov, Phys. Rev. D **74**, 103506 (2006); C. R. Watson, J. F. Beacom, H. Yuksel, and T. P. Walker, Phys. Rev. D **74**, 033009 (2006); A. Boyarsky, J. W. den Herder, A. Neronov, and O. Ruchayskiy, Astropart. Phys. **28**, 303 (2007); K. Abazajian, G. M. Fuller, and W. H. Tucker, Astrophys. J. **562**, 593 (2001); K. Abazajian and S. M. Koushiappas, Phys. Rev. D **74**, 023527 (2006).

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

[25] S. Colombi, S. Dodelson, and L. M. Widrow, Astrophys. J. **458**, 1 (1996); S. H. Hansen, J. Lesgourgues, S. Pastor, and J. Silk, Mon. Not. R. Astron. Soc. **333**, 544 (2002).

[26] A. Kusenko, Phys. Rev. Lett. **97**, 241301 (2006).

[27] M. Viel, G. D. Becker, J. S. Bolton, M. G. Haehnelt, M. Rauch, and W. L. W. Sargent, Phys. Rev. Lett. **100**, 4 (2008).

[28] A. V. Macció, O. Ruchayskiy, A. Boyarsky, and J. C. Munoz-Cuartas, Mon. Not. R. Astron. Soc. **428**, 882 (2012); D. Anderhalden, J. Diemand, G. Bertone, A. V. Macció, and A. Schneider, J. Cosmol. Astropart. Phys. 10 (2012) 047; F. Villaescusa-Navarro, N. Dalal, J. Cosmol. Astropart. Phys. 03 (2011) 024.

[29] A. V. Macció, and F. Fontanot, Mon. Not. R. Astron. Soc. **404**, L16 (2010); D. Anderhalden, A. Schneider, A. V. Macció, J. Diemand, and G. Bertone, J. Cosmol. Astropart. Phys. 03 (2013) 014; A. Schneider, R. E. Smith, and D. Reed, Mon. Not. R. Astron. Soc. **433**, 1573 (2013).

[30] J. Zavala, Y. P. Jing, A. Faltenbacher, G. Yepes, Y. Hoffman, S. Gottlober, and B. Catinella, Astrophys. J. **700**, 1779 (2009).

[31] S. Shao, L. Gao, T. Theuns, and C. S. Frenk, Mon. Not. R. Astron. Soc. **430**, 2346 (2013); D. Boyanovsky, Phys. Rev. D **83**, 103504 (2011).

[32] L. Gao, and T. Theuns, Science **317**, 1527 (2007).

[33] H. Song and J. Lee, Astrophys. J. **703**, L14 (2009).

[34] M. Mapelli, A. Ferrara, and E. Pierpaoli, Mon. Not. R. Astron. Soc. **369**, 1719 (2006).

[35] R. Barkana, Z. Haiman, and J. P. Ostriker, Astrophys. J. **558**, 482 (2001).

[36] R. S. de Souza, A. Mesinger, A. Ferrara, Z. Haiman, R. Perna, and N. Yoshida, Mon. Not. R. Astron. Soc. **432**, 3218 (2013).

[37] N. Menci, F. Fiore, and A. Lamastra, Mon. Not. R. Astron. Soc. **421**, 2384 (2012); A. M. Nierenberg, T. Treu, N. Menci, Y. Lu, and W. Wang, arXiv:1302.3243.

[38] J. Wang and S. D. M. White, Mon. Not. R. Astron. Soc. **380**, 93 (2007).

[39] A. Schneider, R. E. Smith, A. V. Macció, and B. Moore, Mon. Not. R. Astron. Soc. **424**, 684 (2012).

[40] R. E. Smith and K. Markovic, Phys. Rev. D **84**, 063507 (2011).

[41] P. Valageas, Phys. Rev. D **86**, 123501 (2012); D. Boyanovsky and J. Wu, Phys. Rev. D **83**, 043524 (2011).

[42] E. Pierpaoli, S. Borgani, A. Masiero, and M. Yamaguchi, Phys. Rev. D **57**, 2089 (1998).

[43] A. J. Benson, A. Farahi, S. Cole, L. A. Moustakas, A. Jenkins, M. Lovell, R. Kennedy, J. Helly, and C. Frenk, Mon. Not. R. Astron. Soc. **428**, 1774 (2012); R. E. Angulo, O. Hahn, and T. Abel, arXiv:1304.2406; A. Lapi, P. Salucci, and L. Danese, Astrophys. J. **772**, 85 (2013); F. Pacucci, A. Mesinger, and Z. Haiman, arXiv:1306.0009.

[44] S. S. Vogt, S. L. Allen, B. C. Bigelow *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **2198**, 362 (1994).

[45] G. M. Bernstein, A. E. Athey, R. Bernstein, S. M. Gunnels, D. O. Richstone, S. A. Shectman, A. M. Larar, and M. G. Mlynczak, Proc. SPIE Int. Soc. Opt. Eng. **4485**, 453 (2002).

[46] G. D. Becker, M. Rauch, and W. L. W. Sargent, Astrophys. J. **662**, 72 (2007).

[47] G. D. Becker, W. L. W. Sargent, M. Rauch, and A. P. Calverley, Astrophys. J. **735**, 93 (2011).

[48] A. P. Calverley, G. D. Becker, M. G. Haehnelt, and J. S. Bolton, Mon. Not. R. Astron. Soc. **412**, 2543 (2011).

[49] D. D. Kelson, Publ. Astron. Soc. Pac. **115**, 688 (2003).

[50] K. Horne, Publ. Astron. Soc. Pac. **98**, 609 (1986).

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

[52] X. Fan, M. A. Strauss, R. H. Becker, R. L. White, J. E. Gunn, G. R. Knapp, G. T. Richards, D. P. Schneider, J. Brinkmann, and M. Fukugita, Astron. J. **132**, 117 (2006).

[53] E. Rollinde, T. Theuns, J. Schaye, I. Paris, and P. Petitjean, Mon. Not. R. Astron. Soc. **428**, 540 (2012).

[54] V. Springel, Mon. Not. R. Astron. Soc. **364**, 1105 (2005).

[55] M. Viel, M. G. Haehnelt, and V. Springel, Mon. Not. R. Astron. Soc. **354**, 684 (2004).

[56] E. Komatsu *et al.*, Astrophys. J. Suppl. Ser. **192**, 18 (2011).

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

[58] J. S. Bolton, M. Viel, T.-S. Kim, M. G. Haehnelt, and R. F. Carswell, Mon. Not. R. Astron. Soc. **386**, 1131 (2008).

[59] L. Hui and N. Y. Gnedin, Mon. Not. R. Astron. Soc. **292**, 27 (1997).

[60] A. Rorai, J. F. Hennawi, and M. White, arXiv:1305.0210.

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

[62] S. Bird, M. Vogelsberger, D. Sijacki, M. Zaldarriaga, V. Springel, and L. Hernquist, Mon. Not. R. Astron. Soc. **429**, 3341 (2013).

[63] J. A. Regan, M. G. Haehnelt, and M. Viel, Mon. Not. R. Astron. Soc. **374**, 196 (2007).

[64] M. Rauch, J. MiraldaEscude, W. L. W. Sargent, T. A. Barlow, D. H. Weinberg, L. Hernquist, N. Katz, R. Cen, and J. P. Ostriker, Astrophys. J. **489**, 7 (1997).

[65] J. S. Bolton and M. Viel, Mon. Not. R. Astron. Soc. **414**, 241 (2011).

[66] P. McDonald, U. Seljak, R. Cen, P. Bode, and J. P. Ostriker, Mon. Not. R. Astron. Soc. **360**, 1471 (2005); R. A. C. Croft, Astrophys. J. **610**, 642 (2004).

[67] M. Viel and M. G. Haehnelt, Mon. Not. R. Astron. Soc. **365**, 231 (2006).

[68] A. Lewis and S. Bridle, Phys. Rev. D **66**, 103511 (2002); CosmoMC, http://www.cosmologist.info/cosmomc/.

[69] A. Lidz, K. Heitmann, L. Hui, S. Habib, M. Rauch, and W. L. W. Sargent, Astrophys. J. **638**, 27 (2006).

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

[71] J. S. Bolton, G. D. Becker, S. Raskutti, J. S. B. Wyithe, M. G. Haehnelt, and W. L. W. Sargent, Mon. Not. R. Astron. Soc. **419**, 2880 (2012).

[72] P. McDonald *et al.*, Astrophys. J. **635**, 761 (2005).

[73] N. Palanque-Delabrouille *et al.*, arXiv:1306.5896.

[74] P. McDonald *et al.*, Astrophys. J. Suppl. Ser. **163**, 80 (2006).

[75] E. Scannapieco, C. Pichon, B. Aracil, P. Petitjean, R. J. Thacker, D. Pogosyan, J. Bergeron, and H. M. P. Couchman, Mon. Not. R. Astron. Soc. **365**, 615 (2006).

[76] L. Hui, and R. E. Rutledge, Astrophys. J. **517**, 541 (1999).

[77] A. Songaila and L. L. Cowie, Astrophys. J. **721**, 1448 (2010).

[78] T.-S. Kim, R. F. Carswell, S. Cristiani, S. D'Odorico, and E. Giallongo, Mon. Not. R. Astron. Soc. **335**, 555 (2002).

[79] P. F. Hopkins, G. T. Richards, and L. Hernquist, Astrophys. J. **654**, 731 (2007).