Scuola Internazionale Superiore di Studi Avanzati

Constraining Dark Matter properties with the Inter-Galactic Medium and other probes

A thesis submitted in fulfillment of the requirements
for the degree of *Philosophiæ Doctor* in

Astrophysics & Cosmology

Author:
Riccardo Murgia

Supervisor:
Prof. Matteo Viel

September 2019

"What we are doing here is so important,
we better not to take it too seriously."
Shunryū Suzuki Rōshi

Prelude

Mari Pintau.
Sunset.
A glass of Vermentino.
A tasty cigarette.
Some thoughts.
Lots of memories.
Here I am. Here we are.
I have never been really able to separate different aspects of life. Work and Fun. Colleagues and Friends. Science and Art. Sport and Politics. I and we...
That is why this PhD thesis is the result of four years of experiences which somehow go much beyond myself and my scientific research. This thesis would not exist at all if I haven't had the opportunity to cross my path with many fantastic and memorable people, who literally fullfilled these years with fun and crazyness. Papers are just a natural consequence of having been so lucky to find (and grow together) the perfect envinroment to spend most of our time, almost everyday, and sincerely enjoy it (at least I have sincerely enjoyed it!).
First of all, I would like to deeply thank my supervisor Matteo Viel, not only a brilliant and inspiring scientist, but above all, one of the most supportive persons I have ever met. Besides the pleasure of working together, the most important lesson I learnt from him is about the method, rather than about some content: the importance of synergies and cooperation, which, thanks to him, I could constantly experience at first hand. For analogous reasons I am infinitely thankful to Julien Lesgourgues, my advisor in Aachen, and to all the worderful people who really made me feel at home over there (surely not an easy task, I was a Sardinian in Germany!): Maria Archidiacono, Deanna Hooper, and Nils(ixeddu) Schöneberg; but also Thejs Brinckmann, Christian Fidler, Samuel/Salmon Brieden, and all the others with whom I haven't had (yet?) the opportunity to work with.
Cooperation, networking, and teamworking are as valuable in science as in all other aspects of life. Having the chance to collaborate with several excellent physicists from various institutions helped me to grow as a scientist and as a person. I am grateful to all of them: Takeshi Kobayashi (alè, alè), Alex Merle, Max Totzauer, Aurel Schneider, Vid Iršič, Andrea De Simone, Marco Baldi, Alvise Raccanelli,

Mario Ballardini, Fabio Finelli, Nicolao Fornengo, Stefano Gariazzo, Sebastian Bohr, Carisa Miller, Adrienne Erickcek, and of course, *dulcis in fundo*, Matteo Nori. Furthermore, I would like to thank Stefano Camera and Olga Mena, for spending part of their summer to carefully read this thesis.
Another fundamental part of my PhD life has been constituted by outreach: I am thankful to Simona Cerrato and Anna Lisa Cesaro, who gave me the chance to get involved in several science communication projects in Trieste. Moreover, I must spend a few words to thank all my companions in that incredible ongoing adventure we dubbed as IDeAS (*Incontri di Divulgazione e Astrofisica in Sardegna*): Matteo Tuveri, Matteo Cadeddu, Francesca Dordei, Francesca/Frankina Lepori, Mario/Marieddu Cadelano, Fra(ncesca)-Loi, Emmanuele Picciau, my *brother-by-choice* Ale(ssandro)-Loni, my dearest Francipà and Francivè, and the indispensable Brinuzzo. I am also thankful to all the boys and girls involved in our outreach projects, who always repay our efforts with their enthusiasm and passion, making me believe that we can really use art & science to revolutionise the world.
On a more personal side, there is a crew of Torino-Genova-Arenzano-Savona people who have always been present, constantly supporting and encouraging me during these years: as a "representative" of all of them, I deeply thank my *compañero* Lucone, who's one of the sweetest persons I know, even if he pretends not to.
If I will miss Trieste in the future, it's surely because I will miss all my dearest honorary Sardinian citizens, who stayed and/or will be staying at SISSA for a while: Gor(ixeddu) Oganesyan, who'll never stop surprising me; F. Andrej Obuljen, the real *vecchia ciabatta* of my heart; Elias(eddu) Kammoun, who has a solution to every problem; Farida/Farixedda Farsian, one of the best dictators ever; Hasti(xedda) Khoraminezhad, whose surname truly reflects who she is; (Scelfo-)Giulio-Scelfo, with his lovely 5-dimensional smile; Matteo Nuriss(im)o, my designated heir in the Coffice; Tronconi Tommaso Ronconi, who *fa tanto l'anarchico* and that's why I love him; Gabriele (un, due, tre...) Parimbelli, the ideal (C)office mate; Andrea *Dottor*-Oddo, hoping he will survive his next football match; Lumen Boco (*facci un balletto!*); *l'intellectuel belge* Miguel(ingo) Vanvlasselaer (*a.k.a.* Vanvly) & Jon-Pham, *Odder Couple* than Jack Lemmon/Walter Matthau; Titou(an) Lazeyras and *tzia* Pauline Vielzeuf, who don't have excuses to not to come visit me in Occitania soon and often; Mihael Petač, who knows very well that without him I would have never survived winters and post-Sardinia sadness crises.
Last but not least, I have been, I am, and I will always be grateful to Gigi Riva, unintentional metaphor of a better world to be built, and (even more, I must admit) to my family, both the *strictly defined* one and the *extended* one, for being exactly what they are, the primary cause of each of my achievements.

Abstract

According to the standard cosmological (ΛCDM) model, the universe today is mainly composed by a cosmological constant, denoted by Λ, and by Cold Dark Matter (CDM). Whereas this standard paradigm is tremendous agreement with Cosmic Microwave Background (CMB) and Large-Scale Structure (LSS) data, some discrepancies exist, on the cosmological and local determination of the Hubble parameter H_0, and on the measurement of the amplitude of the matter fluctuations, σ_8. Additionally, assuming the ΛCDM model, cosmological N-body simulations predict too many dwarf galaxies and too much (C)DM in the innermost regions of galaxies, with respect to observations. Moreover, the dynamical properties of the most massive Milky Way satellites are not reproduced in simulations. The inclusion of baryon feedback is crucial to give a realistic picture of the aforementioned problems, and it shows that baryons can indeed mitigate this CDM "small-scale crisis". Nevertheless, in the absence of a solution within the ΛCDM framework, and driven by the fact that the fundamental nature of the dark sector is still unrevealed, alternative DM scenarios emerged as a possible way to explain the tensions. In fact, many *non-cold* (nCDM) candidates have been proposed in order to provide a better description of the structure formation and distribution at small scales, with respect to the ΛCDM model.

The effect of the existence of a nCDM particle is a suppression of the matter power spectrum $P(k)$ on small scales, induced, e.g., by its small mass or some non-standard interaction. The suppression in the power spectrum can be described by the so-called transfer function $T(k)$, namely the square root of the ratio of the matter power spectrum in the presence of nCDM with respect to that in the presence of CDM only. Most of the constraints from structure formation data obtained so far, refer to a specific shape of the power suppression, corresponding to the case of thermal Warm DM (WDM), i.e., candidates with a Fermi-Dirac/Bose-Einstein momentum distribution. However, most of the viable particle DM candidates do not feature a thermal momentum distribution, making the oversimplified notion of thermal WDM incapable to describe the shape of their transfer functions. Besides particle DM scenarios, another intriguing possibility that can be tested against small-scale observations is the case where a significant fraction of DM is made by Primordial Black Holes (PBHs), given that Poisson fluctuations in the PBH number density induce a small-scale power enhancement departing from the standard CDM prediction.

In this thesis, we firstly introduce a new analytic parametrisation for the transfer function, simple yet versatile enough to describe the gravitational clustering signal of large classes of non-thermal nCDM models, such as sterile neutrinos, ultra-light scalar DM, mixed DM fluids, and interacting DM. The goal is to systematically test these models against the most constraining data set for small-scale deviations with respect to ΛCDM, i.e., high-resolution and high-redshift measurements of the Lyman-α forest, the absorption line pattern produced by intervening inter-galactic neutral hydrogen in the spectra of distant quasars. We thus illustrate how to exploit such observable to constrain practically any non-standard DM scenarios without the need to run any

specific numerical simulations, due to the novel parametrisation proposed, to a large suite of pre-computed hydrodynamic simulations, and to an advanced scheme efficiently interpolating across different cosmological models. We demonstrate that the shape of the linear matter power spectrum for thermal WDM models is in mild tension ($\sim 2\sigma$ C.L.) with data, compared to non-thermal scenarios, and we probe for the first time the small-scale *shape* of the DM power spectrum for a large set of nCDM models, through extensive Monte Carlo Markov Chain (MCMC) analyses.

We then use the Lyman-α data to update current constraints on ultra-light scalar DM models, and we further investigate the cosmological implications at high and low redshifts. For scalar DM constituting more than 30% of the whole of the DM, we obtain a lower limit $m \gtrsim 10^{-21}\,\mathrm{eV}$ for the scalar DM mass, which implies an upper limit on the initial field displacement of $\phi \lesssim 10^{16}\,\mathrm{GeV}$. We derive limits on the energy scale of cosmic inflation and determine an upper bound on the tensor-to-scalar ratio of $r < 10^{-3}$, in the presence of scalar DM. We also find that there is very little room for scalar DM to solve the CDM small-scale crisis without hitting the Lyman-α bounds. We then focus on quantifying the impact of the Quantum Potential (QP) during the non-linear evolution explored by our hydrodynamic simulations. We improve upon the nearly universally adopted approximation to encode the non-standard nature of the DM candidate in the transfer function used to produce the initial conditions for the simulation, by accurately following the scalar DM evolution in a N-body set-up without approximating its dynamics. Since the new constraints do not depart significantly from the previous ones, this represents the first direct validation of the approximations generally adopted in the literature. Furthermore, we perform a thorough characterisation of the DM halo properties, determining the typical mass scale below which the QP has a significant impact.

We then focus on interacting DM scenarios, specifically on models where the dark sector is composed by two types of relic particles, possibly interacting with each other: non-relativistic DM, and relativistic Dark Radiation (DR). Based on the general parametrisation previously discussed, we introduce a new Lyman-α likelihood, applicable to a wide range of non-standard cosmological models, with complementary scale and redshift coverage with respect to CMB and Baryon Acoustic Oscillation (BAO) data. In fact, for two of the considered interacting scenarios, we find that Lyman-α data strengthen the CMB+BAO bounds on the DM-DR interaction rate by many orders of magnitude. However, models solving the missing satellite problem are still compatible with the new bounds. For the third class of models, Lyman-α data bring no stronger constraints on the interaction rate than CMB+BAO data, except for extremely small values of the DR density. Using a theory-motivated prior on the minimal density of DR, we also find that in this framework the H_0 tension can be reduced from 4.1σ to 2.7σ, while simultaneously accommodating smaller values for σ_8, as hinted by cosmic shear data.

Finally, we present Lyman-α constraints on the PBH mass and abundance, by means of a new grid of high-resolution hydrodynamic simulations. We obtain a marginalised upper limit on the product of the PBH mass and fraction of $f_{\rm PBH} M_{\rm PBH} \sim 60\ M_\odot$ at 2σ C.L., when a Gaussian prior on the reionisation redshift is imposed, preventing its posterior distribution to peak on very high values, which are in disagreement with various recent independent measurements. Such constraint weakens to $f_{\rm PBH} M_{\rm PBH} \sim 170\ M_\odot$, when a more conservative flat prior is instead assumed. Both limits improves previous bounds from the same observable by roughly 2 orders of magnitude. We also extend our predictions to non-monochromatic PBH mass distributions, ruling out large parts of the parameter space for two of the most accredited PBH extended mass functions.

List of Publications

This thesis is a result of the research that I have carried out while being a PhD student in Astrophysics & Cosmology at SISSA, under the supervision of Prof. Matteo Viel, and in collaboration with several people both from SISSA and other institutions. The work presented here is based upon the following publications:

- R. Murgia, A. Merle, M. Viel, M. Totzauer, A. Schneider
  "Non-cold" dark matter at small scales: a general approach
  JCAP 11 046 (2017); arXiv:1704.07838
- T. Kobayashi, R. Murgia, A. De Simone, V. Iršič, M. Viel
  Lyman-α Constraints on Ultralight Scalar Dark Matter: Implications for the Early and Late Universe
  PRD 96, 123514 (2017); arXiv:1708.00015
- R. Murgia, V. Iršič, M. Viel
  Novel constraints on non-cold (non-thermal) Dark Matter from Lyman-α forest data
  PRD 98, 083540 (2018); arXiv:1806.08371
- M. Nori, R. Murgia, V. Iršič, M. Baldi, M. Viel
  Lyman-α forest and non-linear structure characterization in Fuzzy Dark Matter cosmologies
  MNRAS 482, 3, 3227–3243 (2019); arXiv:1809.09619
- R. Murgia, G. Scelfo, M. Viel, A. Raccanelli
  Lyman-α forest constraints on Primordial Black Holes as Dark Matter
  PRL 123, 7, 071102 (2019); arXiv:1903.10509
- M. Archidiacono, D. C. Hooper, R. Murgia, S. Bohr, J. Lesgourgues, M. Viel
  Constraining Dark Matter - Dark Radiation interactions with CMB, BAO, and Lyman-α
  Submitted to JCAP; arXiv:1907.01496

Besides the aforementioned articles, I have also being working on a few other projects, which are not included in this thesis, but have ended up in the following publications:

- R. Murgia, S. Gariazzo, N. Fornengo
  Constraints on the Coupling between Dark Energy and Dark Matter from CMB data
  JCAP 04 014 (2016); arXiv:1602.01765
- C. Miller, A. Erickcek, R. Murgia
  Constraining Nonthermal Dark Matter's Impact on the Matter Power Spectrum
  Submitted to PRD; arXiv:1908.10369
- M. Ballardini, R. Murgia, M. Baldi, F. Finelli, M. Viel
  Imprints of superimposed primordial oscillations on the matter power spectrum down to non-linear scales
  In preparation; arXiv:1909.xxxxx

Contents

Appendices

"Può darsi che il visibile sia nato
da una bagarre di spiriti inferociti.
Ma tempo e spazio erano già creati?
Peccato, dice Crono al suo collega.
Si stava molto meglio disoccupati."
Eugenio Montale

1 Introduction

1.1 The standard cosmological model

The Big Bang cosmology predicts that the universe started in a hot and dense state, and has been expanding over time. The rate of expansion depends on the types of matter and energy present in the universe. The most accredited parametrisation of this theory is the *standard cosmological model* or ΛCDM model, which assumes that General Relativity is the correct theory of gravity on cosmological scales. According to this model, the universe contains a cosmological constant, denoted by Λ, associated with Dark Energy (DE), Cold Dark Matter (CDM), i.e. non-baryonic matter which decoupled while non-relativistic from the primordial plasma, and the small amount of ordinary matter which explains the primordial light element abundances (Alpher et al. 1948). The fundamental nature either of DE or (C)DM still remains undiscovered, however the former is responsible for the accelerating expansion of the universe observed in the light from distant galaxies and supernovae (Hubble 1929; Riess et al. 1998; Perlmutter et al. 1999). The existence of the latter is necessary to explain the shape of the galaxy rotation curves (Corbelli & Salucci 2000), the galaxy cluster dynamics observations (Clowe et al. 2006), and the formation of all the structures present today in the universe, which arose through gravitational instability from the small perturbations originated by quantum mechanical fluctuations in the very early universe.

Many observations are indeed explained and reconciled among each other assuming that we live in an expanding universe, described by the scale factor $a(t)$, which accounts for the time-dependence of physical distances due to the expansion,

and by the current value of the Hubble parameter $H_0 = 100 \cdot h$ km s^{-1} Mpc^{-1} (with $h \simeq 0.7$ (Aghanim et al. 2018; Riess et al. 2019), via the Friedmann equation:

$$H^2 = \frac{8\pi G}{3}\rho_{\rm tot} - \frac{k}{a^2},$$

where G is the universal gravitational constant, the Hubble parameter $H \equiv \frac{1}{a}\frac{da}{dt}$ represents the expansion rate of the universe, $\rho_{\rm tot} \equiv \sum_i \rho_i$ is the total energy density of the universe and k is the spatial curvature, which can assume the values $k = -1, 0, +1$ respectively in the case of open, flat or closed universe. The values of the energy densities of the different components of the universe ρ_i control $a(t)$. They are typically expressed in units of the so-called critical density ρ_c:

$$\Omega_i \equiv \frac{\rho_i}{\rho_c} = \frac{8\pi G \rho_i}{3H^2},$$

where $\rho_c \equiv 3H^2/8\pi G$ is the critical density, namely the total energy density of a flat universe ($k = 0$).

The evolution of each cosmological component is fully specified by its equation of state:

$$w_i = \frac{p_i}{\rho_i},$$

where ρ_i and p_i are the density and pressure of the i-th component, respectively. Whereas DM and baryons have zero pressure ($w_{\rm DM} = w_b = 0$), DE is interpreted as a cosmological constant, so that $w_{\rm DE} \equiv w_\Lambda = -1$. Concerning the radiation density, it accounts for two components, namely $\Omega_r = \Omega_\gamma + \Omega_\nu$: the former is the photon contribution ($w_\gamma = 1/3$), the latter is the neutrino contribution, following from the assumption of 3 neutrino species, with mass $m_\nu \ll 1$ eV. Massive neutrinos have an equation of state $w_\nu = 1/3$ while relativistic, and $w_\nu = 0$ after they become non-relativistic. Neutrinos contribute to the radiation density only until they are relativistic. Since the universe appears to have zero global spatial curvature ($k = 0$), its total density $\rho_{\rm tot}$ has to be equal to the critical density ρ_c, so that $\Omega_{\rm tot} \equiv \rho_{\rm tot}/\rho_c \simeq 1$. In the ΛCDM framework, the current values of the relative densities Ω_i of DE, DM, baryonic matter and radiation are: $\Omega_\Lambda \simeq 68.3\%$; $\Omega_{\rm DM} \simeq 26.8\%$; $\Omega_b \simeq 4.9\%$; $\Omega_r \simeq 10^{-5}$ (Aghanim et al. 2018).

Density fluctuations are determined through the gravitational instability of an initial spectrum of perturbations. The most plausible mechanism to account for the origin of the perturbations in the universe and produce a set of initial conditions is the cosmic inflation, namely a very early phase (presumably 10^{-34} seconds after the Big Bang) in which the universe exponentially expanded with time (Liddle 1999). The inflationary hypothesis predicts that quantum fluctuations are produced in the very early universe, when the relevant scales are casually connected, then

these perturbations are whisked outside the causal horizon by expansion, so that they remain frozen with constant amplitude until they re-enter the horizon at a later time, to serve as initial conditions for the evolution of inhomogeneities and anisotropies (Lyth & Riotto 1999). From inflation one obtains an initial spectrum of perturbations that is a power law with a power law index or tilt $n_s \simeq 1$.

The standard paradigm that we have briefly illustrated predicts the existence of the Cosmic Microwave Background (CMB), a thermal radiation left over when the universe went out of thermal equilibrium and could not be considered as a plasma any longer (Dicke et al. 1965). The CMB, predicted by Gamow in 1948 (Gamow et al. 1948) and discovered by Penzias and Wilson in 1965 (Penzias & Wilson 1965), has an almost isotropic black body spectrum with temperature $T \simeq 2.7$ K. It can be considered as a snapshot of the universe when it was about 300000 years old. However, the CMB is not perfectly uniform: it presents fluctuations in its temperature field across the sky, detected for the first time by the FIRAS experiment in the COBE spacecraft in 1992 (Smoot et al. 1992). Since the end of the 1990's several satellites have been sent in space to improve COBE measurements: the spacecraft WMAP operated nine-year observations both on temperature anisotropies and polarisation until 2010 (Bennett et al. 2013); the ESA satellite Planck was launched in 2009 to take data with extremely high sensitivity (down to an angular scale of a few arc-minutes) both on CMB temperature and polarisation. The Planck Collaboration made three public data releases: in 2013 (Ade et al. 2014), in 2015 (Ade et al. 2016a), and in 2018 (Aghanim et al. 2018), all of them in excellent agreement with ΛCDM predictions. Besides CMB, also the gravitational clustering signal on large scales matches very well the expectation for a CDM scenario. Large-Scale Structure (LSS) surveys such as BOSS/SDSS (Palanque-Delabrouille et al. 2013; Alam et al. 2017) have completed the picture at low and intermediate redshifts, through detection of galaxies ($z < 0.7$) and the Lyman-α forest ($z \simeq 2.4$).

Hence, the possibility to simultaneously explain the early time CMB anisotropies and the LSS of the universe at late times, has solidified the standard CDM paradigm, in which DM is cold and collisionless, as a cornerstone of modern cosmology.

Despite its remarkable success, the standard ΛCDM model has been challenged by possible disagreements in different data sets. The most notable of these is the H_0 tension, where the value of the Hubble Constant inferred from CMB and BAO data (Aghanim et al. 2018), is significantly lower than the value locally measured with supernovae (Riess et al. 2019). Similarly, measurements of σ_8 – the root-mean-square matter fluctuations in a sphere of $8\,\mathrm{Mpc}/h$ radius – also yield a mild discrepancy across different observations; the latest CMB+BAO inferred value (Aghanim et al. 2018) is slightly higher than the value obtained from weak lensing experiments (Hildebrandt et al. 2018; Abbott et al. 2018; Joudaki et al. 2019).

In addition to these tensions, the standard CDM paradigm clearly exhibits some limits at sub-galactic scales, often denoted as the ΛCDM *small-scale crisis*. In fact, assuming the ΛCDM model, cosmological N-body simulations predict too many dwarf galaxies (*missing satellite* problem (Klypin et al. 1999; Moore et al. 1999) within the Milky Way (MW) virial radius, and too much DM in the innermost regions of galaxies (*cusp-core* problem (Donato et al. 2009), as well as showing more diversity in the inner density profile (*diversity* problem (Oman et al. 2015; Tulin & Yu 2018)), with respect to observations. Moreover, the most massive predicted subhalos, which have so much enclosed mass that they should have ignited, remain unseen (the *too-big-to-fail* problem (Boylan-Kolchin et al. 2011).

These small-scale problems could be alleviated by invoking baryon physics, still not perfectly understood and implemented in cosmological simulations (Garrison-Kimmel et al. 2017; Sawala et al. 2016; Pawlowski et al. 2015). For instance, it is known that photo-evaporation from ultraviolet (UV) sources during the reionisation period pushes gas out from small halos, preventing star formation and reducing the number of observed substructures (Okamoto et al. 2008). Moreover, supernova feedback may be able to make the inner parts of halo density profiles significantly shallower (Governato et al. 2012). However, it is currently hard to model and implement these baryon effects in hydrodynamic simulations, without being affected by large uncertainties. Whereas a more accurate modelling of baryon feedback is crucial to provide a realistic picture of these problems and probe to which extent baryons can alleviate the CDM crisis, another possible solution is to modify the nature of DM, by going beyond the standard CDM paradigm.

DM candidates are generally classified according to their velocity dispersion, which defines a free streaming length. On scales smaller than their free streaming length, density fluctuations are erased and gravitational clustering is suppressed. The velocity dispersion of CDM particles is by definition so small that the corresponding free streaming length is negligible for cosmological structure formation. The most accredited CDM candidate, very well motivated by theoretical particle physics, is in the form of a Weakly Interacting Massive Particle (WIMP) (Jungman et al. (1996)). Nonetheless, such particles have so far eluded detection in direct and indirect searches, as well as at colliders (Bertone & Tait 2018).

On the other hand, non-standard DM scenarios featuring small DM particle masses, large momenta and/or non-standard interactions tend to suppress the gravitational clustering at small scales. Such models have thus gained a lot of interest in recent years, as viable alternatives to the standard ΛCDM paradigm. The "golden goal" is to find a theoretically well motivated *non-cold* DM (nCDM) model, providing nearly the same predictions of ΛCDM in terms of CMB and LSS, while simultaneously solving the H_0 and σ_8 tensions, and inducing a small-scale clustering suppression which better fits the astrophysical observations at non-linear scales.

1.2 The matter power spectrum

Even though both matter inhomogeneities and CMB anisotropies originated in the same context, they appear very different today. Whereas matter inhomogeneities have grown because of gravitational attraction, the radiation pressure prevents the same process from occurring in the photons. Hence, while the matter distribution on scales roughly below $\mathcal{O}(10 - 10^2$ Mpc) is organised in the form of non-linear structures like filaments, galaxies and clusters, photon perturbations are linear even on small scales. The linear regime is the easiest to work in from a theoretical standpoint, as linear cosmological perturbation theory is well understood and gives definite predictions.

Though CMB anisotropy observations have nowadays reached incredible levels of precision, they only give us a partial view of the universe, as it was on a thin shell at around redshift $z \simeq 1100$, known as the *last scattering surface*. The main difference between LSS and CMB is in the number of scales that one can measure. While the CMB is coming from a thin spatial shell, the LSS is 3-dimensional. We thus expect the latter to have more constraining power than the former. Nevertheless, only recent constraints coming from LSS have become comparable to the CMB bounds on cosmological parameters, due to the different physical regimes that the two approaches probe.

However, it is already now evident that DM, though dominant in terms of its overall energy density, is only one ingredient in a complex system of matter components that are responsible for forming the structures that we observe. Driven by the tensions between observations and simulations, briefly discussed in the previous Section, our picture of structure formation is indeed becoming progressively more complicated. Ongoing and future LSS surveys have been (or will be soon) providing increasingly accurate measurements of the matter distribution on various scales and redshifts, offering tremendous opportunities to further *stress-test* the standard ΛCDM model and its possible alternatives.

In order to describe and model the observed clustering of matter, galaxies or gas, it is useful to look at the correlation function of the overdensity field of these tracers. Assuming that the overdensity is a homogeneous and isotropic Gaussian random field with zero mean value, all the information is contained in the 2-point (auto-)correlation function[1], $\xi(r)$, defined as an excess probability of finding a pair of galaxies at a given distance $r_{12} = |\mathbf{x}_1 - \mathbf{x}_2|$, namely:

$$dP = n^2 \delta V_1 \delta V_2 [1 + \xi(r_{12})], \tag{1.1}$$

[1]Let us stress, however, that the non-linear structure evolution tends to move part of this information into higher n-point correlation functions, as we will see in the following Subsections.

where n is the average number density and δV_i are the volume elements separated by a distance r_{12}, so that

$$\xi(r_{12}) = \langle \delta(\mathbf{x}_1)\delta(\mathbf{x}_2)\rangle, \tag{1.2}$$

where $\delta(\mathbf{x})$ describes the matter overdensity field at a given time t, i.e.

$$\delta(\mathbf{x}, t) = \frac{\rho(\mathbf{x}, t)}{\bar{\rho}(t)} - 1, \tag{1.3}$$

with $\bar{\rho}$ being the matter field mean density of the universe. Note that ξ only depends on the amplitude r_{12} of $\mathbf{r}_{12}$, following from the fact that the density perturbation field $\delta(\mathbf{x})$ is assumed to be a homogeneous and isotropic random field.

From Equation (1.3) one can notice that the overdensity field must satisfy $-1 \leq \delta(\mathbf{x}, t) < \infty$. As we already stated, the overdensity δ is $\ll 1$ only at large enough scales and/or at early enough times, so that the linear perturbation theory is applicable only in such regimes.

The power spectrum $P(k)$ of the overdensity field can be identified as the Fourier transform of the correlation function, i.e.

$$\xi(r_{12}) = \int \frac{d^3k}{(2\pi)^3}\, P(k)\, e^{-i\mathbf{k}(\mathbf{x}_1 - \mathbf{x}_2)}, \tag{1.4}$$

where the wavenumber k is related to the wavelength λ of the matter fluctuation by $k = 2\pi/\lambda$. It is thus defined as:

$$P(k) = \frac{1}{(2\pi)^3}\langle |\delta(\mathbf{k})|^2\rangle. \tag{1.5}$$

The shape of the power spectrum is sensitive to the values of the cosmological parameters, and different cosmological probes allow us to measure it at different scales. To illustrate this point, in Figure 1.1 we report the reconstructed linear matter power spectrum at redshift $z = 0$, obtained using different observations. Such plot is taken from the latest public release from the Planck Collaboration (Akrami et al. 2018), to which we refer for the details on the data sets under consideration.

Let us note that, at the linear perturbation level, different k-modes evolve independently, and the covariance matrix is purely diagonal. However, as the density field becomes non-linear, different modes couple among each other. In fact, late-time LSS analyses can be roughly split into three regimes:

- large (linear) scales ($\gtrsim 30$ Mpc$/h$), where fluctuations are Gaussian and practically fully characterised by the two-point correlation function (as defined in Equation (1.2));
- intermediate (mildly non-linear) scales ($\sim 1 - 30$ Mpc$/h$);
- small (highly non-linear) scales ($\lesssim 1$ Mpc$/h$), where galaxies are forming and galaxy cluster cosmology becomes important.

In the following Subsections we will briefly discuss these three different regimes, and the impact of non-linear effects on the large-k region of the power spectrum.

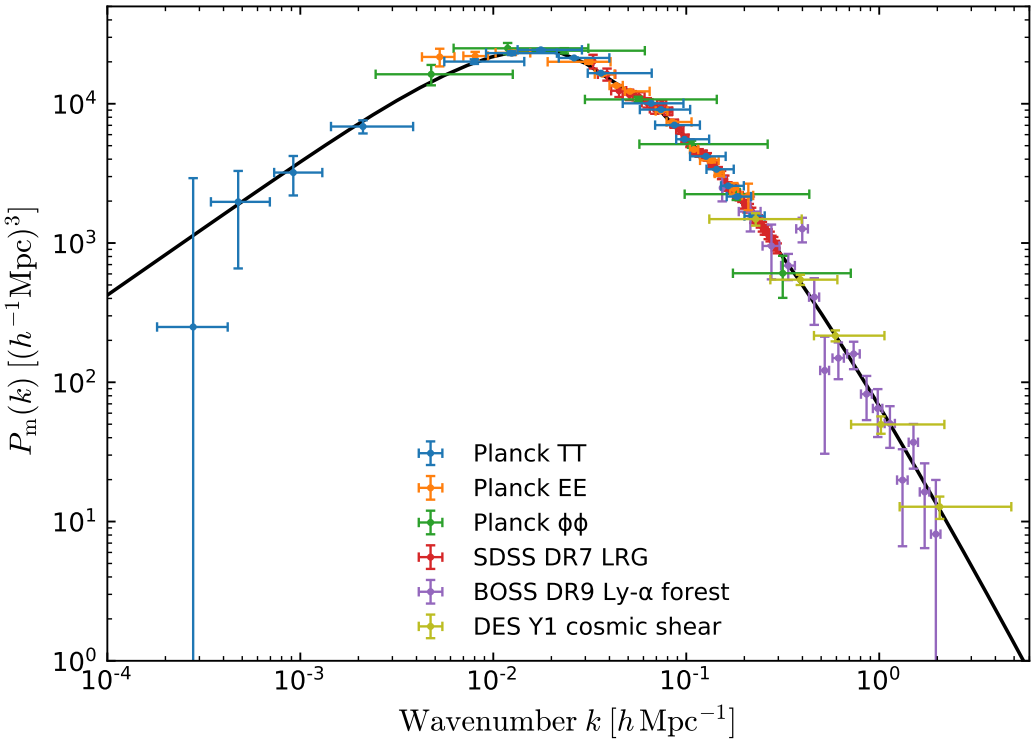

Figure 1.1: The linear matter power spectrum, at redshift $z = 0$, inferred from different cosmological probes. Different colours refer to different cosmological probes; the black solid line represents the standard ΛCDM prediction. This plot is taken from Akrami et al. (2018), to which we refer for the details on the data sets under consideration.

1.2.1 Large scales

As it has been already explained, fluctuations on a sufficiently large scale (and/or at sufficiently high redshift) can be described by linear perturbation theory. Therefore, in this regime, the linear matter power spectrum gives a complete statistical description of the density fluctuations. Even before the formulation of the inflationary paradigm, a scale-invariant power spectrum of the initial perturbations in the gravitational potential Φ – i.e., $P_\Phi(k) \propto \langle|\Phi((k))|^2\rangle$ – was suggested, so that

$$k^3 P_\Phi \propto const. \tag{1.6}$$

It is thus common to express such *primordial* power spectrum in terms of deviations from pure scale-invariance, namely:

$$P_\Phi \propto \frac{k^{n_s-1}}{k^3}, \tag{1.7}$$

where the spectral index n_s has a value close to 1, and the case of $n_s = 1$ is known as the Harrison-Zel'dovich power spectrum. In this case, fluctuations on different length scales correspond to the same amplitude of fluctuation in the gravitational

potential. Inflation theories predict an almost scale-invariant primordial power spectrum, characterised by a tiny deviation from the Harrison-Zel'dovich case (see, e.g., Guth & Pi (1982); Bardeen et al. (1983)).

Let us now relate the power spectrum P_Φ of the perturbations in the gravitational potential to the matter power spectrum, as defined in Equation (1.5). This can be done through the Poisson equation, which relates the gravitational potential Φ to the matter overdensity field δ:

$$\nabla^2\Phi = 4\pi G a^2 \bar{\rho}\delta, \tag{1.8}$$

from which we see that $\delta(k) \propto k^2\Phi(k)$, finally leading to

$$P(k, z) \propto k^{n_s}. \tag{1.9}$$

The value of the primordial tilt was lately measured with great precision by Planck (Aghanim et al. 2018), obtaining $n_s = 0.965 \pm 0.004$, which is in excellent agreement with the predictions from inflation.

The rate at which matter fluctuations grow on different scales is determined by an interplay between self-gravitation, pressure forces and damping processes, yielding a modification of the shape of the primordial power spectrum, i.e. a modification of Equation (1.9). Density fluctuations on scales larger than the size of the Hubble horizon ($\propto H^{-1}$) grow through self-gravity. At very early times, all relevant scales are outside the horizon and do not grow. However, since the horizon grows with time, it progressively incorporates density fluctuations on larger scales, which start growing depending on the epoch when the horizon-crossing occurs. Fluctuations on very small length scales cross the horizon at early times, when the universe is radiation dominated, so that they are effectively frozen until the matter/radiation equality, due to the fact that the radiation pressure prevents their gravitational collapse. After matter/radiation equality ($z \sim 3000$) DM fluctuations, which do not interact electromagnetically with the radiation, can grow. On the other hand, baryons are still coupled to photons, which provide a pressure force that prevents their gravitational collapse until redshift $z \sim 1000$, when baryonic matter decouple from the radiation, and progressively fall into DM gravitational potentials.

The interpretation of the shape of the linear power spectrum, as shown in Figure 1.1, is now clear. Its turn-over corresponds to the horizon scale at matter/radiation equality: high k-modes, which entered the horizon at earlier times, are suppressed, while low k-modes follow the trend given by Equation (1.9).

All of this information is typically encoded in terms of a *transfer function* $\mathcal{T}(k, z)$, which allows us to finally express the power spectrum as a function of scale and time:

$$P(k, z) = A(z)\mathcal{T}^2(k, z)k^{n_s}, \tag{1.10}$$

where the normalisation $A(z)$ has to be determined observationally.

Let us remark that the power spectrum shown in Figure 1.1 and described by Equation (1.10) is the linear one. We must therefore keep in mind that it is supposed to accurately match observations only at sufficiently large scales, i.e. at wavenumbers $k \lesssim 0.2\ h/\mathrm{Mpc}$. Furthermore, in reality we do not directly measure the matter power spectrum. What we do observe are tracers of the underlying matter field, such as galaxies or gas. However, on large scales, where density fluctuations are small, one can safely assume that galaxies follow the distribution of the total matter up to a linear factor, dubbed as the *linear galaxy bias* b_g:

$$P_g(k, z) = b_g^2(z) P(k, z). \tag{1.11}$$

By relating an observable quantity such as the galaxy power spectrum $P_g(k)$ to the matter power spectrum $P(k)$, Equation (1.11) allows one to test the large-scale matter distribution of the universe simply by solving linear perturbation theory equations. To do this, two publicly available numerical codes are are mostly used: `CAMB` (Lewis et al. 2000) and `CLASS` (Lesgourgues 2011; Blas et al. 2011). The latter is the one used in this thesis.

The approach described in this Subsection is completely inadequate on intermediate and small scales, where higher order perturbation methods or numerical N-body/hydrodynamic simulations are needed to efficiently compare theoretical predictions with astrophysical observations, accounting for the non-linear effects affecting such scales.

1.2.2 Intermediate and small scales

In the previous Subsection we have seen that the evolution of a small-amplitude density fluctuation on a given scale can be followed by using linear perturbation theory, independently of fluctuations on other length scales. However, as anticipated above, in the later stages of gravitational collapse, fluctuations on different length scales become coupled and the subsequent evolution is non-linear.

Gravitational instability of matter fluctuations, as already mentioned, somehow drives galaxy formation, which proceeds via a two-step process. First, gravitational instability acting on DM, induces the formation of self-gravitating DM halos. Baryonic gas also takes part in the collapse, and it is heated by shocks to the thermal (or virial) temperature of the DM halos. Second, the hot gas cools radiatively. Although the physical principles governing the gravitational, (hydro-)dynamical and radiative processes involved in galaxy formation are well established, the systems under consideration are too complex to be studied analytically. Therefore, cosmological numerical simulations have to be employed to study the final collapse

of matter fluctuations, and to properly model the shape of the matter power spectrum at intermediate and small scales.

Presently, the most efficient way to investigate the non-linear regime of structure formation is through N-body simulations, which numerically solve the coupled equations of motion of N particles, interacting only through gravity, in an expanding universe. Such approach has been very successfully applied to model the evolution of collisionless CDM. N-body simulations can be supplemented with hydrodynamic routines to follow the evolution of the gas, gravitationally interacting with DM, but also subject to the aforementioned cooling and heating processes. Two different approaches are currently in use: Eulerian methods, and Lagrangian methods, also known as Smoothed Particle Hydrodynamics (SPH). The hydrodynamic simulations used in this thesis were performed adopting the latter approach (see Chapter 2 – Section 2.4 and references therein – for further details). In principle, hydrodynamic simulations can follow the evolution of both the gas and DM without relying on simplified approximations. In practice, however, simulations are limited by numerical resolution and computational power. Consequently, some of the key physical processes, in particular star formation and feedback, have to be modelled approximately, using semi-analytic recipes. The arbitrariness in the implementations of these processes make their validity still uncertain, in particular when analysing very low-redshift and very small-scale regimes, such as the internal structure of DM (sub)halos. When looking for possible deviations from the standard CDM predictions, it is then often useful to focus on relatively smooth overdensities ($\delta \lesssim 20$), intermediate scales, and relatively high redshifts ($z \gtrsim 2$), where the impact of non-linearities and galactic feedback is much less prominent.

1.3 The Inter-Galactic Medium (IGM)

Most of the matter in the universe does not reside in collapsed objects. In fact, by definition, at sufficiently early times, all baryons and DM were part of the the Inter-Galactic Medium (IGM), i.e. the diffuse gas filling the voids between galaxies. Even at present time, the large majority of cosmic matter is still in the IGM. The IGM is indeed a crucial ingredient of any theory of galaxy formation and evolution. Its characteristics can be investigated through its absorption of radiation from background sources or through the radiation that it generates in emission. Studying the IGM properties as a function of redshift can thus provide important information on galaxy formation and evolution, as well as insight into cosmological events occurred since recombination: for the purposes of this thesis, we are mostly interested in the latter possibility. For a comprehensive review on the IGM physics we address the reader to, e.g., Meiksin (2009); McQuinn (2016).

Most of the IGM gas observed at relatively high redshift ($z \gtrsim 2$) has experienced only mild gravitational collapse. Moreover, it has gone through little astrophysical processing other than heating by photo-ionisation and compression, in competition with the adiabatic cooling by Hubble expansion. These facts make it an ideal cosmological laboratory.

At redshifts around $z \sim 1000$, the IGM has recombined, it has then been neutral until the first radiative sources, capable to reionise it, were produced. Once the first galaxies formed, they photo-ionised nearly all the inter-galactic hydrogen, and heated the IGM to temperatures of $\mathcal{O}(10^4 - 10^5 K)$, smoothing the gas distribution and affecting the subsequent formation of galaxies. Whereas at high redshift this ionising background was sourced by stars, by $z \sim 3$ quasars became important, if not dominant. The process of reionisation began as individual sources started to generate expanding bubbles of ionised hydrogen. As more and more UV sources switch on, the ionised volumes grow in size, until they overlap and fill all the inter-galactic space. The IGM is completely reionised when the rate of emission of UV photons per comoving unit volume balances the radiative recombination rate, so that hydrogen atoms are photo-ionised faster than they can recombine. Due to the history that we have just sketched, over $2 \lesssim z \lesssim 5$ there is a wealth of absorption line data for inter-galactic hydrogen, helium, and metals. Being a tracer of the non-linear collapse of cosmological structures, the IGM is indeed known to be a highly inhomogeneous network of filamentary structures, the *cosmic web*. Its main manifestation is the *Lyman-α forest*, namely the absorption pattern of neutral hydrogen along the line of sights to high-redshift quasars. Such observable represent an effective probe of the underlying DM overdensity field on a range of scales ($0.5 \lesssim k \lesssim 50$ h/Mpc) which makes it complementary to estimates of the matter power spectrum from CMB and other LSS probes, such as galaxy surveys or weak gravitational lensing observations. At least in the intermediate redshift regime, where the Lyman-α lines are strong enough to be detected, yet not too strong to be saturated, the absorption spectra of distant quasars are expected to provide a record of the initial conditions of gravitational structure formation.

1.4 The Lyman-α forest

The Lyman-α forest is a region in the spectra of high-redshift quasars characterised by a set of absorption lines arising from the absorption of Lyman-α photons, produced by the distant source, by the residual neutral hydrogen along our line of sight. Due to the expansion of the universe, each neutral hydrogen "cloud" sees the photons at a different wavelength, thus leaving its imprint in the spectrum as an absorption line at $1216 \cdot (1 + z)$ Å, where z is the redshift of the "cloud". The

link between the observable appearance of the Lyman-α forest and cosmology can be described through the *fluctuating Gunn–Peterson approximation*, which yields the following relation for the for Lyman-α optical depth, in velocity space:

$$\tau(z) = 1.3\Delta_b \left(\frac{x_{\rm HI}}{10^{-5}}\right) \left(\frac{1+z}{4}\right)^{3/2} \left(\frac{H(z)/(1+z)}{dv/dx}\right), \tag{1.12}$$

where Δ_b is the baryon density in terms of the cosmic mean, and $x_{\rm HI}$ is the fraction of neutral hydrogen. As long as the gas is in photo-ionisation equilibrium, the value of $x_{\rm HI}$ is given by the balance between photo-ionisation and recombination, except for regions which have been shock heated to temperatures higher than $\mathcal{O}(10^5 - 10^6$ K), where collisional ionisation is not negligible. Note that Equation (1.12) imply that the Lyman-α forest is sensitive to $x_{\rm HI} \sim 10^{-5}$ at $z = 3$, which corresponds to very low neutral hydrogen number densities, i.e. $n_{\rm HI} \sim 10^{-10}$ cm^{-3}, and to overdensities with respect to the mean density of the universe less than about a factor $\delta \sim 10 - 15$.

Inference on cosmological parameters from the Lyman-α forest spectra is complicated by there being no reliable analytic description for the mildly non-linear densities probed by the forest. All of the analyses do require, indeed, to be confronted against cosmological numerical simulations, accounting for all the typical astrophysical contaminants of the signal. Once again, we refer to the review by McQuinn (2016) for a detailed review of the most commonly adopted prescriptions to extract reliable mocks from simulations. Let us point out that the simulations used in this thesis have also been performed following such prescriptions, which are indeed discussed in the next Chapters, when the methods and results of our Lyman-α analyses are presented.

Although the absorption line parameters, appearing in Equation (1.12), are directly related to the physical properties of the absorption systems, estimating such parameters in numerical simulations can be very computationally demanding. Easier comparisons can be made by directly focusing on the statistics of the transmitted Lyman-α flux, $F \propto e^{-\tau}$. In particular, one of the most studied observable quantities is the 1D *flux* power spectrum:

$$P_F(k) = L^{-1}|\delta_F(k)|^2, \tag{1.13}$$

where $\delta_F(k)$ is the Fourier transform of $\delta_F(x)$, namely the overdensity in the transmitted flux at a position x over a line of sight L, i.e.

$$\delta_F(x) = \frac{F(x) - \langle F(x)\rangle}{\langle F(x)\rangle}, \tag{1.14}$$

with $\langle F(x)\rangle$ being the mean flux transmission fraction, obtained by averaging over all the considered lines of sight.

Let us stress that the 1D flux spectrum is, by construction, a projection of the 3D one, i.e.

$$P_{\rm 1D}(k) = \frac{1}{2\pi}\int_k^\infty dk' k' P_{\rm 3D}(k'), \tag{1.15}$$

so that:

$$P_{\rm F}(k) = b^2(k) P_{\rm 1D}(k), \tag{1.16}$$

where $b(k)$ is a scale-dependent bias factor.

Clearly, Equation (1.15) implies that the 1D flux power is much more sensitive to the small scales with respect to the 3D spectrum. This fact, in addition to its robustness to the messiest astrophysical processes – such as spatial variations in the IGM temperature or in the ionising UV background, as well as galactic feedback (see, e.g. Meiksin & White (2004) or the review by McQuinn (2016) and the references therein) – makes this observable an ideal tool to investigate the small-scale DM distribution.

This Chapter is mainly based on:

R. Murgia, A. Merle, M. Viel, M. Totzauer, A. Schneider
"Non-cold" dark matter at small scales: a general approach
JCAP 11 046 (2017); arXiv:1704.07838

R. Murgia, V. Iršič, M. Viel
Novel constraints on non-cold (non-thermal) Dark Matter from Lyman-α forest data
PRD 98, 083540 (2018); arXiv:1806.08371

2

A new & general approach to test non-standard Dark Matter

2.1 Overview

In order to give a better description of the structure formation and distribution at small scales with respect to the standard ΛCDM model, many *non-cold* DM (nCDM) candidates, well motivated by particle physics theories, such as sterile neutrinos (de Vega et al. 2012; König et al. 2016; Schneider 2016; Merle 2017; Yèche et al. 2017; Adhikari et al. 2017) or axion-like particles (Hu et al. 2000; Marsh & Silk 2014; Hui et al. 2017; Iršič et al. 2017a; Armengaud et al. 2017; Kobayashi et al. 2017; Nori et al. 2019), have been recently studied in the literature. In addition, there are other non-standard hypotheses potentially able to induce a small-scale suppression in the matter power spectrum: a mixed (cold and warm) DM fluid (Viel et al. 2005; Schneider et al. 2017; Diamanti et al. 2017; Gariazzo et al. 2017), DM particles coupled to DE (Wang et al. 2016; Murgia et al. 2016; Murgia 2016) or to some relativistic fluid (Boehm et al. 2014; Bringmann et al. 2016; Archidiacono et al. 2019), or Self-Interacting Dark Matter (SIDM) (Cyr-Racine et al. 2016; Vogelsberger et al. 2016).

Different scenarios lead to different shapes in the suppression of the power spectrum, allowing for a direct link between DM models and astrophysical observations (Murgia et al. 2017; Murgia 2018; Murgia et al. 2018; Archidiacono et al. 2019; Miller et al. 2019). However, most of the constraints from structure formation data which have been published so far, refer to a very specific shape of the small-scale power suppression, corresponding to the case of *thermal* Warm Dark Matter (WDM), i.e., candidates with a Fermi-Dirac or Bose-Einstein momentum

distribution (Viel et al. 2013; Baur et al. 2016; Lapi & Danese 2015; Iršič et al. 2017b). Nonetheless, most of viable nCDM candidates do not have thermally distributed momenta (i.e., they feature truly *non-thermal* distribution functions), which may lead to a non-trivial suppression in their power spectra. In particular, given that non-thermal distributions may feature more than one characteristic momentum scale, these settings may in fact provide a whole new approach to resolving the small scale issues of CDM.

The suppressed gravitational clustering characterising nCDM models is usually parametrised through the *transfer function* $T(k)$, i.e., the square root of the ratio of the matter power spectrum predicted by the given model with respect to that in the presence of CDM only, for fixed cosmological parameters. In this Chapter we introduce a novel, general approach to test nCDM models, based on a new analytic fitting formula for $T(k)$, which is simple yet flexible enough to reproduce a large variety of shapes with only three free parameters.

This Chapter is organised as follows: in Section 2.2 we introduce and motivate our general approach; in Section 2.3 we show that the new formula is an useful tool for fitting the behaviour of the most viable DM models provided by particle physics; in Section 2.4 we discuss the results of the cosmological simulations that we have performed in order to investigate different parametrisations of the new transfer function; in Section 2.5 we discuss preliminary astrophysical constraints on its free parameters, based on linear perturbation theory only; in Section 2.6, we present the most up-to-date accurate limits on such parameters, obtained through an extensive analysis the high-resolution and high-redshift Lyman-α forest data, easily translatable to bounds on the fundamental nCDM properties.

2.2 The novel parametrisation

The suppression of gravitational clustering in nCDM scenarios can be described through the transfer function $T(k)$, given by:

$$T^2(k) \equiv \frac{P_{\rm nCDM}(k)}{P_{\rm CDM}(k)}, \tag{2.1}$$

where $P_{\rm CDM}$ and $P_{\rm nCDM}$ are the power spectra of the CDM and the nCDM models, respectively.

For the special case of thermal WDM, the transfer function can be well approximated by the fitting function (Bode et al. (2001)):

$$T(k) = [1 + (\alpha k)^{2\mu}]^{-5/\mu}, \tag{2.2}$$

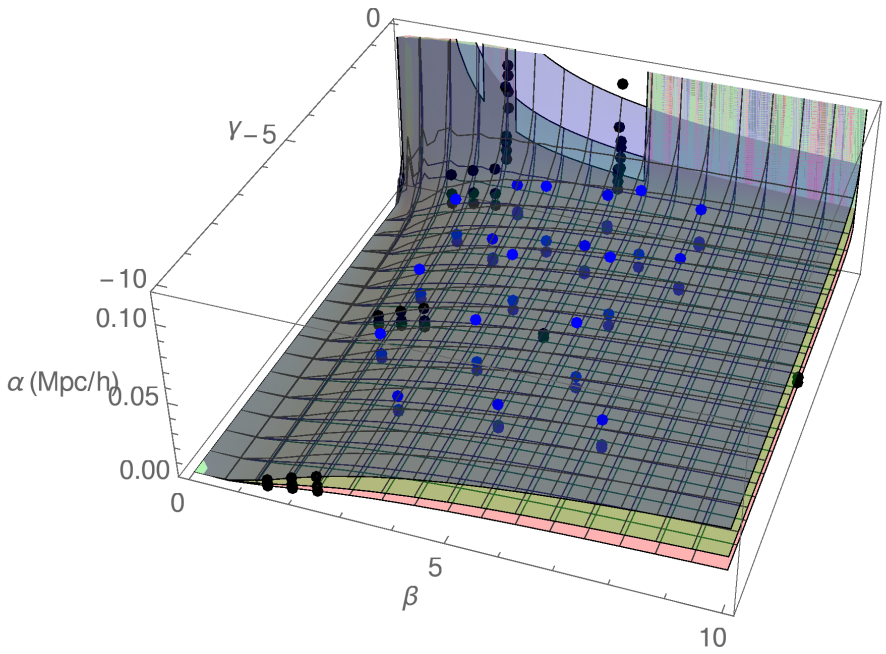

Figure 2.1: The blue, green, and red surfaces represent the regions of the $\{\alpha, \beta, \gamma\}$-space corresponding to thermal WDM masses of 2, 3, and 4 keV, respectively. The dots constitute the non-regular grid that we considered for our analyses: the black ones correspond to the first 55 $\{\alpha, \beta, \gamma\}$-combinations listed in Table 2.1, while the blue ones correspond the the remaining 54 combinations reported in the same Table.

where α is the only free parameter, and $\mu = 1.12$. Therefore, constraints on the mass of the WDM candidate translate into bounds on α, by the following formula (Viel et al. (2005)):

$$\begin{aligned} \alpha &= 0.24 \left(\frac{m_x/T_x}{1\ \mathrm{keV/T_\nu}}\right)^{-0.83} \left(\frac{\omega_x}{0.25\cdot(0.7)^2}\right)^{-0.16} \mathrm{Mpc} \\ &= 0.049 \left(\frac{m_x}{1\ \mathrm{keV}}\right)^{-1.11} \left(\frac{\Omega_x}{0.25}\right)^{0.11} \left(\frac{h}{0.7}\right)^{1.22} h^{-1}\mathrm{Mpc}\ , \end{aligned} \tag{2.3}$$

where the subscripts x and ν refer to WDM and active-neutrino properties, respectively, and the second equation holds only in the case of thermal relics.

Let us now generalise Equation (2.2) and write down the following fitting formula[1]

$$T(k) = [1 + (\alpha k)^\beta]^\gamma, \tag{2.4}$$

which is a function of three free parameters: α, β, and γ. In Section 2.3 we show that the simple function given by Equation (2.4) is generic enough to describe the majority of nCDM models from the literature.

[1]Note that equivalent fitting functions have already been used by, e.g., Barkana et al. (2001); Viel et al. (2012); Destri et al. (2013). However, they have only been applied to special cases, and its general applicability had not been recognised to our knowledge.

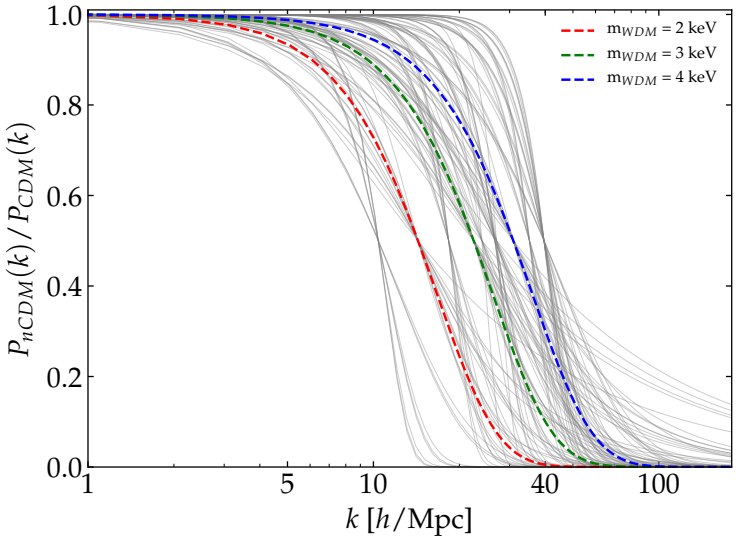

Figure 2.2: Squared transfer functions computed by Equation (2.4), and associated to the 109 $\{\alpha, \beta, \gamma\}$-combinations that we used for our analyses (grey solid lines), each of them corresponding to a different nCDM model (see Table 2.1). We also plot three squared transfer functions computed through Equation (2.2), corresponding to thermal WDM models with masses 2, 3 and 4 keV (red, green and blue dashed lines, respectively).

As a next step, let us define the characteristic half-mode scale $k_{1/2}$, obtained by setting $T^2 \equiv 1/2$. Using Equation (2.4), we therefore have:

$$k_{1/2} = \frac{1}{\alpha}\left[\left(\frac{1}{\sqrt{2}}\right)^{1/\gamma} - 1\right]^{1/\beta} . \tag{2.5}$$

Whereas through Equations (2.2) and (2.3) we had a one-to-one correspondence between m_x and α, constraints on the DM mass are now, by Equations (2.4) and (2.5), mapped to 3D surfaces in the $\{\alpha, \beta, \gamma\}$-space. In other words, given a value of $k_{1/2}$, which corresponds to a certain (thermal) WDM mass, we can easily compute the corresponding surface in the 3D parameter space from Equation (2.5) – but this information alone is not yet sufficient to decide about the validity of the point under consideration. In Figure 2.1 we plot the three surfaces associated to the $k_{1/2}$-values listed below:

$$\begin{aligned} k'_{1/2} &= 14.323\ h/\mathrm{Mpc} \quad (\textit{if thermal}: \longleftrightarrow\ m'_x = 2\ \mathrm{keV}), \\ k''_{1/2} &= 22.463\ h/\mathrm{Mpc} \quad (\textit{if thermal}: \longleftrightarrow\ m''_x = 3\ \mathrm{keV}), \\ k'''_{1/2} &= 30.914\ h/\mathrm{Mpc} \quad (\textit{if thermal}: \longleftrightarrow\ m'''_x = 4\ \mathrm{keV}). \end{aligned} \tag{2.6}$$

It is indeed well established that thermal warm DM candidates with masses of the order of 3 keV can induce a suppression in the corresponding matter power

spectra such that the CDM small-scale crisis vanishes or it is largely reduced (Lovell et al. (2016, 2017)). Hence, it is compelling to investigate the volume of the $\{\alpha, \beta, \gamma\}$-space associated to thermal WDM masses roughly between 2 and 4 keV. We did that by building a 3D grid in the parameter space which samples such volume, with each of the grid points unequivocally identified by a certain $\{\alpha, \beta, \gamma\}$-combination, corresponding to a different nCDM model. The $\{\alpha, \beta, \gamma\}$-combinations which constitute such grid are listed in Table 2.1 and plotted as black and blue dots in Figure 2.1. Note that, although the points marked in Figure 2.1 may appear to be somewhat sparsely distributed at first sight, they in fact cover a large fraction of the relevant parameter space. The reason for this lies in a quasi-degeneracy between the two parameters α and γ, which we discuss in detail in Appendix A.

For each of the models listed in Table 2.1, we computed the corresponding transfer function by using Equation (2.4). We report them in Figure 2.2, where the red, green and blue dashed lines represent the "old" transfer functions, i.e. computed through Equation (2.2), for $m'_x = 2$ keV, $m''_x = 3$ keV, and $m'''_x = 4$ keV, respectively.

Let us now qualitatively describe the role of the different parameters in the generalised fit for the transfer function. The value of α gives the general scale of suppression, i.e., it is the most important parameter for setting the position of $k_{1/2}$. The parameters β and γ are responsible for the slope of the transfer function before and after the half-mode scale $k_{1/2}$, respectively. The parameter β has to be positive in order to have meaningful transfer functions, since negative values for β lead to transfer functions which increase with k and reach 1 at small scales. The larger is β, the flatter is the transfer function before $k_{1/2}$. Analogously, the larger is the absolute value of γ, the sharper is the cut-off.

2.3 Connection to particle physics models

The purpose of this Section is to see to which extent the suggested 3-parameter fitting formula is able to match the transfer functions from different nCDM models. We chose to focus on sterile neutrinos from resonant production (RP), sterile neutrinos from particle decay production, mixed (cold plus warm) DM models, ultra-light scalar DM, and another class of models suggested by the effective theory of structure formation (ETHOS).

2.3.1 Sterile neutrinos by resonant production

Given that keV sterile neutrinos generically mix with the active-neutrino sector, it is a natural idea to use this mixing to produce sterile neutrino DM in the early universe. While it is nowadays known that the production by non-resonant transitions (*Dodelson-Widrow* mechanism) (Langacker (1989); Dodelson & Widrow (1994);

	α [Mpc/h]	β	γ		α [Mpc/h]	β	γ
nCDM1	0.008	1.5	-10.00	nCDM56	0.023	2.0	-6.00
nCDM2	0.005	1.5	-10.00	**nCDM57**	**0.009**	**2.0**	**-6.00**
nCDM3	**0.003**	**1.5**	**-10.00**	**nCDM58**	**0.006**	**2.0**	**-6.00**
nCDM4	0.012	1.5	-5.00	nCDM59	0.029	2.0	-4.00
nCDM5	0.008	1.5	-5.00	**nCDM60**	**0.011**	**2.0**	**-4.00**
nCDM6	**0.006**	**1.5**	**-5.00**	**nCDM61**	**0.008**	**2.0**	**-4.00**
nCDM7	0.039	1.5	-1.00	nCDM62	0.042	2.0	-2.00
nCDM8	0.025	1.5	-1.00	**nCDM63**	**0.016**	**2.0**	**-2.00**
nCDM9	**0.018**	**1.5**	**-1.00**	**nCDM64**	**0.011**	**2.0**	**-2.00**
nCDM10	0.013	2.0	-10.00	nCDM65	0.047	4.0	-6.00
nCDM11	0.008	2.0	-10.00	**nCDM66**	**0.019**	**4.0**	**-6.00**
nCDM12	**0.006**	**2.0**	**-10.00**	**nCDM67**	**0.012**	**4.0**	**-6.00**
nCDM13	0.019	2.0	-5.00	nCDM68	0.053	4.0	-4.00
nCDM14	0.012	2.0	-5.00	**nCDM69**	**0.021**	**4.0**	**-4.00**
nCDM15	**0.009**	**2.0**	**-5.00**	**nCDM70**	**0.014**	**4.0**	**-4.00**
nCDM16	0.045	2.0	-1.00	nCDM71	0.063	4.0	-2.00
nCDM17	0.029	2.0	-1.00	**nCDM72**	**0.025**	**4.0**	**-2.00**
nCDM18	**0.021**	**2.0**	**-1.00**	nCDM73	**0.017**	**4.0**	**-2.00**
nCDM19	0.018	2.5	-10.00	nCDM74	0.060	6.0	-6.00
nCDM20	0.012	2.5	-10.00	**nCDM75**	**0.023**	**6.0**	**-6.00**
nCDM21	**0.008**	**2.5**	**-10.00**	nCDM76	**0.016**	**6.0**	**-6.00**
nCDM22	0.024	2.5	-5.00	nCDM77	0.064	6.0	-4.00
nCDM23	0.016	2.5	-5.00	**nCDM78**	**0.025**	**6.0**	**-4.00**
nCDM24	**0.011**	**2.5**	**-5.00**	**nCDM79**	**0.017**	**6.0**	**-4.00**
nCDM25	0.049	2.5	-1.00	nCDM80	0.073	6.0	-2.00
nCDM26	0.031	2.5	-1.00	**nCDM81**	**0.028**	**6.0**	**-2.00**
nCDM27	**0.023**	**2.5**	**-1.00**	**nCDM82**	**0.019**	**6.0**	**-2.00**
nCDM28	0.011	2.0	-5.00	nCDM83	0.020	3.0	-7.50
nCDM29	**0.010**	**2.0**	**-5.00**	**nCDM84**	**0.010**	**3.0**	**-7.50**
nCDM30	0.015	2.5	-5.00	**nCDM85**	**0.009**	**3.0**	**-7.50**
nCDM31	**0.013**	**2.5**	**-5.00**	nCDM86	0.029	3.0	-2.50
nCDM32	**0.025**	**5.0**	**-5.00**	**nCDM87**	**0.015**	**3.0**	**-2.50**
nCDM33	**0.022**	**5.0**	**-5.00**	**nCDM88**	**0.013**	**3.0**	**-2.50**
nCDM34	**0.032**	**10.0**	**-5.00**	nCDM89	0.041	3.0	-1.00
nCDM35	**0.028**	**10.0**	**-5.00**	**nCDM90**	**0.021**	**3.0**	**-1.00**
nCDM36	0.095	2.5	-0.30	**nCDM91**	**0.019**	**3.0**	**-1.00**
nCDM37	0.169	2.5	-0.15	nCDM92	0.030	5.0	-7.50
nCDM38	0.061	2.5	-0.30	**nCDM93**	**0.015**	**5.0**	**-7.50**
nCDM39	0.108	2.5	-0.15	**nCDM94**	**0.014**	**5.0**	**-7.50**
nCDM40	**0.044**	**2.5**	**-0.30**	nCDM95	0.037	5.0	-2.50
nCDM41	**0.078**	**2.5**	**-0.15**	**nCDM96**	**0.019**	**5.0**	**-2.50**
nCDM42	0.057	2.5	-0.30	**nCDM97**	**0.017**	**5.0**	**-2.50**
nCDM43	0.101	2.5	-0.15	nCDM98	0.046	5.0	-1.00
nCDM44	**0.051**	**2.5**	**-0.30**	**nCDM99**	**0.024**	**5.0**	**-1.00**
nCDM45	**0.090**	**2.5**	**-0.15**	**nCDM100**	**0.021**	**5.0**	**-1.00**
nCDM46	0.082	5.0	-0.30	nCDM101	0.035	7.0	-7.50
nCDM47	0.109	5.0	-0.15	**nCDM102**	**0.018**	**7.0**	**-7.50**
nCDM48	0.052	5.0	-0.30	**nCDM103**	**0.016**	**7.0**	**-7.50**
nCDM49	0.069	5.0	-0.15	nCDM104	0.042	7.0	-2.50
nCDM50	**0.038**	**5.0**	**-0.30**	**nCDM105**	**0.022**	**7.0**	**-2.50**
nCDM51	**0.050**	**5.0**	**-0.15**	**nCDM106**	**0.019**	**7.0**	**-2.50**
nCDM52	**0.049**	**5.0**	**-0.30**	nCDM107	0.048	7.0	-1.00
nCDM53	0.065	5.0	-0.15	**nCDM108**	**0.025**	**7.0**	**-1.00**
nCDM54	**0.043**	**5.0**	**-0.30**	**nCDM109**	**0.022**	**7.0**	**-1.00**
nCDM55	**0.058**	**5.0**	**-0.15**				

Table 2.1: Here we report the 109 $\{\alpha, \beta, \gamma\}$-combinations that we considered for our analyses, each of them associated to a different nCDM model. We used the corresponding transfer functions, computed via Equation (2.4), as initial conditions to perform the cosmological simulations described in Section 2.4. Models highlighted in bold-face are accepted at 2 σ C.L. by our reference Lyman-α forest analysis (see Section 2.6).

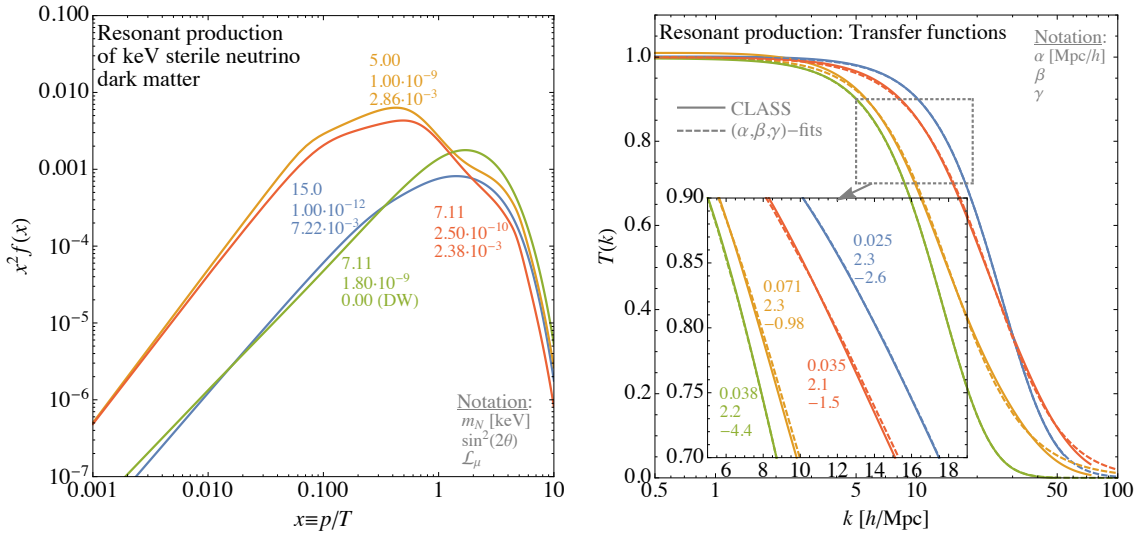

Figure 2.3: Example distributions functions (*left*) and corresponding transfer functions (at $z = 0$; *right*) for resonant (Shi-Fuller) production, including one example for non-resonant (Dodelson-Widrow) production, green curve, for comparison. On the right panel, one can see that the transfer functions are fitted very well by the parametrisation from Equation (2.4).

Merle et al. (2016)[2]) is incompatible with structure formation (Viel et al. (2013); Adhikari et al. (2017)), a suitable lepton number asymmetry in the early universe (whose origin is not necessarily clear, though) can resonantly enhance the active-sterile transitions (*Shi-Fuller* mechanism) and yield spectra that are more likely to be in agreement with data (Enqvist et al. (1990); Shi & Fuller (1999); Abazajian et al. (2001); Canetti et al. (2013); Venumadhav et al. (2016); Ghiglieri & Laine (2015)). Note, however, that also this mechanism is restricted to a small successful region in the parameter space (Cherry & Horiuchi (2017); Adhikari et al. (2017)).

As can be seen from Abazajian et al. (2001); Venumadhav et al. (2016); Ghiglieri & Laine (2015), the distribution functions resulting from resonant production can be highly non-thermal: typically, they feature one or more narrow peaks on top of a continuous spectrum, see Figure 2.3. This Figure shows different momentum distribution functions (*left*) and the corresponding transfer functions (*right*), for a few example values of the sterile neutrino mass m_N, the active-sterile mixing angle $\sin^2(2\theta)$, and the lepton asymmetry $\mathcal{L}_\mu$.[3] Note that the green curve actually

[2]Note that, contrary to previous statements in the literature (Dodelson & Widrow (1994); Colombi et al. (1996)), non-resonantly produced sterile neutrinos also feature a non-thermal distribution (Merle et al. (2016)), rather than a suppressed thermal spectrum. However, while this slightly changes the published numerical values of the bounds on this setting, the basic conclusion of non-resonant production being excluded remains valid (in fact, it is even made stronger) (Merle et al. (2016)).

[3]Due to the current technical limitations of the software developed in conjunction with Venumadhav et al. (2016), the lepton asymmetry can only be placed in the muon sector, if the package `sterile-dm` is used. However, as shown by Ghiglieri & Laine (2015), the results would not be altered dramatically if the lepton asymmetry was present in another sector.

features $\mathcal{L}_\mu \equiv 0$, i.e., a case of non-resonant production. Compared to the red curve, one can see that in this case a larger mixing angle is required to meet the correct abundance, and also the spectrum is different from the resonant cases. The plots in the right panel illustrate the corresponding transfer functions (solid lines), along with the fits obtained from Equation (2.4) using a least squares approach (dashed lines). As can be seen already by eye (and is confirmed by a goodness-of-fit test), our general transfer function, Equation (2.4), provides excellent fits to these cases, with parameter values within the grid that we defined (see Figure 2.1 and Table 2.1). This remains true also for distributions with more than one scale, which can e.g. be seen for the orange curve in the plots.

Given that the distributions that we tested (more than shown here) are very representative for resonant production of sterile neutrino DM, we conclude that our fitting function describes this class of models very well.

2.3.2 Sterile neutrinos from particle decays

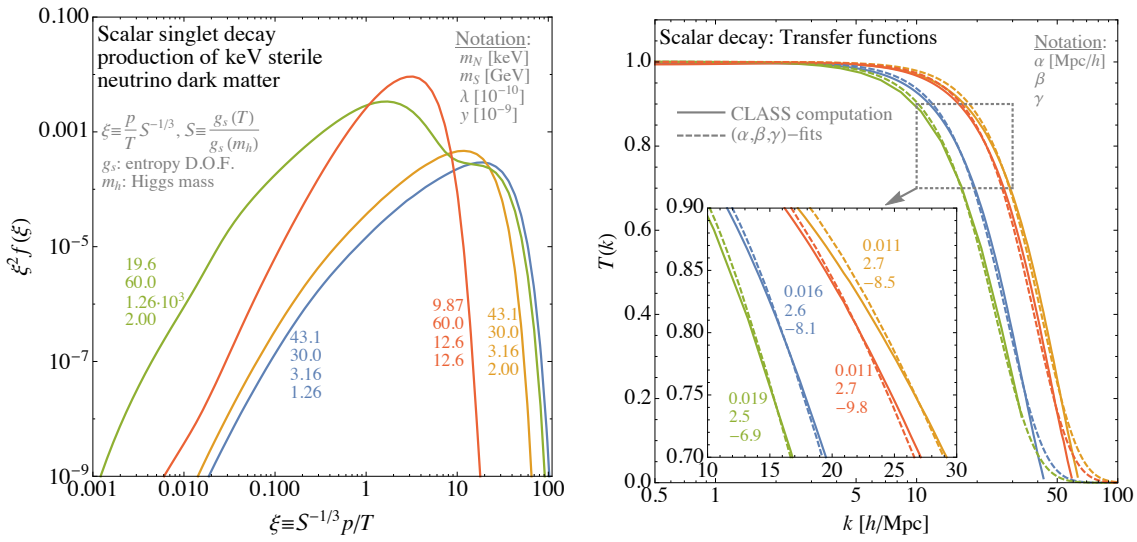

Figure 2.4: Example distributions functions (*left*) and corresponding transfer functions (at $z = 0$; *right*) for scalar decay production. On the right panel, one can see that the transfer functions are fitted very well by the parametrisation from Equation (2.4).

Another potential production mechanism for sterile neutrino DM relies on the decay of a hypothetical parent particle in the early universe, whose properties (in terms of the momentum distribution) translate into those of the resulting keV sterile neutrino. A very simple case discussed frequently in the literature is that of a singlet scalar particle which may via its interactions with the Standard Model Higgs boson either thermalise (and thus be equilibrated and ultimately freeze-out) or not (and thus freeze-in), see Kusenko (2006); Petraki & Kusenko (2008); Merle

et al. (2014); Merle & Totzauer (2015); Shakya (2016); König et al. (2016) for very detailed treatments. Other possibilities for parent particles, like pions (Lello & Boyanovsky (2015)) or electrically charged scalars (Frigerio & Yaguna (2015)), do not exhibit any qualitatively different behaviour.

For decay production, the resulting distribution functions are highly non-thermal, with spectra not only having a shape very different from the thermal one but also featuring, in general, two distinct characteristic scales (Merle & Totzauer (2015); König et al. (2016)), or even three if a subdominant subsequent Dodelson-Widrow modification is taken into account (Merle et al. (2016)). Four example distributions are depicted in the left panel of Figure 2.4, for different values of the sterile neutrino and decaying scalar masses, along with the two parameters λ (Higgs portal) and y (Yukawa coupling), which shape the distributions (see Merle & Totzauer (2015); König et al. (2016) for details). The plots in the right panel illustrate the corresponding transfer functions (solid lines), along with the fits obtained from Equation (2.4) using a least squares approach (dashed lines). As for the case of resonantly produced sterile neutrinos, the general parametrisation of Equation (2.4), provides an excellent fit to the "exact" transfer functions, with parameter values within the grid that we defined (see Figure 2.1). This is true also for distributions with more than one scale, which can be seen for the green curve in the plots.

Given that the distributions that we tested (more than shown here) are very representative for decay production of DM, independently of the details of the particle physics setting under consideration, we conclude that our fitting function describes this class of models very well.

2.3.3 Mixed (cold plus warm) models

In principle, the DM sector may consist of a complicated mixture of different DM particles (Boyarsky et al. (2009)). Here we study a toy model that assumes the presence of both a cold and a warm (thermal) component. This simple model, dubbed as mixed DM, leads to a large variety of shapes in the transfer function, therefore providing an ideal test for the novel parametrisation, i.e., the fitting formula of Equation (2.4).

Mixed DM is characterised by two parameters: the mass of the WDM component and the fraction f of the warm to the total DM abundance, i.e., $f = \Omega_{\rm WDM}/\Omega_{\rm total}$, where $\Omega_{\rm total} = \Omega_{\rm WDM} + \Omega_{\rm CDM}$ denotes the total DM abundance in the universe. The fraction f parametrises the example settings illustrated in Figure 2.5, where we depict both the power spectra (*left*) and the transfer functions (*right*). For mixed DM, it had been pointed out by Boyarsky et al. (2009) that a non-zero plateau can be present in the transfer function at large k's, corresponding to the remaining CDM component once the reduction of small scales by the warm component has

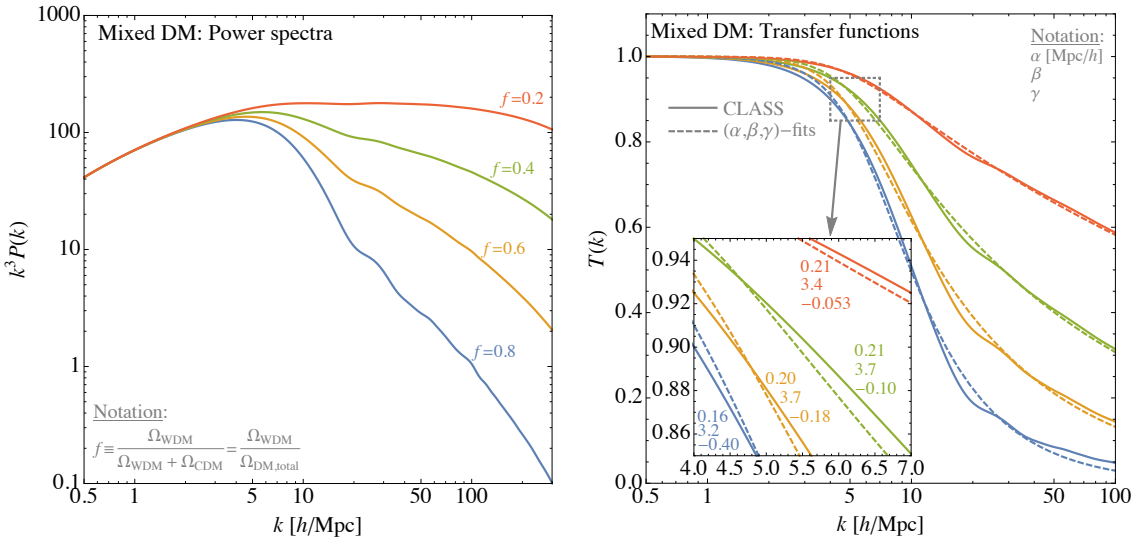

Figure 2.5: Example power spectra (*left*) and corresponding transfer functions (*right*) for mixed DM, derived with `CLASS` (Blas et al. (2011)). On the right panel, one can see that even the transfer functions featuring some type of plateau are fitted well by the parametrisation from Equation (2.4).

died off. However, although our fitting function does not formally feature a plateau, it provides a very good fit to the majority of mixed DM cases, not only by eye but also when performing a goodness-of-fit test. In Chapter 4 (Section 4.4) we will discuss a further generalisation of our fitting formula, which allows much more flexibility to describe scenarios where only a fraction of DM is non-standard.

2.3.4 Fuzzy Dark Matter

A conceptually different class of DM candidates that also affects the small scales of structure formation is the so-called ultra-light scalar DM, or Fuzzy DM (FDM) (Hu et al. (2000); Marsh & Silk (2014); Hui et al. (2017)). This type of DM consists of (initially) condensed scalar particles with tiny masses, $\sim 10^{-22}$ eV, such that their wave-like nature becomes relevant at astrophysical scales. A comprehensive treatment of such scenario is carried out in Chapter 3. However, let us now briefly introduce its main features, before to show that the induced small-scale power suppression is also well described by our $\{\alpha, \beta, \gamma\}$-parametrisation.

These ultra-light scalar particles are assumed to have no self-interactions – quite like axions (Hui et al. (2017)) – which could modify the picture if sufficiently strong (Rindler-Daller & Shapiro (2014)). In the absence of such interactions, the class of FDM models is conveniently described by a single parameter, namely the DM mass $m_{22} \equiv m_\psi/10^{-22}$ eV, where m_ψ denotes the actual physical particle mass. However, note that strong constraints exist on these scenarios, e.g. upper limits $m_{22} < 1.5$ from the kinematics of dwarf galaxies (Marsh & Pop (2015); Calabrese

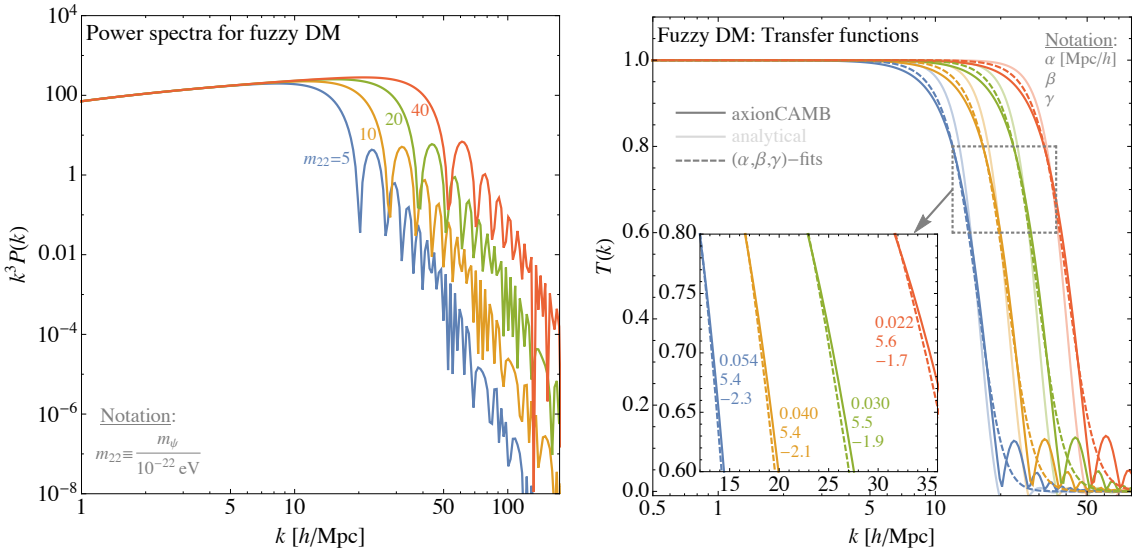

Figure 2.6: Example power spectra (*left*) and corresponding transfer functions (*right*) for FDM, derived with `axionCAMB` (Hlozek et al. (2015)). On the right, we also show the analytic result by Hu et al. (2000). On the right panel, one can see that the transfer functions are fitted very well by the parametrisation from Equation (2.4).

& Spergel (2016)), which were recently complemented by strong lower limits from the abundance of high-z galaxies (Menci et al. (2017)), $m_{22} > 10$, superseeding earlier limits from their luminosity functions ($m_{22} > 1.2$ (Schive et al. (2016))) by nearly one order of magnitude. Even more recently, the IGM provided the tightest limits on the mass of $m_{22} > 20$ (2 σ C.L.) for a very conservative analysis of high-redshift data, while $m_{22} > 37.5$ is obtained for a less conservative scenario where some priors on the IGM thermal history are assumed (Iršič et al. (2017a); Kobayashi et al. (2017); Nori et al. (2019)).

In Figure 2.6, we show a few example power spectra (*left*) and transfer functions (*right*), associated with different values of m_{22}. The point that we want to illustrate is that, even though the known oscillations are present in the FDM power spectra, their transfer functions are still well described by our general parametrisation from Equation (2.4), simply because the oscillations appear only at the smallest scales or, equivalently, at large values of k. To achieve a good fit, we have simply cut off the oscillations, which are very suppressed and therefore negligible for most applications, including the MW satellite counting and the Lyman-α forest methods. Under this approximation, we see that the novel parametrisation in fact provides a very good fit to the FDM transfer functions (see the right panel of Figure 2.6). We show in Section 2.5.3 how well this strategy truly works, when both the actual models and our fits are subjected to a "reality-check". Moreover, we address the reader to Section 2.6.3 and Chapter 4 (Section 4.2.1), where the validity of this procedure is explicitly demonstrated by means of dedicated hydrodynamic simulations.

2.3.5 Effective theory of structure formation (ETHOS)

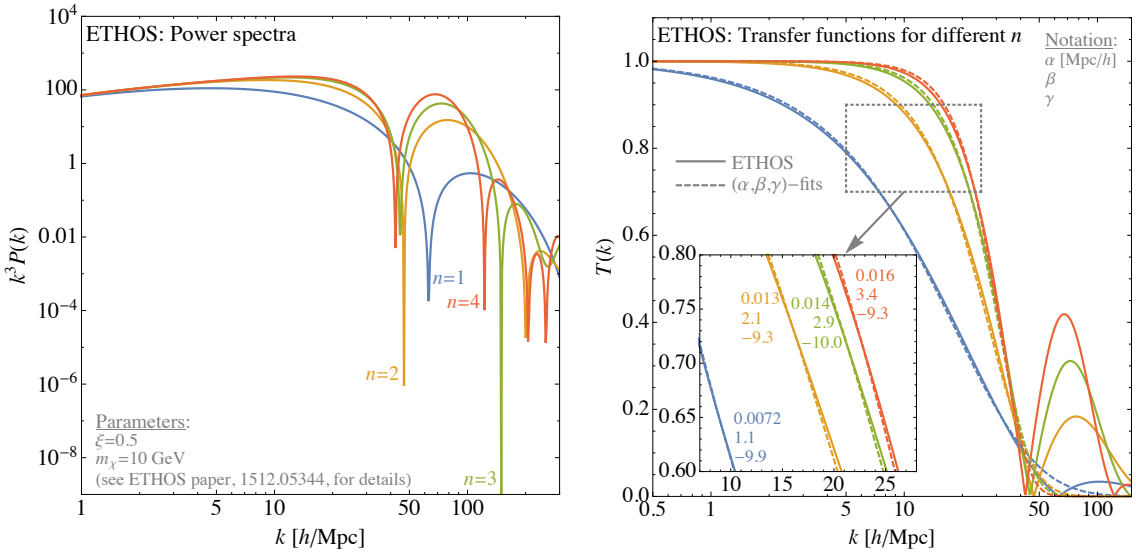

Figure 2.7: Power spectra (*left*) and corresponding transfer functions (*right*) for a few cases of interacting DM as derived with ETHOS, see Figure 1a from Cyr-Racine et al. (2016). On the right panel, one can see that the transfer functions are fitted well by the parametrisation from Equation (2.4). However, note that we have only fitted the part of k left of the first oscillation.

In order to further demonstrate the flexibility of our approach, we now compare it to that of ETHOS (Cyr-Racine et al. (2016); Vogelsberger et al. (2016)), which consists of an attempt to formulate an effective theory of cosmic structure formation, to map virtually any particle physics model to the constraints from astrophysics and cosmology. In Chapter 4 we will carry out a comprehensive analysis of the ETHOS framework, which to some extent has been developed in the same spirit as ours. Here we just want to point out the main differences with respect to our approach:

- *Lagrangian-based vs. transfer function-based approach*:
  While we are simply trying to parametrise the transfer function for nCDM, the approach of ETHOS is to start directly from the particle physics Lagrangian. While at first this seems like a clear disadvantage of our strategy, since a particle physicist would need to compute the matter power spectrum (e.g., using `CLASS` (Blas et al. (2011))) before being able to apply our results to their model, we would like to point out that the mapping of DM models into the quantities relevant for cosmic structure formation is injective. In other words, many models that look quite different from the particle physics point of view lead precisely to the same predictions for structure formation (this can be seen easily for thermal examples, e.g., by re-scaling both temperature and DM mass). With our approach it is therefore possible to summarise most DM models in a

much simpler framework. In fact, it may be a too big effort to start from any possible particle physics Lagrangian, when the key point for comparison lies in the transfer function. We thus consider our approach to be the most economic.

- *Validity for small scales, i.e., for large k*:
  As we pointed out, e.g., in Section 2.3.4, our approach cannot capture the very smallest scales, which are suppressed in the transfer function due to dividing by the CDM power spectrum. However, given that this is in any case the part of the spectrum with the smallest power, we do not expect big effects of this region, unless we find an observable that is truly sensitive to very small scales. The ones that we are using (i.e., satellite counts and Lyman-α) are not sensitive to that extreme region. For instance, as visible in the left panel of Figure 2.7, oscillations in the transfer functions seem very prominent when plotted in log-scales, although they are in fact unimportant for most aspects of structure formation (once again, for further details on this issue, see Sections 2.6.3 and 4.2.1). We thus consider our approach to be safe as long as the regions for (very) large k's play no role.

- *Model coverage*:
  By construction, our approach of fitting the transfer function is much simpler and therefore less versatile than a method starting from the particle Lagrangian. However, we want to point that the original fitting function of Equation 2.4 can be easily generalised if necessary. For example, a plateau in the transfer function can be described by adding one more parameter δ[4], i.e.

$$T(k) = [1 + (\alpha k)^{\beta}]^{\gamma} \rightarrow T(k) = (1 - \delta) \cdot [1 + (\alpha k)^{\beta}]^{\gamma} + \delta. \tag{2.7}$$

  In principle, oscillatory patterns could also be included, for example by simply adding a cosine function of the form

$$T(k) = [1 + (\alpha k)^{\beta}]^{\gamma} \cdot \cos^2(\delta k) \tag{2.8}$$

  where δ is an additional free parameter. We thus consider our approach to be easily extendable to basically cover the same range of models as ETHOS does.

Given that, we can now fit some of the transfer functions obtained by ETHOS, and for this purpose we take the ones given in Figure 1a from Cyr-Racine et al. (2016) as example (see that reference for the details on the data chosen).[5] This allows

[4]This will be the subject of Section 4.4

[5]Note that Cyr-Racine et al. (2016) use an alternative definition of the transfer function of $T_{\rm ETHOS}(k) \equiv P(k)_{\rm nCDM}/P(k)_{\rm CDM}$, instead of the more common definition of $T^2(k) \equiv P(k)_{\rm nCDM}/P(k)_{\rm CDM}$, which is the one we are using. This difference in definitions is what creates the seeming difference between the right panel of our Figure 2.7 and Figure 1a from Cyr-Racine et al. (2016).

us to investigate whether our approach yields a similar result as the more detailed ETHOS analysis. We have again fitted the transfer functions with Equation (2.4), however, note that this time we have only fitted the part for small k's, i.e., left of the first oscillation. In Section 2.5.3 we show that regarding the number of MW satellites as well as the power spectrum from the Lyman-α forest, there is hardly any difference between our simplified fit without oscillations and the full transfer function from ETHOS. In Section 2.6.3 and Section 4.2.1, we further demonstrate the accuracy of this procedure by means of dedicated hydrodynamic simulations.

2.4 Cosmological simulations

In order to put constraints on the nCDM models, we modified the numerical code `2LPTic` (Crocce et al. (2006)) – which generates initial conditions for running cosmological simulations by displacing the DM particles from a cubic Cartesian grid according to second-order Lagrangian Perturbation Theory – by implementing the new, general transfer function: now the code takes as inputs $\{\alpha, \beta, \gamma\}$ instead of the thermal WDM mass, and it computes the corresponding $T(k)$ from Equation (2.4).

Thus, we generated the initial conditions corresponding to the 109 $\{\alpha, \beta, \gamma\}$-combinations listed in Table 2.1, at redshift $z = 99$, and we used these snapshots as inputs to perform two different sets of cosmological simulations, described in the following Subsections. To do that, we used the numerical code `GADGET-3`, a modified version of the publicly available N-body hydrodynamic cosmological code `GADGET-2` (Springel et al. (2001b); Springel (2005)), devised to to reconstruct the density field from the particle distribution through advanced and refined Smoothed Particle Hydrodynamics (SPH) routines.

The general SPH approach relies on the concept that the density field ρ underlying a discrete set of particles can be approximated at particle i position with the weighted sum of the mass m of neighbouring particles NN(i)

$$\rho_i = \sum_{j \in \mathrm{NN}(i)} m_j W_{ij}, \tag{2.9}$$

where the mass is convoluted with a kernel function W_{ij} of choice, characterised by a particle-specific smoothing length h_i, and whose extent is fixed imposing

$$\frac{4}{3}\pi h_i^3 \rho_i = \sum_{j \in \mathrm{NN}(i)} m_j \tag{2.10}$$

so that only a given mass is enclosed within it. Once the density field is reconstructed, every observable is locally computed through weighted sums as

$$O_i = \sum_{j \in \mathrm{NN}(i)} m_j \frac{O_j}{\rho_j} W_{ij} \tag{2.11}$$

and its derivatives are iteratively obtained with

$$\nabla O_i = \sum_{j \in \mathrm{NN}(i)} m_j \frac{O_j}{\rho_j} \nabla W_{ij} \tag{2.12}$$

where the derivative is applied on the window function.

The exact scheme of the SPH algorithm is not fixed, since each observable can be expressed in many analytically equivalent forms that, however, translate into different operative summations. An important consequence of such flexibility is that different but analytically equivalent expressions will map into operative sums that carry different numerical errors. This issue is particularly relevant when one wants to implement non-standard DM features in the non-linear regime explored by cosmological simulations, as we will do in Chapter 3.

2.4.1 Dark Matter–only simulations

The first set of simulations that we used for our analyses is constituted by 55+1 pure N-body (DM-only) simulations (512^3 particles; 20 comoving Mpc/h box length), corresponding to the first 55 $\{\alpha, \beta, \gamma\}$-combinations listed in Table 2.1, plus the standard CDM case, with cosmological parameters $\Omega_m = 0.301$, $\Omega_b = 0.0457$, $n_s = 0.961$, $H_0 = 70.2$ km s^{-1} Mpc^{-1} and $\sigma_8 = 0.829$ (Ade et al. (2016a)), evolved down to redshift $z = 0$. On top of these simulations, we have used a Friends-of-Friends (FoF) algorithm (Davis et al. (1985)) with the standard linking length $b = 0.2$, in order to extract the DM halos. We also run the `SUBFIND` code (Springel et al. (2001a)) for searching for the substructures bound to each main FoF group. Since the constraints presented in Section 2.5 are primarily based on the linear theory, this set of simulations is mainly used as a first cross-check that the results are not significantly altered by non-linearities, with the following purposes:

- present a first assessment of non-linearities for the nCDM models discussed;
- address quantitatively how the DM (sub)halo mass functions based on linear theory predictions compare to the actual results of the N-body simulations (Section 2.5.1 and Appendix B);
- assess whether non-linearities in the matter power spectrum could affect the constraints obtained through the so-called "area criterion" for the Lyman-α forest (Section 2.5.2).

Let us now briefly illustrate the non-linear statistics that we extracted from this suite of simulations, i.e. the matter power spectra and the DM halo mass functions. In Figure 2.8 we report the ratios of the nCDM non-linear matter power spectra with respect to the standard CDM power spectrum. We show the power spectra

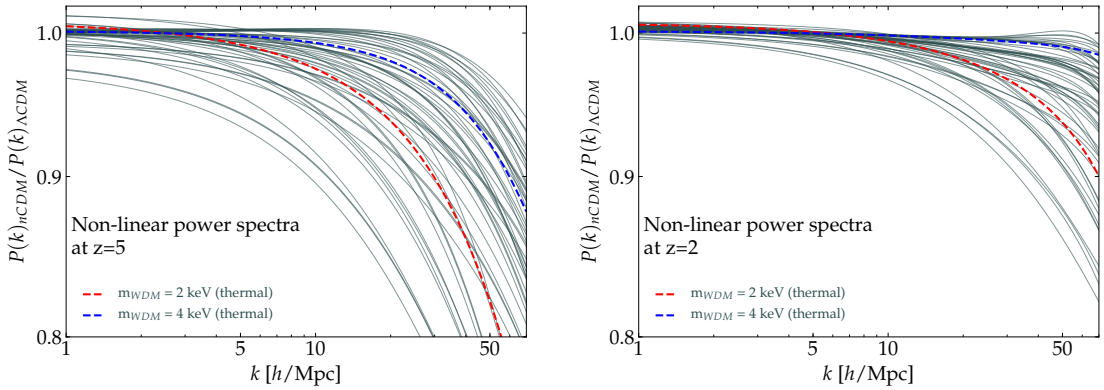

Figure 2.8: Ratios of the nCDM non-linear matter power spectra with respect to the standard CDM power spectrum, at redshifts $z = 5$ and $z = 2$. The black lines correspond to the first 55 models listed in Table 2.1. The red and blue dashed lines refer to thermal WDM models with masses of 2 and 4 keV, respectively.

of all nCDM models at redshift $z = 5$ (left panel) and redshift $z = 2$ (right panel). Additionally, we show the thermal WDM cases with 2 and 4 keV, which are in agreement with a similar study made by Viel et al. (2012). The differences between the models gradually decrease when going to smaller redshifts. Below $z = 2$, the small-scale power enhancement from the non-linear structure evolution starts to dominate the signal from different nCDM models, resulting in a progressive shift of the corresponding half-mode scales towards larger values of k. This is the reason why the Lyman-α data from the highest redshift bins typically provide the strongest limits on the nature of DM, as we show in Section 2.6.

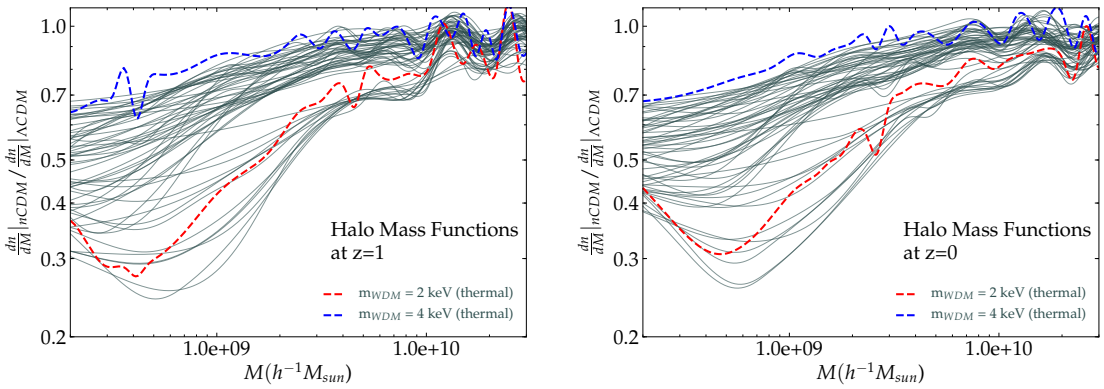

Figure 2.9: Ratios of the nCDM halo mass functions with respect to the standard CDM halo mass function, at redshifts $z = 1$ and $z = 0$. The black lines correspond to the first 55 models listed in Table 2.1. The upturn at low masses is due to artificial clumping, while the oscillation pattern at large masses is due to the cosmic variance. The red and blue dashed lines refer to thermal WDM models with masses of 2 and 4 keV, respectively.

In Figure 2.9 we report the ratios between the nCDM halo mass functions with

respect to the standard CDM halo mass function, at redshift $z = 1$ (*left*) and $z = 0$ (*right*). Note that all the nCDM models predict a lower abundance of DM halos with masses between 10^8 and 10^9 $M_\odot/h$: for some of the $\{\alpha, \beta, \gamma\}$-triplets the number of halos included in this mass range is expected to be 80% smaller with respect to standard CDM predictions.

The visible upturn at low masses is not physical, but a consequence of the phenomenon of *artificial clumping*, which affects models with suppressed initial power spectra (see also Chapter 3 – Section 3.5.1 –), due the limited resolution of the simulations (Angulo et al. 2013; Schneider et al. 2013; Schneider 2015). This is why it is crucial to have good theoretical predictions for the low-mass behaviour of the (sub)halo mass functions. This issue, essential for the analyses presented in Section 2.5.1, is discussed in more detail in Appendix B. The oscillatory pattern which characterises the region corresponding to masses $\gtrsim 10^{10}$ $M_\odot/h$ is due to the cosmic variance, since the size of our simulated box does not allow to have enough large DM halos to have statistically fully meaningful results.

2.4.2 Hydrodynamic simulations

As it has been explained in Chapter 1 (Section 1.4), the physical observable for Lyman-α forest experiments is the flux power spectrum $P_{\rm F}(k, z)$. Therefore, the second set of simulations that we performed is a large suite of full hydrodynamic simulations, where the gas particles are treated with the *Quick-Lyman-Alpha* method, specifically developed for Lyman-α forest analyses (Viel et al. (2004)); it consists in converting into collisionless stars all gas particles which reach densities larger than 1000 times the mean density. Since the absorption systems producing the Lyman-α forest have small overdensity, this criterion has little effect on flux statistics, while significantly speeding up the computations. The goal of the simulations is to provide a reliable template of mock flux power spectra to be compared with observations.

As in the analyses by Iršič et al. (2017b), our reference model simulation has a box length of $20/h$ comoving Mpc with $2 \cdot 768^3$ gas and CDM particles (with gravitational softening $1.04/h$ comoving kpc) in a flat ΛCDM universe with cosmological parameters $\Omega_m = 0.301$, $\Omega_b = 0.0457$, $n_s = 0.961$, $H_0 = 70.2$ km s^{-1} Mpc^{-1} and $\sigma_8 = 0.829$ (Ade et al. (2016a)).

Given that the flux power spectrum is affected both by astrophysical and cosmological parameters, it is important to properly take them into account and accurately quantify their impact. To this end, along the lines of Iršič et al. (2017b), we explored several values of the cosmological parameters σ_8, i.e. the normalisation of the linear matter power spectrum, and $n_{\rm eff}$, namely the slope of the matter power spectrum at the scale of the Lyman-α forest (0.009 s/km). Specifically, we considered five different values for both σ_8 (in the interval $[0.754, 0.904]$) and $n_{\rm eff}$ (in the range

$[-2.3474, -2.2674]$). There have been several studies in the past (e.g., Seljak et al. (2006); McDonald et al. (2006); Arinyo-i Prats et al. (2015)), that have shown that the Lyman-α forest is really measuring the amplitude of the linear matter power spectrum, the slope of the power spectrum, and possibly the effective running, all evaluated at a pivot scale of around 1-10 Mpc/h. Thus σ_8 and $n_{\rm eff}$ used are good tracers of what is actually measured. Given that all our modelling in simulations kept $\Omega_m h^2$ fixed, σ_8 can be directly translated into the amplitude of linear matter power at the pivot scale (similarly to how $n_{\rm eff}$ is used). As pointed by Seljak et al. (2006), these matter power amplitude parameters are equivalent. The linear matter power only weakly depends on $\Omega_m h^2$, and moreover, the effects of Ω_m and H_0 on the linear matter power are already captured in the tracers of the amplitude σ_8 and slope $n_{\rm eff}$. Therefore the constraints are not sensitive to the value of Ω_m nor H_0.

Concerning the astrophysical parameters, we varied the thermal history of the IGM in the form of the amplitude T_0 and the slope $\widetilde{\gamma}$ of its temperature-density relation, generally parametrised as $T = T_0(1+\delta_{\rm IGM})^{\widetilde{\gamma}-1}$, with $\delta_{\rm IGM}$ being the IGM overdensity (we refer to Hui & Gnedin (1997) for the physical motivation of why the IGM is expected to follow the relation above). Specifically, we considered a set of three different temperatures at mean density, $T_0(z=4.2) = 6000, 9200, 12600$ K, which evolve with redshift, as well as a set of three values for the slope of the temperature-density relation, $\widetilde{\gamma}(z=4.2) = 0.88, 1.24, 1.47$. The reference thermal history has been chosen to be defined by $T_0(z=4.2) = 9200$ and $\widetilde{\gamma}(z=4.2) = 1.47$, and it provides a good fit to observations, as demonstrated by Bolton et al. (2017) where several hydrodynamic simulations with the same reference thermal history as the one used here have been carried on.

We have also varied the redshift of the instantaneous reionisation model, for which we considered the three different values $z_{\rm reio} = 7, 9, 15$, with $z_{\rm reio} = 9$ being the reference value. Furthermore, we considered UV fluctuations of the ionising background, that may be particularly important at high redshift. The amplitude of this effect is described by the parameter $f_{\rm UV}$: the corresponding template is built from a set of 3 models with $f_{\rm UV} = 0, 0.5, 1$, where $f_{\rm UV} = 0$ corresponds to a spatially uniform UV background (Iršič et al. (2017b)).

We do not consider IGM temperature fluctuations which have been advocated as potentially mimicking the presence of a cut-off at small scales (Hui et al. (2017)). However, according to sophisticated hydrodynamic simulations such effects appear to happen at large rather than small scales (Cen et al. (2009); D'Aloisio et al. (2015)). A comprehensive treatment of spatial UV and temperature fluctuations would require computationally prohibitive radiative transfer calculations in large volumes, and it is beyond the analysis performed here.

Finally, we varied the mean flux $\langle F(z)\rangle$ by selecting 9 different values for it, namely $(0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4) \cdot \langle F_{\rm REF}\rangle$, with reference values from

SDSS-III/BOSS measurements (Palanque-Delabrouille et al. (2013)). Aiming to have a very fine grid in terms of mean fluxes, we also included 8 additional values, obtained by re-scaling the optical depth $\tau = -\ln\langle F\rangle$, i.e. $(0.6, 0.7, 0.8, 0.9, 1.1, 1.2, 1.3, 1.4) \cdot \tau_{\rm REF}$.

Concerning the nCDM sector, we used the initial conditions corresponding to all the 109 $\{\alpha, \beta, \gamma\}$-combinations listed in Table 2.1, as inputs for running 109 full hydrodynamic simulations (512^3 particles; 20 comoving Mpc/h box length), keeping the astrophysical and cosmological parameters fixed to their reference values.

Before investigating the fully general $\{\alpha, \beta, \gamma\}$-space, we wanted to be able to reproduce the same results obtained by Iršič et al. (2017b), when the analysis is limited to the thermal WDM case. In order to do that, we extended our grid of 109 nCDM points with 8 additional hydrodynamic simulations, in which the values for α correspond to thermal WDM masses of 2,3,4,5,6,7,8,9 keV, β and γ are fixed to their thermal values, i.e. $\beta = 2.24$ and $\gamma = -4.46$, and all the other cosmological and astrophysical parameters are fixed to their reference values. The full nCDM grid, including both thermal and non-thermal simulations, consists thereby in 117 points sampling the $\{\alpha, \beta, \gamma\}$-space.

2.5 Constraints from the linear theory

In this Section we present the first limits on $\{\alpha, \beta, \gamma\}$ from structure formation data, based on linear perturbation theory. We constrained the parameters by using two powerful independent methods for testing the "non-coldness" of DM: MW satellite counts (Subsection 2.5.1), and the so-called "area criterion" for the Lyman-α forest data (Subsection 2.5.2).

Satellite counts rely on the simple fact that any nCDM model must predict a number of substructures within the MW virial radius not smaller than the actual number of MW satellites that we observe. The "area criterion" is instead a simple and effective method, based on linear theory, for testing different nCDM scenarios with Lyman-α forest data, by using an approximate yet intuitive estimator. With such method it is not possible to extract absolute limits on the nCDM parameters, but it allows to look into deviations with respect to a given reference case, which is typically chosen to be the most updated bound from full statistical analyses (Murgia et al. (2017); Kobayashi et al. (2017); Murgia (2018); Murgia et al. (2018)).

Finally, in Subsection 2.5.3, we compare the predictions in terms of MW satellite counts and "area criterion" estimator, by using both the fitted and the "exact" transfer functions of some particle physics models defined in Section 2.3, in order to put our fitting formula, Equation (2.4), to the "reality-check".

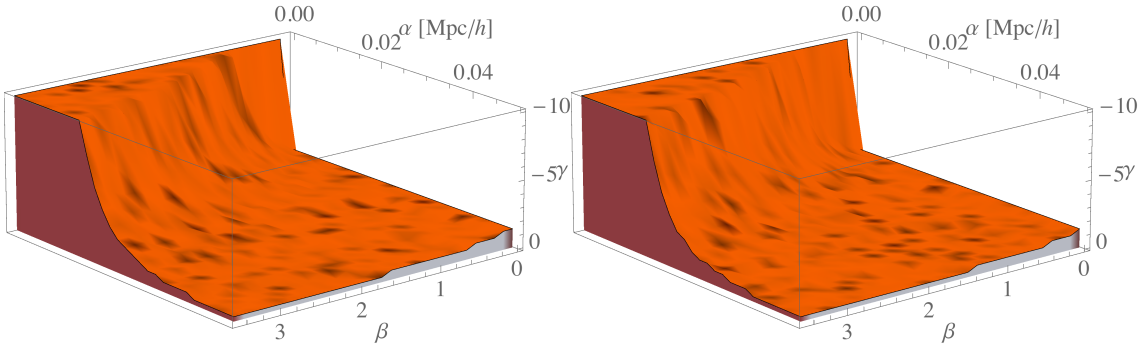

Figure 2.10: The 3D orange contour in the $\{\alpha, \beta, \gamma\}$-space represents the upper bound (i.e., with the largest modulus for γ) on the region of the parameter space which is in agreement with MW satellite counts. The left panel refers to the conservative analysis ($N_{\rm sub} \geq 57$), whereas the right panel refer to the non-conservative case ($N_{\rm sub} \geq 63$).

2.5.1 Milky Way satellite counts

Assuming the standard CDM model, cosmological N-body simulations predict too many dwarf galaxies within the MW virial radius, with respect to observations. Therefore, MW satellite counts represent a very useful tool for constraining DM properties (see, e.g., Jethwa et al. (2018); Vegetti & Koopmans (2009); Vegetti et al. (2012, 2014)). We can estimate the number of MW satellites following the approach of Polisensky & Ricotti (2011); Schneider (2016), i.e., by multiplying the 15 ultra-faint satellites from SDSS by a factor 3.5, in order to account for the limited sky coverage of the survey, and by finally summing the 11 MW classical satellites. We thus obtain $N_{\rm sat} = 63$ as an estimate of the number of observed satellites within the MW viral radius.

	$N_{\rm sub}$		$N_{\rm sub}$		$N_{\rm sub}$		$N_{\rm sub}$		$N_{\rm sub}$
nCDM1	39	**nCDM12**	**105**	**nCDM23**	**75**	nCDM34	42	**nCDM45**	**99**
nCDM2	**78**	nCDM13	27	**nCDM24**	**105**	nCDM35	57	nCDM46	26
nCDM3	**105**	**nCDM14**	**69**	nCDM25	28	nCDM36	52	nCDM47	45
nCDM4	25	**nCDM15**	**104**	**nCDM26**	**68**	**nCDM37**	**71**	**nCDM48**	**63**
nCDM5	**68**	nCDM16	19	**nCDM27**	**104**	**nCDM38**	**81**	**nCDM49**	**76**
nCDM6	**104**	nCDM17	60	**nCDM28**	**76**	**nCDM39**	**90**	**nCDM50**	**105**
nCDM7	18	**nCDM18**	**103**	**nCDM29**	**89**	**nCDM40**	**106**	**nCDM51**	**106**
nCDM8	59	nCDM19	52	**nCDM30**	**69**	**nCDM41**	**106**	**nCDM52**	**71**
nCDM9	**103**	**nCDM20**	**83**	**nCDM31**	**85**	**nCDM42**	**86**	**nCDM53**	**82**
nCDM10	41	**nCDM21**	**106**	nCDM32	49	**nCDM43**	**93**	**nCDM54**	**86**
nCDM11	**78**	nCDM22	38	**nCDM33**	**68**	**nCDM44**	**95**	**nCDM55**	**93**

Table 2.2: Number of subhalos (with mass $M_{\rm sub} \geq 10^8\ M_\odot/h$) within the virial radius of a halo with mass $M_{\rm halo} = 1.7 \cdot 10^{12}\ M_\odot/h$. Each of the 55 models corresponds to a different $\{\alpha, \beta, \gamma\}$-combination, according to Table 2.1. Models that predict a number of subhalos consistent with observations (i.e., at least as many subhalos as the number of observed MW satellites, $N_{\rm sat} = 63$) are written in bold-face.

We can now compare $N_{\rm sat}$ with the number of subhalos $N_{\rm sub}$ predicted by our models, in order to extract some constraints on $\{\alpha, \beta, \gamma\}$. According to Schneider (2015, 2016), we can use the following formula for estimating the number of substructures for a given model:

$$\frac{{\rm d}N}{{\rm d}M_{\rm sub}} = \frac{1}{44.5}\frac{1}{6\pi^2}\frac{M_{\rm halo}}{M_{\rm sub}^2}\frac{P(1/R_{\rm sub})}{R_{\rm sub}^3\sqrt{2\pi(S_{\rm sub}-S_{\rm halo})}}, \tag{2.13}$$

where $M_{\rm sub}$ and $S_{\rm sub}$ are the mass and the variance of a given subhalo, $M_{\rm halo}$ and $S_{\rm halo}$ are the mass and the variance of the main halo, defined as follows:

$$S_i = \frac{1}{2\pi^2}\int_0^{1/R_i} {\rm d}k\ k^2 P(k); \qquad M_i = \frac{4\pi}{3}\Omega_m\rho_c(cR_i)^3; \qquad c = 2.5; \tag{2.14}$$

with $P(k)$ being the linear power spectrum of a given model, computed at redshift $z = 0$. In Appendix B we demonstrate that the mass functions extracted from the grid of N-body simulations described in Section 2.4.1 are in good agreement with the theoretical mass function formalism outlined above.

Under the assumption of a MW halo mass $M_{\rm halo} = 1.7 \cdot 10^{12}\ M_\odot/h$ (Lovell et al. (2014)), and by considering subhalos with masses $M_{\rm sub} \geq 10^8\ M_\odot/h$, we can obtain the number of subhalos $N_{\rm sub}$ predicted by our models, by simply integrating Equation (2.13). The results are reported in Table 2.2, where the models highlighted in bold-face are those in agreement with the number of observed satellites, i.e. with $N_{\rm sub} \geq 63$.

In Figure 2.10 we show a 3D contour plot in the $\{\alpha, \beta, \gamma\}$-space, where each triplet is associated with a different model: the orange contour represents the upper bound on the region of the parameter space which is in agreement with MW satellite counts, according to the method that we have just described. Hence, in the right panel of Figure 2.10, all the $\{\alpha, \beta, \gamma\}$-combinations which sample the orange volume correspond to models that predict a number of substructures at least equal to $N_{\rm sat} = 63$. In the left panel, instead, we show the allowed volume of the parameter space whether we require the number of subhalos predicted by each nCDM model to be equal or larger with respect to a more conservative estimate for the number of MW satellites, $N_{\rm sat} = 57$. This number has been chosen in order to account for a 10% sampling variance in the number of satellites at a given MW halo mass.

By marginalising over β and γ we obtain the following limits on α:

$$\begin{aligned} &\alpha \leq 0.067\ {\rm Mpc}/h\ \ (2\ \sigma\ {\rm C.L.}) && \text{requiring } N_{\rm sub} \geq 63, \\ &\alpha \leq 0.061\ {\rm Mpc}/h\ \ (2\ \sigma\ {\rm C.L.}) && \text{requiring } N_{\rm sub} \geq 57, \end{aligned} \tag{2.15}$$

which would correspond, in the old one-to-one parametrisation, to a thermal WDM particle with mass $m_{\rm WDM} \approx 2$ keV (see Equation (2.2) and Equation (2.3)). These

limits are less constraining than the latest constraints from structure formation data: as expected, modelling the power suppression with three free parameters weakens the constraints on $k_{1/2}$. Within our general approach, due to the dependence of α on β and γ, even lighter DM candidates may provide suppressed power spectra in agreement with MW satellite counts.

Looking at Equation (2.15), it may seem surprising that the constraint on α strengthens when we use a weaker rejection criterion for the nCDM models, i.e. a smaller value for $N_{\rm sat}$. The reason is that, by relaxing the limit on the number of substructures, we accept a larger number of $\{\alpha, \beta, \gamma\}$-triplets characterised by very small values of α. This is a straightforward consequence of the shape of the volume of the parameter space shown in Figure 2.10. By accepting all those models which predict $57 \leq N_{\rm sub} < 63$, we are slightly shifting towards higher values of α the whole orange 3D surface in the plot. As one can see by comparing the two panels of Figure 2.10, due to the geometry of the orange contour, this shift mainly concerns $\{\alpha, \beta, \gamma\}$-combinations with α very close to zero. Therefore, by imposing a smaller value for $N_{\rm sat}$ and marginalising over β and γ, small values of α increase their contribution to its probability distribution with respect to the high-value tail of the distribution, which instead is only minimally affected by the choice of a more conservative value for $N_{\rm sat}$. We are therefore pushing the peak of the probability distribution of α towards zero, yielding a stronger upper bound on it. However, at this approximate level of analysis, the difference between the two bounds reported in Equation (2.15) is negligible, as visible from Figure 2.10.

2.5.2 The "area criterion" for the Lyman-α forest

	δA		δA		δA		δA		δA
nCDM1	0.61	**nCDM12**	**0.34**	nCDM23	0.44	nCDM34	0.39	**nCDM45**	**0.37**
nCDM2	0.45	nCDM13	0.63	**nCDM24**	**0.32**	**nCDM35**	**0.33**	nCDM46	0.61
nCDM3	**0.34**	nCDM14	0.45	nCDM25	0.62	nCDM36	0.57	nCDM47	0.57
nCDM4	0.63	**nCDM15**	**0.32**	nCDM26	0.44	nCDM37	0.54	nCDM48	0.41
nCDM5	0.45	nCDM16	0.64	**nCDM27**	**0.31**	nCDM38	0.42	nCDM49	0.40
nCDM6	**0.32**	nCDM17	0.45	nCDM28	0.42	nCDM39	0.42	**nCDM50**	**0.27**
nCDM7	0.64	**nCDM18**	**0.31**	**nCDM29**	**0.37**	**nCDM40**	**0.31**	**nCDM51**	**0.27**
nCDM8	0.45	nCDM19	0.59	nCDM30	0.42	**nCDM41**	**0.33**	**nCDM52**	**0.38**
nCDM9	**0.31**	nCDM20	0.44	**nCDM31**	**0.37**	nCDM42	0.40	**nCDM53**	**0.37**
nCDM10	0.61	**nCDM21**	**0.34**	nCDM32	0.40	nCDM43	0.40	**nCDM54**	**0.33**
nCDM11	0.45	nCDM22	0.61	**nCDM33**	**0.34**	**nCDM44**	**0.36**	**nCDM55**	**0.33**

Table 2.3: Here we list the 55 models that we have tested, each of them with its corresponding δA, namely the estimator of the small-scale power suppression associated to it. A model is excluded (at 2 σ C.L.) if $\delta A > \delta A_{\rm REF,1}$, i.e., if it shows a power suppression $\gtrsim 38\%$ with respect to the standard CDM power spectrum. Each of the models corresponds to a different $\{\alpha, \beta, \gamma\}$-combination, according to Table 2.1. Accepted models are highlighted in bold-face.

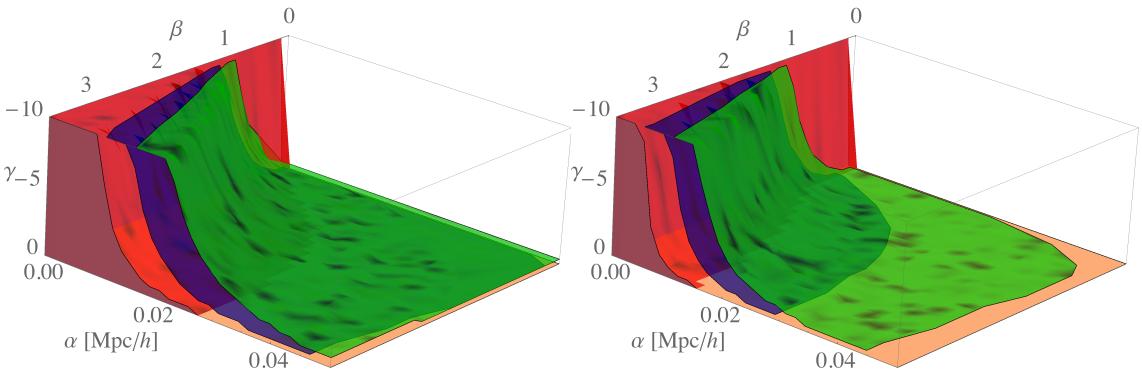

Figure 2.11: The 3D contour plot in the $\{\alpha, \beta, \gamma\}$-space represents the region of the parameter space which contains models in agreement with the Lyman-α forest data, according to the "area criterion". The left panel refers to the conservative analysis ($\delta A < 0.38$), whereas the right panel refer to the non-conservative case ($\delta A < 0.21$). The red contours represent the 1 σ C.L. limits on the $\{\alpha, \beta, \gamma\}$-combinations, while the blue and green contours represent the 2 σ and 3 σ C.L. limits, respectively. All those models associated to $\{\alpha, \beta, \gamma\}$-triplets placed outside of the 3D coloured region are therefore excluded at 3 σ C.L. by the "area criterion".

The Lyman-α forest, i.e., the Lyman-α absorption produced by intervening intergalactic neutral hydrogen in the spectra of distant quasars, provides a powerful tool for constraining the DM properties at small scales. While we present absolute and accurate limits from a comprehensive Monte Carlo Markov Chain (MCMC) analysis of Lyman-α data in Section 2.6, the goal of this Section is to introduce a simpler yet approximate method, dubbed as "area criterion". Instead of computing absolute bounds, the goal of this Section is indeed to investigate deviations of our nCDM models with respect to a thermal WDM reference model, i.e., $m_{\rm WDM} = 3.5$ keV, which is one of the most updated constraints on WDM candidates (at 2 σ C.L.), obtained through extensive Lyman-α analyses by Iršič et al. (2017b).

Firstly, we slightly modify the method developed by Schneider (2016), by parametrising the deviation of a model with respect to the standard CDM scenario through the ratio

$$\zeta(k) = \frac{P_{\rm 1D}(k)}{P_{\rm 1D}^{\Lambda{\rm CDM}}(k)}, \tag{2.16}$$

where $P_{\rm 1D}(k)$ is the 1D power spectrum of the model that we are considering, computed by the following integral on the 3D matter power spectrum:

$$P_{\rm 1D}(k) = \frac{1}{2\pi} \int_k^\infty {\rm d}k' k' P(k'), \tag{2.17}$$

where $P(k')$ is the 3D linear matter power spectrum, computed at redshift $z = 0$.

We are now able to determine whether a model deviates more or less from plain CDM, with respect to the thermal WDM reference model that we have chosen, by adopting the following criterion: a model is excluded (at 2 σ C.L.) if it is characterised by a larger power suppression with respect to the reference model. In order to quantify the suppression in the power spectra, we define the following estimator:

$$\delta A \equiv \frac{A_{\Lambda\mathrm{CDM}} - A}{A_{\Lambda\mathrm{CDM}}}, \tag{2.18}$$

where A is the integral of $\zeta(k)$ over the typical range of scales probed by Lyman-α observations ($0.5\ h/\mathrm{Mpc} \lesssim k \lesssim 20\ h/\mathrm{Mpc}$ for the MIKE/HIRES+XQ-100 combined data set used by Iršič et al. (2017b)), i.e.,

$$A = \int_{k_{\mathrm{min}}}^{k_{\mathrm{max}}} \mathrm{d}k\ \zeta(k), \tag{2.19}$$

such that $A_{\Lambda\mathrm{CDM}} \equiv k_{\mathrm{max}} - k_{\mathrm{min}}$ by construction.

Analogously, by plugging the power spectrum of the thermal WDM reference model into Equations (2.16) and (2.18), we find $\delta A_{\mathrm{REF},1} = 0.38$, which is an estimate of the small-scale power suppression with respect to standard CDM for models that are excluded at 2 σ C.L. by Lyman-α forest data. In Table 2.3 we list the 55 models that we tested, each of them with its corresponding δA: a certain model is excluded (at 2 σ C.L.) if $\delta A > \delta A_{\mathrm{REF},1}$, i.e., if it shows a power suppression $\gtrsim 38\%$ with respect to the standard CDM power spectrum. Accepted models are highlighted in bold-face.

Let us now stress that the constraint on the thermal WDM mass associated to $\delta A_{\mathrm{REF},1}$ (i.e. $m_{\mathrm{WDM}} = 3.5$ keV) has been obtained under very conservative assumptions on the thermal history of the universe (see Iršič et al. (2017b); Murgia et al. (2018) for details). By modifying these assumptions, the lower limit on thermal WDM masses strengthens to $m_{\mathrm{WDM}} = 5.3$ keV (at 2 σ C.L.), which represents the tightest bound from Lyman-α forest data up to date. By taking this limit as reference, we find indeed a corresponding small-scale suppression $\delta A_{\mathrm{REF},2} = 0.21$, with respect to the standard CDM power spectrum.

In Figure 2.11 we show a 3D contour plot in the $\{\alpha, \beta, \gamma\}$-space, which represents the region of the parameter space that contains models in agreement with Lyman-α forest data, according to the "area criterion". The left panel refers to the conservative analysis, whereas the right panel refers to the less conservative case. The red contours represent the 1 σ C.L. limit on the $\{\alpha, \beta, \gamma\}$-combinations, while the blue and green contours represent the 2 σ and 3 σ C.L. limits, respectively. All those models associated to $\{\alpha, \beta, \gamma\}$-triplets placed outside of the 3D coloured region are therefore excluded at 3 σ C.L. by our analysis. By marginalising over β and γ we obtain the following limits on α:

$$\begin{array}{ll} \alpha \leq 0.058\ \mathrm{Mpc}/h\ \ (2\ \sigma\ \mathrm{C.L.}), & \text{conservative analysis,} \\ \alpha \leq 0.044\ \mathrm{Mpc}/h\ \ (2\ \sigma\ \mathrm{C.L.}), & \text{non-conservative analysis.} \end{array} \tag{2.20}$$

These limits would correspond, in the old one-to-one parametrisation, to a thermal WDM particle with a mass of $m_{\rm WDM} \approx 3$ keV (see Equations (2.2) and (2.3)). Even at this approximate level of analysis, the Lyman-α forest tends to provide more stringent constraints than MW satellite counts. Note that, however, this is not true for all models: e.g., the point nCDM35 is allowed by the Lyman-α forest, while rejected by satellite counts. This is because, after all, the two methods probe slightly different scales, and thus are in reality complementary when it comes to constraining DM models. Note that this is a qualitative difference of non-thermal settings compared to thermal WDM: indeed, using different methods to constrain DM models is paramount to obtain a clear picture of what is allowed and what is not. Still, the limit shown is weaker compared to the most updated constraints on thermal WDM masses. As we discussed before, this is primarily due to the new general parametrisation of $T(k)$. With our approach, thanks to the mutual dependence among α, β, and γ, it is possible to model nCDM scenarios with non-trivially suppressed power spectra. Therefore, models with shallower transfer functions may be found to be in agreement with Lyman-α forest data even if the corresponding DM candidate mass lies below the current constraints for thermal WDM masses, given that those constraints refer to a very specific shape of the power suppression (i.e., a very specific $\{\beta, \gamma\}$-combination).

As a double-check, we applied our method to the non-linear power spectra extracted from the N-body simulations described in Section 2.4.1, finding consistent results with respect to the linear analysis. All the models rejected when comparing their linear power spectra are also rejected when comparing the non-linear ones.

Finally, let us recall that the actual physical observable in Lyman-α forest experiments is the flux power spectrum $P_{\rm F}(k, z)$, and not the 1D or 3D linear matter power. However, two different key aspects of the Lyman-α forest physics suggest that the analysis that we just presented could be also quantitatively correct. Firstly, the relation between linear matter and flux power can be modelled as $P_{\rm F} = b^2(k)P(k)$ (see Chapter 1), with a bias factor $b(k)$ which differs only very little between CDM and nCDM models, at least for models reasonably close to the standard case (see, e.g., Croft et al. (2002); Viel et al. (2005)): this justifies the application of Equation (2.16) to flux power spectra as well. Secondly, the "area criterion" is motivated by the fact that IGM peculiar velocities (typically < 100 km/s) tend to redistribute the small-scale power within a relatively wide range of wavenumbers in the probed region of the flux power (Gnedin & Hamilton (2002)).

It is thus important to check the bounds derived in this Section against full hydrodynamic simulations which can provide a forward modelling of the flux power spectrum. This has been done for the first time by Kobayashi et al. (2017), in the context of ultra-light scalar DM, finding a relatively good agreement between the predictions of the two methods, as we will see in Chapter 3. Furthermore, in

Section 2.6.4 we perform an even more explicit comparison, by testing all the nCDM models listed in Table 2.1 with a full statistical analysis, probing the effectiveness of the "area criterion" as a tool to perform preliminary tests on nCDM scenarios in an immediate and simplified way.

2.5.3 Constraints on particle physics models

The goal of this Section is to compare the predictions in terms of MW satellite counts and power in the range probed by the Lyman-α forest (estimated through the "area criterion") for the fitted and the "exact" transfer functions of some particle physics models defined in Section 2.3. Hence, we want to put our fitting formula, Equation (2.4), to the "reality-check".

	α	β	γ	$k_{1/2}$ [h/Mpc]	$N_{\rm sub}^{\rm fit}$ ($N_{\rm sub}^{\rm true}$) [%]	Agree?	$\delta A_{\rm fit}$ ($\delta A_{\rm true}$) [%]	Agree?
RP neutrinos	0.025	2.3	−2.6	17.276	38 (39) [−2.6%]	✓	0.555 (0.571) [−2.8%]	✓
	0.071	2.3	−1.0	9.828	15 (14) [+7.1%]	✓	0.743 (0.754) [−1.5%]	✓
	0.038	2.3	−4.4	8.604	5 (5) [±0.0%]	✓	0.799 (0.810) [−1.4%]	✓
	0.035	2.1	−1.5	15.073	35 (37) [−5.4%]	✓	0.599 (0.613) [−2.3%]	✓
Neutrinos from particle decay	0.016	2.6	−8.1	19.012	38 (42) [−9.5%]	✓	0.521 (0.535) [−2.6%]	✓
	0.011	2.7	−8.5	28.647	**91** (**97**) [−6.2%]	✓	*0.339* (*0.360*) [−5.8%]	✓
	0.019	2.5	−6.9	16.478	27 (28) [−3.6%]	✓	0.582 (0.576) [+1.0%]	✓
	0.011	2.7	−9.8	26.31	**79** (**87**) [−9.2%]	✓	0.375 (*0.390*) [−3.8%]	✗
Mixed models	0.16	3.2	−0.4	6.743	9 (9) [±0.0%]	✓	0.823 (0.834) [−1.3%]	✓
	0.20	3.7	−0.18	7.931	28 (27) [+3.7%]	✓	0.738 (0.752) [−1.9%]	✓
	0.21	3.7	−0.1	11.36	*60* (*62*) [−3.2%]	✓	0.596 (0.610) [−2.3%]	✓
	0.21	3.4	−0.053	33.251	**110** (**114**) [−3.5%]	✓	*0.365* (*0.377*) [−3.2%]	✓
Fuzzy DM	0.054	5.4	−2.3	13.116	8 (9) [−11.1%]	✓	0.691 (0.708) [−2.4%]	✓
	0.040	5.4	−2.1	18.106	21 (23) [−8.7%]	✓	0.543 (0.565) [−3.9%]	✓
	0.030	5.5	−1.9	25.016	56 (60) [−6.7%]	✓	*0.376* (0.399) [−5.8%]	✗
	0.022	5.6	−1.7	34.590	**121** (**126**) [−4.0%]	✓	*0.228* (*0.250*) [−8.8%]	✓
ETHOS models	0.0072	1.1	−9.9	7.274	18 (19) [−5.3%]	✓	0.780 (0.788) [−1.0%]	✓
	0.013	2.1	−9.3	16.880	36 (39) [−7.7%]	✓	0.568 (0.581) [−2.2%]	✓
	0.014	2.9	−10.0	21.584	50 (53) [−5.7%]	✓	0.463 (0.477) [−2.9%]	✓
	0.016	3.4	−9.3	23.045	53 (56) [−5.4%]	✓	0.430 (0.439) [−2.1%]	✓

Table 2.4: Here we list 20 $\{\alpha, \beta, \gamma\}$-triplets, with the corresponding values of $k_{1/2}$, which represent the real model examples presented in Section 2.3, split into five groups. For each case, we have confronted the $\{\alpha, \beta, \gamma\}$-fit (the real model) with both satellite counting, $N_{\rm sub}^{\rm fit}$ ($N_{\rm sub}^{\rm true}$), and the Lyman-α forest area estimator, $\delta A_{\rm fit}$ ($\delta A_{\rm true}$), where for each case we also indicate the percentage (in purple for better visibility) by which the value predicted from the fit differs from the "true" value predicted by the model point. Bold-faced numbers indicate that both the restrictive and conservative bounds are met, while Italic numbers indicate that only the conservative bound was met. For each case we have indicated whether the conclusion drawn from the fit – i.e., whether or not a certain choice of parameters is allowed by the data – does (✓) or does not (✗) agree with the one drawn from the data for the real model.

Note that we do not aim to give a full account of the validity (or invalidity) of the different nCDM models presented in Section 2.3, since in any case a few example points will not be able to give us a clear answer. Instead, we would like to find out whether the $\{\alpha, \beta, \gamma\}$-fit to a certain nCDM setting would lead us to the same conclusion about its validity when confronted with MW satellite counts and Lyman-α bounds, while we do not care very much at this stage about whether a certain point in the model parameter space is now marked as rejected or allowed – we only want to check whether our conclusion about the points under consideration changes when we look at the fits instead of looking at the actual points.

To do so, we depict in Table 2.4 first of all the fit parameters for the example model points discussed in Section 2.3, which should be rather simple to grasp. For example, the second line for RP corresponds to the green curves in Figure 2.3, while the first line for particle decay corresponds to the blue curves in Figure 2.4. For each point, we computed the number of satellites for both the fit and the corresponding transfer function of the real model, $N_{\rm sub}^{\rm fit}$ ($N_{\rm sub}^{\rm true}$), as well as the Lyman-α area estimator, $\delta A_{\rm fit}$ ($\delta A_{\rm true}$) – with the difference of the fit to the real point indicated by the percentages in square brackets – which are in both cases matched to the respective conservative and non-conservative observational constraints. In all cases, no matter if the resulting number corresponds to a real model or to a fitted point, we use bold-faced scripts/Italic scripts/Roman scripts to indicate that a certain number is in agreement with both the conservative and non-conservative bounds/only with the conservative bound/with none of the bounds. As can be seen from the distribution of bold-faced or Italic numbers, many of the example points shown here are not in agreement with the bounds. However, what we are interested in is whether or not the fitted points would have brought us to the same conclusion. Indeed, this is the case for the vast majority of cases. In fact, given that the predictions from the fits deviate from the real model predictions only by a few percent at most, we would expect agreement of the conclusions drawn from both versions of the transfer function (i.e., fitted and "exact") in all cases up to a few percent.[6]

Thus, except for a tiny amount of borderline cases, the fitted points always yield the same conclusion as the actual model points. We can thereby conclude that our fitting formula reproduces the true results to a very high degree. Hence, whenever is it desired to match a nCDM setting to observational data, there is no need to do the whole computation. Instead, it is perfectly sufficient to match the resulting transfer functions to our Equation (2.4) and to check whether the fitted points are allowed by data.

[6]Two fails in forty comparisons, i.e., an empirical failure rate of 5% seems to support this prediction rather well – even though, of course, we have not selected the examples shown completely arbitrarily, but rather we have picked them to somehow reflect some of the variation possible for the different production mechanisms.

In light of all of that, in the next Section we will go beyond the simple and approximate data analyses just presented, by discussing a more comprehensive approach, not limited to linear theory, which provides accurate absolute bounds on α, β, γ, easily translatable to constraints on the fundamental nCDM properties through the scheme that we have illustrated.

2.6 Accurate limits from the Lyman-α forest

In this Section we discuss the results of the comprehensive MCMC analyses of the high-resolution and high-redshift Lyman-α forest data from the MIKE/HIRES spectrographs, based on the large suite of cosmological hydrodynamic simulations illustrated in Section 2.4.2.

2.6.1 Data set and methods

The data set that we used for our analyses is constituted by the high-resolution and high-redshift HIRES/MIKE samples of quasar spectra. It has been obtained with the HIRES/KECK and the MIKE/Magellan spectrographs, at redshift bins $z = 4.2, 4.6, 5.0, 5.4$ and in 10 k-bins in the interval 0.001-0.08 s/km, with spectral resolution of 6.7 and 13.6 km/s, for HIRES and MIKE, respectively (Viel et al. (2013)). As in the analyses by Viel et al. (2013); Iršič et al. (2017b), we imposed a conservative cut on the flux power spectra obtained from MIKE/HIRES data, and only the measurements with $k > 0.005$ s/km have been used, in order to avoid possible systematic uncertainties on large scales due to continuum fitting. Furthermore, we did not consider the highest redshift bin for MIKE data, for which the error bars on the flux power spectra are very large (see Viel et al. (2013) for details). We have therefore used a total of of 49 (k, z) data points for the MIKE/HIRES data set, which has the advantage with respect to other surveys of exploring small scales and high redshifts, being thereby the most constraining up-to-date, for the models that we have considered.

Note that low-resolution surveys such as SDSS-III/BOSS (Palanque-Delabrouille et al. (2013)) can be used for constraining nCDM scenarios in the quasi-linear regime which characterises larger scales with respect to the ones that we are focusing on. An interesting attempt of modelling the relevant features of the flux power spectrum in order to obtain an approximate estimator for testing such relatively large scales was recently done by Garny et al. (2018). Let us stress that the two approaches are complementary, since the different scale and redshift coverage may lead to different constraints and degeneracies, with the common goal of developing an effective framework which does not require to run specific numerical simulations per each nCDM model.

With the models of the flux power spectra obtained from the large suite of hydrodynamic simulations presented in Section 2.4.2, we set a sparse grid of points in the multidimensional parameter space of $\{\langle F(z)\rangle, T_0(z), \widetilde{\gamma}(z), \sigma_8, z_{\rm reio}, n_{\rm eff}, f_{\rm UV}, \alpha, \beta, \gamma\}$. For the interpolation between different grid points, we adopted an improved method with respect to the linear interpolation scheme used by Iršič et al. (2017b,a), i.e., the *Ordinary Kriging* method, which is widely used in very different fields from cosmology, such as geo-statistics or environmental science, since it is particularly effective for dealing with sparse and non-regular grids (see, e.g., Webster & Oliver (2007)). Such method basically consists in predicting the value of the flux power at a given point by computing a weighted average of all its known values, with weights inversely proportional to the distance from the considered point. The interpolation is done in terms of ratios between the flux power spectra of the nCDM models and the reference one. We first interpolate in the astrophysical and cosmological parameter space for the pure CDM case, i.e., in the $\alpha = 0$ plane. We then correct all the $\{\alpha, \beta, \gamma\}$-grid points accordingly, and we finally interpolate in the $\{\alpha, \beta, \gamma\}$-space. This procedure relies on the assumption that the corrections due to non-reference astrophysical or cosmological parameters are universal, so that we can apply the same corrections computed for the standard CDM case ($\alpha = 0$) to all the nCDM models described by our parametrisation.

We carefully tested the new interpolation scheme, by iteratively predicting the value of the flux power spectrum at a given grid point without using that grid point, as well as by reproducing the bounds on the thermal WDM mass obtained by Iršič et al. (2017b). Let us note that, in doing the latter, rather than using the full nCDM grid of 117 points, we only used the 8 thermal WDM simulations mentioned in Section 2.4.2, and hence applied the interpolation procedure to the same parameter space investigated in the previous analyses (Viel et al. (2005, 2013); Iršič et al. (2017b)). The results of such comparison are discussed in details in Appendix C.

Another difference with respect to previous analyses is that we did not use cross-simulations between the nCDM parameters and the astrophysical and cosmological ones. Notice, however, that the expected degeneracies, e.g., between the IGM temperature and α, emerged consistently with respect to the results published so far (see Appendix C). On the other hand, we noticed that our current interpolation scheme is not fully accurate for reproducing power spectra which are very far from the reference cases. This issue does not affect our final results on the nCDM parameters, having the only consequence of a further weakening of the bounds for those cosmological parameters which were not tightly constrained even in previous IGM studies (see Section 2.6.2). Regarding this aspect, in Chapter 4 we will combine Lyman-α data with other complementary cosmological observables, such as CMB and Baryon Acoustic Oscillations (BAOs), obtaining a remarkable improvement in the constraining power.

In light of the aforementioned caveats, we determined the constraints on both astrophysical, cosmological, and nCDM parameters, by maximising a Gaussian likelihood with a MCMC approach, using the publicly available affine-invariant MCMC sampler `emcee` (Foreman-Mackey et al. (2013)).

Regarding the IGM thermal history, we carried out two different analyses: following the approach of Iršič et al. (2017b), we modelled the redshift evolution of both T_0 and $\widetilde{\gamma}$ as power laws, such that $T_0(z) = T_0^A[(1+z)/(1+z_p)]^{T_0^S}$ and $\widetilde{\gamma}(z) = \widetilde{\gamma}^A[(1+z)/(1+z_p)]^{\widetilde{\gamma}^S}$, where the pivot redshift z_p is the redshift at which most of the Lyman-α forest pixels are coming from (i.e. $z_p = 4.5$ for MIKE/HIRES). We adopted weak priors ($T^A \in [0, 20000]$ K and $T^S \in [-5, 5]$) on the slope and amplitude of those power law relations. We refer to this power law parametrisation as our reference MCMC analysis.

However, in order to be conservative, we repeated the same analysis by letting the amplitude $T_0(z)$ free to vary in each bin, only requiring to forbid differences greater than 5000 K between adjacent redshift bins (Viel et al. (2013)). Furthermore, in order to prevent unreasonably cold values for the IGM temperatures, which would hardly be physically motivated, we adopted broad Gaussian priors centred on $T_0(z)$ reference values, with standard deviation $\sigma = 3000$ K. As it has been thoroughly discussed by Iršič et al. (2017b), different choices of the thermal history priors sensibly affect the results, due to the degeneracy between the IGM temperature evolution and the nCDM parameters (see also Garzilli et al. (2017); Garzilli et al. (2018) for different analyses on the impact of very different thermal histories). When the power law evolution for $T_0(z)$ is not assumed, the constraints on the small-scale power suppression associated to nCDM are expected to be weaker. For these reasons, in Appendix C we compare different prior choices on the IGM thermal history, both in the standard and the new framework, showing and discussing how they do influence the final results.

For the mean fluxes $\langle F(z) \rangle$ we have chosen Gaussian priors with standard deviation $\sigma = 0.04$, approximately corresponding to the normalisation uncertainties given by different observations. For all other free parameters we adopted flat priors within the boundaries given by the grid of simulations.

2.6.2 Results and discussion

Let us now discuss the results of the comprehensive MCMC analyses that we performed for the MIKE/HIRES data set. Firstly, we focused on the thermal WDM case, by switching off the parameters responsible of non-trivial features in the shape of the small-scale power cut-off (i.e. β and γ), and thus by constraining the same parameter space studied by Viel et al. (2013); Iršič et al. (2017b), where α plays the role of the thermal WDM mass parameter. In Appendix C we compare

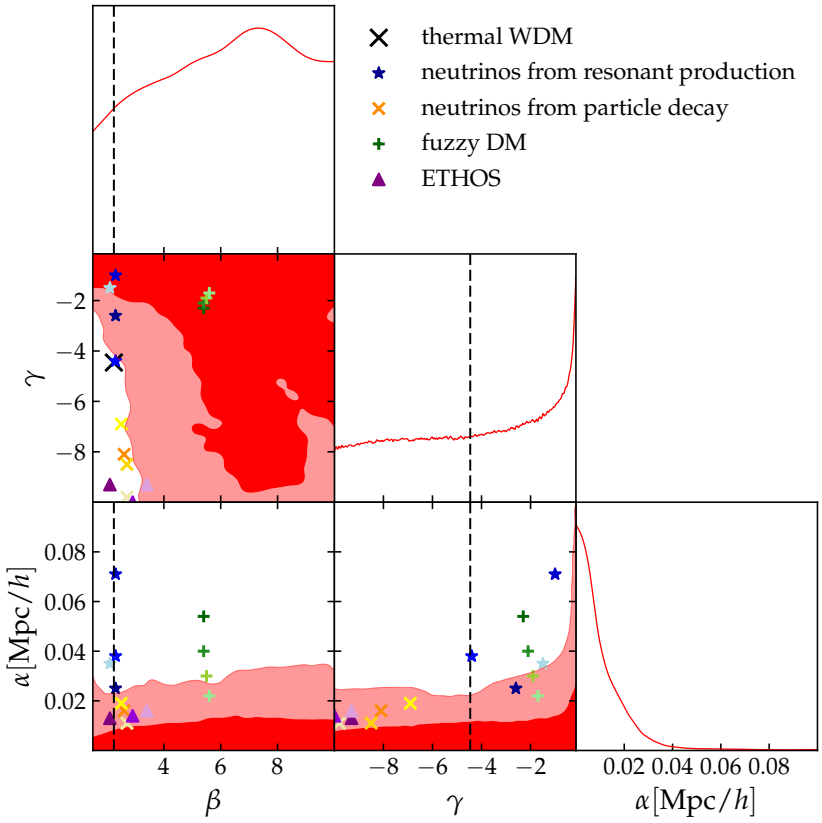

Figure 2.12: 1 σ and 2 σ C.L. contour plots for α, β and γ, obtained by assuming a IGM temperature power law evolution. The dashed vertical lines and the black cross stand for the thermal WDM case, i.e. $\beta = 2.24$ and $\gamma = -4.46$. The other symbols shown in the legend correspond to the $\{\alpha, \beta, \gamma\}$-combinations associated to the nCDM models discussed in Section 2.3, and listed in Table 2.5. Different colour gradients are used for distinguishing between different models belonging to the same class of nCDM scenarios. For each class, the darkest tonality corresponds to the first model listed in Table 2.5, the lightest one corresponds to the last model evaluated.

our results for this specific class of nCDM models with the constraints previously published, probing the full consistency between them.

In this Section, let us then focus on our main goal, i.e., putting limits on the $\{\alpha, \beta, \gamma\}$-space. In Figure 2.12 we condense our main results, namely the 1 σ and 2 σ C.L. exclusion plots showing the bounds on the three parameters describing the power suppression induced by nCDM. We chose to focus on the analysis based on the assumption of a power law evolution for both the amplitude and the slope of the IGM temperature, which is also the case adopted as reference by Iršič et al. (2017b), due to its robust physical motivations. However, as it is manifestly shown in Figure 2.13, the more conservative assumption of a thermal history with freely floating $T_0(z)$ does not change at all our conclusions.

From Figure 2.12 we note that, even in our new general framework, the parameter α, responsible of the position of the cut-off in the power spectrum is well constrained

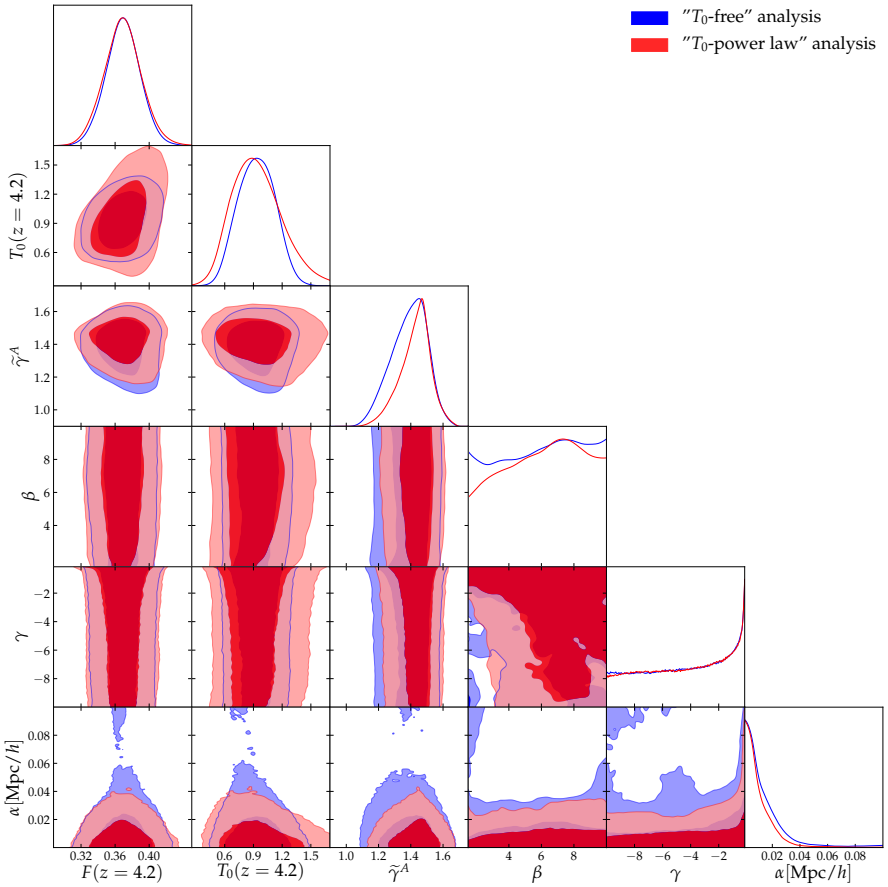

Figure 2.13: 1 σ and 2 σ C.L. contour plots for both α, β, γ, and the main astrophysical free parameters. The blue contours refer to the freely floating IGM temperature analysis, while the red contours refer to our reference analysis, i.e., when a power law evolution is assumed. The values of the temperature are expressed in 10^4 K units.

by current data. On the other hand, both β and γ are quite unconstrained. It is interesting to notice, however, that the 1D posterior distribution of the former shows a peak around $\beta = 7$, which is far from its thermal value, i.e., $\beta = 2.24$. The natural interpretation is that standard thermal WDM models are not favoured by Lyman-α data with respect to non-thermal scenarios. In the plots, the thermal values for β and γ are highlighted by the dashed vertical lines and the black cross. More specifically, the black cross, which corresponds thereby to $\beta = 2.24$ and $\gamma = -4.46$, lies slightly outside of the 2 σ C.L. contour. The peak in the 1D distribution of α corresponds to the standard CDM model, and the mild degeneracy between large values of α and small values of $|\gamma|$ is extensively discussed in Section 2.2 and Appendix A.

In Section 2.3 we described several viable classes of nCDM models motivated by particle physics (i.e., sterile neutrinos both from resonant production and particle decay, FDM models, mixed DM fluids, DM-Dark Radiation interaction models), and we analysed some examples from each of the families. The symbols reported

in Figure 2.12 correspond to the $\{\alpha, \beta, \gamma\}$-combinations which have been shown to provide a good fit for the transfer functions associated to such example models. In order to quantify their viability, we list all of them, with the corresponding α, β, γ and $k_{1/2}$, in Table 2.5, where we report the χ^2-values determined through our reference analysis. Clearly, the χ^2-values were computed only for those models associated to $\{\alpha, \beta, \gamma\}$-combinations sampling a parameter region which is covered by our grid of simulations (see Table 2.1). In Figure 2.12, different colour gradients are used for distinguishing between different models belonging to the same class of nCDM scenarios. For each group of models, the darkest tonality corresponds to the first model listed in Table 2.5, while the lightest one corresponds to the last model evaluated.

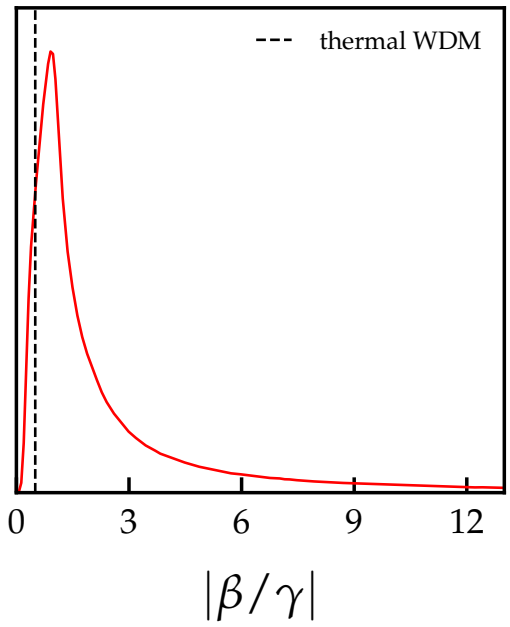

Figure 2.14: Marginalised 1D distribution of $|\beta/\gamma|$, which is a useful estimator for condensing the constraints on the two parameters governing the slope of the power suppression. The vertical dashed line corresponds to the thermal WDM $\{\beta, \gamma\}$-combination, i.e. $\beta = 2.24$ and $\gamma = -4.46$.

Let us shortly recall that the first group of models corresponds to few example values for RP sterile neutrino mass ($m_{\rm S} = 5, 7, 15$ keV), the active-sterile neutrino mixing angle, and the lepton asymmetry. Concerning sterile neutrinos from particle decay, each of the model is characterised by different values of the sterile neutrino and decaying scalar particle masses, along with different Higgs portal and Yukawa coupling parameters. FDM scenarios rely on the assumption that all of DM is constituted by an ultra-light scalar particle: the first two models belonging to this class correspond to DM masses of $5 \cdot 10^{-22}$ and $10 \cdot 10^{-22}$ eV, and they are rejected

by our analysis; the latter two FDM scenarios correspond to masses $20 \cdot 10^{-22}$ and $40 \cdot 10^{-22}$ eV, and they are accepted.

Let us stress that, thanks to the new general parametrisation, it has not been necessary to run any specific numerical simulations in order to test such nCDM scenarios. Whenever one wants to constrain any model belonging to one of these families, it is sufficient to fit the corresponding linear power spectrum in terms of $\{\alpha, \beta, \gamma\}$, and interpolate in the parameter space delineated by our full grid of simulations, i.e., in a refined χ^2-table having a similar structure to Table 2.5, but also including the astrophysical and cosmological parameters involved.

By looking at the positions of the various symbols in Figure 2.12, one cannot give a definitive answer about the viability of the whole classes of candidates. This aspect is particularly relevant for the models suggested by ETHOS, which often feature oscillations at very small scales, that our parametrisation cannot capture. For the few examples considered here, all referring to a weak Dark Acoustic Oscillation (DAO) regime, the presence of such oscillations is totally negligible for the data analyses, as it is explicitly displayed in Section 2.6.3. In Chapter 4 we will carry out a thorough investigation for quantifying to which extent our fitting procedure is able to cover the whole class of ETHOS models, both in the weak and in the strong DAO regime (see also Cyr-Racine et al. (2016); Vogelsberger et al. (2016) for further details). A similar issue concerns FDM models, which are expected to modify the dynamics during the non-linear phase of structure formation, due to Quantum Pressure (QP) effects. Concerning this point, we instead address the reader to Chapter 3, where we will demonstrate that such effects do not affect predictions obtained under the standard approximation of treating ultra-light scalars as standard collisionless DM, at least for models where FDM constitutes the whole of the DM amount.

Taking into consideration all the points above, it is now possible to see which nCDM models are excluded by current data and which ones are not. The FDM model examples that we have considered (green crosses) are associated to values for β and γ which are in perfect agreement with data. However, from Figure 2.12 it is manifest that only the values of α relative to the last two models reported in Table 2.5, i.e., those featuring a power suppression at the smallest scales, are allowed by our analysis. This means that a DM fluid fully composed by ultra-light scalar particles is characterised by a shape for the power suppression which can always accommodate Lyman-α data, provided that the mass of the FDM particle is sufficiently large. Analogous arguments apply to the four RP sterile neutrino models that we have tested (blue stars). Conversely, by looking at the positions of the orange and yellow crosses in Figure 2.12 we conclude that, even though the power suppression due to sterile neutrinos from particle decay occurs at scales allowed by data, the corresponding $\{\beta, \gamma\}$-combinations lie at the border of the

2 σ contour, displaying that it is the peculiar shape of the power suppression to slightly disfavour this class of models. Concerning the ETHOS example models that we have analysed (purple triangles), it is interesting to note that, whereas all the corresponding values of α are in agreement with data, the $\{\beta, \gamma\}$-combinations associated to the first two scenarios are rejected. Such combinations lead indeed to a power suppression at relatively large scales, as it is quantified by the corresponding $k_{1/2}$ values listed in Table 2.5.

To summarise the take-home message from this Chapter, due to the new parametrisation it is now possible to test a wide variety of nCDM models with Lyman-α data, by simply fitting their linear power spectra in terms of $\{\alpha, \beta, \gamma\}$, and confronting the corresponding $\{\alpha, \beta, \gamma\}$-combinations against exclusion plots like the ones shown in Figure 2.12. In other words, it is sufficient to look at which region of the parameter space they sample, without the need of running any numerical simulations.

Now we would like to stress that none of the mixed (cold plus warm) DM scenarios discussed in Section 2.3 is over-plotted in Figure 2.12. This is due to the fact that the values of α which are needed in order to fit the transfer functions associated to such models must necessarily be greater than 0.1, well beyond the constraint that we obtained. By marginalising over all the other parameters, we obtain indeed an upper limit on $\alpha < 0.03$ Mpc/h (2 σ C.L.), which could be interpreted as the largest possible scale at which a power suppression induced by any nCDM scenario can be present, in order to be in agreement with Lyman-α data, provided that such nCDM scenario is captured by our parametrisation. Such constraint constitutes a strong hint that DM fluids composed by a mixture of a cold and a warm (thermal) component are disfavoured by structure formation data. A more comprehensive and systematic study focused on this particular class of models is needed before claiming that they are completely ruled out. Nevertheless, in light of our analyses it is clear that only scenarios with large masses and/or tiny abundances for the warm (thermal) component may accommodate Lyman-α data, i.e., scenarios which are practically indistinguishable with respect to the standard CDM model, with current structure formation observations.

In Figure 2.13 we show the 1D and 2D distributions for both α, β, γ, and the main astrophysical free parameters. The blue contours refer to the freely floating IGM temperature analysis, while the red contours refer to the case where a power law evolution is assumed. Both in Section 2.6.1 and in Appendix C we extensively discuss how different prior choices on the IGM temperature evolution influence the constraints on the WDM mass, when the analysis is limited to thermal models. Interestingly, the consequences of such different choices on α are mitigated in the generalised $\{\alpha, \beta, \gamma\}$-analysis, where the effects of different

assumptions on the thermal history are somehow spread on the distributions of the three parameters associated to the nCDM nature. By examining the contour plots shown in Figure 2.13, the relative stability of the 1D distribution of α is visibly evident. However, by dropping the assumption of a temperature power law evolution and letting $T_0(z)$ free to vary bin by bin, the marginalised 2 σ C.L. limit on α is weakened, becoming $\alpha < 0.05$ Mpc/h.

	α [Mpc/h]	β	γ	$k_{1/2}$ [h/Mpc]	χ^2
RP sterile neutrinos	0.025	2.3	−2.6	17.276	101
	0.071	2.3	−1.0	9.828	266
	0.038	2.3	−4.4	8.604	283
	0.035	2.1	−1.5	15.073	149
Neutrinos from particle decay	0.016	2.6	−8.1	19.012	104
	0.011	**2.7**	**-8.5**	**28.647**	**38**
	0.019	2.5	−6.9	16.478	105
	0.011	**2.7**	**-9.8**	**26.31**	**45**
Mixed models	0.16	3.2	−0.4	6.743	229
	0.20	3.7	−0.18	7.931	-
	0.21	3.7	−0.1	11.36	-
	0.21	3.4	−0.053	33.251	-
Fuzzy DM	0.054	5.4	−2.3	13.116	169
	0.040	5.4	−2.1	18.106	104
	0.030	**5.5**	**-1.9**	**25.016**	**40**
	0.022	**5.6**	**-1.7**	**34.590**	**30**
ETHOS models	0.0072	1.1	−9.9	7.274	-
	0.013	2.1	−9.3	16.880	153
	0.014	2.9	−10.0	21.584	70
	0.016	3.4	−9.3	23.045	60

Table 2.5: Here we list 20 $\{\alpha, \beta, \gamma\}$-combinations, with the corresponding value for $k_{1/2}$, each of them referring to one of the nCDM particle models examined in Section 2.3 and over-plotted in Figure 2.12. They are split into five groups, which represent some of the most viable classes of nCDM scenarios up-to-date. In the last column we report the corresponding χ^2-values from our reference MCMC analysis. Those models for which the χ^2 is not shown are associated with $\{\alpha, \beta, \gamma\}$-combinations sampling a parameter region not covered by our grid of simulations (see Table 2.1). Models highlighted in bold-face are accepted (at 2 σ C.L.) by our analysis.

It is also informative to look at the marginalised 1D distribution for the quantity $|\beta/\gamma|$, reported in Figure 2.14, which somehow compresses the information about the slope of the power suppression, and appears to be well constrained by our analysis. The corresponding 2 σ C.L. upper limit, obtained by marginalising over α, is the following: $|\beta/\gamma| < 14$. The vertical dashed line refers to the position of the $\{\beta, \gamma\}$-combination associated to thermal WDM models. Consistently with what

we have already pointed out, its position shows that the particular shape of the power suppression induced by a thermal WDM candidate is not favoured by data.

As it was expected, the bounds on the astrophysical and cosmological parameters do not present dramatic differences with respect to the thermal WDM case, apart from a mild overall weakening due to the addition of two more free parameters to the analysis. The effects on the flux power spectra induced by variations of most of the astrophysical and cosmological parameters, in fact, mainly occur at larger scales with respect to the ones influenced by the nCDM parameters, with the remarkable exception, e.g., of the IGM temperature $T_0(z)$.

Lastly, in Figure 2.15 we explicitly show, among the 109 $\{\alpha, \beta, \gamma\}$-combinations listed in Table 2.1, the squared transfer functions corresponding to nCDM models which are accepted at 2 σ C.L. by our reference MCMC analysis (red lines) and the ones associated to models that are rejected (blue lines). The nCDM models corresponding to the red curves are associated to the $\{\alpha, \beta, \gamma\}$-triplets highlighted in bold-face in Table 2.1.

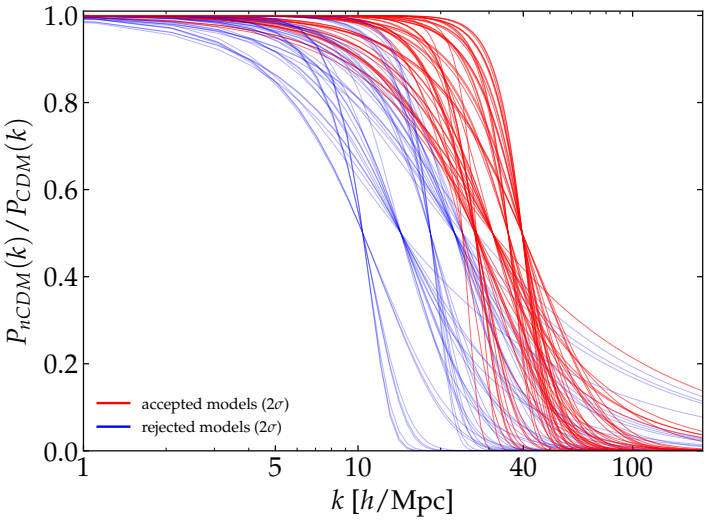

Figure 2.15: Here we show the squared transfer functions associated to the 109 $\{\alpha, \beta, \gamma\}$-combinations listed in Table 2.1, and we highlight which ones are accepted by our reference MCMC analysis at 2 σ C.L. (red lines), and which ones are rejected (blue lines).

In Table 2.6 we report the best fit parameters and their 1 σ and 2 σ C.L. intervals, corresponding to our reference MCMC analysis, namely to the assumption of a power law evolution for both the IGM temperature amplitude and slope. By looking at Table 2.6, one might notice that the lower limits on σ_8 and $n_{\rm eff}$ are sensibly underestimated with respect to the previous results. This effect cannot be fully addressed to the overall weakening of the constraints due to the presence of two additional free parameters in the MCMC analysis. As it was already mentioned in Section 2.6.1, this issue is partly due to intrinsic difficulties of the new interpolation

Parameter	(1 σ)	(2 σ)	Best Fit
$\bar{F}(z=4.2)$	[0.35, 0.39]	[0.33, 0.41]	0.35
$\bar{F}(z=4.6)$	[0.27, 0.32]	[0.25, 0.35]	0.26
$\bar{F}(z=5.0)$	[0.16, 0.21]	[0.15, 0.24]	0.18
$\bar{F}(z=5.4)$	[0.05, 0.09]	[0.02, 0.11]	0.07
T_0^A [10^4 K]	[0.55, 0.95]	[0.41, 1.23]	0.74
T_0^S	[-5, -2.72]	[-5, 1.34]	-4.38
$\widetilde{\gamma}^A$	[1.35, 1.53]	[1.21, 1.60]	1.45
$\widetilde{\gamma}^S$	[-2.16, -1.32]	[-2.41, 1.07]	-1.93
σ_8	[0.67, 0.99]	[0.53, 1.11]	0.84
$z_{\rm reio}$	[7.73, 10.32]	[7, 12.30]	9.16
$n_{\rm eff}$	[-2.6, -2.35]	[-2.6, -2.20]	-2.46
$f_{\rm UV}$	[0, 1]	[0, 1]	0.02
β	[1.5, 10]	[1.5, 10]	3.2
γ	[-6.24, -0.15]	[-10, -0.15]	-4.8
α [Mpc/h]	[0, 0.01]	[0, 0.03]	0.005

Table 2.6: Marginalised constraints at 1 σ and 2 σ C.L., and best fit values for all the free parameters of our reference MCMC analysis (see the text for further details). Our best fit model has a $\chi^2/d.o.f. = 29/38$.

scheme when sampling regions that are very far from the range covered by our simulations, a situation which can only occur for σ_8 and $n_{\rm eff}$. These two parameters are indeed the only ones for which we are scanning an interval of values spread significantly beyond the range covered by our simulations. Nevertheless, such parameters were nearly unconstrained even in previous IGM analyses, and our results are not biased by this problem. We made sure of that by performing several times each of our analyses, for both the thermal case and the general one, imposing Gaussian priors centred around Planck values for σ_8 and $n_{\rm eff}$ (Ade et al. (2016a)), both individually and in combination, with various values for the standard deviations. None of the runs provided sensibly different bounds on the other parameters. Thus, we can conclude that not being able to perfectly constrain extremely low values for σ_8 and $n_{\rm eff}$ does not affect any of our predictions. Let us remark, in this context, the importance of combining Lyman-α data with other cosmological probes, e.g., CMB and BAO measurements, to achieve tighter constraints. Such an extended data set will be used in Chapter 4 to test interacting DM models.

2.6.3 Robustness of the method

This Section is devoted to test both the $\{\alpha, \beta, \gamma\}$-fitting procedure and the interpolation scheme previously illustrated, in order to demonstrate the robustness and accuracy of the novel approach. To this end, we focus on three specific nCDM models belonging to the ETHOS framework, i.e., associated to DM-Dark Radiation interaction scenarios with three different strengths. For the purposes of this Section,

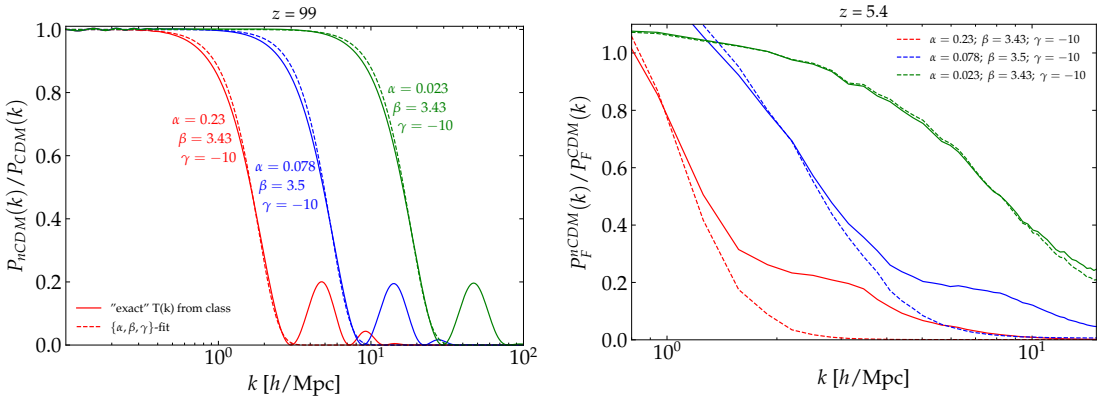

(a) **(b)**

Figure 2.16: (a) In the left panel we plot the "exact" squared transfer functions associated to three different ETHOS models (solid lines), all of them featuring small-scale oscillations. We also plot the corresponding fitted transfer functions (dashed lines), obtained by neglecting such oscillations when fitting the "exact" ones. (b) In the right panel we plot the ratios with respect to a pure CDM model (at redshift $z = 5.4$) for the flux power spectra extracted from two sets of simulations: the solid lines by using as initial conditions the "exact" transfer functions, the dashed lines by using as initial conditions the fitted $T(k)$. The colour code of the two panels is the same.

it is not necessary to go into the particle physics details of the models, we just need to notice that the corresponding transfer functions are characterised by a cut-off at the scales of interest for Lyman-α forest observations. Such squared transfer functions are plotted as solid lines in Figure 2.16a, together with the corresponding $\{\alpha, \beta, \gamma\}$-fits (dashed lines). The solid lines are dubbed as "exact" $T(k)$, since they have been produced by a specific version (Archidiacono et al. (2019)) of the numerical Boltzmann solver `CLASS` (Blas et al. (2011)), where the non-standard interactions characterising the considered models are fully implemented. That is why they feature small-scale oscillations which our $\{\alpha, \beta, \gamma\}$-parametrisation is not able to capture. For each squared transfer function plotted in Figure 2.16a, we also reported the corresponding $\{\alpha, \beta, \gamma\}$-values, obtained by fitting the solid curves down to their first minima, completely neglecting the oscillations. In Figure 2.16b we have plotted the ratios with respect to a pure CDM model (at redshift $z = 5.4$) for the flux power spectra extracted from two sets of simulations: the solid lines by using as initial conditions the aforementioned "exact" transfer functions, the dashed lines by using as initial conditions the fitted ones. By analysing the two panels, for which we adopted the same colour code, it is clear that differences between the "exact" flux power spectra and the $\{\alpha, \beta, \gamma\}$-predictions appear only when the power suppression with respect to the standard CDM case is more than 50%. Furthermore, this is true only for nCDM models characterised by suppression at significantly large scales, which are indeed associated with α-values well above the marginalised

2 σ upper limit presented in Section 2.6.2, i.e., $\alpha < 0.03$ Mpc/h. Let us highlight, in fact, that the fitted transfer function of the most viable model that we considered (green dashed line), is practically superposed with the corresponding "exact" $T(k)$ (green solid line). For both these reasons, we conclude that the differences due to our inability to capture the small-scale oscillations appear only for flux power spectra which lie anyway very far from the Lyman-α forest data points. This completely justifies ignoring such oscillations when applying our fitting procedure (regarding this subject, see also Chapter 4, where further applicability checks of the method are successfully carried out).

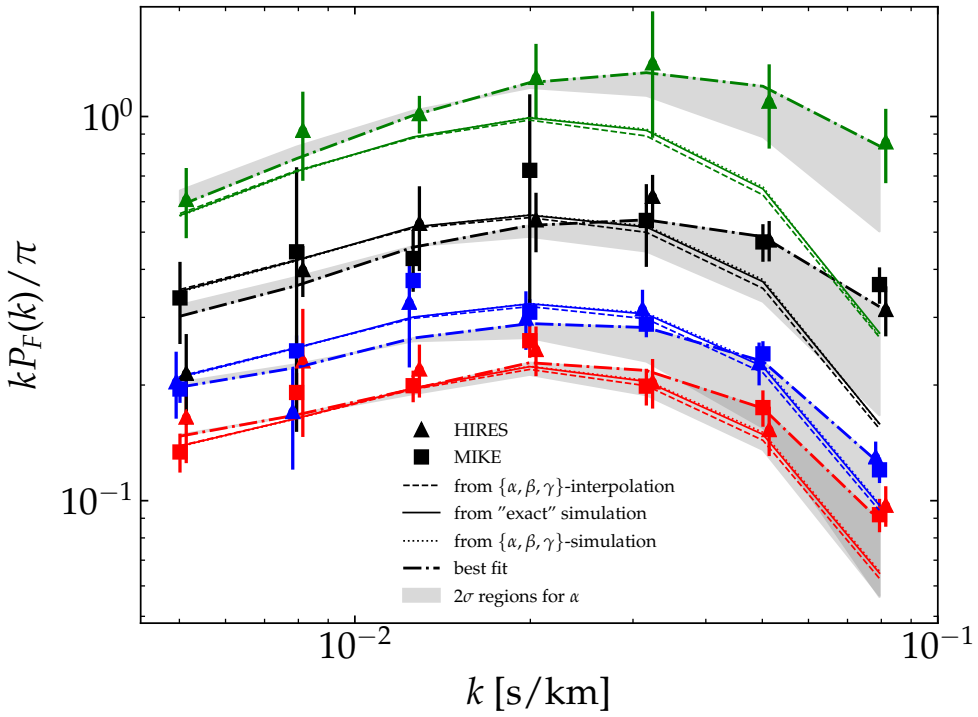

Figure 2.17: Flux power spectra for the most viable model shown in Figure 2.16, together with MIKE/HIRES data points and error bars. Different colours stand for different redshifts, i.e., from the bottom to the top, $z = 4.2, 4.6, 5.0, 5.4$. See the text for the detailed discussion about the different power spectra listed in the legend.

Let us now look at Figure 2.17, where we focus on the flux power spectra of the model described by the green curves in Figure 2.16, given that its α-value is the only one accepted at 2 σ C.L. by our analysis. Different colours stand for different redshifts, i.e., from the bottom to the top, $z = 4.2, 4.6, 5.0, 5.4$. For each redshift bin we plot both MIKE and HIRES data points and error bars, as triangles and squares, respectively. Firstly, we note again the excellent agreement between the flux power obtained by using the "exact" initial conditions (solid lines) with respect to the fitted ones (dashed lines). Most importantly, the dotted lines correspond to

flux power spectra computed by simply interpolating in our coarse grid, without running any dedicated simulation, and they nicely coincide with the solid ones, at each redshift. It is thus evident that models associated with $\{\alpha, \beta, \gamma\}$-combinations sampling our grid of simulations are perfectly reproduced by the novel interpolation scheme. As a reference, we also plot our best fit flux power spectra (dot-dashed lines), associated to the parameter values reported in Table 2.6. The grey dashed areas represent instead the region spanned by flux power spectra with values of α varying up to its 2 σ C.L. marginalised upper bound.

2.6.4 Comparison with the "area criterion"

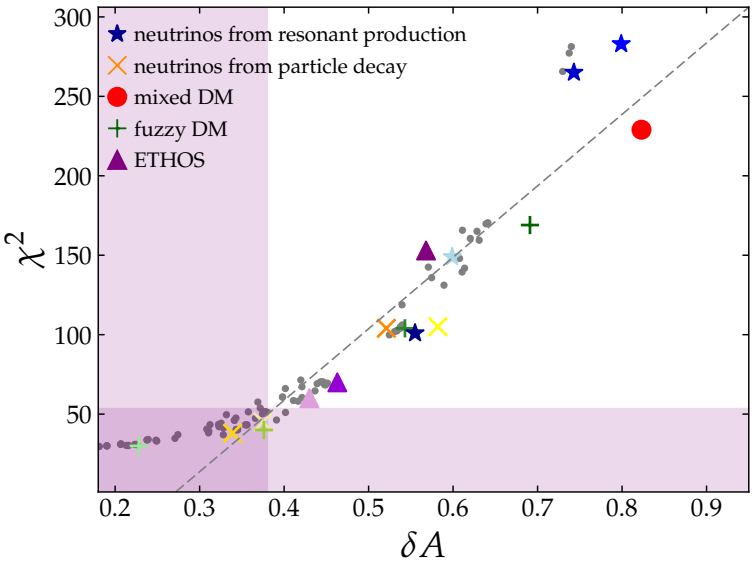

Figure 2.18: Correlation between the area estimators δA relative to each of the nCDM models examined in Section 2.3, and the corresponding χ^2-values, obtained with our reference MCMC analysis and reported in Table 2.5. Different symbols refer to points belonging to different groups of nCDM models, consistently recalling the convention adopted in Figure 2.12. The dashed line represents the linear regression fit to the results. The grey dots correspond to the $\{\alpha, \beta, \gamma\}$-combinations listed in Table 2.1, each of them associated to a different nCDM simulation. The vertical and horizontal shaded bands represent the regions corresponding to the 2 σ C.L. for the "area criterion" and the MCMC analysis, respectively.

In Section 2.5.2 we introduced a simple method, based on linear perturbation theory, for testing different nCDM scenarios with Lyman-α forest data by using an approximate yet intuitive estimator, and we named it as "area criterion". We have shown that with such method it is not possible to extract absolute limits on the nCDM parameters, but it allows to look into deviations with respect to a given

reference case. Let us stress that the choice of the reference thermal WDM model for calibrating the "area criterion" is crucial for establishing the threshold which defines which models to accept/reject. In Section 2.5.2, we performed an analysis in the $\{\alpha, \beta, \gamma\}$-space by calibrating the method with $m_{\rm WDM} = 3.5$ keV (i.e., the 2 σ C.L. limit from the conservative MIKE/HIRES+XQ-100 analysis by Iršič et al. (2017b)). We quantified the reference power suppression through the area estimator δA, and we rejected (at 2 σ C.L.) all those nCDM models that feature a larger power suppression with respect to the reference one. Since the aforementioned 2 σ reference limit roughly coincides with the constraint quoted in Appendix C for thermal WDM masses (power law analysis), we can consistently compare the results presented in Section 2.6.2 with the approximate conclusions reported in Section 2.5.2. It is interesting, indeed, to quantify the precision of the simple "area criterion" with respect to the full statistical data analysis that we just illustrated.

Firstly, let us note that by marginalising the "area criterion" results over β and γ, we obtained the following upper limit on $\alpha < 0.058$ Mpc/h (2 σ C.L.), which is weaker with respect to the bounds quoted in Section 2.6.2. This could mostly be due to the unavoidably prominent tail at large values in the 1D α-distribution, when the analysis is done with the approximate area method. As we pointed out in Section 2.6.2, such tail corresponds to extreme values for both β and γ, as well as very cold IGM temperatures, unlikely to be physically motivated. By simply applying the "area criterion" it is intrinsically impossible to account for this astrophysical information.

Let us now compare the results listed in Table 2.5 with the conclusions reported in Table 2.4, which have been obtained by applying the "area criterion" to the same nCDM particle model examples. In Figure 2.18 we show the correlation between the area estimators δA relative to each of the examined nCDM models and the corresponding χ^2-values, obtained with our reference MCMC analysis and reported in Table 2.5. The two different sets of predictions are visibly correlated, with correlation coefficient $r = 0.94$. The dashed line represents the linear regression fit to the results. Different symbols are used to identify points belonging to different groups of nCDM models, consistently recalling the convention adopted in Figure 2.12. The vertical and horizontal shaded bands represent the regions which are included at 2 σ C.L. by the "area criterion" and the MCMC analysis, respectively. All those models sampling the lower left intersection between the two bands are thus accepted (at 2 σ C.L.) by both the analyses, whereas models which sample the white region are excluded by both of them. Interestingly, none of the particle model examples examined is rejected by one analysis and accepted by the other.

When considering also the grey dots, which correspond to the $\{\alpha, \beta, \gamma\}$-triplets listed in Table 2.1, the correlation between the conclusions drawn by the two

methods is even more evident. Since the grey dots which sample the vertical shaded band refer to models accepted at 2 σ C.L., they are also associated to the red transfer functions shown in Figure 2.15. Note that the departure from the linear correlation occurring for small values of δA is an intrinsic feature of the method. Whereas the χ^2-distribution saturates when approaching the best fit χ^2-value, the area estimator can assume arbitrarily small (positive) values. It is remarkable that, among 109 models thoroughly sampling the $\{\alpha, \beta, \gamma\}$-space, only two of them are accepted by the "area criterion" while rejected by the MCMC analysis. Conversely, it is worthwhile to notice that only one borderline $\{\alpha, \beta, \gamma\}$-triplet is rejected by the "area criterion" while accepted by the MCMC analysis, that is a confirmation of the suitability of the former as an approximate yet effective and conservative method to perform preliminary tests on non-standard DM scenarios.

This Chapter is mainly based on:

T. Kobayashi, R. Murgia, A. De Simone, V. Iršič, M. Viel
Lyman-α Constraints on Ultralight Scalar Dark Matter: Implications for the Early and Late Universe
PRD 96, 123514 (2017); arXiv:1708.00015

M. Nori, R. Murgia, V. Iršič, M. Baldi, M. Viel
Lyman-α forest and non-linear structure characterization in Fuzzy Dark Matter cosmologies
MNRAS 482, 3, 3227–3243 (2019); arXiv:1809.09619

3 Ultra-light scalar Dark Matter

3.1 Overview

Ultra-light scalar DM – also denoted as Fuzzy DM (FDM) – scenarios describe the DM sector as made up of very light bosonic particles (see, e.g., Hui et al. (2017) for a review on the topic), so light that their quantum nature becomes relevant also on cosmological scales. This requires a description of the DM dynamics in terms of the Schrödinger equation, in order to take into account for quantum corrections, and can be mapped in a fluid-like description where a Quantum Potential (QP) enters the classical Navier–Stokes equation (Hu et al. (2000)).

The typical wave-like quantum behaviour adds to the standard CDM dynamics a repulsive effective interaction that, along with creating oscillating interference patterns, actively smooths matter overdensities below a redshift-dependent scale that decreases with the cosmic evolution – as confirmed by FDM linear simulations (see, e.g., Marsh & Ferreira (2010); Hlozek et al. (2015)) – thus potentially easing some of the small-scale tensions of the standard CDM paradigm.

The lack of density perturbations at small scales induced by the QP is represented, in Fourier space, by a sharp suppression of the matter power spectrum, that persists – at any given scale – until the action range of the QP shrinks below such scale and cannot balance any longer the effect of the gravitational potential (see, e.g., Marsh (2016a) for another detailed review on the subject). As a matter of fact, while linear theory predicts that perturbations at scales smaller than the cut-off never catch up with those at larger scales – untouched by FDM peculiar dynamics –, non-linear cosmological simulations have shown that gravity is able to restore intermediate scales to the unsuppressed level, in a sort of healing process (Marsh (2016b); Nori & Baldi (2018)).

FDM non-linear cosmological simulations have been performed over the years either with highly numerically intensive high-resolution Adaptive Mesh Refinement (AMR) algorithms able to solve the Schrödinger-Poisson equations over a grid (Schive et al. (2010, 2018)) or with standard N-body codes that, however, include the (linear) suppression only in the initial conditions, but neglect the integrated effect of the FDM interaction during the subsequent dynamical evolution (Schive et al. (2016); Iršič et al. (2017a); Armengaud et al. (2017)) – basically treating FDM as standard collisionless CDM with a suppressed primordial power spectrum, as it is routinely done in nCDM simulations, like the ones presented in Chapter 2. The former approach led to impressive results in terms of resolution (Woo & Chiueh (2009); Schive et al. (2014); Veltmaat et al. (2018)) but is extremely computationally demanding, thereby hindering the possibility of adding a full hydrodynamic description of gas and star formation for cosmologically representative simulation domains. On the other hand, the latter allows for such possibility because of its reduced computational cost which is, however, gained at the price of the substantial approximation of neglecting QP effects during the simulation (Schive et al. (2016)).

For these reasons, following the idea proposed by Mocz & Succi (2015), Nori & Baldi (2018) devised `AX-GADGET`, a modified version of the numerical code `GADGET` (Springel et al. (2001b); Springel (2005)), to include the dynamical effect of QP through SPH numerical methods. The explicit approximation of the dependence on neighbouring particles results in a less numerically demanding code with respect to full-wave AMR solvers, without compromising cosmological results, with the additional ability to exploit the gas and star physics already implemented in `GADGET`. Given that gravity, as mentioned above, can restore the suppressed power at intermediate scales in the non-linear regime, major observables related to the LSS at such scales may appear similar in both FDM and CDM cosmologies at sufficiently low redshifts. For this reason, Lyman-α forest observations could play a crucial role in distinguishing such radically different models of DM, being one of the most far reaching direct astrophysical probes in terms of redshift of the LSS observables, sampling the redshift range $z \sim 2 - 5$.

This Chapter is organised as follows: in Section 3.2 we briefly describe the FDM models under consideration; in Section 3.3 we present the most updated FDM constraints obtained from Lyman-α forest data, we discuss their implications for the CDM small-scale crisis, and we quantify for the first time the impact of the full non-linear treatment of the QP; in Section 3.4 we address the implications of such results for the early universe: in particular, on the initial value of the scalar field, isocurvature perturbations, and the tensor-to-scalar ratio; in Section 3.5 we study the statistical properties of DM halos, and extract information on how FDM affects the abundance, the shape, and density profiles of DM (sub)structures.

3.2 Theoretical framework

The idea of describing DM and its key role in the LSS formation in terms of ultra-light scalar particles, i.e., particles with mass $\sim 10^{-22} eV/c^2$, was introduced by Hu et al. (2000), in which the term *Fuzzy* DM was used for the first time and the cosmological implications induced by the quantum behaviour of an ultra-light scalar DM field on linear cosmological perturbations were outlined. In this Section we recall the main properties of such light bosonic field in a cosmological framework, and how it affects the growth of LSS.

3.2.1 Fuzzy Dark Matter as a fluid

Let us start by describing FDM as a cosmological fluid, using the formalism discussed by, e.g., Arvanitaki et al. (2010); Hu et al. (2000); Hui et al. (2017); Widrow & Kaiser (1993). A large number of scalar particles with mass m behaves as a classical field obeying the Klein–Gordon and Einstein's equations,

$$\nabla_\mu \nabla^\mu \phi = m^2 \phi, \quad G_{\mu\nu} = 8\pi G\, T_{\mu\nu}, \tag{3.1}$$

where the scalar field contributes to the energy-momentum tensor as

$$T^\phi_{\mu\nu} = g_{\mu\nu}\left(-\frac{1}{2}\partial_\rho\phi\partial^\rho\phi - \frac{1}{2}m^2\phi^2\right) + \partial_\mu\phi\partial_\nu\phi. \tag{3.2}$$

Regarding the metric of the Friedmann-Robertson-Walker (FRW) universe with perturbations, we take the Newtonian gauge and ignore anisotropic stress,

$$ds^2 = -(1+2\Phi)dt^2 + a(t)^2(1-2\Phi)d\boldsymbol{x}^2. \tag{3.3}$$

The Hubble rate is defined as $H = \dot{a}/a$, with a dot denoting a derivative with respect to cosmic time t. When $H \ll m$, and hence the scalar field is harmonically oscillating, it is convenient to re-write ϕ as

$$\phi = \frac{1}{\sqrt{2m}}\left(\varphi e^{-imt} + \varphi^* e^{imt}\right), \tag{3.4}$$

in terms of a complex field φ describing the oscillation amplitude whose time-dependence is slow compared to the oscillation period. We re-write the Klein–Gordon equation in terms of φ under the assumption of tiny perturbations $|\Phi| \ll 1$, and thus ignoring terms that contain quadratic or higher orders of Φ. Furthermore, let us focus on non-relativistic modes ($k/a \ll m$), and suppose the time scales for the variations of Φ, φ, a, and their derivatives to be much longer than the oscillation period, i.e., $|\dot{\Phi}| \ll m|\Phi|$, $|\ddot{\varphi}| \ll m|\dot{\varphi}|$, etc. This, in particular, allows us to drop second time-derivatives in the equation. After taking an average over

the oscillation period, the Klein–Gordon equation reduces to a Schrödinger-type equation (Widrow & Kaiser (1993)),

$$i\left(\dot{\varphi}+\frac{3}{2}H\varphi\right)=-\frac{\partial^2\varphi}{2a^2m}+m\Phi\varphi, \tag{3.5}$$

where $\partial^2 \equiv \partial_i\partial_i$, and sum over repeated spatial indices is implied irrespective of their positions.

In terms of the amplitude and the phase of φ, we now define

$$\rho_\phi \equiv m\varphi\varphi^*, \quad v_i \equiv \frac{\partial_i\{\arg(\varphi)\}}{am}=-\frac{i}{2am}\left(\frac{\partial_i\varphi}{\varphi}-\frac{\partial_i\varphi^*}{\varphi^*}\right), \tag{3.6}$$

whose meaning will soon become clear. Multiplying both sides of the Schrödinger equation (Equation (3.5)) by φ^*, its real and imaginary parts respectively lead to the following equations:

$$\dot{v}_i+Hv_i+\frac{v_j\partial_j v_i}{a}=-\frac{\partial_i\Phi}{a}+\frac{1}{2a^3m^2}\partial_i\left(\frac{\partial^2\sqrt{\rho_\phi}}{\sqrt{\rho_\phi}}\right), \tag{3.7}$$

$$\dot{\rho}_\phi+3H\rho_\phi+\frac{\partial_i(\rho_\phi v_i)}{a}=0. \tag{3.8}$$

One can also re-write the Einstein's equation in a similar fashion. Focusing on sub-horizon and non-relativistic modes, i.e., $H \ll k/a \ll m$, it can be checked that the $(0,0)$ component of the Einstein's equation yields

$$\frac{\partial^2\Phi}{a^2}=4\pi G\left(\rho_\phi+T_{00}^{\text{others}}\right)-\frac{3}{2}H^2, \tag{3.9}$$

where T_{00}^{others} denotes the contributions to the energy-momentum tensor from components other than FDM. Interpreting ρ_θ and v_i as the density and velocity fields, the set of Equations (3.7), (3.8), and (3.9) are seen to correspond respectively to the Euler, continuity, and Poisson equations; thus we have arrived at a fluid description of the scalar DM. The only difference with the familiar CDM fluid is the existence of the last term in the right hand side of Equation (3.7), which represents a pressure due to the wave-like nature of the scalar field on small scales.

3.2.2 Evolution of the density perturbations

Let us now study the evolution of density perturbations in a matter dominated universe that is filled with FDM and CDM. We also describe CDM as a fluid obeying the Euler and continuity equations as in Equation (3.7) and (3.8), except for that there is no pressure term, and replace T_{00}^{others} in the Poisson equation (Equation (3.9))

by the CDM density $\rho_{\rm c}$. Then one immediately sees from that the homogeneous and isotropic background satisfies

$$\frac{\dot{\bar{\rho}}_\phi}{\bar{\rho}_\phi} = \frac{\dot{\bar{\rho}}_{\rm c}}{\bar{\rho}_{\rm c}} = -3H, \quad \bar{\rho}_\phi + \bar{\rho}_{\rm c} = \frac{3H^2}{8\pi G}, \tag{3.10}$$

where a bar is used to denote unperturbed values. We discuss the density fluctuations around the background in terms of the density contrast,

$$\delta_n = \frac{\rho_n - \bar{\rho}_n}{\bar{\rho}_n}, \tag{3.11}$$

where $n = \phi, {\rm c}, {\rm m}$, with $\rho_{\rm m} = \rho_\phi + \rho_{\rm c}$ being the total DM density. Expressing the ratio between the unperturbed densities of the scalar and total DM as

$$F = \frac{\bar{\rho}_\phi}{\bar{\rho}_{\rm m}}, \tag{3.12}$$

which is a constant in the range $0 \leq F \leq 1$, the density contrast of the total DM is written as

$$\delta_{\rm m} = F\delta_\phi + (1-F)\delta_{\rm c}. \tag{3.13}$$

Expanding Equations (3.7), (3.8), and (3.9) up to linear order in δ and v_i, and then combining them to eliminate v_i and Φ, one arrives at the evolution equations for the density contrasts,

$$\ddot{\delta}_{\phi\boldsymbol{k}} + 2H\dot{\delta}_{\phi\boldsymbol{k}} + \frac{c_s^2 k^2}{a^2}\delta_{\phi\boldsymbol{k}} - \frac{3}{2}H^2\delta_{{\rm m}\boldsymbol{k}} = 0, \tag{3.14}$$

$$\ddot{\delta}_{{\rm c}\boldsymbol{k}} + 2H\dot{\delta}_{{\rm c}\boldsymbol{k}} - \frac{3}{2}H^2\delta_{{\rm m}\boldsymbol{k}} = 0. \tag{3.15}$$

Here we have written the linearised equations in terms of the Fourier components ($\boldsymbol{k}$ is a comoving wavenumber, with $k = |\boldsymbol{k}|$), and the sound speed of the FDM component is

$$c_s^2 \equiv \frac{k^2}{4a^2m^2}. \tag{3.16}$$

From their derivations, it should be noted that Equations (3.14) and (3.15) are valid during the matter dominated epoch, and for wavenumbers that are sub-horizon and non-relativistic, i.e., $H \ll k/a \ll m$.

In the case where FDM constitutes the entire DM, i.e., $F = 1$, one can read off its Jeans wavenumber from the last two terms in Equation (3.14) as

$$\frac{k_{\rm J}}{a} = \sqrt{Hm}, \tag{3.17}$$

where we have ignored numerical factors. As we are interested in FDM masses larger than the Hubble rate at matter/radiation equality, the Jeans length is smaller than the Hubble length throughout the matter dominated epoch. On length scales larger than the Jeans length ($k < k_{\rm J}$), the pressure term is negligible and FDM behaves similarly to CDM; thus the density fluctuation possesses the usual growing mode $\delta_{\phi \boldsymbol{k}} \propto a$. However, below the Jeans length ($k > k_{\rm J}$), $\delta_{\phi \boldsymbol{k}}$ undergoes oscillations of $\propto \exp(\pm 2ic_s k/aH)$ and thus does not grow. Here note that in a matter dominated universe, the comoving Jeans wavenumber grows as $k_{\rm J} \propto a^{1/4}$. Hence for modes that are sub-Jeans at the time of matter/radiation equality, i.e., $k > k_{\rm Jeq}$, the FDM perturbation is prevented from growing until the mode k crosses the Jeans scale at $a = a_k$, where

$$a_k = a_{\rm eq} \left(\frac{k}{k_{\rm Jeq}} \right)^4 . \tag{3.18}$$

With a slight abuse of language, we will refer to $k_{\rm J}$ as defined in Equation (3.17) as the Jeans wavenumber, even in cases with $F \neq 1$.

If DM is a mixture of FDM and CDM, i.e., $0 < F < 1$, then the matter fluctuations grow even on small scales, albeit slowly. This is seen by dropping $\delta_{\phi \boldsymbol{k}}$ in Equation (3.15) by taking an average over $\delta_{\phi \boldsymbol{k}}$'s oscillation period, giving

$$\frac{d^2 \delta_{{\rm c}\boldsymbol{k}}}{da^2} + \frac{3}{2a} \frac{d\delta_{{\rm c}\boldsymbol{k}}}{da} - \frac{3(1-F)\delta_{{\rm c}\boldsymbol{k}}}{2a^2} = 0. \tag{3.19}$$

Here we have re-written the derivatives in terms of time by those of a, using $H \propto a^{-3/2}$ (see Equation (3.10)). This equation has a general solution of

$$\delta_{{\rm c}\boldsymbol{k}} = C_+ a^{n_+} + C_- a^{n_-}, \quad \text{with} \quad n_\pm = \frac{-1 \pm \sqrt{25 - 24F}}{4}, \tag{3.20}$$

whose first term is a growing mode, and the second a decaying mode. The growing CDM perturbation drags the FDM component and thus $\delta_{\phi \boldsymbol{k}}$ starts to grow even before crossing the Jeans scale.

On the other hand, wave modes with $k < k_{\rm Jeq}$ stay super-Jeans throughout the matter dominated epoch. We should also note that during radiation domination, there is no significant growth for both FDM and CDM perturbations on sub-horizon scales.[1] Therefore the difference in the matter perturbations between cases with and without FDM becomes prominent on wavenumbers $k > k_{\rm Jeq}$, mainly due to the difference in the evolution during the matter dominated epoch. For reference, the Jeans wavenumber at the matter/radiation equality is

$$\frac{k_{\rm Jeq}}{a_0} = \frac{a_{\rm eq}}{a_0} \sqrt{H_{\rm eq} m} \approx 7\,{\rm Mpc}^{-1} \left(\frac{m}{10^{-22}\,{\rm eV}} \right)^{1/2}, \tag{3.21}$$

where the subscript "0" denotes quantities today.

[1] The evolution of the matter fluctuations prior to matter/radiation equality can be described by further including a homogeneous radiation component $\bar{\rho}_{\rm r} \propto a^{-4}$ to $T_{00}^{\rm others}$ in Equation (3.9) (hence ignoring radiation perturbations on sub-horizon scales); this amounts to multiplying the term $(3/2)H^2 \delta_{{\rm m}\boldsymbol{k}}$ in Equations (3.14) and (3.15) by $\bar{\rho}_{\rm m}/(\bar{\rho}_{\rm m} + \bar{\rho}_{\rm r})$. During radiation domination ($\bar{\rho}_{\rm m}/\bar{\rho}_{\rm r} \to 0$), $\delta_{\phi \boldsymbol{k}}$ has general solutions of $\exp\{\pm i(c_s k/aH) \ln a\}$.

3.2.3 Impact on the linear matter power spectrum

In order to evaluate the growth of the perturbations for arbitrary F, we suppose an adiabatic initial condition and consider δ_ϕ to behave similarly to δ_c during radiation domination, and on super-Jeans scales. For wavenumbers $k > k_\mathrm{Jeq}$, let us make a rough approximation that the total matter perturbation follows $\delta_{\mathrm{m}\boldsymbol{k}} \propto a^{n_+}$ since matter/radiation equality until crossing the Jeans scale at $a = a_k$, then subsequently grows as the usual $\delta_{\mathrm{m}\boldsymbol{k}} \propto a$. This approximation is crude, but allows us to understand the overall spectral shape of the suppression and how it is determined by the FDM parameters.

By comparison with the pure CDM ($F = 0$) case where $\delta_{\mathrm{m}\boldsymbol{k}} \propto a$ throughout matter domination, one can estimate the suppression of the linear matter power spectrum due to FDM. For wave modes that have crossed the Jeans scale, the matter perturbation containing FDM is suppressed relative to that with pure CDM by

$$\left|\frac{\delta^{(\phi+\mathrm{c})}_{\mathrm{m}\boldsymbol{k}}}{\delta^{(\mathrm{c})}_{\mathrm{m}\boldsymbol{k}}}\right|_{k_\mathrm{Jeq}<k<k_\mathrm{J}} = \left(\frac{a_\mathrm{eq}}{a_k}\right)^{1-n_+} = \left(\frac{k_\mathrm{Jeq}}{k}\right)^{4(1-n_+)}. \tag{3.22}$$

On the other hand, for modes that are still sub-Jeans at the time the matter perturbations are measured, the suppression saturates to a k-independent value of

$$\left|\frac{\delta^{(\phi+\mathrm{c})}_{\mathrm{m}\boldsymbol{k}}}{\delta^{(\mathrm{c})}_{\mathrm{m}\boldsymbol{k}}}\right|_{k>k_\mathrm{J}} = \left(\frac{a_\mathrm{eq}}{a}\right)^{1-n_+} = \left(\frac{k_\mathrm{Jeq}}{k_\mathrm{J}}\right)^{4(1-n_+)}. \tag{3.23}$$

See also Arvanitaki et al. (2010), where similar results were obtained.

The relative suppression of the linear matter power spectrum $\mathcal{P}_\mathrm{m}(k) \propto |\delta_{\mathrm{m}\boldsymbol{k}}|^2$ due to FDM, obtained from squaring Equations (3.22) and (3.23), is sketched in the left panel of Figure 3.1. The actual spectrum with FDM can be oscillatory, and so what is illustrated in the figure should be considered as the envelope. We stress that while the FDM mass m determines the Jeans wavenumber at equality k_Jeq above which the suppression appears, the spectral index of the suppression is set by the fraction F. Moreover, the suppression saturates at the Jeans wavenumber k_J at the time of the measurement, where the saturated suppression factor is independent of m.[2]

Provided that the square of Equation (3.23) is smaller than 1/2, we can define the wavenumber $k_{1/2}$ at which the power spectrum $\mathcal{P}_\mathrm{m}(k)$ is suppressed by 1/2, as in Chapter 2. From Equation (3.22), $k_{1/2}$ is obtained in terms of the Jeans scale at equality (Equation (3.21)) as

$$\frac{k_{1/2}}{k_\mathrm{Jeq}} = f(F), \quad \text{where} \quad f(F) = 2^{\frac{1}{10-2\sqrt{25-24F}}}. \tag{3.24}$$

[2]The suppression function presented in Equation (8) of Hu et al. (2000) behaves quite differently from what we discussed in the region $k \gg k_\mathrm{Jeq}$. This is because the function of Hu et al. (2000) was obtained to describe the first few efolds of suppression at around k_Jeq, instead of the asymptotic behaviour. We thank Wayne Hu for private communication on this point.

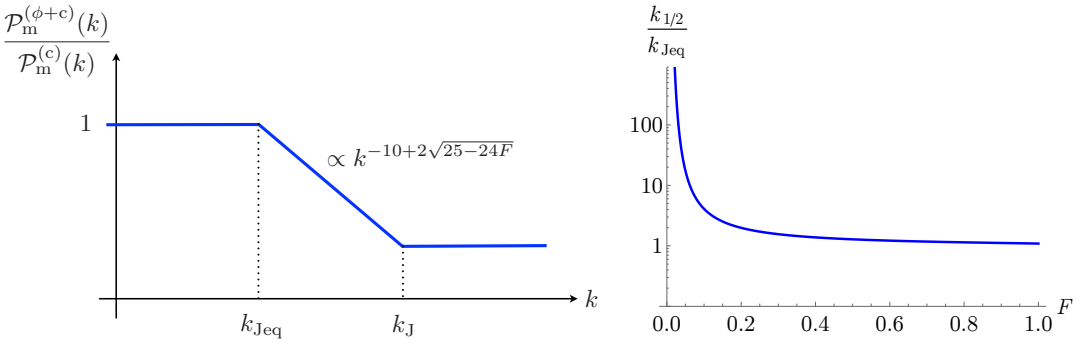

Figure 3.1: *Left*: A sketch of the suppression of the linear matter power spectrum containing FDM relative to that with pure CDM. *Right*: The wavenumber $k_{1/2}$ at which the linear matter power spectrum is suppressed by 1/2, in units of the Jeans wavenumber at matter/radiation equality $k_{\rm Jeq}$, as a function of the FDM fraction F.

This ratio as a function of F describes how the suppression effect is diluted for a smaller FDM fraction. We plot this in the right panel of Figure 3.1, where one sees that $k_{1/2} \sim k_{\rm Jeq}$ for $F > 0.1$, whereas $k_{1/2}$ becomes exponentially larger than $k_{\rm Jeq}$ for $F \lesssim 0.1$. In particular if $F \lesssim 0.07$, then $k_{1/2}$ even exceeds the present-day Jeans scale $k_{\rm J0} = (a0/a_{\rm eq})^{1/4} k_{\rm Jeq} \sim (3000)^{1/4} k_{\rm Jeq}$ (here we are ignoring DE for simplicity); namely, the linear power spectrum today does not fall below 1/2 of that from pure CDM. This indicates that FDM, no matter how light its mass is, would not impact structure formation as long as its fraction is below $\sim 10\%$ of the total DM.

We have also numerically solved the coupled evolution Equations (3.14) and (3.15). In order to incorporate the slow growth of the CDM perturbation towards the end of the radiation dominated epoch, we further included a background radiation component as explained in Footnote 1, and started the computations at $a = a_{\rm eq}/10$ with a simplified adiabatic initial condition $\delta_{\phi \boldsymbol{k}} = \delta_{{\rm c}\boldsymbol{k}}$, $\dot{\delta}_{\phi \boldsymbol{k}} = \dot{\delta}_{{\rm c}\boldsymbol{k}} = 0$.[3] The ratio of the resulting linear matter power spectrum today between cases with FDM and pure CDM is shown in Figure 3.2, for various values of the FDM fraction F. The FDM mass is fixed to $m = 10^{-22}\,{\rm eV}$, and the spectrum is shown for wavenumbers that are sub-horizon and non-relativistic ($H < k/a < m$) at the initial time of the calculation. As was indicated by the analytic arguments, the suppression appears at around the Jeans scale at equality $k_{\rm Jeq}/a_0 \sim 7\,{\rm Mpc}^{-1}$, and saturates at around the Jeans scale today $k_{\rm J0}/a_0 \sim 50\,{\rm Mpc}^{-1}$. It is also clearly seen that $k_{1/2} \sim k_{\rm Jeq}$ for $F > 0.1$, whereas for $F \lesssim 0.1$ the suppression does not fall much below 50% on any scale.

[3] A more rigorous treatment would involve solving the relativistic perturbation equations until the fluid description becomes valid (Khlopov et al. (1985); Nambu & Sasaki (1990); Ratra (1991); Hwang (1997)), and including other components such as baryons and DE. However, our simplified treatment should suffice to estimate the relative suppression of FDM perturbations at the order-of-magnitude level.

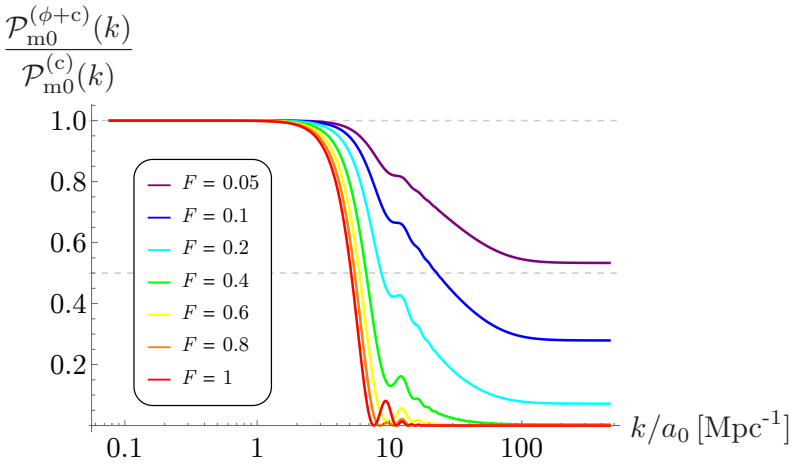

Figure 3.2: Suppression of the linear matter power spectrum today due to FDM with mass $m = 10^{-22}\,\mathrm{eV}$, for various values of the FDM fraction F. The results are obtained by numerically solving Equations (3.14) and (3.15).

The scale $k_{1/2}$ also offers a rough guide to estimate the mass of halos whose formation are suppressed. The mass contained in a sphere of diameter $2\pi a/k_{1/2}$ is

$$M_{1/2} = \frac{H^2}{2G}\left(\frac{\pi a}{k_{1/2}}\right)^3 \sim 10^{10} M_\odot\, f(F)^{-3}\left(\frac{m}{10^{-22}\,\mathrm{eV}}\right)^{-3/2}, \tag{3.25}$$

where in the far right hand side we substituted Equation (3.24) to $k_{1/2}$, used the fact that $M_{1/2}$ is time-independent during matter domination, and thus estimated its value at matter/radiation equality. From linear perturbation theory, we can thereby infer that FDM suppresses the number of halos with masses below $M_{1/2}$.

3.3 Lyman-α forest constraints

The FDM mass was recently constrained from the Lyman-α forest by Iršič et al. (2017a); Armengaud et al. (2017), for the case where FDM constitutes the entire DM, and under the universally adopted approximation that ultra-light scalars behave as standard pressureless CDM at redshift $z < 99$. All the limits available in the literature have been indeed established by comparing Lyman-α data with flux power spectra obtained from standard SPH cosmological simulations, which completely neglected the QP effects during the non-linear structure evolution. In other words, the non-standard nature of the DM candidate is simply encoded in the suppressed initial conditions, computed at $z < 99$, and used as inputs for performing the hydrodynamic simulations. In this Section we extend such analyses in two different directions:

- we study the case where DM consists of both FDM and CDM, which is crucial for constraining general theories with light scalars which are not necessarily designed to explain the DM of our universe (Subsection 3.3.1);

- we use the aforementioned numerical code `AX-GADGET` (Nori & Baldi (2018)) to provide the first fully accurate constraints on the FDM mass, by going beyond the standard dynamical approximation of ignoring the time-integrated QP effect, and perform a meticulous comparison with the bounds determined under such approximation, in order to exactly quantify its validity (Subsection 3.3.4).

For both the two analyses, we rely on a sample of 100 medium-resolution, high-signal-to-noise quasar spectra of the XQ-100 survey (López et al. (2016)), with emission redshifts $3.5 < z < 4.5$. A detailed description of the data and the power spectrum measurements of the XQ-100 survey has been discussed by Iršič et al. (2017c). The spectral resolution of the X-shooter spectrograph is 30-50 km/s, depending on wavelength. The flux power spectrum $P_{\rm F}(k, z)$ has been calculated for a total of 133 (k, z) data points in the ranges $z = 3, 3.2, 3.4, 3.6, 3.8, 4, 4.2$ and 19 bins in $k-$space in the range 0.003-0.057 s/km. We also use the measurements of the flux power spectrum by Viel et al. (2013), at redshift bins $z = 4.2, 4.6, 5.0, 5.4$ and in 10 $k-$bins in the range 0.001-0.08 s/km. In this second sample the spectral resolution of the quasar absorption spectra obtained with the MIKE and HIRES spectrographs are about 13.6 and 6.7 km/s, respectively (see Chapter 2 – Section 2.6.1 – for a more comprehensive description of this data set).

Compared to XQ-100, the HIRES/MIKE sample has the advantage of probing smaller scales and higher redshift, where the primordial power spectrum is more linear and thereby more constraining for the models considered here. Since the thermal broadening (measured in s/km) of the Lyman-α forest lines is approximately constant with redshift, the presence of a cut-off in the matter power spectrum due to the wave-like nature of FDM becomes more prominent in velocity space at high redshift, due to the $H(z)/(1 + z)$ scaling between the fixed comoving length scale set by the FDM properties (see Equations (3.21) and (3.24)) and the corresponding velocity scale.

To obtain the results presented in Subsection 3.3.1, we simulated 9 different FDM models with `GADGET-3` (Springel et al. (2001b); Springel (2005)), with FDM masses m of 1, 4 and $15.7{\cdot}10^{-22}$ eV, and density ratios F between the FDM and total DM (i.e., FDM plus CDM) of 1, 0.75 and 0.25. We also simulated the corresponding ΛCDM model, and since the interpolation is done in terms of $\alpha = 10^{-22}$ eV$/m$, the entire range from $\alpha = 0$ (ΛCDM case) to $\alpha = 1$ ($m = 10^{-22}$ eV) is covered by interpolation alone. For larger values of α, or equivalently $m < 10^{-22}$ eV, linear extrapolation is used. In the plane of the DM ratio and the FDM mass,

we assume that the ΛCDM model is exact on the axis of $F = 0$ and any m. The initial conditions have been produced at redshift $z = 99$ using the `2LPTic` code (Crocce et al. (2006)), according to a random realisation of the suppressed linear power spectrum as calculated by `axionCAMB` (Hlozek et al. (2015)), for the different FDM models under investigation.

Concerning the analysis presented in Subsection 3.3.4, we performed a set of 6+1 hydrodynamic simulations with `AX-GADGET` (Nori & Baldi (2018)), consisting in three pairs of simulations (plus the ΛCDM case), one pair for each considered FDM mass, evolved either including or neglecting the effect of the QP in the non-linear dynamics – labelling these two cases as *FDM* and *FDMnoQP*, respectively. As already stated, the goal of this set is to assess and quantify the entity of the approximation typically employed in the literature. Each of the simulations follows the evolution of 512^3 DM particles in a comoving periodic box with side length of 15 Mpc/h. The initial conditions have been generated at redshift $z = 99$ using the `2LPTic` code (Crocce et al. (2006)), according to a random realisation of the suppressed linear power spectrum as calculated by `axionCAMB` (Hlozek et al. (2015)) for the different FDM masses under investigation. To ensure a coherent comparison between simulations, we used the same random phases to set up the initial conditions. In particular, the FDM masses m_χ considered here are $2.5 \cdot 10^{-22}$, $5 \cdot 10^{-22}$ and $2.5 \cdot 10^{-21}\mathrm{eV}/c^2$, in order to sample the mass range preferred by the first Lyman-α constraints in the literature (Iršič et al. (2017a); Armengaud et al. (2017); Kobayashi et al. (2017)).

Regarding the astrophysical parameters, we vary the IGM thermal history parameters in the form of the amplitude T_0 and the slope $\widetilde{\gamma}$ of the IGM temperature density relation, usually parametrised as $T = T_0(1 + \delta_{\mathrm{IGM}})^{\widetilde{\gamma}-1}$, with δ_{IGM} the IGM overdensity. The thermal parameters $(T_0, \widetilde{\gamma})$ were assumed to follow a power law redshift evolution (e.g., $T_0(z) = T_0^A(1+z)^{T_0^S}$), However, even if a more conservative temperature evolution with redshift was allowed in the MCMC runs, the MCMC constraints are expected to become weaker by only an order of unity (as was the case for $F = 1$ in the study by Iršič et al. (2017a)). The conservative approach allowed $T_0(z)$ to vary independently in each redshift bin, but prevented non-physical jumps in temperature (jumps with > 5000 K were not allowed between consecutive redshift bins). Regarding this aspect, see also Section 2.6.2 and Appendix C.

We also vary the timing of the instantaneous reionisation model z_{reio}. As in Chapter 2, three values for each of these parameters are considered, in the regime based on recent observational results. As we already discussed, the IGM thermal history is supposed to be the most important contaminant since a hotter medium in general tends to produce a smoother flux distribution and a flux power with less substructure at small scales, like nCDM models do. However, the redshift evolution of thermal and cosmological effects is very different and the wide range

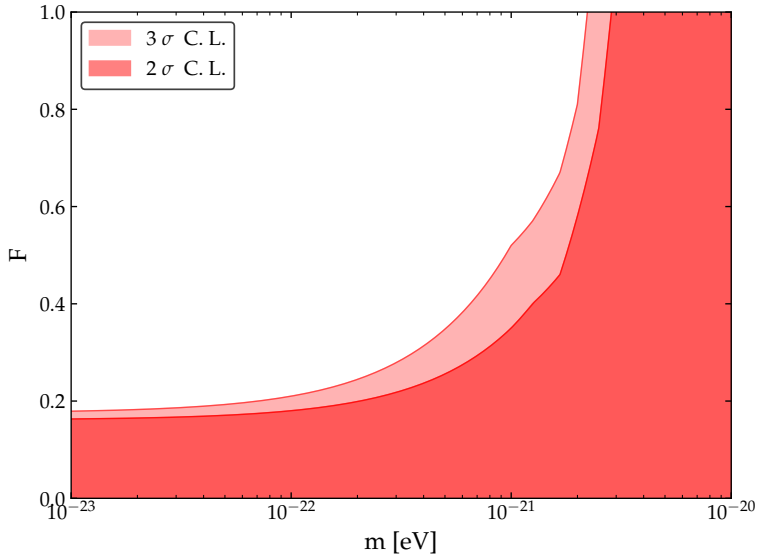

Figure 3.3: Constraints on the FDM mass m and fraction F of the total DM density in FDM obtained from Lyman-α forest data; the two different areas indicate 2 and 3 σ C.L. These results have been obtained for the reference combination of data sets described by Iršič et al. (2017a), with a physically motivated conservative prior on the IGM thermal evolution. The regime of $m < 10^{-22}$ eV has been extrapolated.

explored by our data allows to break the degeneracies between the parameters in a very effective way. We further consider ultraviolet (UV) fluctuations of the ionising background, which could be particularly important at high redshift, and build a refined power spectrum template that incorporates this effect. The amplitude of this effect is let free and is described by the parameter $\mathrm{f_{UV}} = [0, 1]$, which we marginalise over in the final constraints. All other cosmological parameters are fixed: $\Omega_\mathrm{m} = 0.317$, $\Omega_\Lambda = 0.683$, $\Omega_\mathrm{b} = 0.0492$ and $H_0 = 67.27$ km/s/Mpc, $A_s = 2.20652 \cdot 10^{-9}$ and $n_s = 0.9645$ (Ade et al. (2016a)).

3.3.1 Constraints on mixed (cold plus fuzzy) models

With the models of the flux power spectra obtained from the hydrodynamic simulations that we just described, we established a sparse grid of points in the parameter space and by using linear interpolation between the grid points we obtained predictions for the quantity $P_\mathrm{F}(k, z, \mathbf{p})$, with $\mathbf{p}$ a vector containing all the parameters described in the analysis, by performing a Taylor expansion for the desired models in a much finer grid for the highly multidimensional parameter space. We refer to Viel & Haehnelt (2006) for a more detailed description of the basic idea of this approach. We used an MCMC code in order to estimate the parameter constraints, and the results in terms of FDM mass and fraction have been obtained

by marginalising over the whole set of other parameters. The MCMC results for the mixed models are shown in Figure 3.3, using the reference analysis of Iršič et al. (2017a) which relies on all the data sets, and the assumption that the IGM thermal evolution follows a power law, with flat priors on the cosmological parameters.

When FDM constitutes the entire DM, i.e., $F = 1$, the Lyman-α forest data yields a lower bound on the FDM mass of $m \gtrsim 10^{-21}$ eV, as was also shown by Iršič et al. (2017a). On the other hand, the Lyman-α forest becomes insensitive to FDM at $F \lesssim 0.2$, which reflects the fact that the matter power spectrum is only mildly suppressed by FDM with such small fraction, no matter how light it is (as it was explicitly shown in Section 3.2.3 through analytic computations of the linear matter power spectrum). Although the regime of $m < 10^{-22}$ eV is only explored through extrapolation of our simulation models, we do not expect the bound in this regime to change significantly even with actual simulations, as the finite size of the error bars on the measured flux power spectrum makes it difficult to detect the mild suppression of the matter power.

Let us stress that the results presented in this Subsection have been obtained by completely neglecting the effect of QP during the non-linear evolution. However, this should not impact our results for FDM masses down to $m \sim 10^{-22}$ eV, as we explicitly demonstrate in Subsection 3.3.4 (see also, e.g., Iršič et al. (2017a); Schive et al. (2016)). On the other hand, we have just shown that for even smaller masses the FDM fraction becomes small, hence the QP is also expected to be negligible there. Anyhow, if the QP at the non-linear level is actually non-negligible, it should lead to further suppression of structure formation; hence the bounds presented here for the FDM parameters can be considered as conservative.

3.3.2 Comparison with the "area criterion"

We want now compare the Lyman-α forest constraints determined through the full MCMC analysis discussed in detail in Section 3.3.1, against the bounds which can be obtained by applying the simple and intuitive "area criterion", introduced in Chapter 2 (Section 2.5.2).

In order to use the "area criterion" for constraining the FDM mass m and abundance F, we calibrate the method by taking as references the 2 and 3 σ C.L. limits on the FDM mass where it constitutes all the DM content of the universe, namely the values given by the intercepts between the 2 and 3 σ C.L. contours with the $F = 1$ line in Figure 3.3. We can then compute the corresponding linear power spectra with `axionCAMB` (Hlozek et al. (2015)) and plug them into Equations (2.16) and (2.18). The resulting $\delta A_{\mathrm{REF}_{2\sigma}}$ and $\delta A_{\mathrm{REF}_{3\sigma}}$ are the estimate of the small-scale power suppression with respect to standard CDM, for FDM models that are excluded at 2 and 3 σ C.L., respectively. Thereafter, we build a grid in the $\{m, F\}$-space,

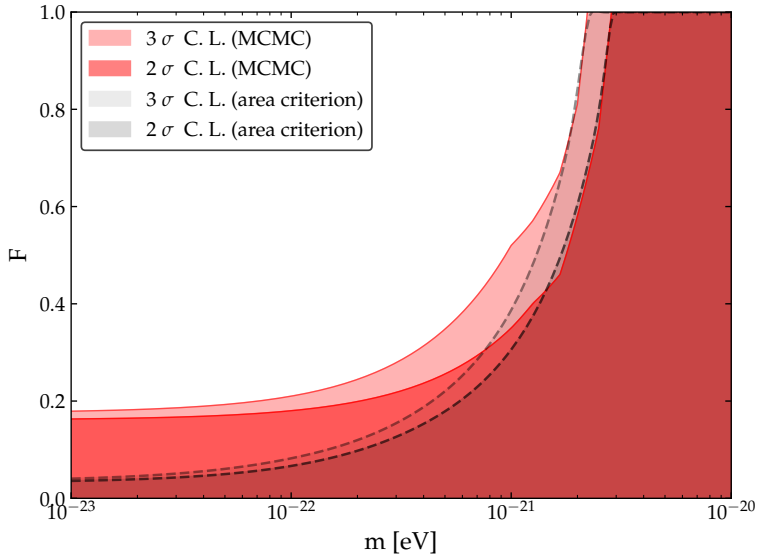

Figure 3.4: Here we compare the 2 and 3 σ C.L. limits from the MCMC analysis (red regions), with the 2 and 3 σ C.L. limits determined through the "area criterion" (superposed grey regions bounded by dashed lines). All the combinations of FDM mass and fraction which sample the region outside the grey contours are thus rejected by the "area criterion".

where each grid point is associated to a different combination of FDM mass and fraction, in order to compare all the corresponding δA with $\delta A_{\rm REF}$, and accept only those combinations which exhibit a smaller power suppression, i.e., $\delta A < \delta A_{\rm REF}$.

The results are reported in Figure 3.4, where the red contours refer to the 2 and 3 σ C.L. contours from the MCMC analysis, while the superposed grey areas bounded by dashed lines correspond to the 2 and 3 σ C.L. contours determined through the "area criterion". Hence, all the combinations of FDM mass and fraction which sample the region outside the grey contours are not allowed by the "area criterion".

Firstly we notice that for the case $F = 1$ the numbers returned by the area criterion are by definition in agreement with the more exact and comprehensive MCMC analysis discussed in Section 3.3.1. However, below FDM abundances of around 30%, the contours predicted by applying the "area criterion" clearly depart from the results of the full statistical analysis that we performed. This is due to the fact that when we apply the "area criterion" to models with small masses, even the practically negligible suppressions associated with small fractions (see Section 3.2.3) correspond to larger estimators δA with respect to the reference one. In other words, shifting the position of the cut-off at lower wavenumbers (i.e., investigating small FDM masses) unavoidably leads to a suppression which, quantified through the area estimator, is larger than the reference one. Consequently, although the Lyman-α bound is expected to be insensitive to FDM with small fractions, the bound obtained using the approximate area estimator improves towards smaller mass even in the regime $m \lesssim 10^{-22}$ eV.

3.3.3 Implications for the Milky Way satellites

In this Section we discuss the astrophysical implications of an ultra-light scalar DM scenario: the possibility, in this framework, of alleviating the "small-scale crisis" of the standard CDM paradigm. In particular, we focus on the well known *missing satellite* problem (Klypin et al. (1999); Moore et al. (1999)).

As we already discussed in Chapter 2, it is nowadays well established that DM models described by suppressed matter power spectra may be able to relax the tensions present in the standard CDM context at sub-galactic scales, i.e., the discrepancy between the observed number of dwarf galaxies within the MW virial radius and the number of MW substructures predicted by cosmological N-body simulations, assuming the standard CDM model. It is thus interesting to investigate the implications of a FDM scenario at sub-galactic scales, in order to check if the $\{m, F\}$-combinations which have been found to be in agreement with Lyman-α forest data are also capable of solving/alleviating such tension.

An accurate calculation of the number of substructures $N_{\rm sub}$ with FDM would require very high-resolution numerical simulations. Here we instead make a rough estimate of $N_{\rm sub}$, following the approach described in Section 2.5.1, i.e., using the analytic expression for the number of MW subhalos, proposed by Schneider (2015, 2016), given by Equation (2.13). Let us recall that the use of this procedure is supported by the analysis performed by Murgia et al. (2017), where the accuracy of the theoretical predictions has been checked against a large suite of N-body simulations (see Appendix B).

For a given MW halo mass, e.g., $M_{\rm halo} = 1.7 \cdot 10^{12}\ M_\odot/h$ (Lovell et al. (2014)), we can now obtain the number of subhalos $N_{\rm sub}$ with masses $M_{\rm sub} \geq 10^8\ M_\odot/h$ predicted by different parametrisations of the FDM scenario (i.e., by different combinations of m and F). This is done simply by integrating Equation (2.13). Note that a different choice of $M_{\rm halo}$ mainly leads to an overall shift in $N_{\rm sub}$ for all DM scenarios, as is seen from the $M_{\rm halo}$-dependence in Equation (2.13). Hence, instead of studying the absolute value of $N_{\rm sub}$, which is sensitive to the MW halo mass, we firstly would like to focus on the relative suppression of $N_{\rm sub}$, i.e., the ratio between $N_{\rm sub}$ for cases with and without FDM. As a benchmark value for the relative suppression of $N_{\rm sub}$, we take the thermal WDM with masses in the interval between 2 and 3 keV as reference models, and consider the mixed FDM and CDM scenarios that yield similar relative suppressions to be able to solve the missing satellite problem (Lovell et al. (2016, 2017)).

The parameter window for $\{m, F\}$ where the number of subhalos lies within the "solving" range is shown in Figure 3.5 as a cyan shaded area bounded by dashed lines. For reference, when using $M_{\rm halo} = 1.7 \cdot 10^{12}\ M_\odot/h$, the aforementioned computation gives the number of subhalos with CDM only as $N_{\rm sub} = 158$, while

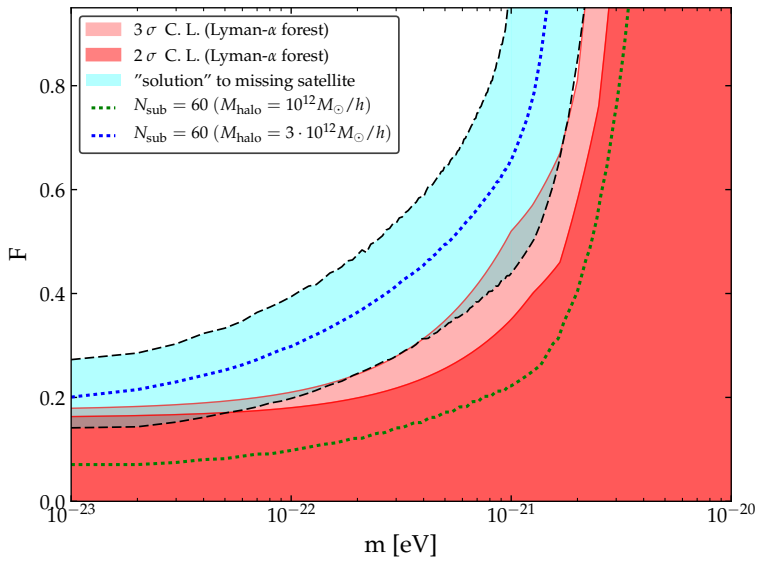

Figure 3.5: Comparison between the constraints on the FDM parameter space from the Lyman-α forest data analysis at 2 and 3σ C.L. (red regions, see Figure 3.3), and the region capable of "solving" the missing satellite problem (cyan region bounded by dashed lines). The green and blue dotted lines refer to models which predict $N_{\rm sub} = 60$, when choosing $M_{\rm halo} = 10^{12} M_\odot/h$ and $M_{\rm halo} = 3 \cdot 10^{12} M_\odot/h$, respectively.

$20 \leq N_{\rm sub} \leq 60$ for the reference WDM models. However, as we explained, these absolute values are irrelevant when focusing on the relative suppression of $N_{\rm sub}$. The red areas in Figure 3.5 represent the 2 and 3σ C.L. contours from the Lyman-α forest data analysis, discussed in Section 3.3.1. As one can easily see from the plot, there is very little room for simultaneously satisfying these constraints and solving the missing satellite problem.

Let us also discuss the effects of the observational uncertainties in the MW halo mass; for instance, a recent comprehensive dynamical analysis of redshifts and distances of 64 dwarf galaxies around the MW has led to $M_{\rm halo} = 2.8 \cdot 10^{12}\ M_\odot$ (Peebles (2017)). The detailed value of $M_{\rm halo}$ does matter when focusing instead on the absolute value of $N_{\rm sub}$, and past studies such as the analysis by Kennedy et al. (2014) pointed out the degeneracy between the MW halo mass and the DM parameters required for relaxing the missing satellite problem.

In order to take into account these issues, we have iterated the same analysis with different input values for $M_{\rm halo}$, and compared the corresponding predictions to a fixed satellite number $N_{\rm sub} = 60$; as in Chapter 2, this value is chosen as a sum of the 11 MW classical satellites and the 15 ultra-faint satellites from SDSS, with the latter value multiplied by a numerical factor which accounts for the limited sky coverage of the survey (Polisensky & Ricotti (2011); Schneider (2016); Murgia et al. (2017)). For $M_{\rm halo} = 1.7 \cdot 10^{12}\ M_\odot/h$, the number $N_{\rm sub} = 60$ is realised

at the lower boundary of the cyan band. The green and blue dotted lines in Figure 3.5 respectively indicate where $N_{\rm sub} = 60$ is realised for $M_{\rm halo} = 10^{12} M_\odot/h$ and $M_{\rm halo} = 3 \cdot 10^{12} M_\odot/h$; these values for the MW halo mass roughly correspond to the current observational limits (Cautun et al. (2014); Wang et al. (2015)).

For $M_{\rm halo} = 10^{12} M_\odot/h$, even in the pure CDM case the satellite number is as low as $N_{\rm sub} = 94$, indicating that the missing satellite problem itself is ameliorated if the MW halo mass takes a value close to its lower bound. Consequently, the green line lies inside the region allowed by the Lyman-α forest. On the other hand, a larger MW halo mass makes the problem worse ($N_{\rm sub} = 274$ for $M_{\rm halo} = 3 \cdot 10^{12} M_\odot/h$ with pure CDM), and thus further reduces the FDM parameter space for satisfying the Lyman-α forest constraint and solving the missing satellite problem at the same time. To summarise, unless the MW halo mass is close to its current lower bound and thus the satellite number is suppressed, the Lyman-α constraint leaves very little room for FDM to serve as a solution to the missing satellite problem. However, we should also remark that we have used rather simple analytic approximations for estimating the satellite number, hence it would be important to verify this conclusion with numerical simulations.

3.3.4 Quantum Potential in the non-linear evolution

This Section is devoted to a scrupulous comparison between the Lyman-α constraints previously discussed, obtained neglecting the time-integrated QP effect during the non-linear evolution, with the first constraints extracted from the same astrophysical observable accurately accounting for the full non-linear treatment of the QP (Nori et al. (2019)), thanks to the numerical module `AX-GADGET` (Nori & Baldi (2018)).

In Section 2.4 we gave a brief description of the SPH approach – which all `GADGET` modules rely on – that we exploited in order to perform all the hydrodynamic simulations used for our analyses. Let us remark that, with such method, physical observables can be expressed in many analytically equivalent forms that translate into different operative summations, carrying thereby different numerical errors. `AX-GADGET` features a new type of particle in the system – i.e., ultra-light-axion (ULA) – whose strongly non-linear quantum dynamics is solved through the following equations:

$$\nabla \rho_i = \sum_{j \in \mathrm{NN}(i)} m_j \nabla W_{ij} \frac{\rho_j - \rho_i}{\sqrt{\rho_i \rho_j}} \tag{3.26}$$

$$\nabla^2 \rho_i = \sum_{j \in \mathrm{NN}(i)} m_j \nabla^2 W_{ij} \frac{\rho_j - \rho_i}{\sqrt{\rho_i \rho_j}} - \frac{|\nabla \rho_i|^2}{\rho_i} \tag{3.27}$$

$$\nabla Q_i = \frac{\hbar^2}{2 m_\chi^2} \sum_{j \in \mathrm{NN}(i)} \frac{m_j}{f_j \rho_j} \nabla W_{ij} \left(\frac{\nabla^2 \rho_j}{2 \rho_j} - \frac{|\nabla \rho_j|^2}{4 \rho_j^2} \right). \tag{3.28}$$

Among the several strategies that have been employed in the literature to reduce the residual numerical errors (Brookshaw (1985); Cleary & Monaghan (1999); Colin et al. (2006)), Equation (3.26) has indeed proven to be the more stable and accurate for the QP case (for a comparison between different implementations, see Nori & Baldi (2018)).

The implementation of FDM physics in `AX-GADGET` includes the possibility to simulate universes with multiple CDM and FDM species, or FDM particles with self- or external interactions, due to the merging of the `AX-GADGET` module with the C-Gadget module for Coupled DM models (Baldi et al. (2010)). Furthermore, `AX-GADGET` inherits automatically all the large collection of physical implementations – ranging from gas cooling and star formation routines to DE and modified gravity implementations – that have been already developed for `GADGET-3`.

Among the different Lyman-α flux statistics that can be considered, we focused on the 1D flux power spectrum, and on the flux Probability Distribution Function (PDF), by extracting them from the three pairs of *FDM* and *FDMnoQP* simulations described in Section 3.3. Unless otherwise stated, we normalise the extracted flux arrays in order to have the same observed mean flux over the whole sample considered, and for all the simulations. In any case, we do find that the scaling factor for the optical depth arrays over the whole simulated volume is 1.6, 1.4 and 1.1 times higher than in the ΛCDM case in order to achieve the same mean flux for the $m_{22} = 2.5$, 5 and 25 FDM cases with negligible ($\sim 1 - 2\%$) differences between the *FDM* and *FDMnoQP* cases.

In Figure 3.6 we show the flux (*left*) and gas (*right*) PDF ratios between the simulations that include the QP and those that do not include it – *FDM* and *FDMnoQP*, respectively – at $z = 5.4$, one of the highest redshift bins in which Lyman-α data are available.

It is possible to see that there is a $2 - 6\%$ peak at flux $\sim 0.6 - 0.8$, i.e., in regions of low transmissivity that are expected to trace voids. The fact that *FDM* simulations display a more peaked PDF compared to *FDMnoQP* ones for this range of fluxes means that, on average, in those models it is more likely to sample such void environments. In fact, the different PDFs should reflect the underlying different gas PDFs at the same redshifts and along the same lines of sight. In the right panel of Figure 3.6, showing the corresponding gas PDF, it is indeed apparent that in models with QP the gas PDF is more skewed towards less dense regions, that are typically associated to high transmission. The effect due to the QP is thus to increase the volume filling factor of regions below the mean density with respect to the corresponding *FDMnoQP* case.

In Figure 3.7 we plot the percentage difference in terms of flux power spectrum at three different redshifts for the *FDM* simulations (i.e., with QP), compared both

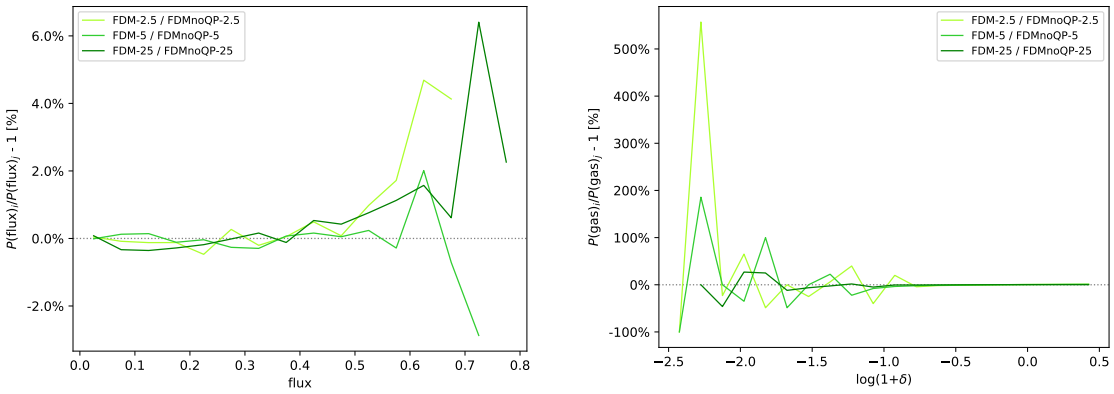

Figure 3.6: Relative differences of the flux PDF (*left*) and gas PDF (*right*) for FDM models with respect to their corresponding FDMnoQP counterparts, at redshift $z = 5.4$.

to pure CDM (right panels) and to the corresponding *FDMnoQP* case (left panels). The increase of power at $z = 5.4$ in the largest scales – compared to the standard CDM case – is due to the imposed normalisation at the same mean flux, while the evident suppression at small scales is related to the lack of structures at those scales. The comparison with the *FDMnoQP* set-ups, instead, reveals an additional suppression which is always below the 5% level for all the masses considered.

Since the Lyman-α constraint are calculated by weighting the contribution from all the scales, we expect the bound on m_{22} found by Iršič et al. (2017a) to change comparably to the additional suppression introduced, that in our case is about of $2 - 3\%$.

This is exactly what can be seen in in Figure 3.8, where the marginalised posterior distribution of m_χ obtained by our analysis (with QP) is plotted and compared with the results presented by Iršič et al. (2017a). The red line refers to our MCMC analysis, whereas the green line corresponds to the results obtained by Iršič et al. (2017a). The corresponding vertical lines show the marginalised 2 σ C.L. limit on the FDM mass. Such constraint on the FDM mass changes from $20.45 \cdot 10^{-22}$ eV to $21.08 \cdot 10^{-22}$ eV, which matches with our expectation and confirms that the universally adopted approximation of neglecting the QP dynamical effects is legitimate to investigate the Lyman-α typical scales. The agreement between the sets of results obtained with and without the dynamical QP implementation is evident and is not sensibly affected by varying the assumptions on the IGM thermal history. Such agreement implies that the non-linear evolution of the cosmic structures at intermediate scales, and the non-linear mapping between flux and density, effectively make up for the additional suppression introduced.

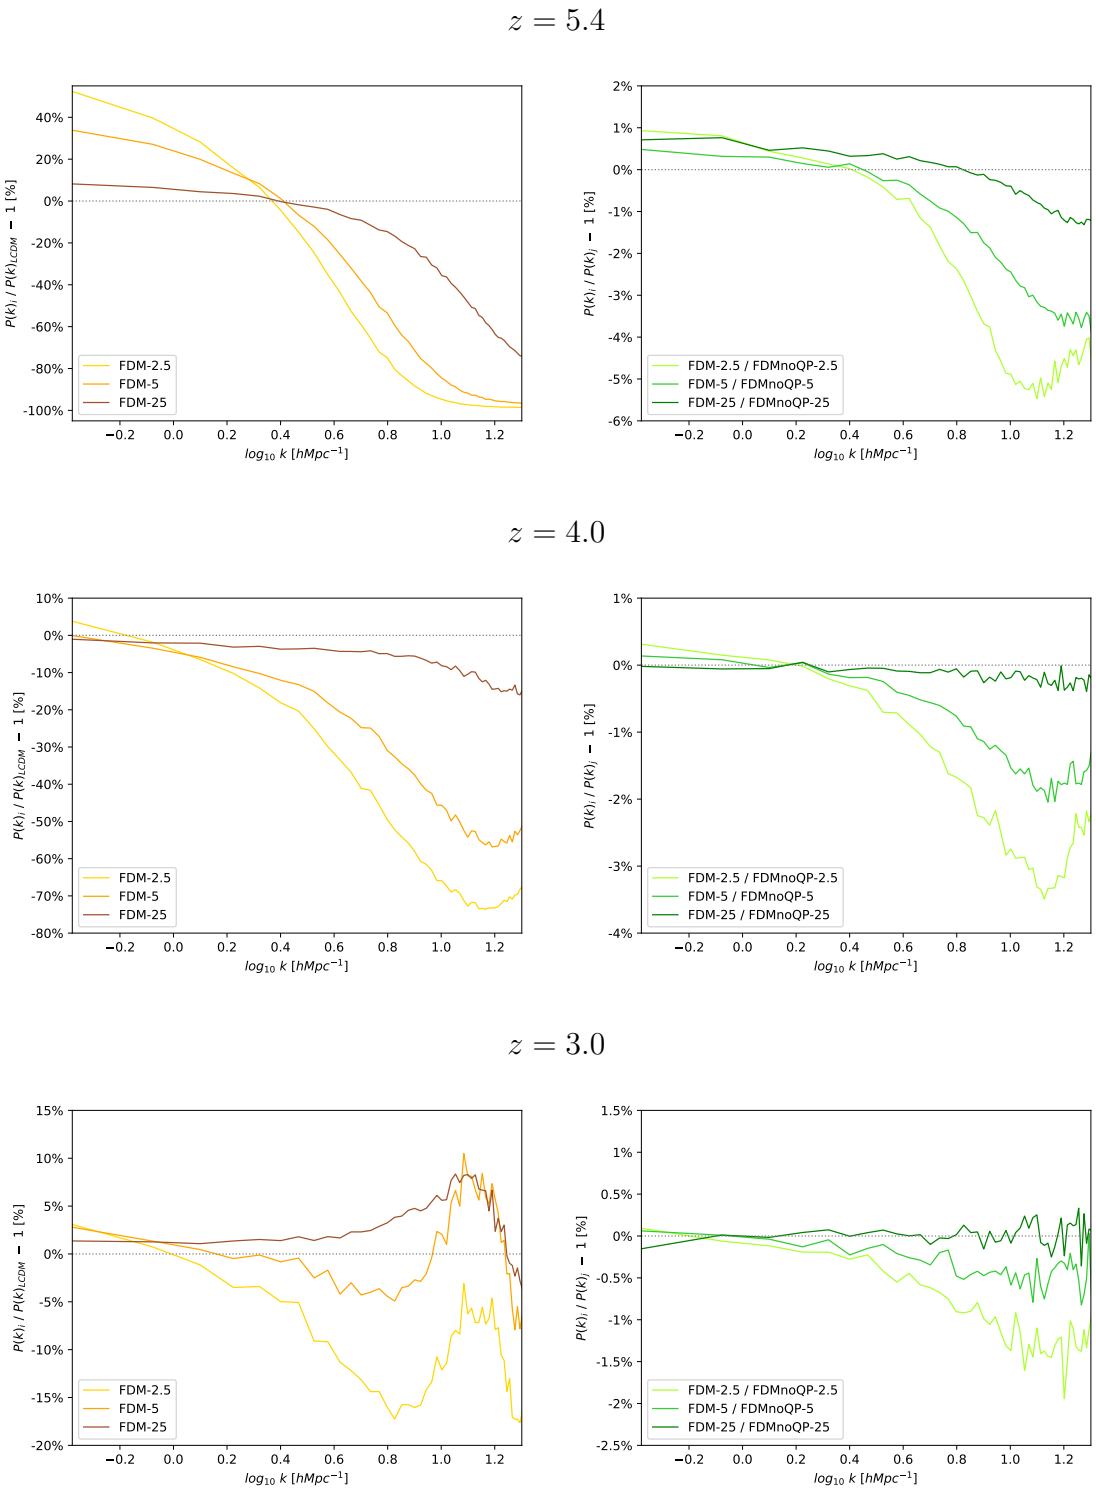

Figure 3.7: Flux power spectrum comparison between all simulations and ΛCDM (left panels), and between *FDM* simulation and their *FDMnoQP* counterparts (right panels) at different redshifts.

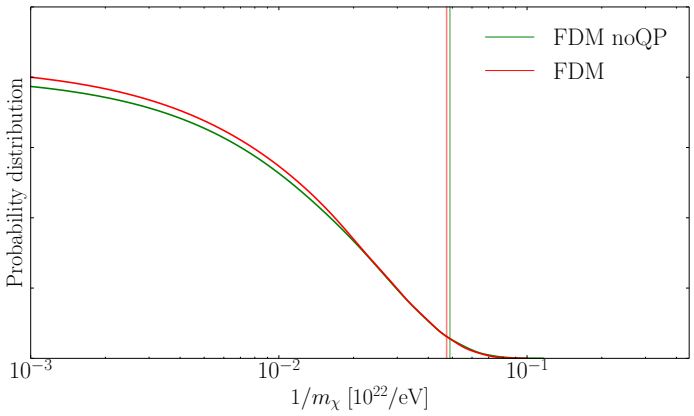

Figure 3.8: Here we plot the marginalised posterior distribution of $1/m_\chi$ from both the analyses performed by Iršič et al. (2017a) (green lines, without QP) and ours (red lines, with QP). The vertical lines stand for the marginalised 2 σ C.L. limits.

3.4 Cosmological implications

We now move on to discuss the implications of the Lyman-α constraints for the scalar field, and for the early universe. After evaluating the field range of the scalar (Subsection 3.4.1), we compute the isocurvature perturbations in the scalar DM density sourced during cosmic inflation. Combining with CMB constraints on DM isocurvature perturbations, we also derive bounds on the inflation scale (Subsection 3.4.2). All the results presented in this Section are based on the constraints reported in Section 3.3.1, hence they have been obtained neglecting the QP effects during the non-linear phase of structure formation and evolution. Nevertheless, as we explicitly showed in Section 3.3.4, such approximation is expected to have a marginal effect at the scales probed by the Lyman-α forest.

3.4.1 Initial displacement of the vacuum

A light scalar field stays frozen at its initial field value in the early universe. Hence, any initial displacement from the potential minimum gives rise to a FDM density in the later universe. We consider such an initial vacuum misalignment to be the main source of the density, and also suppose the FDM mass to be time-independent (unlike the QCD axion whose mass depends on the cosmic temperature). Then the scalar can collectively be described by a homogeneous Klein–Gordon equation in a FRW background universe,

$$\ddot{\phi} + 3H\dot{\phi} + m^2\phi = 0, \tag{3.29}$$

where an dot denotes a derivative in terms of the cosmic time t, and $H = \dot{a}/a$. The homogeneous scalar field forms a perfect fluid with an energy density and pressure of

$$\rho_\phi = \frac{1}{2}\left(\dot{\phi}^2 + m^2\phi^2\right), \quad p_\phi = \frac{1}{2}\left(\dot{\phi}^2 - m^2\phi^2\right). \tag{3.30}$$

We denote the initial displacement of the scalar field from its potential minimum by $\phi_\star$. In the early universe, when $H \gg m$, the scalar field is frozen at $\phi_\star$ due to the Hubble friction, and thus contributes to the vacuum energy. On the other hand in the later universe, when $H \ll m$, the scalar undergoes harmonic oscillations along the quadratic potential and behaves as pressureless matter. Thus the scalar densities in the two epochs are written as

$$\rho_\phi = \begin{cases} \frac{1}{2}m^2\phi_\star^2 & \text{when } H \gg m, \quad (3.31) \\ \frac{1}{2}m^2\phi_\star^2 \left(\frac{a_{\rm osc}}{a}\right)^3 & \text{when } H \ll m. \quad (3.32) \end{cases}$$

These asymptotic behaviours smoothly connect to each other at around $H \sim m$. Here $a_{\rm osc}$ represents the scale factor at the onset of the scalar oscillation; the explicit value of $a_{\rm osc}$ is chosen such that the scalar density in the asymptotic future matches with Equation (3.32). We also denote quantities measured at $a = a_{\rm osc}$ by the subscript "osc".

We are interested in ultra-light scalars that start oscillating in the radiation dominated epoch, instead of during times prior to reheating. The exact solution of the Klein–Gordon equation in a radiation dominated background is given in Appendix D, where the ratio between the mass and Hubble rate at $a = a_{\rm osc}$ is shown to take the value of

$$\frac{m^2}{H_{\rm osc}^2} = \left(\frac{8}{\pi}\right)^{4/3} \left[\Gamma\left(\frac{5}{4}\right)\right]^{8/3} \approx 2.68. \quad (3.33)$$

Since the total energy density, and hence the Hubble rate, of a radiation dominated universe are related to the cosmic temperature by

$$\rho_{\rm r} = 3M_p^2H^2 = \frac{\pi^2}{30}g_*T^4, \quad (3.34)$$

with $M_p = (8\pi G)^{-1/2}$ being the reduced Planck mass, the temperature at $a = a_{\rm osc}$ is obtained as

$$T_{\rm osc} \approx 0.5\,{\rm keV} \left(\frac{g_{*\rm osc}}{3.36}\right)^{-1/4} \left(\frac{m}{10^{-22}\,{\rm eV}}\right)^{1/2}. \quad (3.35)$$

Thus for instance, a scalar with $m = 10^{-22}\,{\rm eV}$ starts oscillating when the cosmic temperature drops to $T_{\rm osc} \approx 0.5\,{\rm keV}$. Moreover, the scalar would be oscillating at matter/radiation equality as long as $m \gg 10^{-28}\,{\rm eV}$.

Using that the entropy of the universe is conserved since the radiation dominated epoch, the scalar density today can be expressed in terms of the entropy density s as

$$\rho_{\phi 0} = \frac{1}{2}m^2\phi_\star^2 \frac{s_0}{s_{\rm osc}}, \quad (3.36)$$

where the subscript "0" denotes quantities in the present universe. Here, $s_{\rm osc}$ is written in terms of $H_{\rm osc}$ using Equation (3.34) as

$$s_{\rm osc} = \frac{2\pi^2}{45} g_{s*\rm osc} T_{\rm osc}^3 = \frac{2\pi^2}{45} g_{s*\rm osc} \left(\frac{90}{\pi^2} \frac{M_p^2 H_{\rm osc}^2}{g_{*\rm osc}} \right)^{3/4} . \tag{3.37}$$

Thus, the present FDM abundance is obtained by combining Equations (3.33), (3.36), and (3.37). Expressing it in terms of the ratio to the CDM density measured by Planck (Ade et al. (2016a)), $\Omega_c h^2 = 0.1186 \pm 0.0020$ (1 σ C.L., TT+lowP+lensing), one finds[4]

$$F \equiv \frac{\Omega_\phi}{\Omega_{\rm c}} \approx 0.6 \left(\frac{g_{*\rm osc}}{3.36} \right)^{3/4} \left(\frac{g_{s*\rm osc}}{3.91} \right)^{-1} \left(\frac{\phi_\star}{10^{17}\,\mathrm{GeV}} \right)^2 \left(\frac{m}{10^{-22}\,\mathrm{eV}} \right)^{1/2} . \tag{3.38}$$

Using this relation, the Lyman-α constraints on $\{m, F\}$, reported in Figure 3.3, are translated into bounds on the scalar parameters $\{m, \phi_\star\}$. In Figure 3.9 we plot the 2 and 3σ C.L. limits on $\{m, \phi_\star\}$ from Lyman-α forest data, where the shaded regions indicate the allowed parameter space. Here we have set $g_{*\rm osc} = 3.36$, $g_{s*\rm osc} = 3.91$, since $T_{\rm osc} \ll 1\,\mathrm{MeV}$ for the displayed masses (see Equation (3.35)). The dashed lines in Figure 3.9 indicate the contours of constant fraction F. We find that the initial displacement is bounded from above as $|\phi_\star| \lesssim 10^{16}\,\mathrm{GeV}$ for most values of the mass. For masses of $m \gtrsim 10^{-21}\,\mathrm{eV}$, $\phi_\star$ is constrained mainly by the requirement that the scalar should not lead to over-abundance of DM (i.e., $F \leq 1$). On the other hand for $m \lesssim 10^{-21}\,\mathrm{eV}$, the Lyman-$\alpha$ forest gives the strongest constraint. As one goes to even smaller masses $m \lesssim 10^{-22}\,\mathrm{eV}$, the Lyman-$\alpha$ bound on $\phi_\star$ weakens since the FDM density F decreases. There the Lyman-α bound closely follows the $F = 0.2$ contour, allowing $\phi_\star$ to take larger values.

3.4.2 Isocurvature perturbation and inflation scale

The ultra-light scalar acquires super-horizon field fluctuations during cosmic inflation.[5] As a consequence, the initial field value $\phi_\star$ possesses fluctuations with a power spectrum of

$$\mathcal{P}_{\delta\phi_\star}(k) = \left(\frac{H_k}{2\pi} \right)^2 , \tag{3.39}$$

[4]Note that we defined F as the density ratio between FDM and total DM (i.e., FDM plus CDM). Here we are identifying the measured CDM density with the total DM density.

[5]This is not necessarily the case if the scalar arises after inflation as a pseudo Nambu–Goldstone boson of a spontaneously broken global U(1) symmetry. However in such cases, topological defects are produced, which would over-close the universe (unless the number of degenerate vacua along the bottom of the Mexican hat potential is $N = 1$) (Vilenkin & Everett (1982); Sikivie (1982); Linde & Lyth (1990)). For a recent discussion, see also Visinelli (2017). Thus in this Section we suppose the scalar field to have existed already during inflation.

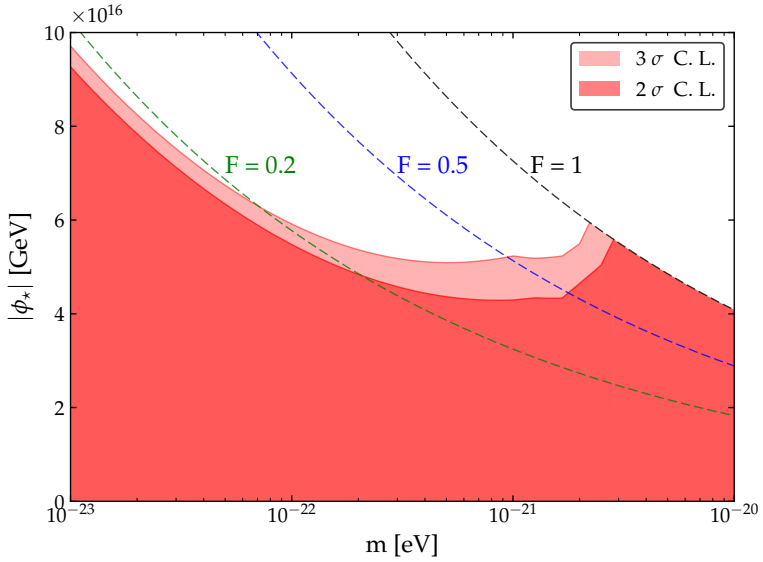

Figure 3.9: Upper limits on the initial displacement of the scalar field $\phi_\star$ from Lyman-α forest data (2 and 3 σ C.L.) as a function of the FDM mass m. The coloured dashed lines show contours of constant FDM fraction F.

where H_k represents the Hubble rate during inflation when the comoving wavenumber k exits the horizon. Since the scalar does not dominate the universe until matter/radiation equality, its field fluctuations lead to isocurvature perturbations.

From Equation (3.36) the FDM density depends on the initial field value as $\rho_\phi \propto \phi_\star^2$, therefore the density fluctuates as

$$\frac{\delta\rho_\phi}{\rho_\phi} = 2\frac{\delta\phi_\star}{\phi_\star}, \tag{3.40}$$

up to linear order in the field fluctuations. Identifying this with the FDM isocurvature perturbation $S_{\phi\gamma}$ using Equation (3.39), and further multiplying with the DM fraction yields the effective CDM isocurvature power spectrum,

$$\mathcal{P}_{c\gamma}(k) = F^2\mathcal{P}_{\phi\gamma}(k) = \left(\frac{FH_k}{\pi\phi_\star}\right)^2. \tag{3.41}$$

Given that the Hubble rate during inflation is nearly constant, the isocurvature spectrum is nearly scale-invariant. Moreover, the scalar ϕ does not contribute to curvature perturbations and hence there is no correlation between the isocurvature and curvature perturbations.

Since the FDM compatible with the Lyman-α analysis behaves similarly to CDM on large scales, the CMB constraints on CDM isocurvature perturbations

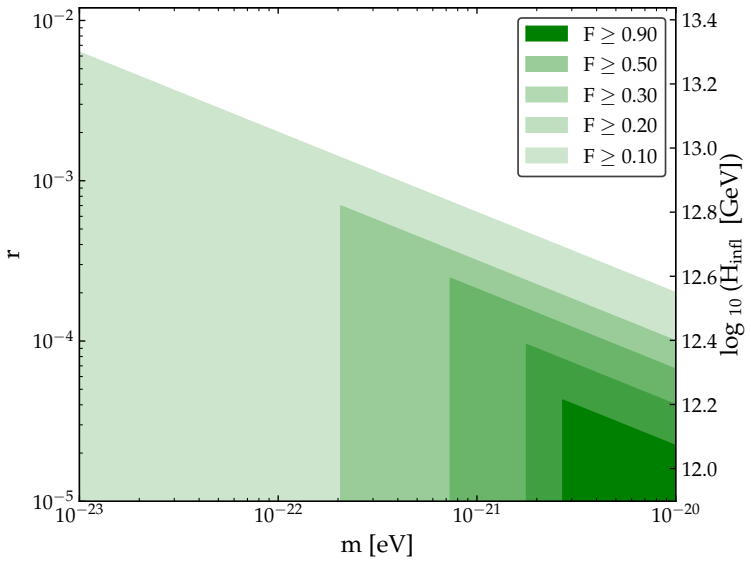

Figure 3.10: Upper bound on the inflation scale $H_{\rm inf}$ and tensor-to-scalar ratio r at the pivot scale $k_{\rm piv}/a_0 = 0.05\,{\rm Mpc}^{-1}$, as a function of the FDM mass m (2σ C.L.). Differently coloured regions represent the allowed parameter space when the FDM constitutes a certain fraction F of the total DM.

also apply to FDM. Parametrising the isocurvature power spectrum in terms of the curvature power as

$$\mathcal{P}_{c\gamma}(k) = \frac{\beta_{\rm iso}(k)}{1-\beta_{\rm iso}(k)}\mathcal{P}_\zeta(k), \tag{3.42}$$

uncorrelated and scale-invariant CDM isocurvature is constrained by Planck (Ade et al. (2016c)) at the pivot scale $k_{\rm piv}/a_0 = 0.05\,{\rm Mpc}^{-1}$ as

$$\beta_{\rm iso}(k_{\rm piv}) < 0.038 \quad (95\%\ \text{C.L., TT, TE, EE+lowP}), \tag{3.43}$$

with $\mathcal{P}_\zeta(k_{\rm piv}) \approx 2.2\cdot 10^{-9}$.

The CMB upper bound on the isocurvature translates into a bound on the inflation scale through Equation (3.41); eliminating $\phi_\star$ using Equation (3.38), we obtain an upper limit on the Hubble rate when the pivot scale leaves the horizon as

$$H_{k_{\rm piv}} < 4\cdot 10^{12}\,{\rm GeV}\left(\frac{g_{*\rm osc}}{3.36}\right)^{-3/8}\left(\frac{g_{s*\rm osc}}{3.91}\right)^{1/2} F^{-1/2}\left(\frac{m}{10^{-22}\,{\rm eV}}\right)^{-1/4}. \tag{3.44}$$

This can also be expressed as a bound on the tensor-to-scalar ratio,[6]

$$r(k) = \frac{\mathcal{P}_T(k)}{\mathcal{P}_\zeta(k)} = \frac{1}{\mathcal{P}_\zeta(k)}\frac{2H_k^2}{\pi^2 M_p^2}, \tag{3.45}$$

[6]Here we assume that the sound speed of the tensor fluctuations is unity.

as

$$r(k_{\rm piv}) < 2 \cdot 10^{-4} \left(\frac{g_{*\rm osc}}{3.36}\right)^{-3/4} \left(\frac{g_{s*\rm osc}}{3.91}\right) F^{-1} \left(\frac{m}{10^{-22}\,\rm eV}\right)^{-1/2} . \tag{3.46}$$

Alternatively, in terms of m and $\phi_\star$, the bound is written as:

$$r(k_{\rm piv}) < 4 \cdot 10^{-4} \left(\frac{g_{*\rm osc}}{3.36}\right)^{-3/2} \left(\frac{g_{s*\rm osc}}{3.91}\right)^{2} \left(\frac{m}{10^{-22}\,\rm eV}\right)^{-1} \left(\frac{\phi_\star}{10^{17}\,\rm GeV}\right)^{-2} . \tag{3.47}$$

These constraints become weaker for a smaller m. On the other hand, the Lyman-α forest sets a lower bound on m. Thus, by combining the Lyman-α and CMB constraints, an upper bound on the inflation scale can be obtained. This is presented in Figure 3.10, where each coloured region represents the values allowed for the FDM mass m and the tensor-to-scalar ratio r, or the inflation scale $H_{\rm inf}$, when the FDM constitutes a certain fraction F of the total DM. Here we combined the 2σ C.L. limit on FDM from the Lyman-α forest analysis (see Figure 3.3) with the CMB 2σ C.L. limit on isocurvature perturbations (i.e., Equation (3.46) with $g_{*\rm osc} = 3.36$, $g_{s*\rm osc} = 3.91$). The former sets the left boundaries of each region, and the latter the upper boundaries.[7] One clearly sees that FDM is incompatible with an observably large r, with the upper limits on r becoming stronger for a larger F. In particular if FDM constitutes more than 20% of the total DM, the tensor-to-scalar ratio would be as low as $r < 10^{-3}$. This in turn suggests that any detection of primordial Gravitational Waves (GWs) in the near future would rule out FDM produced from a vacuum misalignment as the main component of DM.[8]

We also illustrate this point in Figure 3.11, where contours of the upper bounds of r (Equation (3.47)) are shown on the $\{m,\, \phi_\star\}$-plane; regions above the contour would be excluded if r is detected at the displayed value. Figure 3.11 actually summarises all the main results discussed in this Section, showing the allowed values for the mass m and initial displacement of the scalar field $\phi_\star$. The field displacement is generically bounded as $|\phi_\star| \lesssim 10^{16}\,\rm GeV$; otherwise the scalar would either lead to too much DM in the universe, or suppress structure formation and contradict the Lyman-α forest measurements.[9] By combining the Lyman-α constraints with the CMB bounds on DM isocurvature perturbations, we further derived upper limits on

[7] Isocurvature perturbations can also impact the Lyman-α forest (Beltran et al. (2005)), thus for a rigorous treatment, the isocurvature should also be included in the Lyman-α analyses. However, since the FDM isocurvature is nearly scale-invariant, its effect on the Lyman-α should be tiny; hence here we simply combine the result of Section 3.3.1 with the CMB limit.

[8] However, we should also remark that there have been attempts to make light FDM consistent with high-scale inflation by adding further ingredients. One such example is an axion-like field with a time-dependent decay constant (Linde & Lyth (1990); Linde (1991); Higaki et al. (2014); Chun (2014); Fairbairn et al. (2015); Kobayashi & Takahashi (2016)).

[9] In Subsection 3.4.3 we show that if the scalar is an axion-like field, the bound on $\phi_\star$ corresponds to that on the product of the axion decay constant and the initial misalignment angle, $f_a\theta_\star$, when anharmonic effects are negligible.

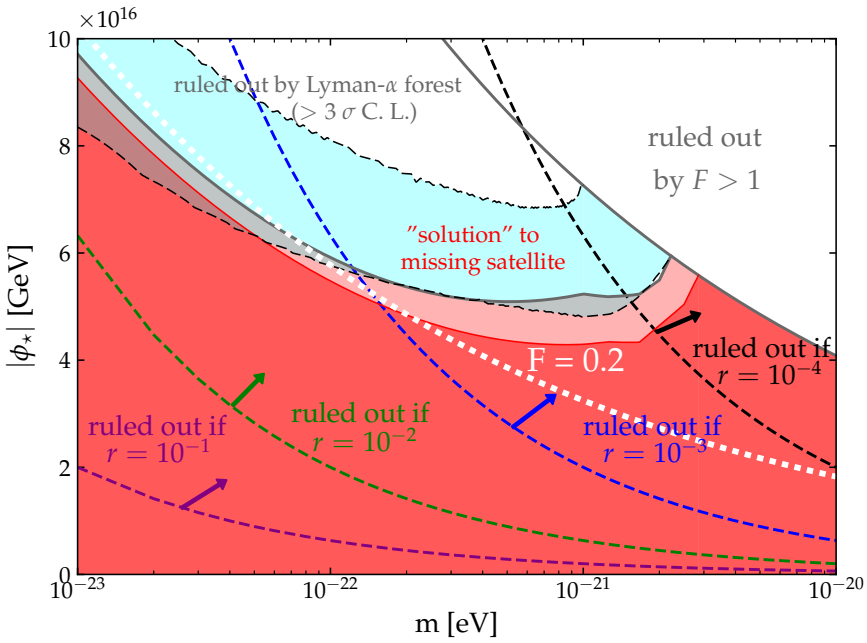

Figure 3.11: Summary of constraints on the scalar field mass m and the initial displacement $\phi_\star$ (for an axion-like field, $\phi_\star = f_a\theta_\star$). The 2 and 3σ C.L. regions allowed by Lyman-α forest data are shown in red. The upper-right corner is excluded by the over-abundance of DM. The dashed lines indicate the parameter regions that will be ruled out by a detection of a tensor-to-scalar ratio r. The cyan band bounded by dashed lines shows where the missing satellite problem can be solved. On the white dotted contour, FDM constitutes 20% of the total DM.

the scale of cosmic inflation in the presence of FDM. These are shown in Figure 3.11 as dashed lines, indicating the parameter regions that will be ruled out if primordial GWs are detected in the future. A dotted white line is also overlaid to indicate where the fraction of the DM in FDM is 20%; this value serves as the fraction threshold below which the Lyman-α forest becomes insensitive to the presence of FDM.

3.4.3 Comments on axion-like fields

Scalar fields with approximate continuous shift symmetries, often referred to as axion-like fields, have been studied as an ultra-light DM candidate (Arvanitaki et al. (2010); Ringwald (2012); Hui et al. (2017)). These models are typically described by an action with a periodic potential of

$$S = \int d^4x\sqrt{-g}\left[-\frac{1}{2}f_a^2 g^{\mu\nu}\partial_\mu\theta\partial_\nu\theta - m^2 f_a^2\left(1-\cos\theta\right)\right], \tag{3.48}$$

with f_a being an "axion decay constant". When focusing on the vicinity of a minimum $\theta = 0$, the potential is expanded as

$$V(\theta) = m^2 f_a^2\left(1-\cos\theta\right) = \frac{1}{2}m^2 f_a^2\theta^2 + \mathcal{O}(\theta)^4. \tag{3.49}$$

Thus our analyses apply to axion-like fields by replacing ϕ with a product of the decay constant and the angle, i.e.,

$$\phi \to f_a\,\theta. \tag{3.50}$$

In particular, if the initial misalignment angle $\theta_\star$ is of order unity, then the bounds on $\phi_\star$ directly translates into bounds on the decay constant f_a.

However, we should also remark that the discussions are modified if the initial angle is as large as $|\theta_\star| > 1$; then the expansion (Equation (3.49)) breaks down and the axion potential can no longer be approximated as quadratic. The non-quadratic nature of the axion potential would lead to (*i*) an increase in the final axion DM density due to a delayed onset of the scalar oscillation (Turner (1986); Bae et al. (2008)), and (*ii*) a significant enhancement of the axion isocurvature due to a non-uniform onset of the oscillation, giving much stronger bounds on the inflation scale (Lyth (1992); Strobl & Weiler (1994); Kobayashi et al. (2013)). We also note that, besides such anharmonic initial conditions, a time-dependent m (e.g., of the QCD axion) or f_a (e.g., Linde & Lyth (1990); Linde (1991); Higaki et al. (2014); Chun (2014); Fairbairn et al. (2015); Kobayashi & Takahashi (2016)) might also modify the cosmological evolution of the axion.

3.5 Non-linear structure characterisation

This Section is dedicated to an extended analysis of the statistical and structural properties of the DM halos in FDM cosmologies, with and without QP.

By exploiting the large statistical sample at our disposal, we have extracted valuable information about how FDM affects, among others, the halo mass function as well as the shape and density distribution of DM halos. To this end, we performed a set of 6+1 DM-only simulations, evolved down to $z = 0$, and we have used it to characterise the small-scale structures at low redshift. As before, this set consists

Model	QP	m_χ $[10^{-22}\mathrm{eV}/c^2]$	N halos	N genuine halos	M_{CUT} $[10^{10}M_\odot]$
ΛCDM	×	-	57666	56842	-
FDM-25	✓	25	25051	13387	0.04056
FDM-5	✓	5	10058	2736	0.1645
FDM-2.5	✓	2.5	8504	1301	0.3151
FDMnoQP-25	×	25	25432	13571	0.04056
FDMnoQP-5	×	5	10376	2856	0.1645
FDMnoQP-2.5	×	2.5	8819	1374	0.3151

Table 3.1: Summary of the properties of the set of simulations used for structure characterisation (see the text for the details).

in three pairs of simulations (plus the ΛCDM case), one pair for each considered FDM mass, evolved either including or neglecting the effect of QP in the dynamics (once again, labelling these two cases as *FDM* and *FDMnoQP*, respectively). All simulations follow the evolution of 512^3 DM particles in a comoving periodic box with side length of 15 Mpc/h. The mass resolution is $2.2124 \cdot 10^6 M_\odot$. A summary of their properties can be found in Table 3.1.

The structure of this Section is the following: in Subsection 3.5.1 and 3.5.2 we outline the procedures to deal with numerical fragmentation, and to match halos across different simulations, respectively; in Subsection 3.5.3 we present and discuss our results.

3.5.1 Numerical fragmentation

For cosmological models whose LSS properties depart sensibly from ΛCDM only at small scales, as FDM models, the thorough analysis of the statistical overall properties and the specific inner structures of halos represents the most relevant and often largely unexploited source of information. In N-body simulations, this implies the use of a suitable clustering algorithm to build a halo catalogue in order to identify gravitationally bound structures that can then be studied in their inner structural properties.

As in Chapter 2, in this Section we rely on the `SUBFIND` routine (Springel et al. (2001a)), already implemented in `GADGET-3`, namely a two step halo-finder which combines a FoF algorithm (Davis et al. (1985)) to find particle clusters – that defines the primary structures of our halo sample – with an unbinding procedure to identify gravitationally bound substructures within the primary halos. Hereafter, we use the terms *primary structures* to identify the substructures of each FoF group containing the most gravitationally bound particle, *subhalos* for the non-primary structures and *halos* when we generally consider the whole collection of structures found.

As we already discussed in Chapter 2, a long-standing problem that affects N-body simulations, when characterised by a sharp and resolved cut-off of the matter power spectrum, has to be taken into account in the process of building a reliable halo sample. This is the *artificial clumping*, often referred to as *numerical fragmentation*, i.e., the formation of artificial low-mass spurious clumps within filaments (see, e.g., Wang & White (2007); Schneider et al. (2012); Lovell et al. (2014); Angulo et al. (2013); Schive et al. (2016)).

While it has been initially debated whether the nature of such fragmentation was to be considered physical or numerical, the detailed analysis by Wang & White (2007) showed that in nCDM simulations (as, e.g., Bode et al. (2001)) characterised by a highly suppressed matter power spectrum, the formation of low-mass subhalos

was resolution-dependent, and related to the large difference between force resolution and mean particle separation (as already suggested by Melott & Shandarin (1989)).

To identify spurious halos in simulations, and select a clean sample to study and characterise the FDM structures halos, we take cue from the procedure outlined by Lovell et al. (2014): in particular, we use the mass at low redshift and the spatial distribution of particles as traced back in the initial conditions as proxies for the artificial nature of halos, as described below.

In fact, the more the initial power spectrum is suppressed at small scales, the more neighbouring particles are coherently homogeneously distributed, thus facilitating the onset of artificially bounded and small ensembles that eventually outnumber the physical ones. As already shown by Wang & White (2007), the dimensionless power spectrum peak scale k_{peak} and the resolution of the simulation – expressed by the mean inter-particle distance d – can be related together to get the empirical estimate

$$M_{\mathrm{lim}} = 10.1 \ \rho_b \ d \ / \ k_{\mathrm{peak}}^2 \tag{3.51}$$

describing the mass at which most of the halos have a numerical rather than a physical origin. In Lovell et al. (2014), this mass is used as a pivotal value for the mass M_{CUT} used to discriminate genuine and spurious halos – lying above and below such threshold, respectively – which is set as $M_{\mathrm{CUT}} = 0.5 M_{\mathrm{lim}}$.

In addition to the mass discriminating criterion, Lovell et al. (2014) showed that particles that generate spurious halos belong to degenerate regions in the initial conditions and are more likely to lie within filaments, stating that the reconstructed shape of the halo particles ensemble in the initial conditions can be used to identify spurious structures. N-body initial conditions are generally designed as regularly distributed particles on a grid from which are displaced in order to match the desired initial power spectrum. Hence, numerical fragmentation originates mostly from particles lying in small planar configurations, belonging to the same row/column domain, or a few adjacent ones.

Therefore, we need a method to quantitatively describe the shape of subhalos and of the distribution of their member particles once traced back to the initial conditions of the simulation. To this end, we resort to the inertia tensor of the particle ensemble

$$I_{ij} = \sum_{\mathrm{particles}} m \ (\hat{e}_i \cdot \hat{e}_j) \ |r|^2 - (\mathbf{r} \cdot \hat{e}_i) \ (\mathbf{r} \cdot \hat{e}_j) \tag{3.52}$$

where m and $\mathbf{r}$ are the particle mass and position, respectively, and $\hat{e}$ are the unit vectors of the reference orthonormal base. The eigenvalues and the eigenvectors of the inertia tensor represent the square moduli and unitary vectors of the three axes of the equivalent triaxial ellipsoid with uniform mass distribution. We define as $a \geq b \geq c$ the moduli of the three axes, and the sphericity, $s = c/a$, as the

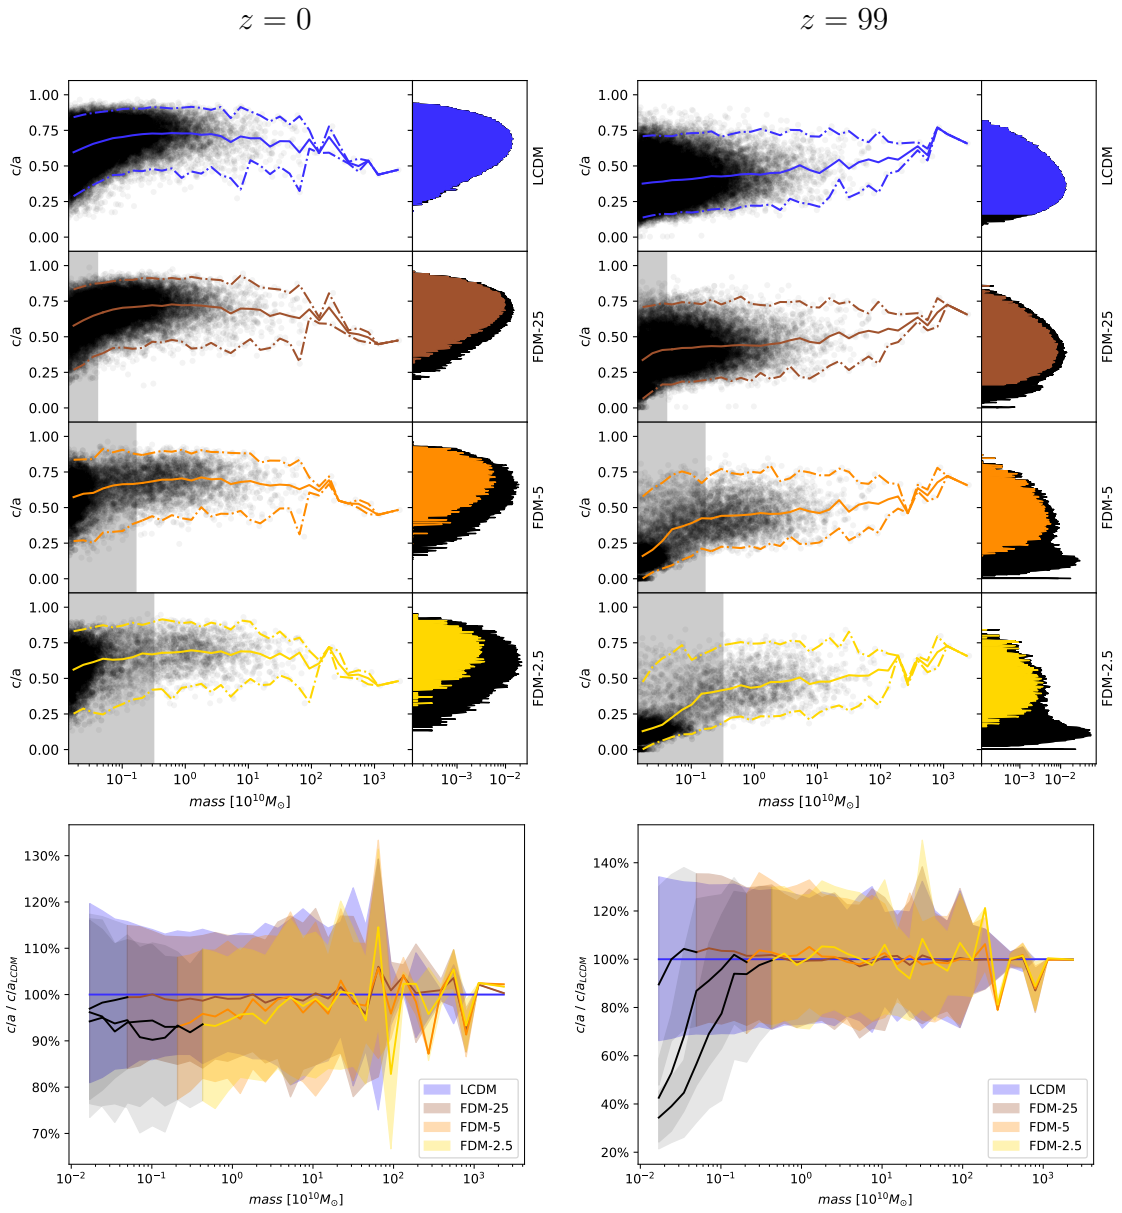

Figure 3.12: Sphericities of all DM particle ensembles as found by `SUBFIND`, as function of their mass (upper panels), at redshifts $z = 0$ and $z = 99$ (*left* and *right* panels, respectively). The black shaded area represents the discarded region below the different mass cuts $M_{\rm CUT}$, corresponding to each model. Each black dot represents a subhalo and the solid (dot-dashed) lines describe the median (99-th percentile) of the total distribution, which are all gathered and contrasted with ΛCDM in the lower panel. The total sphericity distribution – integrated in mass – is represented in the side panels, where the contribution of the discarded sample to medians and distributions are portrayed in black. Lower panels feature the median of the mass-sphericity distributions, presented as the ratio with respect to ΛCDM. The shaded areas, corresponding to the ± 1 σ of the distribution, are colour coded as in the upper panels. The blackened median and shaded areas represent the excluded portion of the sphericity distributions below the corresponding $M_{\rm CUT}$.

ratio between the minor and the major ones: a very low sphericity will characterise the typical degenerate domains of numerical fragmentation. For these reasons, we use the combined information carried by the mass and the sphericity in the initial condition to clean the halo catalogues from spurious ones by applying independent cuts on both quantities, as it is be detailed below.

In Figure 3.12, the mass-sphericity distributions of the different simulations are plotted at redshift $z = 0$ (upper left panel) and at $z = 99$ (upper right panel) where each point represents a halo identified by `SUBFIND`, without applying any selection. Solid and dot-dashed lines denote the median and the 99-th percentile of the distribution; in the side panels we display the cumulative distributions, where the contribution of spurious halos is highlighted in black.

By looking at the two panels, it is possible to notice that the total cumulative sphericity distribution at low redshift is fairly model-independent, so that distinguishing spurious halos from genuine ones is impossible. However, if we trace the particle ensembles of each halo found at $z = 0$ back to the initial conditions at $z = 99$, using particles ID, and we study the resulting reconstructed mass-sphericity relation, the anomalous component of the distribution associated with spurious halos clearly emerges as a low-sphericity peak, which is more pronounced for smaller values of the FDM particle mass.

In fact, as the mass m_χ decreases, the smoothing action of QP becomes more efficient, inducing homogeneity at larger and larger scales in the initial conditions and increasing, consequently, the contamination of numerical fragmentation. It clearly appears that the population of halos in the initial conditions is homogeneously distributed in a pure CDM scenario, while a bimodal structure emerges at lower and lower FDM mass. In particular, an increasing number of halos are located in a small region characterised by low mass ($M \lesssim 10^9 M_\odot$) and low sphericity ($s \lesssim 0.20$).

As there is no theoretical reason why QP should favour the collapse of ensembles with very low sphericities in the initial conditions with respect to the ΛCDM case, we consider this second population as the result of numerical fragmentation.

As in Lovell et al. (2014), we chose to compute $M_{\rm CUT} = 0.5 M_{\rm lim}$ using Equation (3.51) – one $M_{\rm CUT}$ for each value of the FDM mass, as reported in Table 3.1 –, that defines the upper bound of the discarded mass regions, i.e., the black shaded areas in all panels of Figure 3.12. It is interesting to notice that the masses $M_{\rm CUT}$ appear to be very close to the values at which the sphericity medians of the simulation sample – in the initial conditions – depart from the ones of ΛCDM, as can be seen in the lower right panel of Figure 3.12. As the $M_{\rm CUT}$-values we obtain are slightly larger compared to these departing values, we confirm the choice of the former over the latter, as a more conservative option for the mass thresholds dividing spurious from genuine halos.

m_χ $[10^{-22}\mathrm{eV}/c^2]$	N halos	N genuine halos
25	101.6%	101.4%
5	103.5%	104.4%
2.5	103.1%	105.6%

Table 3.2: The total and genuine number of halos, presented as the ratio between the simulations neglecting (*FDMnoQP*) and considering (*FDM*) the QP dynamical effects.

In Lovell et al. (2014), the selection in terms of initial sphericity was operationally performed discarding every halo with sphericity lower than $s_{\mathrm{CUT}} = 0.16$, equal to the 99-th percentile of the distribution of halos with more than 100 particles in the ΛCDM simulation. In our set of simulations, a similar value denotes the 99-th percentile as measured at the M_{CUT} mass in each simulations, so we adopt it as our own threshold in sphericity. Let us stress that the halos that are discarded through sphericity selection in the initial conditions have sphericities at $z = 0$ that are statistically consistent with the genuine sample, making their numerical origin impossible to notice based only on the sphericity distribution at $z = 0$. However, the mass constraint is far more rigid than the sphericity one in all models but ΛCDM, where no mass limit is imposed.

Finally, in Table 3.2 we summarise the comparison of the number of halos in the *FDMnoQP* set-up with respect to the corresponding *FDM* set-up, presented as the ratio of the total number of halos found by SUBFIND and the number of genuine halos remaining after the exclusion of spurious ones. It is possible to see that in the *FDMnoQP* simulations, for the three FDM masses considered, the total number of halos is overestimated by a factor $\sim 2.5\%$ on average while the genuine halos excess becomes more important as the FDM mass decreases, up to 5.6% for $m_{22} = 2.5$. This means that neglecting the QP effects during the simulation leads to the formation of halos which are not present when the full QP dynamics is taken into account and that, using our *à la* Lovell et al. (2014) spurious detection selection, such halos pass the numerical fragmentation test and contaminate any halo statistical property characterisation.

3.5.2 Inter-simulations halo matching

In FDM models, as we have extensively discussed, not only the initial power spectrum of matter perturbation is suppressed at small scales, thereby preventing the formation of low-mass structures, but the dynamical evolution of density perturbations changes due to the effect of QP, intimately affecting the development of structures during the whole cosmological evolution by opposing gravitational collapse. The implementation of such effect in `AX-GADGET` breaks the one-to-one correspondence of the spatial position of collapsed structures in simulations with

different FDM masses – especially for smaller objects –, despite the identical random phases used to set up the initial conditions.

We indeed expect bigger halos not to change dramatically their position at low redshift across different simulation, while this is not the case for lighter subhalos which are more affected by the evolving local non-linear balance between gravity and QP of the environment. This makes it more difficult to identify matching collapsed objects of common origin across the simulations, to study how FDM models affect the inner structure of halos on a halo-to-halo basis.

To this end, we devised an iterative matching procedure, to be repeated until no more couples are found, as the following: given a halo i at position $\mathbf{r}_i$ and total mass m_i in simulation A,

1. select all halos j belonging to simulation B as potential counterparts if $|\mathbf{r}_i - \mathbf{r}_j|/(a_i + a_j) < \tilde{R}$ where a_i and a_j are the major axes of the halos computed through the inertia tensor of all their member particles.

2. within the ensemble selected at the previous point, retain only the halos $k \subseteq j$ whose masses satisfy the condition $|m_i - m_k|/(m_i + m_k) < \tilde{M}$

3. if more than one halo $l \subseteq k$ is left, then choose the one for which $|\mathbf{r}_i - \mathbf{r}_l|/(a_i + a_l)$ is minimum.

4. after having considered all the halos in A, if more than one are linked to the same halo l belonging to B, choose the couple (i, l) that minimises $[|\mathbf{r}_i - \mathbf{r}_l|/(a_i + a_l)]^2 + [|m_i - m_l|/(m_i + m_l)]^2$, in order give the same weight to the two criteria.

This method is flexible enough to account for the shift in mass and position we expect from simulations with different FDM mass models, yet conservative enough to ensure the common origin of the subhalo couples. Moreover, using the combination of position and mass filters, we are able to discriminate couples in all mass ranges: position filtering is a weaker constraint in the case of bigger halos – since they occupy a big portion of a simulation – where instead the mass filter is very strict; vice versa, it is more powerful for smaller halos for which the mass filter select a large number of candidates.

Operatively, we used the previous procedure to match halos in each simulation with the ΛCDM one and we refer to the subset of halos that share the same ΛCDM companion across all the simulations as the *common sample*.

For geometrical reasons, we set the limit value for $\tilde{R}$ to be 0.5: this represents the case in which two halos with the same major axis a have centres separated exactly by the same amount a. The configurations that are selected by point (i) are the ones for which the distance between the halo centres is less or equal the smallest

$\tilde{M}$	m_{min}/m_{max}	N matches
1/39	95%	53
1/19	90%	162
3/37	85%	234
1/9	80%	279
1/7	75%	304
1/3	50%	346
3/5	25%	361
1	0%	389

Table 3.3: Number of common matches across ΛCDM and FDMs simulations, using different values of the parameter $\tilde{M}$ representing the minimum allowed ratio between the minimum and maximum masses for each candidate couple.

major axis between the two. A higher value for $\tilde{R}$ would include genuine small halos that have been more subject to dynamical QP drifting but would also result in a spurious match of bigger halos. For these reasons, we adopt $\tilde{R} = 0.5$, checking that the selected sample gains or loses $\sim 5\%$ of components if values 0.45 and 0.55 are used, without modifying the overall statistical properties of the sample itself.

With respect to $\tilde{M}$ at point (ii), instead, we applied the matching algorithm using several values, each denoting a specific threshold of the minimal value allowed for the mass ratio of halo couples in order to be consider as a match. As reported in Table 3.3, more than 60% of all the matching halos across ΛCDM and FDM simulations without mass selection – $\tilde{M} = 1$ case – have a mass ratio in the $100-85\%$ ratio range and almost 80% in the $100 - 75\%$ range. In order not to spoil our matching catalogue, especially with very close but highly different in mass halo couples, we chose the limiting value of $\tilde{M} = 1/7$.

3.5.3 Results and discussion

The statistical properties of the genuine halos belonging to each simulation are summarised in Figure 3.13, where we display the cumulative halo mass function (top right panel), the halo mass outside R_{200} (top left panel) – where R_{200} identifies the distance from the halo centre where the density is 200 times the critical density of the universe and M_{200} the mass contained within a R_{200} radius sphere –, the subhalo mass function (bottom left panel), and the subhalo radial distribution (bottom right panel).

In order to highlight the impact of numerical fragmentation and simplify the comparison of the different models to ΛCDM, relative ratios are displayed in the bottom panels, and shaded lines represent the distribution of the full halo sample, i.e., including also spurious halos.

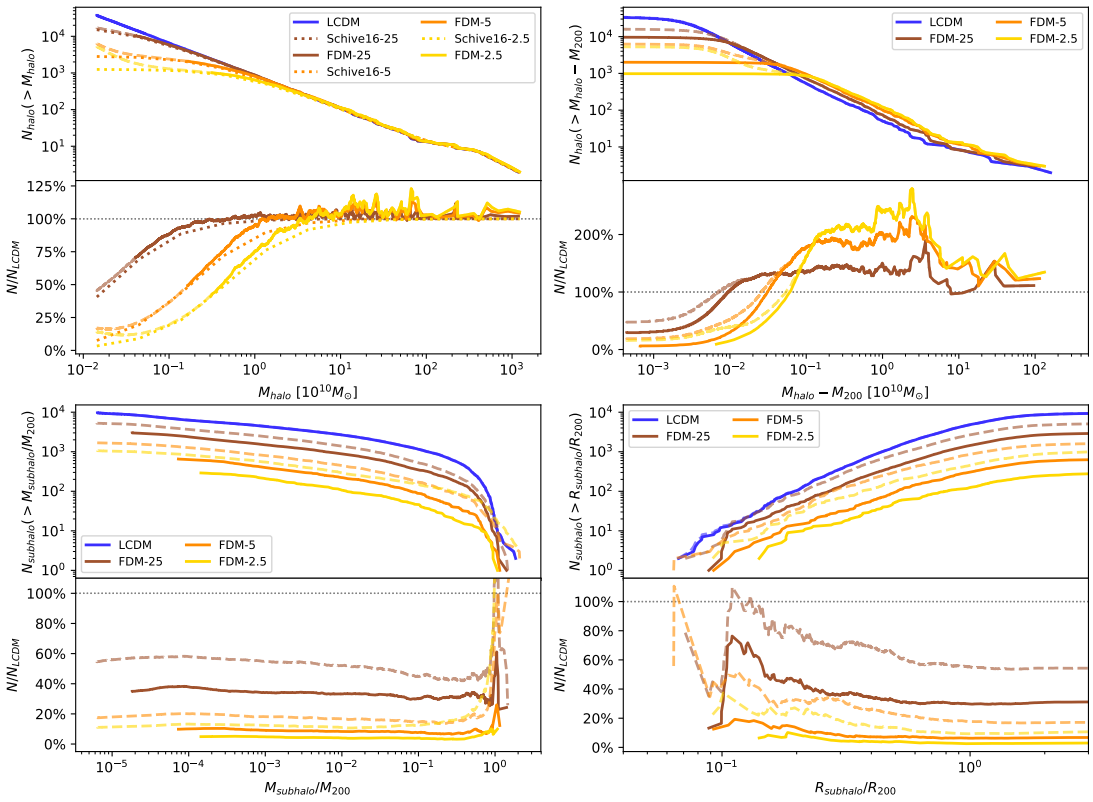

Figure 3.13: Properties of the halo and subhalo samples at $z = 0$, with (dashed lines) and without (solid lines) including the halos marked as spurious as described in Section 3.5.1. In particular, the cumulative distributions of halo mass (top left panel), the halo mass outside R_{200} (top right panel), the subhalo-halo relative mass (bottom left panel) and the subhalo-halo distance (bottom right panel) are displayed. The fitting functions of the cumulative halo mass distribution from Schive et al. (2016) (Equation (3.53)) are plotted for reference (dotted line in the top left panel).

The analytic fit used by Schive et al. (2016) to parametrise the cumulative halo mass function drop of the FDM models with respect to ΛCDM

$$N(>M)_{\rm FDM} = \int_M^{+\infty} \partial_M N_{\rm CDM} \left[1 + \left(\frac{M}{M_0}\right)^{-1.1}\right]^{-2.2} dM \tag{3.53}$$

with $M_0 = 1.6 \cdot 10^{10} m_{22}^{-4/3}\ M_\odot$, are plotted as reference – one for each FDM mass – in the top left panel of Figure 3.13 (dotted lines).

As expected, we find that the number of low-mass subhalos is drastically reduced in the FDM models, and the cumulative distributions depart from ΛCDM at higher and higher masses, as the m_χ mass decreases. The values at which the drop occurs are approximately $5 \cdot 10^{10} M_\odot$, $2.5 \cdot 10^{10} M_\odot$ and $5 \cdot 10^9 M_\odot$ for values of m_{22} of 2.5, 5 and 25, respectively: this suggest a linear trend of the threshold mass

$$M_t \simeq 5 \cdot 10^{10} M_\odot \left(\frac{2.5}{m_{22}}\right) \tag{3.54}$$

describing the approximate mass below which the number of halos starts decreasing with respect to ΛCDM.

Looking at the distribution of subhalo masses as compared to their associated primary halo M_{200} and the radial distribution to R_{200}, it is evident how the numerous small subhalos in ΛCDM, far from the gravitational centre of the main halo, are the ones that were not able to form in a FDM universe.

The halos that have masses above M_t not only have been able to survive the disrupting QP action up to redshift $z = 0$, but the cumulative distribution shows how they also gained extra mass, at the smallest (sub)halo expenses. This is confirmed by the cumulative distribution of the primary structures $N(> M_{\rm tot} - M_{200})$, representing the mass accumulated outside the R_{200} radius, which is systematically higher with respect to the ΛCDM case as the FDM mass lowers – up to peaks of 200% ratio for the lowest m_{22} –: this is consistent with the picture of bigger primary halos accreting the mass of uncollapsed smaller subhalos that did not form.

The fitting function expressed by Equation (3.53) is consistent with the scale of the drop of the halo mass function, which is indeed expected to be almost redshift-independent, since it is predominantly given by the initial cut-off in the power spectrum (Hu et al. 2000). However, it fails to reproduce the data on two levels: on one hand it does not recover the slope of the cumulative distribution – especially in the mass range close to M_t where the halo mass function departs from ΛCDM – and, on the other hand, does not account for the mass transfer from smaller halos, unable to collapse due to QP repulsive interaction, to bigger ones, that accrete the more abundant available matter from their surroundings. The discrepancies between the Schive et al. (2016) fitting function and our results are probably due to the fact that the former is based on simulations with approximated FDM dynamics and evolved only to redshifts $z = 4$, thus representing a different collection of halos that are, moreover, in an earlier stage of evolution.

Therefore, the analysis of the aggregated data of cumulative distributions of genuine halos in each simulation lead us to conclude that formation, evolution and properties of a FDM halo subject to the real effect of QP – as compared to the *FDMnoQP* approximation – can follow three general paths depending on its own mass and on the mass of the FDM boson: if the halo mass is $M \ll M_t$, there is high chance that the halo does not form at all since gravitational collapse is prevented by QP; if $M \gtrsim M_t$, the halo can be massive enough to form but its properties will be affected by QP – especially on its internal structure, as we will see below –, while for $M \gg M_t$ the halo is not severely affected by QP, and will simply accrete more easily uncollapsed mass available in its surroundings.

In order to study in more detail the impact of QP on the halo properties and structures, we divided our common sample, that by construction collects the halos

across all the simulations that share the same ΛCDM match (as described in detail in Section 3.5.2), in three contiguous mass ranges. Let us remind that matching halos have similar but not necessarily equal mass, so mass intervals are to be referred to the ΛCDM halo mass; the other matching halos belonging to the FDM simulations are free to have lower and higher mass, compatibly with the limit imposed by the $\tilde{M}$ parameter of the common sample selection procedure. The common sample low-mass end is clearly limited by the FDM-2.5 model, since it is the one with higher M_t, below which halos have statistically lower chance to form. The three mass ranges are $[0.5-4], [4-100], [100-4000] \cdot 10^{10} M_\odot$, in order to be compatible with the three halo categories described in the previous paragraph for the FDM-2.5 model, being $M_t(m_{22}=2.5) \sim 5 \cdot 10^{10} M_\odot$

For all the matching halos considered, we have tested the sphericity distribution, the halo volume and the total halo mass with respect to ΛCDM, as well as the radial density profiles.

Properties of inter-simulation matching halos are gathered in Figure 3.14, where the total sample is divided column-wise in the three mass ranges. The sphericity, the volume occupied and the total mass of the halos – contrasted with the corresponding ΛCDM match – are shown in the first row (left panels), together with related distribution functions (right panels). The second and the third row represent the overall density profiles, stacked in fractional spherical shells of R_{200} and ellipsoidal shells of the major axis a – identified with the vertical dashed lines –, respectively. Density profiles are divided by the value of the density calculated within the R_{200}- and a-shells, and are shown both in absolute value (top panels) and relatively to ΛCDM (bottom panels).

The sphericity distributions confirm that, in the mass range considered, there is no statistical deviation from ΛCDM, except for a mild deviation towards less spherical configurations of the less massive halos, especially in the $m_{22}=2.5$ model. This is consistent with the analysis of the sphericity distributions of the genuine samples (see lower panels in Figure 3.12) that reveals that halos appear to be statistically less spherical with respect to ΛCDM at $z=0$ when lower FDM masses are considered, down to a maximum of $\sim 10\%$ decrease in sphericity for $m_{22}=2.5$ and halo mass of $\sim 5 \cdot 10^9 M_\odot$.

For all FDM models, the volume occupied by the halos is systematically larger, consistently with a delayed dynamical collapse of the halos. All mass ranges show such property and it is emphasised by lower m_{22} mass – i.e., stronger QP force –; however, while bigger halos occupy almost systematically 20% more volume for $m_{22}=2.5$, smaller halos can reach even twice the volume occupied by their ΛCDM counterparts, when the same model is considered.

Comparing the mass of the halos in the various models with the one in ΛCDM, it is possible to see that small halos are less massive and big ones, on the contrary,

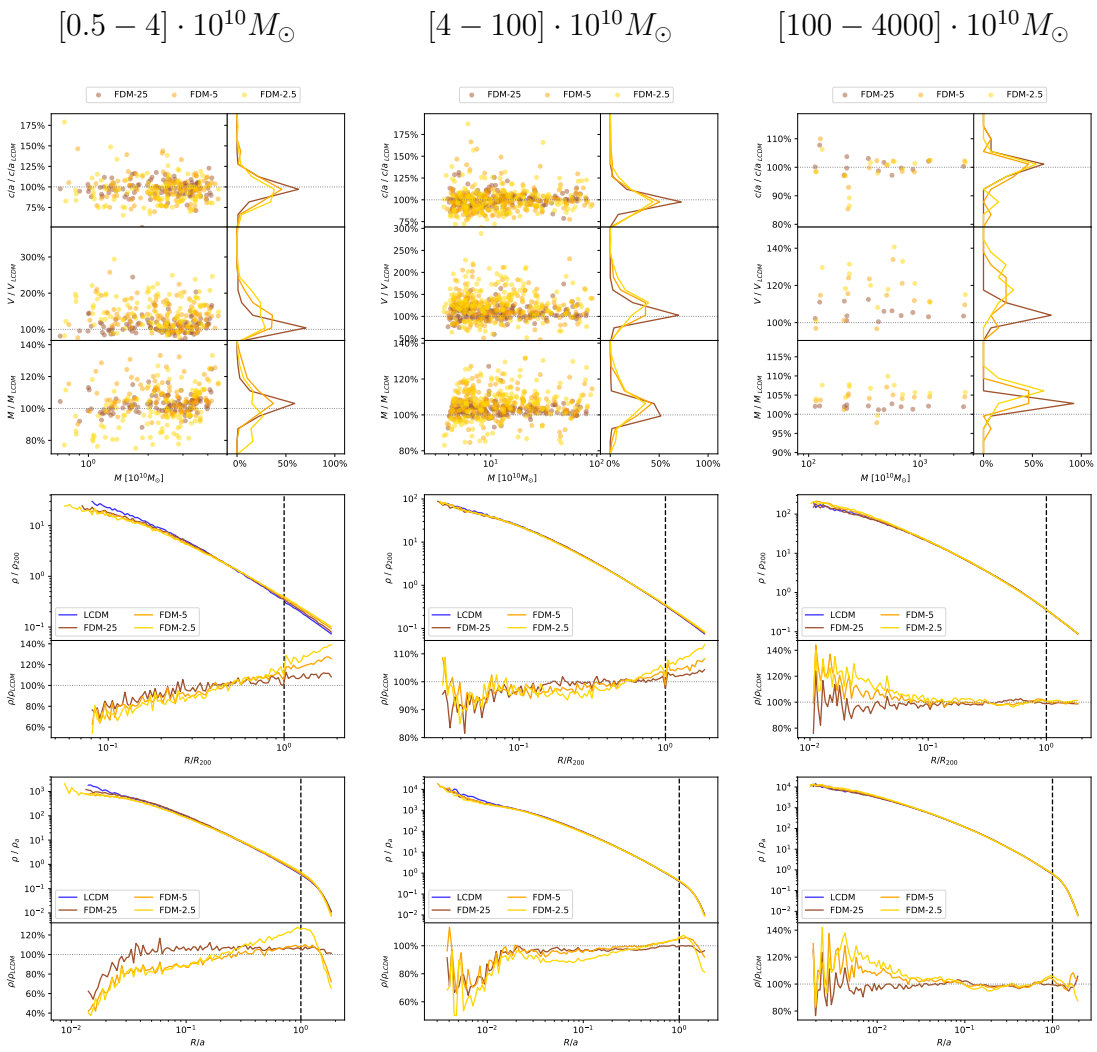

Figure 3.14: Properties of inter-simulation matching halos. The total sample is divided column-wise in three mass ranges. The sphericity, the volume occupied and the total mass of the halos – contrasted with the corresponding ΛCDM match – are shown in the first row (left panels), together with related distribution functions (right panels). The second and third row represent the overall density profiles, stacked in fractional spherical shells of R_{200} and ellipsoidal shells of the major axis a – identified with the vertical dashed lines –, respectively. Density profiles are divided by the value of the density calculated within R_{200} and a, and are shown both in absolute value (top panels) and relatively to ΛCDM (bottom panels).

become even more massive, confirming our hypothesis of mass transfer from substructures towards main structures.

The stacked density profiles provide even more insight on the underlying different behaviour between the chosen mass ranges. Starting from the less massive one, the stacked profiles look very differently if plotted using the spherical R_{200}-based or the ellipsoidal a-based binning. This is due to two concurrent reasons related to the properties of this mass range: first of all, as we said before, the sphericity is m_χ-dependent, and thus it is not constant with respect to ΛCDM, so that the geometrical difference in the bin shape becomes important when different models are considered; secondly, since the FDM halos have lower mass but occupy larger volumes, the two lengths are different from each other – being R_{200} related to density and a purely to geometry – so that the actual volume sampled is different. Nevertheless, it is possible to see that in FDM models there is an excess of mass in the outskirts of the halo – seemingly peaking exactly at distance a – and less mass in the centre.

The intermediate mass range shows also a suppression in the innermost regions but a less pronounced overdensity around a as expected, since the effectiveness of the repulsive force induced by the QP in tilting the density distribution decreases as its typical scale becomes a smaller fraction of the size of the considered objects. In fact, stacked density profiles of the most massive halos are very similar in the two binning strategies, being $R_{200} \sim a$ and sphericity constant among the various models, and consistent with no major deviation from ΛCDM, except for a central overdensity. It is our opinion, however, that such feature in the very centre of most massive halos could be a numerical artefact, since its extension is comparable with the spatial resolution used.

The results presented in this Section have been obtained through the detailed analysis of the statistical properties of halos found at $z = 0$ in the *FDM* simulations. The same analysis, repeated at $z = 0$, of the *FDMnoQP* simulations yield very similar results which are, therefore, not shown. Such consistency suggests that the properties of halos at low redshift are – at the investigated scales – not sensible to modifications induced by the dynamical QP repulsive effect, which are expected to appear more prominently at scales of ~ 1 kpc with the formation of solitonic cores. Higher resolution simulations will soon allow us to explore such even smaller scales.

This Chapter is mainly based on:

M. Archidiacono, D. C. Hooper, R. Murgia, S. Bohr, J. Lesgourgues, M. Viel
Constraining Dark Matter - Dark Radiation interactions with CMB, BAO, and Lyman-α
Submitted to JCAP; arXiv:1907.01496

4 Interacting Dark Matter

4.1 Overview

A class of nCDM models that has gained a lot of interest in recent years is that in which DM couples to an additional relativistic dark sector, known as Dark Radiation (DR) (Bringmann et al. 2016; Buckley et al. 2014; Boddy et al. 2014; Buckley & Fox 2010). Some of these models have been proposed to solve the missing satellite problem (Archidiacono et al. 2017; Vogelsberger et al. 2016; Schewtschenko et al. 2016) or to delay reionisation (Das et al. 2018) by means of a cut-off in the matter power spectrum. Other classes of DM-DR models, with a smooth damping of the matter power spectrum, can alleviate the cosmological tensions on H_0 and/or σ_8, as proposed by Buen-Abad et al. (2015); Lesgourgues et al. (2016); Buen-Abad et al. (2018).

Given the suppression that these interacting DM-DR models can have on small-scale structure growth, the matter power spectrum is an essential tool to study these dark sector interactions (Gluscevic et al. 2019). The main goal of this Chapter is to apply the novel parametrisation for suppressed matter power spectra, described in Chapter 2, to test DM-DR interaction scenarios with small-scale structure formation data. As it has been already discussed, in the framework presented in Chapter 2, the full shape of the cut-off is captured by a three-parameter analytic fitting formula. High-resolution simulations are performed on a grid of nodes given by several combinations of these parameters, allowing one to interpolate at a later stage on the pre-existing grid to derive constraints on a specific model.

To this end, we developed a new likelihood for the parameter inference code `MontePython` (Audren et al. 2013; Brinckmann & Lesgourgues 2018), making use of the method discussed in Chapter 2. This has been used together with the

implementation in the Boltzmann code `CLASS` (Blas et al. 2011) of the generic ETHOS parametrisation (Cyr-Racine et al. 2016) for DM-DR interactions, also introduced in Chapter 2 (Section 2.3.5).

Note that there have been several studies in the past that used the Lyman-α forest data to constrain DM-DR and DM-baryons interactions. In particular Dvorkin et al. (2014); Krall et al. (2017); Xu et al. (2018) used measurements of the linear matter power spectrum amplitude and slope obtained from SDSS-II low-resolution low-signal-to-noise quasar spectra. However, these measurements were obtained assuming a vanilla ΛCDM cosmology or small departures from it (McDonald et al. 2006, 2005). More recently, Garny et al. (2018) compared interacting dark sector models to the 1D flux power spectrum derived from SDSS-III data. They proposed a new modelling of the flux power spectrum in which the non-linear evolution is calculated analytically using viscous two-loop perturbation theory (Blas et al. 2015), while uncertainties on the flux power spectrum modelling are accounted for by marginalising over several nuisance parameters. None of these investigations rely on a forward modelling of the flux power spectrum based on high-resolution hydrodynamic simulations, and this is the approach that we discuss here.

This Chapter is organised as follows: in Section 4.2 we present the interacting DM-DR model, we describe its impact on structure formation and cosmological observables, we then describe the implementation of our Lyman-α likelihood (Subsection 4.2.1), as well as our modifications to the `CLASS` code (Subsection 4.2.2); in Section 4.3 we present and discuss our results; in Section 4.4 we briefly describe a further generalisation of our methods, and its future applications.

4.2 Theoretical framework and methods

In order to solve the missing satellite problem, besides self-interactions, DM has to scatter off a relativistic particle. The Standard Model particles (neutrinos and photons) cannot play this role both because of model building issues (Bringmann et al. 2016) and because of cosmological consequences (e.g., bounds on free streaming neutrinos (Archidiacono & Hannestad 2014; Brust et al. 2017)). Therefore, we need to invoke the existence of an extra DR component, which requires an extension of the Standard Model of particles. Here, we will not focus on one specific particle model, because our aim is to devise a general phenomenological approach that can be applied to several models.

In order to keep the discussion as general as possible, we assume the ETHOS parametrisation of DM-DR interactions (Cyr-Racine et al. 2016), already introduced in Chapter 2 (Section 2.3.5). The ETHOS framework was first implemented in `CLASS` by Archidiacono et al. (2017) (see Section 4.2.2 for more details), where

the effects on cosmological observables at large scales due to DM-DR interactions through a massive mediator were investigated, determining that, given the ETHOS parametrisation, the impact of the specific particle physics model (e.g., the DM mass, the presence of DR self-interactions, the vector/scalar nature of the mediator (Binder et al. 2016)) have no (or negligible) impact. The only relevant physical quantities are:

- the temperature dependence of the comoving interaction rate $\Gamma_{\rm DR-DM} \propto T^n$, where $\Gamma_{\rm DR-DM}$ can be seen as the DR drag opacity, i.e., the scattering rate of DR off DM,
- the strength of the interaction $a_{\rm dark}$ ($\Gamma_{\rm DR-DM} \propto a_{\rm dark} T^n$),
- the amount of DR parametrised through the temperature ratio $\xi = T_{\rm DR}/T_\gamma$, where T_γ is the temperature of CMB photons,
- the nature of DR (i.e., free streaming or not).

The effects of DM-DR interactions on the CMB were already discussed in detail by Cyr-Racine et al. (2014) and Archidiacono et al. (2017). Here we summarise the most important differences between such models and a ΛCDM model with an equivalent number of extra neutrino-like particles $N_{\rm eff}$ (i.e., with the same background density of radiation):

- Non-free streaming DR: due to its self-interactions and/or its coupling with DM, DR does not lead to additional anisotropic stress, and thus, does not induce the damping and phase-shift of the CMB acoustic peaks that is typically expected in presence of additional relativistic degrees of freedom;
- Non-growing DM fluctuations: the momentum exchange between DM and DR particles reduces the growth rate of DM perturbations compared to the ΛCDM model; this can lead to a fast mode in the DM perturbation evolution (Weinberg 2002; Voruz et al. 2014) and thus to a gravitational coupling between DM and photons that suppresses the odd (compression) CMB peaks.

Since we want to derive Lyman-α bounds on DM-DR models, we are interested in the behaviour of the matter power spectrum $P(k)$ on small scales. As described by Buckley et al. (2014); Archidiacono et al. (2017), the effect of the coupling between DM and DR is twofold:

- The late kinetic decoupling induced by DM-DR scattering yields a collisional damping of the matter power spectrum. The damping translates into a cut-off in the halo mass function, possibly alleviating/solving the missing satellite problem.

- Besides the exponential damping (only apparently similar to thermal WDM), the opposite forces of DM gravitational clustering and DR relativistic pressure may lead to a series of so-called Dark Acoustic Oscillations (DAOs) typical of models of DM-DR interactions mediated by a new light mediator (Cyr-Racine et al. 2014).

A special comment has to be dedicated to a class of models discussed by Buen-Abad et al. (2015); Lesgourgues et al. (2016); Buen-Abad et al. (2018), like for instance Non-Abelian Dark Matter (NADM), in which the momentum transfer rate from DM to DR, related to the ETHOS rate by $\Gamma = -\Gamma_{\rm DM-DR}/a$, scales like a^{-2}. In this case the suppression of the matter power spectrum is smooth, as the temperature dependence of the interaction rate ($n = 0$ in the ETHOS parametrisation) is the same as the temperature dependence of the expansion rate during the radiation dominated epoch. Moreover, DR particles tend to have strong self-interactions caused either by their charge under the new gauge group of the dark sector, or by the fact that they are the gauge bosons of this group. These models are described by the parameters:

- $\Delta N_{\rm fluid} \equiv \rho_{\rm dr}/\rho_{1\nu}$, which gives the amount of self-interacting DR, parametrised as the effective number of extra neutrino families;
- $\Gamma_0 \equiv \Gamma\,(a/a_0)^2$, which gives the momentum transfer rate from DM to DR at redshift $z = 0$.

They can thus be described with the ETHOS parametrisation in the $n = 0$ case, provided that DR is treated as a perfect fluid.

4.2.1 The Lyman-α likelihood

In order to provide limits on the properties of interacting DM-DR scenarios from the Lyman-α forest, we have devised a new `MontePython` (Audren et al. 2013; Brinckmann & Lesgourgues 2018) likelihood, based on the general parametrisation introduced in Chapter 2. The corresponding data set is the HIRES/MIKE sample of quasar spectra, described in detail in Section 2.6.1.

The new likelihood takes advantage of the scheme illustrated in Chapter 2, which allows to interpolate between different cosmological models without the need of running dedicated numerical simulations. Such procedure relies in fact on the large set of pre-computed hydrodynamic simulations described in Section 2.4.2, and on the advanced interpolation method discussed in Section 2.6.1.

Let us just remark that the non-standard nature of the dark sector is accurately followed only at the linear level, i.e., its impact is assumed to be fully encoded in the suppressed initial power spectra produced by our modified version of `CLASS` and used

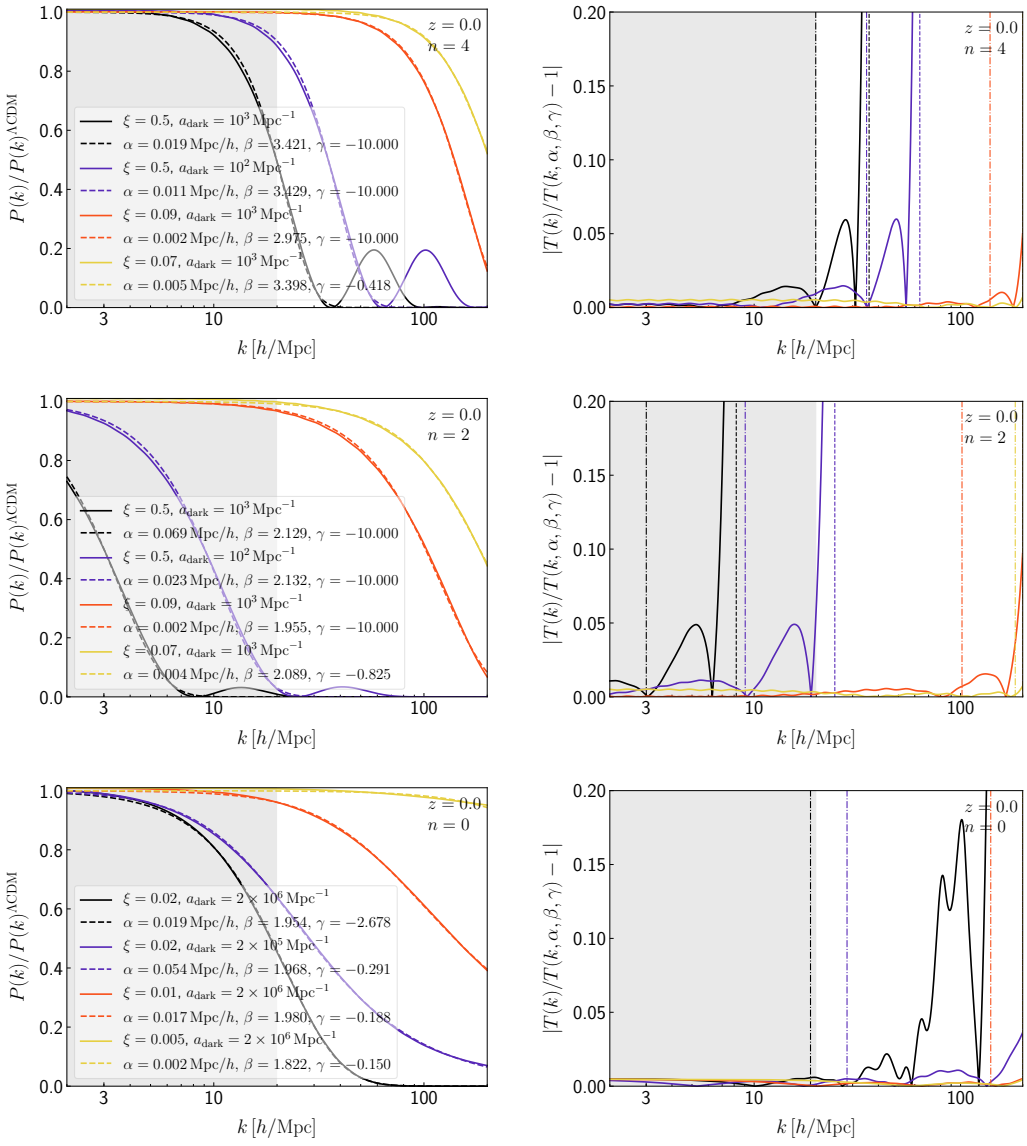

Figure 4.1: *Left:* Linear transfer functions $T^2(k) = P(k)/P(k)^{\Lambda\mathrm{CDM}}$ at $z = 0$, for $n = 4$ (top row), $n = 2$ (second row), $n = 0$ (bottom row). The different colours correspond to different values of the amount of dark radiation ξ and of the strength of the interaction a_{dark}. Solid lines depict the true $T^2(k)$, while dashed lines of the same colour show the corresponding $\{\alpha, \beta, \gamma\}$-fit. *Right:* Relative deviation of the $\{\alpha, \beta, \gamma\}$-fit from the true $T^2(k)$ (solid lines) for the same models (colours) as in the left panel. The vertical lines show $k_{1/2}$ (dot-dashed lines) and k_{fit} (dashed lines - for $n = 0$, $k_{\mathrm{fit}} = k_{\mathrm{max}}$). The grey shaded region approximately represents the k range probed by Lyman-α data.

as inputs for `2LPTic`; while during the non-linear structure evolution investigated by our numerical simulations, DM is treated as standard, pressureless CDM. The motivation for this treatment is twofold. First, DM-DR interactions have significant effects only at earlier times ($z > 99$) (Cyr-Racine et al. 2016). Second, DM self-interactions – which are expected to be relevant at late times – can also be safely

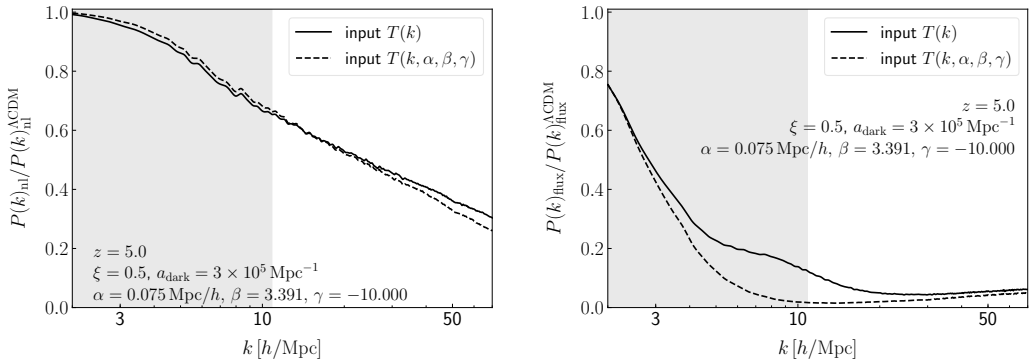

Figure 4.2: Ratio between the non-linear matter power spectra (left panel) and the corresponding ratio of 1D flux power spectrum (right panel) at $z = 5$. The spectra are obtained from simulations with the linear input given either by the true $T(k)$ (solid lines) or by the fit $T(k, \alpha, \beta, \gamma)$. The theoretical model is $n = 4$ and it has $\xi = 0.5$ and $a_{\rm dark} = 3 \times 10^5\,{\rm Mpc}^{-1}$. The grey shaded region defines the k range of MIKE/HIRES data.

neglected during the non-linear evolution, since the scales probed by Lyman-α are somewhat too large to be affected by such exotic DM properties (Vogelsberger et al. 2016) (but see also Chapter 3, where we explicitly demonstrated it in the analogous context of small-scale power suppression induced by ultra-light scalar DM).

The main advantage of the $\{\alpha, \beta, \gamma\}$-parametrisation is that it allows to systematically explore the parameter space of any non-standard DM cosmological model, provided that the corresponding linear power spectrum can be fitted in terms of the three aforementioned parameters. The new likelihood directly translates the limits on α, β, and γ obtained through MIKE/HIRES data into constraints on the fundamental particle physics parameters. The scheme is the following:

- The linear matter power spectrum associated to a given combination of cosmological parameters (six ΛCDM parameters plus additional non-standard DM parameters) is produced by `CLASS` up to a maximum wavenumber chosen to be $k_{\rm max} = 200\,h/{\rm Mpc}$. Corresponding values of the derived parameters $\{\sigma_8, n_{\rm eff}, z_{\rm reio}\}$ used to define the N-body simulations are computed. When these values fall outside of the conservative range assumed in the simulations, the model can safely be rejected, given that such models would be very bad fits to the Planck data (this will be further cross-checked in some dedicated runs called "Lyman-α prior", discussed in Section 4.3). The only exception is the case of models with a low reionisation redshift. The prior used in our grid of simulations, $7 < z_{\rm reio} < 15$, was motivated by Planck 2013 results. Instead Planck 2015 + BAO data are compatible with $z_{\rm reio} = 8.7 \pm 1.1$ ($1\,\sigma$ C.L.), such that in our runs, models with $6 < z_{\rm reio} < 7$ might still be acceptable fits and should not be systematically rejected. In practice, within our Lyman-α

likelihood, we re-map any $6 < z_{\rm reio} < 7$ to $z_{\rm reio} = 7$. This is a satisfactory approximation given that the value of the reionisation redshift has a small impact on the flux power spectrum.

- The linear matter power spectrum of the "equivalent" ΛCDM model is also produced. Note that our grid of hydrodynamic simulations for ΛCDM models assumes a fixed standard value of the ultra-relativistic relic density, corresponding to $N_{\rm eff} = 3.046$. In general, for nCDM models with the same $N_{\rm eff}$, computing the spectrum of the "equivalent" ΛCDM model would be very straightforward: we would just need to re-run `CLASS` with an infinite DM mass and/or zero interaction rates. However, all models discussed in this Chapter include DR and an enhanced value of $N_{\rm eff}$. To deal with this, we use the accurate procedure described by Rossi et al. (2015), which allows to re-map a ΛCDM model with $N_{\rm eff} > 3.046$ to another one sharing the same matter power spectrum up to some scale, but with $N_{\rm eff} = 3.046$: this can be achieved by adjusting the value of other cosmological parameters according to some analytic relations. In other words, for each model with $N_{\rm eff} > 3.046$, our Lyman-α likelihood automatically reformulates the problem in terms of an equivalent ΛCDM model with $N_{\rm eff} = 3.046$, for which we study the effect of a suppression in the small-scale matter power spectrum caused only by non-standard DM effects.

- The transfer function, i.e., the square root of the ratio between the two power spectra, is fitted in terms of $\{\alpha, \beta, \gamma\}$ with a simple least squares method. The fitting algorithm only includes points until a finite value $k_{\rm fit}$ which is set by default to $k_{\rm max} = 200\,h/{\rm Mpc}$. However, for transfer functions with oscillations within the range $[0, k_{\rm max}]$, $k_{\rm fit}$ is reduced to the first zero of the function. The fit is also restricted to values of $\{\alpha, \beta, \gamma\}$ within the region covered by the grid of simulations: $0 \leq \alpha \leq 0.17$, $1.5 \leq \beta \leq 10$, and $-10 \leq \gamma \leq -0.15$. Furthermore, if the difference between the "exact" transfer function and the fitted one is too large in a region in which the power spectrum is not strongly suppressed, our method cannot be considered accurate and reliable enough. Thus we need to implement a conservative "applicability check" rejecting models giving bad $\{\alpha, \beta, \gamma\}$-fits, but such that the Lyman-α data remain more constraining than the applicability check itself. If this condition is met, this check is just a technical step, not biasing our final results, because rejected models would anyway conflict the data. In practice, our likelihood requires that the $\{\alpha, \beta, \gamma\}$-fit to the transfer function is accurate to better than 10% in the whole region where this function is larger than 0.2. The 10% accuracy is sufficient for data with statistical uncertainties of $\sim 10\%$ such as in Iršič et al. (2017c). In Section 4.3 we will describe some dedicated runs probing that this applicability check is much less constraining than the data and has no impact on our final results.

- At this point, if the considered model has passed the aforementioned applicability checks, its flux power spectrum is produced by performing the interpolation procedure described in Section 2.6.1. By confronting such flux power spectrum against Lyman-α forest data, a χ^2-value is associated to the corresponding combination of parameters.

- The procedure is iterated per each MCMC step, until convergence is reached, i.e., until accurate constraints on the cosmological and astrophysical parameters of the model are determined.

- At each step, the fitted values of $\{\alpha, \beta, \gamma\}$ are kept in the MCMC chains as derived parameters, to check a posteriori the range of power spectrum shapes covered by a given cosmological model.

In order to show how the pipeline described above works, we plot in the left panels of Figure 4.1 the square of the linear transfer function of a few selected DM-DR models, or in other words, their linear matter power spectrum divided by that of the ΛCDM equivalent model. By construction, the transfer function always has an asymptote of one in the small-k limit. For each model, we compare it with its best fit using the $\{\alpha, \beta, \gamma\}$-parametrisation. In the right panels we show the relative error of the fit. Notice that the $\{\alpha, \beta, \gamma\}$ cannot reproduce the oscillations in $T(k)$ after the first zero (for $k > k_{\rm fit}$). However, the power of the subsequent oscillations is small. Figure 4.2 demonstrates that the impact of these oscillations on structure formation is negligible and located at scales smaller than those probed by Lyman-α, even for a rather large interaction strength. As it has already been shown in Chapter 2 (Section 2.6.3), significant differences between the "exact" flux power spectrum and the $\{\alpha, \beta, \gamma\}$-prediction appear only when the power suppression with respect to the standard CDM case is more than 50%, i.e., for models whose power spectra lie very far from the Lyman-α forest data points. This fully justifies ignoring such oscillations when applying our fit. For $k \lesssim k_{\rm fit}$, the $\{\alpha, \beta, \gamma\}$-parametrisation works rather well in reproducing the real transfer function (see Figure 4.1). The relative error features a small bump at scales $k_{1/2} < k < k_{\rm fit}$, with an amplitude related to the DR content. Then it diverges at $k \longrightarrow k_{\rm fit}$, which is unavoidable given that the reference transfer function goes to zero, but harmless due to the small power of those scales. This is not a problem for our applicability check, which only applies up to the wavenumber at which the transfer function crosses 0.2.

We stress one important point here. Our Lyman-α forest likelihood significantly improves over previous likelihood analyses present in the literature addressing DM-DR interactions (e.g., Krall et al. (2017)). In previous works, the likelihood was based on an estimate of the linear matter power spectrum amplitude slope and curvature obtained from low-resolution and low-signal-to-noise SDSS-II data (McDonald et al.

2006). These measurements were derived only in the standard ΛCDM model, and are expected to be valid only for small deviations around this model. There exists no explicit proof that models with interacting DM-DR fall inside the range of validity of this method. We are instead fitting the observed quantity, the 1D flux power, using a set of dedicated simulations that take in input the linear power spectra of the $\{\alpha, \beta, \gamma\}$-parametrisation, thus fully taking into account the cosmological signature of DM-DR interactions. Such a dedicated investigation of the non-linear evolution of structures in interacting DM-DR scenarios using N-body/hydro simulations has only been performed by Bose et al. (2019). However, the focus of the latter reference was actually on the survival of oscillatory features in the flux power spectrum, and not on a full MCMC analysis of the flux power. Furthermore, the data set that we used is expected to be the most updated and constraining one for models with a small-scale suppression. Indeed, let us stress again that high-resolution high-signal-to-noise quasar spectra can go down to the smallest scales probed by IGM structures (Viel et al. 2013; Boera et al. 2018). Conversely, low-resolution data from surveys like SDSS have smaller statistical errors, but they are limited to larger scales. They are thereby more appropriate for constraining neutrino masses and/or cosmological parameters (Palanque-Delabrouille et al. 2013), and less constraining for small-scale features.

4.2.2 Dark Matter–Dark Radiation interaction in `CLASS`

`CLASS` already incorporated several DM species and related input parameters: the CDM sector (including the effects of energy release from DM annihilation or decay into electromagnetic particles (Poulin et al. 2017a)); the decaying DM sector into DR (Poulin et al. 2016); and the nCDM sector featuring an arbitrary number of nCDM species (Lesgourgues & Tram 2011).

In Archidiacono et al. (2017), the ETHOS framework for an effective description of DM-DR interactions (Cyr-Racine et al. (2016)) was implemented in `CLASS` as a set of modifications to the existing CDM equations. For our purposes (and for the public release of the code (Archidiacono et al. 2019)), we re-implemented the same set of equations, but for a new "interacting DM species" coexisting with the plain CDM species and enlarging the total number of DM sectors available. This allows to investigate mixed DM models, and it has an appropriate structure for accommodating in future versions more types of DM interactions (e.g., DM-baryon or DM-photon), either separately or at the same time. For the moment, the new sector includes parameters like the DM mass (`m_dm`), the fraction of the total CDM density (`f_idm_dr`), as well as other parameters related to the ETHOS model, fully described in `explanatory.ini`, and appearing here in typefaces.

Here we only recall the main equations of DM and DR perturbations in the ETHOS model, and we refer to Cyr-Racine et al. (2016); Archidiacono et al. (2017) for further details.

As already mentioned in Section 4.2, the amount of DR is set by the temperature ratio $\xi = T_{\rm DR}/T_\gamma$ (`xi_idr`), and its physical density is:

$$\omega_{\rm DR} = \left(\frac{g_{\rm DR}}{2}\right) f_{\rm DR}\xi^4\omega_\gamma, \tag{4.1}$$

where the statistical factor $f_{\rm DR}$ (`stat_f_idr`) is 7/8 for fermionic DR and 1 for bosonic, $g_{\rm DR}$ is the DR number of internal degrees of freedom and it is assumed to be 2. The input parameter `idr_nature` describes the DR nature, i.e., free streaming or fluid: in the former case the DR hierarchy is evolved up to $\ell_{\rm dark}$ (`l_max_idr`) (set by default to 17), while in the latter only the modified continuity and Euler equations are present. In the free streaming case, the DR hierarchy in Newtonian gauge is:

$$\dot\delta_{\rm DR} + \frac{4}{3}\theta_{\rm DR} - 4\dot\phi = 0, \tag{4.2}$$

$$\dot\theta_{\rm DR} + k^2\left(\sigma_{\rm DR} - \frac{1}{4}\delta_{\rm DR}\right) - k^2\psi = \Gamma_{\rm DR-DM}\left(\theta_{\rm DR} - \theta_{\rm DM}\right), \tag{4.3}$$

$$\dot\pi_{{\rm DR},\ell} + \frac{k}{2\ell+1}\left((\ell+1)\pi_{{\rm DR},\ell+1} - \ell\pi_{{\rm DR},\ell-1}\right) = \left(\alpha_\ell\Gamma_{\rm DR-DM} + \beta_\ell\Gamma_{\rm DR-DR}\right)\pi_{{\rm DR},\ell}, \quad 2 \le \ell \le \ell_{\rm dark}. \tag{4.4}$$

The density and velocity dispersion perturbations are labelled as δ and θ, respectively, the DR shear perturbation is $\pi_{\rm DR} = 2\sigma_{\rm DR}$, ϕ and ψ are the gravitational potentials. The specifications related to the DR-DM interactions are embedded into $\Gamma_{\rm DR-DM}$, which is the comoving interaction rate (see the formula and the discussion below), and α_ℓ (`alpha_dark`) is the array of the interaction angular coefficients for $\ell = 2, ..., \ell_{\rm dark}$. The DR self-interactions are encoded in the comoving rate $\Gamma_{\rm DR-DR}$, whose strength is `b_dark`, and whose angular coefficients are β_ℓ (`beta_dark`).

The DM perturbation equations are:

$$\dot\delta_{\rm DM} + \theta_{\rm DM} - 3\dot\phi = 0, \tag{4.5}$$

$$\dot\theta_{\rm DM} - k^2c^2_{\rm DM}\delta_{\rm DM} + \mathcal{H}\theta_{\rm DM} - k^2\psi = \Gamma_{\rm DM-DR}\left(\theta_{\rm DM} - \theta_{\rm DR}\right), \tag{4.6}$$

where $c^2_{\rm DM}$ is the dark sound speed. The interactions are embedded in the right hand side of DR and DM dipole equations, and of DR higher order momenta. The effective comoving scattering rate of DR off DM can be parametrised as:

$$\Gamma_{\rm DR-DM} = -\Omega_{\rm DM}h^2a_{\rm dark}\left(\frac{1+z}{1+z_d}\right)^n, \tag{4.7}$$

where $z_d = 10^7$ is a normalisation factor, n (`nindex_dark`) is the temperature dependence, and $a_{\rm dark}$ (`a_dark`) is the interaction strength. Applying energy-momentum conservation, we obtain:

$$\begin{aligned}\Gamma_{\rm DM-DR} &= R_{\rm dark}\Gamma_{\rm DR-DM} \\ &= \left(\frac{4}{3}\frac{\rho_{\rm DR}}{\rho_{\rm DM}}\right)\Gamma_{\rm DR-DM}. \end{aligned} \tag{4.8}$$

With respect to Archidiacono et al. (2017), the present version of the code implements the tight coupling regime between DM and DR. By default, `CLASS` uses a stiff integrator (`ndf15`) (Blas et al. 2011) for the perturbation equations, which means that rather large values of the interaction rate can be reached while using the default equations and keeping the code fast. However, in order to investigate the very small scales probed by Lyman-α, the tight coupling is required. This regime is fully operational in our released `CLASS` version (Archidiacono et al. 2019). The tightly-coupled equations are switched on automatically whenever the ratio between the conformal interaction rate and Hubble times, $\mathcal{H}/\Gamma_{\rm DR-DM}$, falls below a threshold set by default to 0.005, and the ratio between the conformal interaction and acoustic oscillation times, $k/\Gamma_{\rm DR-DM}$, falls below 0.01[1]. At first order in $\Gamma^{-1}_{\rm DR-DM}$ the DM-DR slip is:

$$\begin{aligned}\dot{\Theta}^{\rm TCA}_{\rm DM-DR} &= \dot{\theta}_{\rm DR} - \dot{\theta}_{\rm DM} \\ &= \left(n - \frac{2}{1+R_{\rm dark}}\right)\frac{\dot{a}}{a}\left(\theta_{\rm DM} - \theta_{\rm DR}\right) + \frac{1}{1+R_{\rm dark}}\frac{1}{\Gamma_{\rm DR-DM}}\times \\ &\times\left[-\frac{\ddot{a}}{a}\theta_{\rm DM} - \frac{\dot{a}}{a}\left(k^2\frac{1}{2}\delta_{\rm DR} + k^2\psi\right) + k^2\left(c^2_{\rm DM}\dot{\delta}_{\rm DM} - \frac{1}{4}\dot{\delta}_{\rm DR}\right)\right],\end{aligned}$$

where a dot denotes the derivative with respect to conformal time. The slip is then plugged into the exact equations for the DM and DR dipole moments $\dot{\theta}_{DM}$ and $\dot{\theta}_{DR}$.

4.3 Results and discussion

With the method described in the previous Section, we have used the parameter inference code `MontePython` (Audren et al. 2013; Brinckmann & Lesgourgues 2018), interfaced with our modified `CLASS` version, in its default Metropolis Hastings mode, to scan the combination of $\{\omega_{\rm b},\ \omega_{\rm cdm},\ \log(10^{10}A_{\rm s}),\ n_{\rm s},\ \tau_{\rm reio},\ H_0,\ \xi,\ a_{\rm dark}\}$, for the ETHOS models with $n = 4$, $n = 2$, and $n = 0$ (corresponding to different powers of the temperature dependence of the comoving interaction rate $\Gamma \propto T^n$).

[1]The two thresholds are defined as the precision parameters `dark_tight_coupling_trigger_tau_c_over_tau_h` and `dark_tight_coupling_trigger_tau_c_over_tau_k`.

For $n = 4$ and $n = 2$ we assume DR to be free streaming and we neglect the impact of DR self-interactions (Archidiacono et al. 2017) ($\Gamma_{\rm DR-DR} = 0$, `b_dark`$= 0$), while for $n = 0$ we assume DR to behave like a fluid ($\Gamma_{\rm DR-DR} \to \infty$, `idr_nature`=`fluid`). For the final case, we have also studied the impact of a different parameter choice to match the NADM model by Lesgourgues et al. (2016); Buen-Abad et al. (2018), thus giving us $\{\omega_{\rm b},\ \omega_{\rm cdm},\ \log(10^{10}A_{\rm s}),\ n_{\rm s},\ \tau_{\rm reio},\ H_0,\ \Delta N_{\rm fluid},\ \Gamma_0\}$.

For each of the studied ETHOS models, we performed MCMC analysis for two different data combinations:

- **Planck + BAO**: This is the combination of Planck 2015 high-ℓ TT+TE+EE, low-ℓ data (Aghanim et al. 2016) and Planck 2015 lensing data (Ade et al. 2016b). We further add BAO data, using measurements of $D_V/r_{\rm drag}$ by 6dFGS at $z = 0.106$ (Beutler et al. 2011) by SDSS from the MGS galaxy sample at $z = 0.15$ (Ross et al. 2015), and additionally by BOSS from the CMASS and LOWZ galaxy samples of SDSS-III DR12 at $z = 0.2 - 0.75$ (Alam et al. 2017).
- **Planck + BAO + Lyman-α**: Same as above, with the additional Lyman-α likelihood discussed in Section 4.2.1.

The results for the different cases are discussed in the following Subsections.

4.3.1 ETHOS n = 4 model

The underlying particle physics model that leads to the $n = 4$ temperature dependence of the comoving interaction rate is represented by fermionic relativistic particles (DR), e.g., sterile neutrinos, interacting with DM particles through a new massive boson mediator of a new $U(1)$ broken symmetry. Given the negligible impact of DR self-interactions induced by these processes on the matter power spectrum (Archidiacono et al. 2017), we set $\Gamma_{\rm DR-DR} = 0$. The results of our MCMC runs for this model are shown in Figure 4.3 and Table 4.1, for both of the data set combinations mentioned above.

CMB constraints. We expect a clear degeneracy between the amount of DR ξ and the interaction strength $a_{\rm dark}$, because the data should remain compatible with DM interacting either strongly with a small amount of DR or barely with a large amount of DR. To capture this behaviour, we chose to use a flat prior on $\log_{10}(a_{\rm dark})$ in the range $[-3, 20]$. Indeed, a linear prior on $a_{\rm dark}$ would only give weight to the region with a high interaction rate and thus a tiny DR density. This would lead to very strong bounds on ξ that would not reflect the fact that the data is perfectly compatible with values up to $\xi \sim 0.40$.

In the middle plot of the left panel of Figure 4.3, we can see the expected degeneracy between ξ and $\log_{10}(a_{\rm dark})$. The results of MCMC runs are usually

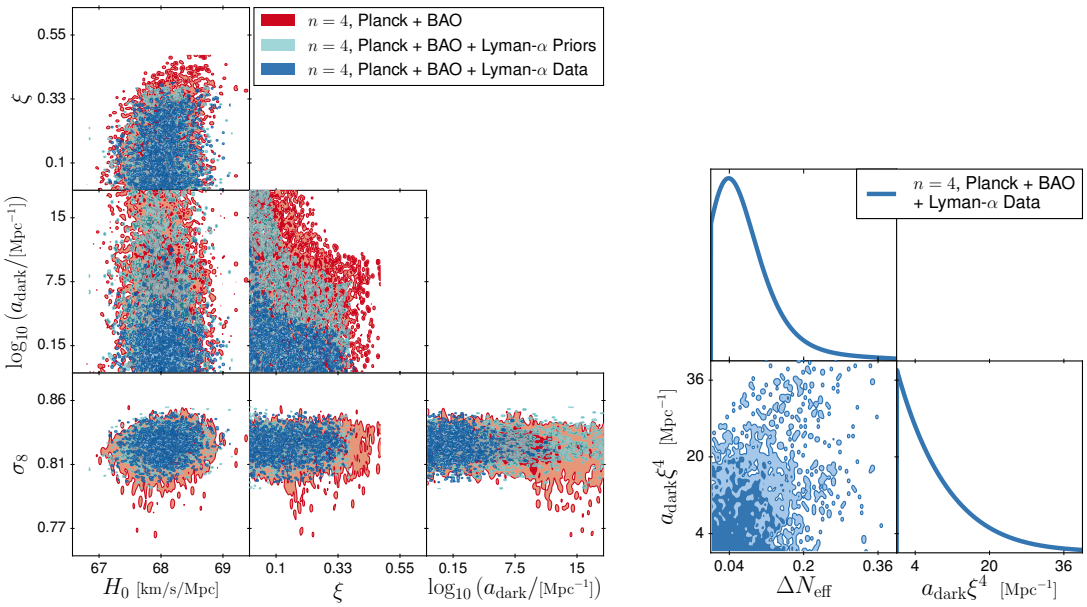

Figure 4.3: *Left*: 2D posterior distributions for all main parameters for the $n = 4$ case, with Planck + BAO (red), Planck + BAO + Lyman-α Data (dark blue), and the Lyman-α Prior check run explained in the text (light blue), when running with a flat prior on ξ and logarithmic prior on $a_{\rm dark}$. The smoothing has deliberately been turned off to show the sharp boundaries of the preferred regions more clearly. *Right*: Posterior distributions when using linear priors on $\Delta N_{\rm eff}$ and $a_{\rm dark}\xi^4$.

plotted as smoothed contour plots. In this particular context, we choose instead to plot the non-smoothed density of points in the chains[2], in order to precisely visualise the edges of the region preferred by the data. The Planck+BAO allowed region has two sharp edges set by the data rather than the priors:

- a vertical line corresponding to the maximum allowed value of ξ (and therefore $\Delta N_{\rm eff}$) in the ETHOS $n = 4$ model. We find $\xi < 0.40$ (95% C.L.), which is consistent within 1 σ with the bound obtained by Archidiacono et al. (2017), with our bounds being slightly tighter. This small difference can be attributed to the inclusion of the lensing and BAO likelihoods, which were not used in the previous study. This can be translated into $\Delta N_{\rm eff} < 0.10$, but the latter result must be taken with a grain of salt because it derives from a flat prior on ξ. Later in this Subsection we will report another bound obtained with a flat prior on $\Delta N_{\rm eff}$. The physical interpretation of this boundary is that CMB data are incompatible with too much DR, even when the latter is self-interacting. This is caused by various effects, the dominant one being the influence of the amount of extra radiation on the CMB damping tail (Hou et al. 2013). DR

[2]In practice this is achieved by analysing the chains with a high number of bins (one hundred).

Parameter	ΛCDM	ETHOS $n = 4$	
	Planck+BAO	Planck + BAO	+ Lyman-α
$100\ \omega_{\rm b}$	$2.219^{+0.013}_{-0.014}$	$2.221^{+0.015}_{-0.015}$	$2.222^{+0.017}_{-0.014}$
$\omega_{\rm cdm}$	$0.1192^{+0.0011}_{-0.0010}$	$0.1195^{+0.0011}_{-0.0014}$	$0.1192^{+0.0011}_{-0.0010}$
$\log(10^{10} A_{\rm s})$	$3.050^{+0.023}_{-0.023}$	$3.053^{+0.025}_{-0.023}$	$3.057^{+0.024}_{-0.024}$
$n_{\rm s}$	$0.9618^{+0.0042}_{-0.0041}$	$0.9622^{+0.0044}_{-0.0045}$	$0.9626^{+0.0044}_{-0.0037}$
$\tau_{\rm reio}$	$0.060^{+0.012}_{-0.012}$	$0.061^{+0.013}_{-0.013}$	$0.063^{+0.013}_{-0.013}$
H_0 / [km/(s Mpc)]	$67.94^{+0.46}_{-0.49}$	$68.06^{+0.52}_{-0.54}$	$68.09^{+0.46}_{-0.48}$
σ_8	$0.8234^{+0.0085}_{-0.0090}$	$0.823^{+0.024}_{-0.013}$	$0.826^{+0.010}_{-0.009}$
$n_{\rm eff}$	$-2.3080^{+0.0034}_{-0.0035}$	$-2.9^{+4.3}_{-22.1}$	$-2.3070^{+0.0039}_{-0.0035}$
ξ	–	< 0.40	< 0.38
$\log_{10}(a_{\rm dark}\ /\ [{\rm Mpc}^{-1}])$	–	n.l.	< 6.8
$\Delta\chi^2$	–	0	-3.62
$\Delta N_{\rm eff}$	–	–	< 0.23
$a_{\rm dark}\xi^4 / [{\rm Mpc}^{-1}]$	–	–	< 30

Table 4.1: Preferred regions at the 68% C.L. (or at the 95% C.L. in the case of upper bounds) for the parameters of the ETHOS $n = 4$ case, both with Planck + BAO and Planck + BAO + Lyman-α. With the first data set, the interaction parameter is not bounded within the prior range. The $\Delta\chi^2$ is given with respect to ΛCDM with the same data sets. The last two rows show the results obtained with linear priors on $\Delta N_{\rm eff}$ and $a_{\rm dark}\xi^4$ using the second data set. Entries with "n.l." means that there is no upper limit within the prior range, while "–" means that the parameter is not present.

has other effects on the scale and amplitude of the acoustic peaks that depend on the rate of DR self-interactions and DR-DM interaction (Archidiacono et al. 2017): thus the bound found in this case is specific to the ETHOS $n = 4$ model, and in principle different from what one would obtain in a plain ΛCDM+$N_{\rm eff}$ fit with only free streaming relativistic relics.

- a roughly hyperbolic boundary, corresponding physically to the limit set by the CMB on the effect of the DM-DR interaction. In particular, a too large rate $\Gamma_{\rm DM-DR}$ implies that DM develops a fast mode (Weinberg 2002; Voruz et al. 2014) that influences the CMB power spectrum, with a suppression of the clustering of the baryon-photon fluid (Archidiacono et al. 2017; Cyr-Racine et al. 2016, 2014).

We obtain no upper bound on $\log_{10}(a_{\rm dark})$, since in the limit of small DR density the DM-DR and DR-DM interaction rates can be arbitrarily high. Thus the allowed region extends up to our upper prior boundary $\log_{10}(a_{\rm dark}) \leq 20$.

For the other cosmological parameters, error bars are slightly larger than for the ΛCDM model with the same data combination, but smaller than for the ΛCDM+$N_{\rm eff}$ model. This arises from several reasons: our flat prior on ξ gives more weight to small values of $\Delta N_{\rm eff}$; we only allow $N_{\rm eff}$ to increase beyond 3.046, while a run with a flat prior on $N_{\rm eff}$ would return $N_{\rm eff} = 2.98 \pm 0.18$ (68% C.L.) (Ade et al. 2016a); and in our model, increasing $\Delta N_{\rm eff}$ comes at the price of introducing DM-DR interaction effects not favoured by the data. In any case, we see that the ETHOS $n = 4$ model offers no clear opportunities to accommodate the high value of H_0 (Riess et al. 2019) and/or the low value of σ_8 hinted by some data sets (Hildebrandt et al. 2018; Abbott et al. 2018; Joudaki et al. 2019).

Lyman-α constraints. With the addition of the Lyman-α likelihood, we obtain approximately the same bound on $\xi < 0.38$ (95% C.L.), as the number of additional relativistic degrees of freedom is already well constrained by CMB data. Instead the upper limit on the interaction rate shrinks by about ten orders of magnitude, because DM-DR interactions result in a suppression of the small-scale matter power spectrum strongly constrained by Lyman-α data. Quantifying this effect is our main goal. Figure 4.1 already showed that a larger value of $a_{\rm dark}$ could potentially be compensated by a smaller value of ξ leading to the same cut-off scale. Indeed, we checked explicitly that the edge of the allowed region is a curve of constant $a_{\rm dark}\xi^4$. This behaviour was expected because the term that accounts for interactions in the DM Euler equation has a coefficient $\Gamma_{\rm DM-DR} \propto \rho_{\rm DR}\,\Gamma_{\rm DR-DM} \propto a_{\rm dark}\xi^4$.

This run gives an upper bound $\log_{10}(a_{\rm dark}/{\rm Mpc}^{-1}) < 6.8$ (95% CL) that is strongly prior-dependent. Indeed, since $a_{\rm dark}$ is compatible with zero, upper bounds on $\log_{10}(a_{\rm dark})$ are inevitably influenced by the choice of a lower prior boundary on this parameter. Moreover, the data are compatible with arbitrarily large values of $a_{\rm dark}$ for arbitrarily small ξ's, such that the bound would entirely disappear if we had chosen a logarithmic prior on ξ.

The analysis with flat priors on ξ and $\log_{10}(a_{\rm dark})$ is particularly useful for identifying the physical mechanisms responsible for the various bounds. It allowed us to check that the data are mostly sensitive to the effect of the density of extra radiation, proportional to $\Delta N_{\rm eff}$, and of the DM-DR rate $\Gamma_{\rm DM-DR}$, parametrised by $a_{\rm dark}\xi^4$. Therefore, the most informative and robust way to formulate our final results is to quote bounds on $\{\Delta N_{\rm eff},\, a_{\rm dark}\xi^4\}$, assuming flat priors on these parameters.

We thus performed another MCMC run with such a choice of priors. The results are shown in the right panel of Figure 4.3. Our final results for the $n = 4$ ETHOS model are summarised by the 95% C.L. upper bounds $\Delta N_{\rm eff} < 0.23$ and $a_{\rm dark}\xi^4 < 30\,{\rm Mpc}^{-1}$. The upper limit of the Bayesian confidence interval for $\Delta N_{\rm eff}$ is slightly stronger than for a ΛCDM+$N_{\rm eff}$ model with extra free streaming relativistic

relics and Planck+BAO data, $\Delta N_{\rm eff} < 0.28$ (95% C.L., see Ade et al. (2016a)), because in our case models with $\Delta N_{\rm eff} > 0$ also come with DM-DR interaction effects that are not favoured by the data.

Knowing the upper bound on $a_{\rm dark}\xi^4$ is also convenient for model building. A typical particle-physics-motivated model would predict a given value of ξ (related to the physics of the dark sector and to its interactions with the visible sector). In such a case one can immediately conclude that the Lyman-α data impose a maximum value on the scattering rate $a_{\rm dark}$ given by $30\,\xi^{-4}\,{\rm Mpc}^{-1}$.

It is important to check that our results are actually driven by the Lyman-α data, and not by the restrictions imposed on the small-scale matter power spectrum by the method implemented in our likelihood, described in Section 4.2.1. For this purpose, we also performed a run with the Planck+BAO likelihoods combined with a modified version of the Lyman-α likelihood that returns a constant value if the power spectrum passes all of the applicability checks, and a zero likelihood otherwise. Thus this run relies on the Planck+BAO data and on the Lyman-α likelihood prior, but not on the Lyman-α data. It allows us to derive regions of validity for our implementation. We call it "Planck+BAO+Lyman-α Prior" and its results are also shown in the left panel of Figure 4.3. If the edge of the allowed region was similar in the Lyman-α Prior and Lyman-α Data runs, we would know that our bounds are driven by the applicability of the method and not by the data. This is not the case, as we can clearly see when comparing the dark and light blue regions in Figure 4.3. As such, we conclude that the applicability checks of our implementation impose no further restriction besides the region that is already excluded by other means.

Furthermore, when adding the Lyman-α likelihood, our error bars on $n_{\rm eff}$, which is the slope of the Lyman-α spectrum, are greatly reduced. This comes mainly from our improved bound on $a_{\rm dark}$; when the interaction strength is allowed to vary over many orders of magnitude, our $P(k)$ is not monotonic, and thus $n_{\rm eff}$ can assume any value, both negative and positive (if the corresponding k value is, for example, just after the first oscillation in $P(k)$).

The inclusion of Lyman-α data tightens the error bars on σ_8, while the mean value is not significantly affected. The mean value and error bars of H_0 are not impacted by the addition of Lyman-α data for this model. The bounds for both parameters are in very close agreement with those obtained for a standard ΛCDM model with the same data sets.

Finally, the χ^2 obtained in the Planck+BAO case is not any better than for the vanilla ΛCDM model, while the addition of Lyman-α data brings it down by $\Delta\chi^2 = -3.6$. Considering that the model features two additional parameters, we conclude that interacting DM-DR models provide a fit of Planck+BAO+Lyman-α as good as ΛCDM.

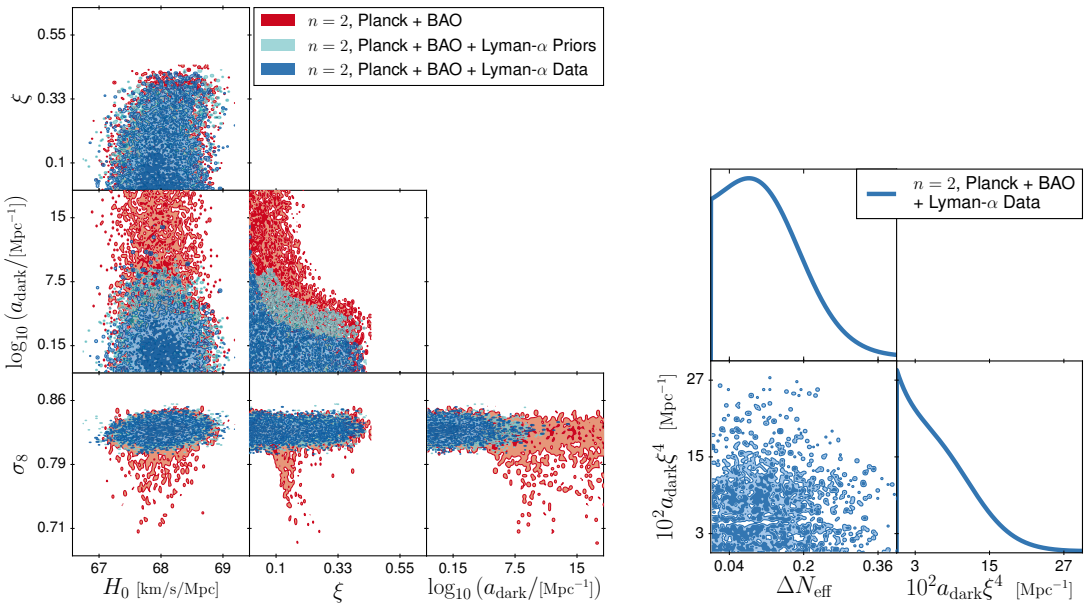

Figure 4.4: *Left*: 2D posterior distributions for all main parameters for the $n = 2$ case, with Planck+BAO (red), Planck+BAO+Lyman-α Data (dark blue), and the Lyman-α Prior check run explained in the text (light blue), when running with a flat prior on ξ and logarithmic prior on $a_{\rm dark}$. The smoothing has deliberately been turned off to show the sharp boundaries of the preferred regions more clearly. *Right*: Posterior distributions when using linear priors on $\Delta N_{\rm eff}$ and $10^2 a_{\rm dark}\xi^4$.

4.3.2 ETHOS n = 2 model

The scenario where the comoving scattering rate of DR off DM scales like T^2 can be realised, e.g., with 4-point contact-only interaction. As for $n = 4$, we neglect the subdominant contribution of DR self-interactions.

CMB constraints. The results of our MCMC run with Planck+BAO data for the $n = 2$ case are shown in Figure 4.4 and Table 4.2. Once more, the middle plot in the left panel of Figure 4.4 shows that the data impose two limitations on the ETHOS parameter: an upper bound $\xi < 0.43$ at the 95% C.L. and a hyperbolic-shaped limit on $\{\xi,\ a_{\rm dark}\}$.

For the other parameters, the preferred intervals only widen moderately with respect to the ΛCDM model, excepted for σ_8 which is compatible with much smaller values. The contour plot for $\{\xi,\ \sigma_8\}$ shows a degeneracy allowing to reach such small values for specific values of ξ and a large interaction rate $a_{\rm dark} > 10^7$. The degeneracy is captured by the relation $\sigma_8 \simeq 0.823 - 210\,\xi^4$, and stretches down to $\sigma_8 = 0.75$ for $\xi \simeq 0.14$. It is potentially interesting to explain the low value of σ_8 returned by several data on cosmic shear and cluster counts, but we will not investigate it in details because this region will be excluded in the next paragraph

by Lyman-α bounds on the interaction rate. Once this region is ignored, we find that the ETHOS $n = 2$ model does not offer opportunities to accommodate larger H_0 or smaller σ_8 values than ΛCDM.

Lyman-α constraints. Like for $n = 4$, the inclusion of Lyman-α data marginally affects the bound on ξ, but considerably strengthens the upper limit on the interaction rate, which is given once more by a line of constant $a_{\rm dark}\xi^4$. This limit is stronger than in the $n = 4$ case by about two orders of magnitude. We checked explicitly that the suppression in the matter power spectrum takes place roughly at the same scale when we change n and keep the same $10^{-n}a_{\rm dark}\xi^4$. This is consistent with the fact that the scales constrained by our Lyman-α data crossed the Hubble scale roughly around $z \simeq 10^6$, and have been suppressed according to the rate $\Gamma_{\rm DM-DR}(z)$ evaluated at that time. Equations (4.7) and (4.8) show that up to constant numbers,

$$\Gamma_{\rm DM-DR}(z) \propto (1+z)\left(\frac{1+z}{1+z_d}\right)^n a_{\rm dark}\xi^4 \tag{4.9}$$

with $z_d = 10^7$, implying

$$\Gamma_{\rm DM-DR}(10^6) \propto 10^{6-n} a_{\rm dark}\xi^4 \ . \tag{4.10}$$

Thus it is normal that the Lyman-α data set provides comparable limits on the combination $(10^{-n}a_{\rm dark}\xi^4)$ for all n's, and that limits on $a_{\rm dark}\xi^4$ become one hundred times stronger when n decreases by two.

We find a bound $\xi < 40$ (95% C.L.) very similar to that in the $n = 4$ case, while the bound $\log_{10}(a_{\rm dark}/[{\rm Mpc}^{-1}]) < 8.4$ (95% C.L.) should again be taken with great care due to its strong dependence on the choice of a linear prior for ξ and on the lower prior edge for $\log_{10}(a_{\rm dark})$. We thus switch to linear priors on the parameters directly related to the physical effects probed by the data, and obtain our final results for the ETHOS $n = 2$ model: $\Delta N_{\rm eff} < 0.29$ and $10^2 a_{\rm dark}\xi^4 < 18\,{\rm Mpc}^{-1}$ (95% C.L.). The first bound is identical to what is obtained when fitting Planck+BAO with a ΛCDM+$N_{\rm eff}$ model.

Once again we performed a "Planck+BAO+Lyman-α Prior" run to check that our bounds do not come from the limitations of the method. In this case, if we compare the $\xi - \log(a_{\rm dark})$ posteriors for the Lyman-α Prior and Lyman-α Data runs in the left panel of Figure 4.4, we see that for $\xi > 0.05$ our constraints are really derived from the data rather than from the range of validity of our method. This is not true anymore in a very small region with $\xi < 0.05$, where the two contours overlap. This is because for these models, the $\{\alpha, \beta, \gamma\}$-parametrisation is not accurate. However $\xi < 0.05$ implies a tiny DR density $\Delta N_{\rm eff} < 2 \cdot 10^{-5}$. This small region is not very interesting for model building, because such tiny values

Parameter	ΛCDM	ETHOS $n = 2$	
	Planck+BAO	Planck+BAO	+Lyman-α
$100\ \omega_{\rm b}$	$2.219^{+0.013}_{-0.014}$	$2.220^{+0.014}_{-0.014}$	$2.220^{+0.014}_{-0.016}$
$\omega_{\rm cdm}$	$0.1192^{+0.0011}_{-0.0010}$	$0.1195^{+0.0011}_{-0.0013}$	$0.1194^{+0.0011}_{-0.0011}$
$\log(10^{10} A_{\rm s})$	$3.050^{+0.023}_{-0.023}$	$3.053^{+0.025}_{-0.025}$	$3.051^{+0.023}_{-0.024}$
$n_{\rm s}$	$0.9618^{+0.0042}_{-0.0041}$	$0.9621^{+0.0044}_{-0.0043}$	$0.9618^{+0.0039}_{-0.0043}$
$\tau_{\rm reio}$	$0.060^{+0.012}_{-0.012}$	$0.061^{+0.013}_{-0.013}$	$0.059^{+0.013}_{-0.013}$
H_0 / [km/(s Mpc)]	$67.94^{+0.46}_{-0.49}$	$68.02^{+0.51}_{-0.51}$	$67.99^{+0.51}_{-0.51}$
σ_8	$0.8234^{+0.0085}_{-0.0090}$	$0.819^{+0.021}_{-0.017}$	$0.8244^{+0.0088}_{-0.0095}$
$n_{\rm eff}$	$-2.3080^{+0.0034}_{-0.0035}$	$-2.9^{+7.0}_{-3.5}$	$-2.3080^{+0.0034}_{-0.0037}$
ξ	–	< 0.43	< 0.40
$\log_{10}(a_{\rm dark}\ /\ [{\rm Mpc}^{-1}])$	–	n.l.	< 8.4
$\Delta\chi^2$	–	0	-0.12
$\Delta N_{\rm eff}$	–	–	< 0.29
$10^2 a_{\rm dark}\xi^4 / [{\rm Mpc}^{-1}]$	–	–	< 18

Table 4.2: Preferred ranges at the 68% C.L. (or 95% upper bound in some cases) for all relevant parameters for the ETHOS $n = 2$ case, both with Planck+BAO and Planck+BAO+Lyman-α. With the first data set, the interaction parameter is not bounded within the prior range. The $\Delta\chi^2$ is given with respect to ΛCDM with the same data sets. The last two rows show the results obtained with linear priors on $\Delta N_{\rm eff}$ and $10^2 a_{\rm dark}\xi^4$ using the second data set.

are difficult to motivate theoretically (for instance, they may derive from a DR particle decoupling from thermal equilibrium with standard model particles when the number of relativistic degrees of freedom is unusually large, $g_* \sim \mathcal{O}(10^4)$). Also, even if our method was improved in order to deal correctly with this corner of the parameter space, there would be no reason for the 95% C.L. upper bound on $\{\xi,\ a_{\rm dark}\}$ to be different from $10^2 a_{\rm dark}\xi^4 = 18$, since the shape of this limit can be inferred from simple analytic arguments. Thus we can safely extrapolate it below $\xi = 0.05$. Finally, we should note that this minor issue is irrelevant when running with a flat prior on $\Delta N_{\rm eff}$, since with such a prior it affects a completely negligible fraction of the preferred region volume.

Like for the $n = 4$ case, we obtain a significantly tighter bound on $n_{\rm eff}$, while the mean value and error bars of H_0 are not impacted by the addition of Lyman-α data. The preferred intervals for H_0 and σ_8 are very close to those of the ΛCDM model. For both data combinations, the difference obtained in the $\Delta\chi^2$ with respect to the base ΛCDM are negligible, thus we once again find no preference for the interacting DM-DR models.

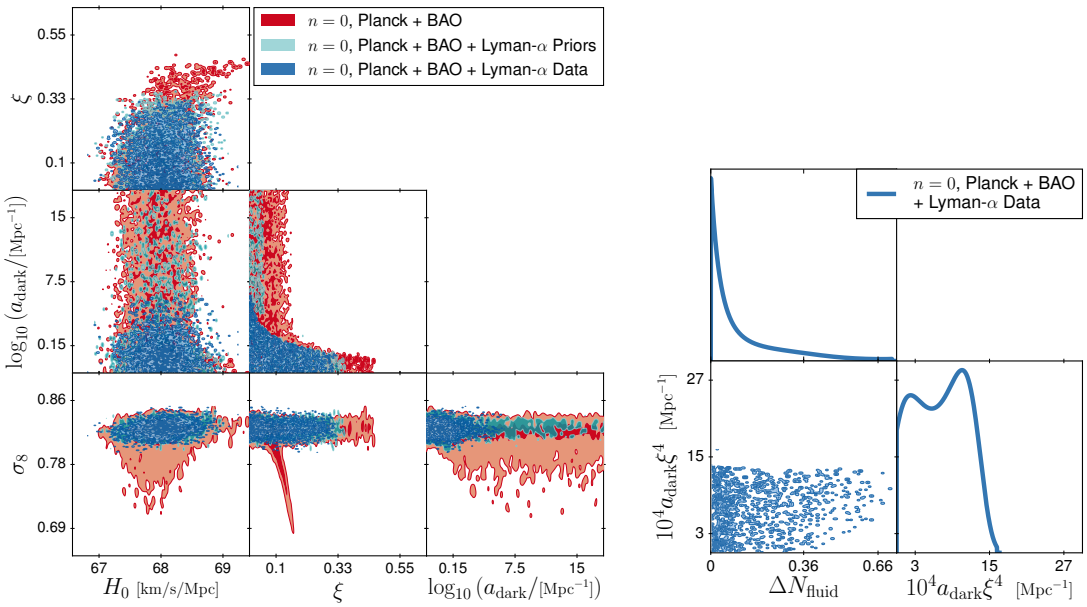

Figure 4.5: *Left*: Two-dimensional posterior distributions for all main parameters for the $n = 0$ case, with Planck+BAO (red), Planck+BAO+Lyman-α Data (dark blue), and the Lyman-α Prior check run explained in the text (light blue), when running with a flat prior on ξ and logarithmic prior on $a_{\rm dark}$. The smoothing has deliberately been turned off to show the sharp boundaries of the preferred regions more clearly. *Right*: Posterior distributions when using linear priors on $\Delta N_{\rm fluid}$ and $10^4 a_{\rm dark}\xi^4$.

4.3.3 ETHOS n = 0 model

The $n = 0$ case is motivated by several particle physics set-ups in which the DM-DR momentum transfer rate with respect to proper time scales like T^2, meaning that the ETHOS rate $\Gamma_{\rm DR-DM}$ is constant. This occurs for instance in the NADM scenario, in which DM particles are charged under a dark non-Abelian symmetry whose dark gluons play the role of DR (see, e.g., Buen-Abad et al. (2015); Lesgourgues et al. (2016); Buen-Abad et al. (2018); Krall et al. (2017); Pan et al. (2018)). Since these models tend to predict strong self-interactions in the DR sector, in this Subsection we always assume that DR is a relativistic perfect fluid described by one continuity and one Euler equation (unlike for $n = 4$ and $n = 2$). To stress this point, we will denote the DR density (in units of effective neutrino number) $\Delta N_{\rm fluid}$ instead of $\Delta N_{\rm eff}$.

For the $n = 0$ model, we can use different parametrisations and priors corresponding to different approaches discussed in the previous literature – either in the ETHOS general framework, or for specific models like the NADM one. We first look at the standard ETHOS parametrisation, with the same choice of priors as in previous cases. Then, to compare our results with the ones from Lesgourgues et al. (2016); Buen-Abad et al. (2018), we will switch to linear priors on the interaction

rate, combined with either linear or logarithmic priors on the parameter $\Delta N_{\rm fluid}$. This will also allow us to see the influence of the choice of priors on our results.

CMB constraints with ETHOS $n = 0$ parametrisation. Our results for this case, assuming a flat prior on $\xi \geq 0$ and on $-3 \leq \log_{10}(a_{\rm dark}/{\rm Mpc}^{-1}) \leq 20$, are shown in Figure 4.5 and Table 4.3. In this case, the general behaviour is similar to the previous cases: we obtain an upper bound on ξ and a hyperbolic-shaped upper limit on $\{\xi, a_{\rm dark}\}$. CMB bounds are much stronger in this model than in previous cases, which is consistent with the fact that the rate $\Gamma_{\rm DM-DR}(z)$ evaluated near photon decoupling, when $z \sim \mathcal{O}(10^3)$, is much larger for the same value of $a_{\rm dark}\xi^4$ when n decreases (as shown by Equation (4.9)). We shall see that for $n = 0$, CMB bounds dominate over Lyman-α bounds at least for some values of ξ. Thus it is worth quantifying these bounds precisely. In the space $\{\xi, \log_{10}(a_{\rm dark})\}$ and within our prior range, the 95% C.L. preferred region is defined in good approximation by:

- either $\xi < 0.13$,
- or $\xi < 0.38$ and $10^4 a_{\rm dark}\xi^4 < 14\,{\rm Mpc}^{-1}$.

This means that the CMB excludes all ETHOS $n = 0$ models with either a too large DR density ($\xi > 0.38$) or a too large $\Gamma_{\rm DM-DR}$ rate ($10^4 a_{\rm dark}\xi^4 > 14\,{\rm Mpc}^{-1}$), but looses sensitivity to these parameters when the DR density is very small ($\xi < 0.13$).

Another interesting aspect of these results is that the ETHOS $n = 0$ model allows to reach larger values of H_0 or lower values of σ_8 than the ΛCDM model. By looking at the 2D contour plots in the left panel of Figure 4.5, we see that:

- high values of H_0 require a large DR density, $\xi > 0.4$ (i.e., $\Delta N_{\rm fluid} > 0.1$): indeed this is a consequence of the well known $H_0 - \Delta N_{\rm fluid}$ degeneracy, that works particularly well in this case because DR is self-interacting, and thus less constrained by CMB observables than extra free streaming relics (Lesgourgues et al. 2016). Our 95% C.L. preferred region reaches values up to $H_0 \simeq 70\ {\rm km\,s^{-1}Mpc^{-1}}$ for $\xi \simeq 0.38$ ($\Delta N_{\rm fluid} \simeq 0.08$). With our choice of priors, this part of the allowed parameter space has little weight, and the 68% C.L. preferred interval for H_0 is still nearly the same as for ΛCDM. Later in this Subsection, runs with different priors will give more weight to this degeneracy;
- a close inspection of the $\{\xi, \sigma_8\}$ contour plot of Figure 4.5 shows that in this plane, the marginalised posterior is bimodal, i.e., made of the superposition of two separate categories of models. The first one has $\sigma_8 = 0.823 \pm 0.017$ (95% C.L.) for any allowed value of the DR density parameter (in the range $0 < \xi < 0.38$). The second one corresponds to a strongly degenerate direction in $\{\xi, \sigma_8\}$, captured by the relation $\sigma_8 \simeq 0.823 - 210\,\xi^4$ (like for the ETHOS

Parameter	ΛCDM	ETHOS $n = 0$	
	Planck+BAO	Planck+BAO	+Lyman-α
$100\ \omega_{\rm b}$	$2.219^{+0.013}_{-0.014}$	$2.220^{+0.015}_{-0.015}$	$2.221^{+0.015}_{-0.015}$
$\omega_{\rm cdm}$	$0.1192^{+0.0011}_{-0.0010}$	$0.1195^{+0.0011}_{-0.0014}$	$0.1192^{+0.001}_{-0.001}$
$\log(10^{10} A_{\rm s})$	$3.050^{+0.023}_{-0.023}$	$3.053^{+0.025}_{-0.024}$	$3.054^{+0.025}_{-0.024}$
$n_{\rm s}$	$0.9618^{+0.0042}_{-0.0041}$	$0.9621^{+0.0042}_{-0.0045}$	$0.9624^{+0.0044}_{-0.0041}$
$\tau_{\rm reio}$	$0.060^{+0.012}_{-0.012}$	$0.061^{+0.013}_{-0.012}$	$0.061^{+0.013}_{-0.014}$
H_0 / [km/(s Mpc)]	$67.94^{+0.46}_{-0.49}$	$68.04^{+0.50}_{-0.60}$	$68.03^{+0.47}_{-0.49}$
σ_8	$0.8234^{+0.0085}_{-0.0090}$	$0.815^{+0.044}_{-0.009}$	$0.8237^{+0.0097}_{-0.0093}$
$n_{\rm eff}$	$-2.308^{+0.0034}_{-0.0035}$	$-3.4^{+9.5}_{-4.2}$	$-2.3100^{+0.0071}_{-0.0079}$
ξ	–	< 0.38	< 0.33
$\log_{10}(a_{\rm dark}\ /\ [{\rm Mpc}^{-1}])$	–	n.l.	< 3.3
$\Delta\chi^2$	–	0	-0.70
$\Delta N_{\rm fluid}$	–	–	< 0.47
$10^4 a_{\rm dark}\xi^4 / [{\rm Mpc}^{-1}]$	–	–	< 14

Table 4.3: 68% C.L. constraints (or 95% upper bound in some cases) for all relevant parameters for the ETHOS $n = 0$ case, both with Planck+BAO and Planck+BAO+Lyman-α. With the first data set, the interaction parameter is not bounded within the prior range. The $\Delta\chi^2$ is given with respect to ΛCDM with the same data sets. The last two rows show the results obtained with linear priors on $\Delta N_{\rm fluid}$ and $10^4 a_{\rm dark}\xi^4$ using the second data set.

$n = 2$ model), and requires a large interaction rate $a_{\rm dark} \geq 1$. It stretches down to $\sigma_8 = 0.68$ for $\xi \simeq 0.16$. This part of the parameter space will also play an enhanced role in some of the runs that we will discuss later with different physical motivations and priors.

Lyman-α constraints with ETHOS $n = 0$ parametrisation. At first sight, the discussion of the Lyman-α constraints seems very similar to that for $n = 2$ or 4. We expect that Lyman-α data will slightly tighten the bound on ξ and put a strong limit on $10^4 a_{\rm dark}\xi^4 < \mathcal{O}(10)$. This is indeed what happens in our run with a linear prior on ξ and a logarithmic prior on $a_{\rm dark}$: we get $\xi < 0.33$ and $10^4 a_{\rm dark}\xi^4 < 14\,{\rm Mpc}^{-1}$ (95% C.L.). Doing a second run with flat priors on ($\Delta N_{\rm fluid}$, $10^4 a_{\rm dark}\xi^4$), we find $\Delta N_{\rm fluid} < 0.47$ and a confirmation of $10^4 a_{\rm dark}\xi^4 < 14\,{\rm Mpc}^{-1}$ (95% C.L.).

However, a run with the "Planck+BAO+Lyman-α Prior" combination shows that the previous results must be taken with great care. Looking at the middle plot of the left panel of Figure 4.5, we see that:

- the different checks performed inside our Lyman-α likelihood induce a cut at $\xi < 0.33$. Thus the previous bound on ξ did not come from the Lyman-α data

but from our methodology, i.e., from the fact that ETHOS $n = 0$ models with $\xi > 0.33$ do not yield a power spectrum that can be accurately represented by the $\{\alpha, \beta, \gamma\}$-parametrisation. Thus we should not trust any bound on ξ apart from the one obtained with Planck+BAO alone, namely $\xi < 0.38$;

- for $\xi > 0.13$, the upper bound on $a_{\rm dark}\xi^4$ is nearly the same in the three ETHOS $n = 0$ runs (without Lyman-α likelihood, with Lyman-α Prior and with Lyman-α Data), suggesting that CMB data alone provide the strongest bounds in this case: $10^4 a_{\rm dark}\xi^4 < 14\,{\rm Mpc}^{-1}$ (95% C.L.). Given the impact of this model on CMB and LSS observables, already discussed in the previous works by Lesgourgues et al. (2016); Buen-Abad et al. (2018), this is not a surprise: for parameter values leading to significant effects in the CMB temperature and polarisation spectrum, this model only generates a very smooth and progressive suppression in the small-scale matter power spectrum, much more difficult to constrain with Lyman-α data than the sharp exponential cut-off observed for $n = 2, 4$;

- for $\xi < 0.02$, the Lyman-α Prior run sets no upper limit on the interaction rate, while the Lyman-α Data run returns $10^4 a_{\rm dark}\xi^4 < 14$ (95% C.L.): thus we can trust this bound which really comes from the data;

- there is a problematic range $0.02 < \xi < 0.13$ in which the Lyman-α Prior run also sets an upper limit $10^4 a_{\rm dark}\xi^4 < 14\,{\rm Mpc}^{-1}$. The reason is that for $n = 0$ and $\xi > 0.02$, the $\{\alpha, \beta, \gamma\}$ parametric function cannot provide an accurate fit of the suppression in the matter power spectrum. This indicates that for this class of models, the bounds are driven by the limitations of the method, in particular by the flexibility of the parametric fitting function, and not by the data. We could search for a better method, but we believe that this is not well motivated, for two reasons. First, $0.02 < \xi < 0.13$ means $6 \cdot 10^{-7} < \Delta N_{\rm fluid} < 10^{-3}$. The weight of this region would be negligible if we would run with a flat prior on $\Delta N_{\rm fluid}$, so we may simply ignore it. Second, the analytic argument suggesting that the Lyman-α bound on the DM-DR interaction takes the form of an upper limit on $a_{\rm dark}\xi^4$ worked very well for $n = 2$ and $n = 4$, and still works very well in the present case for $\xi < 0.02$ and $\xi > 0.13$. We have no reason to believe that this would not be the case in the intermediate range. Thus it is reasonable to expect that a better method would return $10^4 a_{\rm dark}\xi^4 < 14$ (95% C.L.) throughout the range of allowed values $0 < \xi < 0.38$.

In summary, we should retain from this analysis that for $\xi > 0.13$, Lyman-α data, at least with our approach, cannot improve over Planck+BAO bounds, which give $10^4 a_{\rm dark}\xi^4 < 14\,{\rm Mpc}^{-1}$ (95% C.L.). For $0 < \xi < 0.02$, the Lyman-α data give the same bound. In the intermediate range, a different approach would be needed, but there are some hints that the Lyman-α data would give again the same bound.

CMB constraints with a particle-physics-motivated flat prior on $\Delta N_{\rm fluid} \geq 0.07$. Several works have presented particle physics models that can be effectively described by the ETHOS $n = 0$ parametrisation, with weakly interacting DM-DR, and strongly self-interacting DR. In the NADM model by Buen-Abad et al. (2015), the DR is made up of the dark gluons of a non-Abelian gauge symmetry $SU(N)$. Its density is parametrised by $\Delta N_{\rm fluid} = 0.07(N^2 - 1)$ with $N \geq 2$. Lesgourgues et al. (2016) presents a second set-up leading to approximately the same cosmological signature, in which the DR has two components: the dark photon of a dark $U(1)$ gauge symmetry, plus N_f massless fermions with a dark charge q. For $q \geq 1/3$ the DR density is parametrised by $\Delta N_{\rm fluid} = 0.07(1 + \frac{7}{4}N_f)$, but for smaller charges one gets $\Delta N_{\rm fluid} = 0.07$. These models motivate dedicated runs with a flat prior on $\Delta N_{\rm fluid} \geq 0.07$. To compare our results with previous works, we will also adopt a flat prior on the DM-DR momentum exchange rate evaluated today, Γ_0, related to the ETHOS parameters through $a_0\Gamma_0 = \Gamma_{\rm DM-DR}(z = 0) = \frac{4}{3}\omega_{\rm DR}\, a_{\rm dark} = \frac{4}{3}\omega_\gamma\, a_{\rm dark}\xi^4$. With such a correspondence, we checked that we could accurately reproduce Figures 3-6 from Buen-Abad et al. (2018): thus our version of `CLASS` modified for the ETHOS parametrisation does agree perfectly with the version of `CLASS` modified specifically for the NADM model by Buen-Abad et al. (2018).

The prior $\Delta N_{\rm fluid} \geq 0.07$ translates in the ETHOS parametrisation to $\xi \geq 0.367$. Looking at our previous results, we see that this clearly corresponds to the region in which the CMB bounds are at least as strong as the Lyman-α bounds: thus for this case it is sufficient to run with Planck+BAO data only. Note that with such a prior, we avoid the bimodality of the posterior found in the results of the previous run (corresponding to a degeneracy between σ_8 and ξ for $\xi \leq 0.16$). Thus the theoretical prior $\Delta N_{\rm fluid} \geq 0.07$ offers a technical advantage: it limits the exploration of the model parameter space to a region where the posterior is smooth and unimodal, leading to more robust MCMC results.

Our findings, presented in the left panel of Figure 4.6 and middle column of Table 4.4, are consistent with those from Lesgourgues et al. (2016); Krall et al. (2017) when using Planck 2015+BAO 2011 data. Our bounds are however slightly stronger and more up-to-date, because we include Planck lensing data and more recent BAO data. We do not compare directly our results with those from Buen-Abad et al. (2018), as the latter always included direct H_0 measurements, as well as Planck data on Sunyaev-Zel'dovich cluster counts.

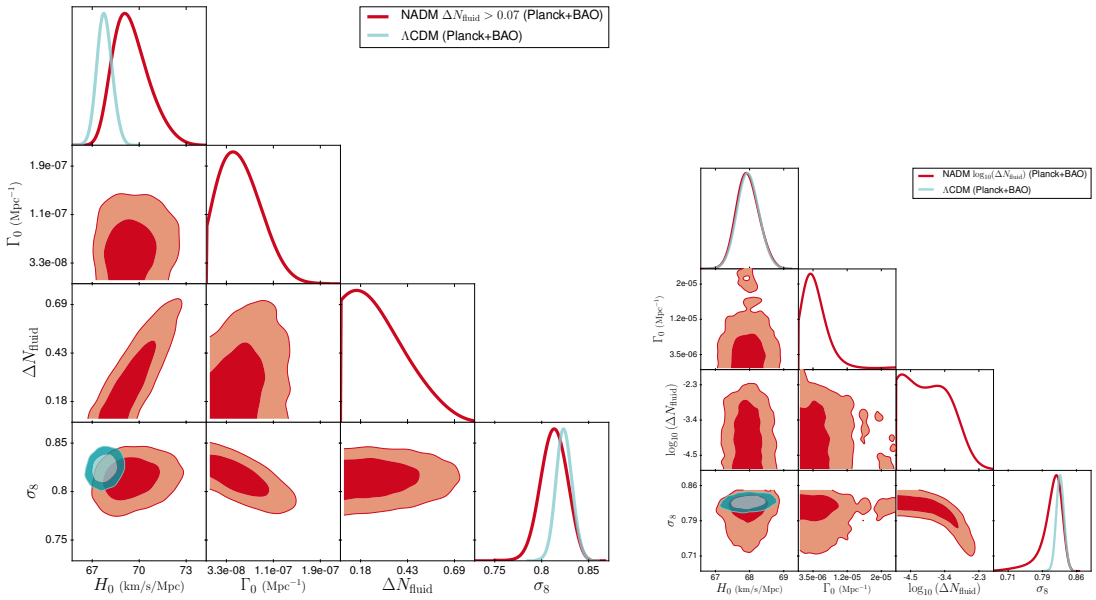

Figure 4.6: *Left*: 2D posterior distributions for all main parameters using Planck+BAO, for the NADM case (red) and for ΛCDM (blue), with the lower prior $\Delta N_{\rm fluid} > 0.07$. *Right*: Same as left but with the log prior $-5 \leq \log_{10}(\Delta N_{\rm fluid}) \leq 0$.

We find $0.07 \leq \Delta N_{\rm fluid} \leq 0.59$ (95% C.L.), corresponding to $0.367 < \xi < 0.626$ with a non-flat prior on ξ, and $\Gamma_0 < 1.2 \cdot 10^{-7}{\rm Mpc}^{-1}$ (95% C.L.), corresponding to $10^4 a_{\rm dark}\xi^4 < 36$ (95% C.L.). We see that the lower prior edge on $\Delta N_{\rm fluid}$ and the linear prior on both $\Delta N_{\rm fluid}$ and Γ_0 have pushed the MCMC to explore regions that were not reached with the previous ETHOS $n = 0$ prior: the previous preferred region only stretched up to twice smaller values of ξ and $10^4 a_{\rm dark}\xi^4$. However, the current run is not forced to explore a region in tension with the data, since the best-fit χ^2 only increases marginally (by 1.9) with respect to the best-fit ΛCDM χ^2.

Even if this model is not preferred by Planck+BAO data, it remains very interesting as a possible way to reconcile CMB+BAO data with high values of H_0 and low values of σ_8 (Lesgourgues et al. 2016; Buen-Abad et al. 2018). Indeed, we find that this model can accommodate a large $H_0 = 69.6^{+0.8}_{-1.3}$ (68% C.L.) reducing the tension with the most recent SH0ES data (Riess et al. 2019) from 4.1 σ to 2.7 σ, and a low $\sigma_8 = 0.813^{+0.015}_{-0.012}$ (68% C.L.). It also allows for smaller values of the parameter combination $S_8 \equiv \sigma_8\sqrt{\Omega_{\rm m}/0.3} = 0.813^{+0.015}_{-0.012}$ (68% C.L.)[3] than the ΛCDM model which gives $S_8 = 0.8235^{+0.0088}_{-0.0091}$ (68% C.L.) for the same data set. Thus it increases the compatibility with the combined KiDS+VIKING-450+DES-Y1 measurement

[3]For this model, we find exactly the same bounds on σ_8 and S_8, because $\Omega_{\rm m}$ remains very close to 0.3.

Parameter	ΛCDM	$\Gamma_0 > 0$, $\Delta N_{\rm fluid} > 0.07$	$\Gamma_0 > 0$, $-5 \leq \log(\Delta N_{\rm fluid}) \leq 0$
$100\ \omega_{\rm b}$	$2.219^{+0.013}_{-0.014}$	$2.232^{+0.017}_{-0.019}$	$2.219^{+0.014}_{-0.016}$
$\omega_{\rm cdm}$	$0.1192^{+0.0011}_{-0.0010}$	$0.1249^{+0.0023}_{-0.0037}$	$0.1192^{+0.0011}_{-0.0011}$
$\log(10^{10} A_{\rm s})$	$3.050^{+0.023}_{-0.023}$	$3.069^{+0.026}_{-0.025}$	$3.054^{+0.025}_{-0.026}$
n_s	$0.9618^{+0.0042}_{-0.0041}$	$0.9653^{+0.0042}_{-0.0045}$	$0.9617^{+0.0042}_{-0.0045}$
$\tau_{\rm reio}$	$0.060^{+0.012}_{-0.012}$	$0.0696^{+0.013}_{-0.013}$	$0.06181^{+0.013}_{-0.014}$
H_0 / [km/(s Mpc)]	$67.94^{+0.46}_{-0.49}$	$69.55^{+0.84}_{-1.3}$	$67.94^{+0.48}_{-0.50}$
σ_8	$0.8234^{+0.0085}_{-0.0090}$	$0.813^{+0.015}_{-0.012}$	$0.806^{+0.029}_{-0.011}$
$n_{\rm eff}$	$-2.308^{+0.0034}_{-0.0035}$	$-2.332^{+0.018}_{-0.011}$	$-3.261^{+0.96}_{-0.36}$
Γ_0 / $\left[\mathrm{Mpc}^{-1}\right]$	–	$< 1.2 \cdot 10^{-7}$	$< 1.5 \cdot 10^{-5}$
$\Delta N_{\rm fluid}$	–	< 0.59	–
$\log_{10}(\Delta N_{\rm fluid})$	–	–	< -2.66
$\Delta\chi^2$	–	1.90	2.34

Table 4.4: 68% C.L. constraints (or 95% upper bound in some cases) for all relevant parameters for the NADM case with two different prior choices, and using Planck+BAO. The $\Delta\chi^2$ is given with respect to ΛCDM with the same data sets.

by Joudaki et al. (2019) from 2.3 σ to 1.8 σ level[4]. The physical explanation is that this model is able to exploit the $H_0 - \Delta N_{\rm fluid}$ degeneracy thanks to its self-interacting DR component, while reducing at the same time the small-scale matter power spectrum amplitude thanks to the effect of DR dragging DM perturbations.

These results are consistent with those based on the previous ETHOS $n = 0$ parametrisation (with flat priors on ξ and $\log_{10}(a_{\rm dark})$ and the same data set), although the comparison is not straightforward since the new run explores a different region of the parameter space. The previous results did show the trend to accommodate a larger H_0 when $\Delta N_{\rm fluid}$ increases. This is even clearer in this run that reaches higher values of $\Delta N_{\rm fluid}$. The previous results also showed that when the interaction rate increases from $\log_{10}(a_{\rm dark}) \simeq -2$ to $\log_{10}(a_{\rm dark}) \simeq 0$, smaller values of σ_8 can be reached. This is confirmed in the new run by the clear correlation between the interaction rate and σ_8 in the left panel of Figure 4.6.

The comparison between the two runs alerts us on the fact that the ability of this model to reconcile data sets depends on the priors: the model would appear less effective in this respect with a lower prior edge $\Delta N_{\rm fluid} \geq 0$ (or with logarithmic priors on $\Delta N_{\rm fluid}$ or Γ_0). This prior-dependence of the conclusions applies anyway to most of the models attempting to resolve the tensions, and would only go away if we included the anomalous H_0 and σ_8 data in the analysis: then, even with

[4]Measurements of S_8 from weak lensing surveys are still very debated and potentially affected by poorly known systematics; for instance, the independent analysis by Fluri et al. (2019) gives a result compatible with our ΛCDM S_8 bounds at the 1.2 σ level.

different priors, some non-zero values of $\Delta N_{\rm fluid}$ and Γ_0 would be preferred with a statistical significance of a few sigmas.

CMB constraints with a logarithmic prior on $\Delta N_{\rm fluid}$. Buen-Abad et al. (2018) explored the same model with a flat prior on $-5 \leq \log_{10}(\Delta N_{\rm fluid}) \leq 0$ and on $\Gamma_0 \geq 0$. The motivation for this choice was to provide complementary results to the previous case, exploring very small values of the DR density which can always be motivated by specific particle physics constructions. We now update these results with our Planck+BAO data set, still not using Lyman-α data here, as we have seen that our method cannot provide accurate constraints for these models.

Our results are presented in the right panel of Figure 4.6 and right column of Table 4.4. We find $-5 \leq \log_{10} \Delta N_{\rm fluid} \leq -2.66$ (95% C.L.), corresponding to $0.04 < \xi < 0.15$ with a non-flat prior on ξ, and $\Gamma_0 < 1.5 \cdot 10^{-5}{\rm Mpc}^{-1}$ (95% C.L.), corresponding to $a_{\rm dark}\xi^4 < 0.45$ (95% C.L.). With this prior choice, we no longer allow for larger H_0, which is in agreement with our ETHOS $n = 0$ results. This can be understood in the following way: the flat prior on ξ (and indeed the logarithmic prior on $\Delta N_{\rm fluid}$) gives less weight to large amounts of DR, and thus the possibility to relax the H_0 tension goes away. However, we can accommodate a lower σ_8, thanks to a degeneracy between σ_8 and $\log_{10} \Delta N_{\rm fluid}$ that is clearly visible in the right panel of Figure 4.6. This degeneracy is equivalent to the $\sigma_8 - \xi$ degeneracy previously observed in the ETHOS $n = 0$ results, and could in principle reconcile the Planck+BAO data with values as low as $\sigma_8 \sim 0.7$. The model predicts $S_8 = 0.8058^{+0.0088}_{-0.0085}$ (68% C.L.), which is compatible with KiDS+VIKING-450+DES-Y1 (Joudaki et al. 2019) at 1.7 σ level.

4.4 Towards a more flexible parametrisation

All the results discussed in this Chapter have been obtained under the assumption that the whole of DM is interacting. Therefore, it is paramount to extend our analyses to scenarios in which only a fraction of DM is non-standard.

In Chapter 2 we have already stressed that our $\{\alpha, \beta, \gamma\}$-parametrisation is capable to capture models where the DM sector is made by a mixture of standard CDM and thermal WDM (Section 2.3.3), as well as being easy to be further generalised by adding an extra free parameter (Section 2.3.5). Indeed, the goal of this Section is to briefly outline a further generalisation of our approach, which will enable one to test in full generality any mixture of CDM and nCDM. In order to do so, we introduce the following new analytic fitting formula (to be compared with Equation 2.4):

$$T(k) = (1 - \delta) \cdot [1 + (\alpha k)^\beta]^{-1.5 \cdot \beta} + \delta, \tag{4.11}$$

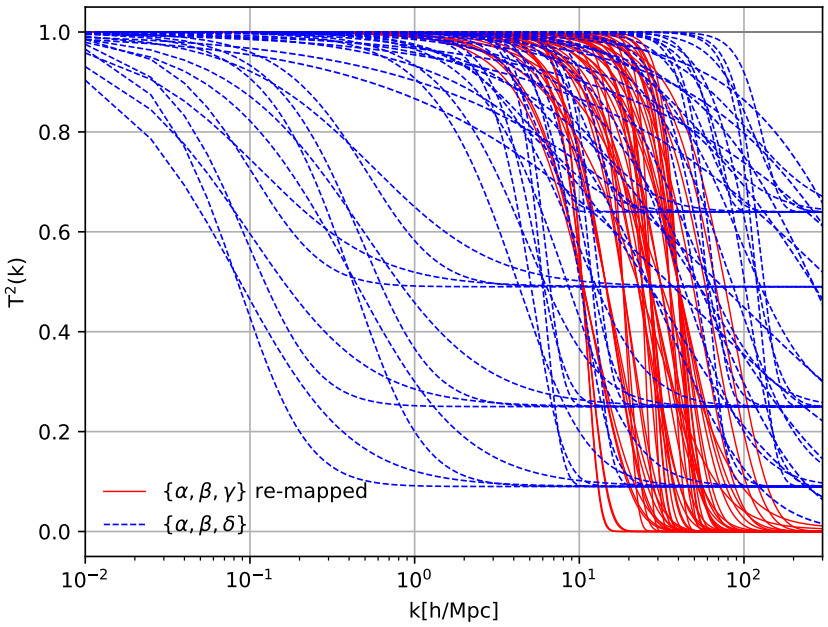

Figure 4.7: Here we show a set of squared transfer functions, each of them corresponding to a different $\{\alpha, \beta, \delta\}$-combination. The red solid curves coincide with $\{\alpha, \beta, \gamma\}$-triplets re-mapped to the new parametrisation, whereas the dashed blue lines have been chosen in order to fully exploit the flexibility given by the introduction of δ, i.e., to investigate in full generality mixtures of CDM and nCDM.

where the slope of the power suppression is primarily governed by the value of β (i.e., given a value for β, γ is now fixed). The motivation for this choice is threefold: firstly, the quasi-degeneracy between α and γ, discussed in Appendix A, allows to trade changes in the latter parameter for changes in the former one; secondly, as it has been shown in Chapter 2 (Section 2.6), current structure formation data do not have enough resolution to put tight constraints on the very large-k region of the matter power spectrum (see Figure 2.12, where γ is practically unconstrained); lastly, models where the whole of the DM is exotic imply relatively sharp power suppressions, allowed by Lyman-α data only if the nCDM candidate is so massive (or so "weakly" interacting) to be almost indistinguishable from pure CDM. On the other hand, scenarios where just a DM fraction is non-standard induce shallower transfer functions, featuring a large-k plateau corresponding to the remaining CDM component once the small-scale suppression by the non-standard component has died off. For this reason, the more flexible formula (Equation 4.11) is characterised by a new free parameter δ, directly responsible for the height of the plateau, i.e. for the relative abundances of the different DM components.

Notice that these arguments apply to any of the nCDM models that we analysed previously: see, e.g., Figure 3.3 for the Fuzzy DM case, from which the dependence of the bound on the DM particle mass with respect to its abundance is manifest. In fact, we are currently working on implementing another `MontePython` likelihood, devised along the lines of Section 4.2.1, but based on the $\{\alpha, \beta, \delta\}$-parametrisation. The goal is to use it to systematically scan the whole parameter space of both DM-DR and DM-baryon interaction scenarios, testing models with arbitrary fractions of interacting DM.

To visually illustrate the versatility of the new formula, in Figure 4.7 we show a set of squared transfer functions, each of them corresponding to a different $\{\alpha, \beta, \delta\}$-combination. The red solid curves coincide with $\{\alpha, \beta, \gamma\}$-triplets re-mapped to the new parametrisation, whereas the dashed blue lines have been chosen in order to fully exploit the flexibility given by the introduction of δ. Indeed, another nice feature of the newer framework is that its flexibility allowed us to re-map in terms of the new parameters the vast majority of the simulations that we had already performed (and used in this Chapter and in Chapter 2).

This Chapter is mainly based on:

R. Murgia, G. Scelfo, M. Viel, A. Raccanelli
Lyman-α forest constraints on Primordial Black Holes as Dark Matter
PRL 123, 7, 071102 (2019); arXiv:1903.10509

5 Primordial Black Holes as Dark Matter

5.1 Overview

The standard CDM scenario is characterised by two assumptions: the DM is made by heavy ($\gtrsim$ GeV), weakly interacting particles; and cosmic structures are formed from initial inhomogeneities through a hierarchical *bottom-up* process. Both these assumptions may be questioned and tested against astrophysical and cosmological observations: in the previous Chapters we focused on the possibility that the DM sector comprises some *non-cold* particles and/or is characterised by some non-standard interaction; in this Chapter we will instead investigate the intriguing possibility that (part of) the DM might be constituted by a population of Primordial Black Holes (PBHs) formed in the very early universe.

PBHs were first theorised decades ago (Hawking 1971) from the collapse of overdense regions originated by large curvature perturbations in the early universe. Many proposals have been made so far for the specific formation mechanism, such as collapsing large fluctuations produced during inflation (Ivanov et al. 1994; García-Bellido et al. 1996; Ivanov 1998), collapsing cosmic string loops (Hawking 1989; Polnarev & Zembowicz 1991; Wichoski et al. 1998), domain walls (Berezin et al. 1983; Ipser & Sikivie 1984), bubble collisions (Crawford & Schramm 1982; La & Steinhardt 1989), or collapse of exotic DM clumps (Shandera et al. 2018).

After the first GW detection revealed merging Black Hole (BH) binaries of masses $\mathcal{O}(10\ M_\odot)$ (Abbott et al. 2016a,b), the interest toward PBHs as DM candidates has been revived (Bird et al. 2016). Since then, several proposals to determine the nature of the merging BH progenitors have been made, involving methods such as GW×LSS cross-correlations (see, e.g., Raccanelli et al. (2016); Scelfo et al. (2018)),

BH binary eccentricities (Cholis et al. 2016), BH mass function studies (Kovetz et al. 2017; Kovetz 2017), and lensing of fast radio bursts (Muñoz et al. 2016).

Several constraints on the possibility that PBHs could compose a DM fraction have been determined through different astrophysical observables, such as gravitational lensing (Barnacka et al. 2012; Katz et al. 2018; Griest et al. 2014; Niikura et al. 2017; Tisserand et al. 2007; Calchi Novati et al. 2013; Alcock et al. 2001; Mediavilla et al. 2009; Wilkinson et al. 2001; Zumalacarregui & Seljak 2018), dynamical (Graham et al. 2015; Capela et al. 2013; Quinn et al. 2009; Brandt 2016; Ali-Haïmoud et al. 2017; Magee et al. 2018), and accretion effects (Gaggero et al. 2017; Ricotti et al. 2008; Ali-Haïmoud & Kamionkowski 2017; Poulin et al. 2017b; Bernal et al. 2017), exploring all possible PBH mass ranges. Nevertheless, varying the numerous assumptions involved might significantly alter these limits (see, e.g., Aloni et al. (2017); Bellomo et al. (2018); Nakama et al. (2017)). Therefore, the investigation towards PBHs as DM candidates is still fully open. In particular, two mass "windows" are currently of large interest: one around $\mathcal{O}(10^{-10} M_{\odot})$, and another around $\mathcal{O}(10\ M_{\odot})$; see Sasaki et al. (2018); Carr & Silk (2018); Kashlinsky et al. (2019) for details on the state-of-the-art development of the field.

Another suitable and mostly unexplored method for constraining the mass and abundance of PBHs is offered by the Lyman-α forest. As we extensively discussed in the previous Chapters, such observable represents a very powerful tool for tracing the DM distribution at (sub-)galactic scales. Lyman-α forest data were used almost two decades ago to set an upper limit of few $10^4 M_{\odot}$ on PBH masses, in the simple case in which all DM is made by PBHs with the same mass (Afshordi et al. 2003). In this Chapter, based on Murgia et al. (2019), we update and improve such limit, by using the highest-resolution Lyman-α forest data set up-to-date (Viel et al. 2013), and a new set of high-resolution hydrodynamic simulations that allows a more precise modelling of the full 1D flux power. Furthermore, we generalise our results to different PBH fractions with respect to the total DM amount, and to non-monochromatic mass distributions.

This Chapter is organised as follows: in Section 5.2 we discuss the impact on the linear matter power spectrum due to the existence of PBHs; in Section 5.3 we extend the discussion to non-monochromatic PBH mass distributions; in Section 5.4 we present the data set and the methods used for our analyses; in Section 5.5 we report and discuss our results.

5.2 Impact on the linear matter power spectrum

Stellar-mass PBHs would cause observable effects on the matter power spectrum; due to discreetness, a small-scale plateau in the linear matter power spectrum is induced by a Poisson noise contribution (Meszaros 1975; Afshordi et al. 2003; Carr & Silk 2018; Gong & Kitajima 2017).

Let us firstly consider the simple case where PBHs are characterised by a Monochromatic Mass Distribution (MMD), i.e., they all have the same mass $M_{\rm PBH}$. In this scenario, PBHs are parametrised only by their mass and abundance, so that the fraction parameter $f_{\rm PBH} \equiv \Omega_{\rm PBH}/\Omega_{\rm DM} = 1$ in the case where the whole of the DM is made of PBHs.

Assuming that PBHs are randomly distributed, their number follows a Poisson distribution, and each wavenumber k can be associated to an overdensity $\delta_{\rm PBH}(k)$, due to Poisson noise. The PBH contribution to the power spectrum is thus defined as

$$P_{\rm PBH}(k) = \langle|\delta_{\rm PBH}(k)|^2\rangle = \frac{1}{n_{\rm PBH}}, \tag{5.1}$$

where $n_{\rm PBH}$ is the comoving PBH number density, i.e.

$$n_{\rm PBH} = \frac{\Omega_{\rm DM}\rho_{\rm cr}f_{\rm PBH}}{M_{\rm PBH}}, \tag{5.2}$$

with $\rho_{\rm cr}$ being the critical density of the universe. Since $n_{\rm PBH}$ is a k-independent quantity, $P_{\rm PBH}$ is scale-invariant. In other words, as anticipated above, the existence of PBHs induces a small-scale plateau departing from the standard ΛCDM prediction.

A crucial concept pointed out by Afshordi et al. (2003); Gong & Kitajima (2017) is that the PBH overdensity $\delta_{\rm PBH}$ can be interpreted as an isocurvature perturbation (see also Gong & Kitajima (2018) for the case of a Lognormal spatial distribution of PBHs). Hence, the CDM power spectrum can be written as:

$$P_{\rm CDM}(k,z) = D^2(z)\left(T^2_{\rm ad}(k)P_{\rm ad} + T^2_{\rm iso}(k)P_{\rm iso}\right), \tag{5.3}$$

where $D(z)$ is the growth factor, $P_{\rm iso}$ is the isocurvature power spectrum, and $P_{\rm ad} \propto A_s k^{n_s}$ is the primordial adiabatic power spectrum. $T_{\rm ad}$ and $T_{\rm iso}$ are the adiabatic and isocurvature transfer functions, respectively (see, e.g., Bardeen et al. (1985) for their analytic expressions). The PBH linear power spectrum is thus defined by:

$$P_{\rm iso} = f^2_{\rm PBH}P_{\rm PBH} = \frac{2\pi^2}{k^3}A_{\rm iso}\left(\frac{k}{k_*}\right)^{n_{\rm iso}-1}, \tag{5.4}$$

where we set the pivot scale $k_* = 0.05/\mathrm{Mpc}$, and the primordial isocurvature tilt $n_{\rm iso} = 4$ in order to ensure the scale-invariance of the power spectrum. Given that the

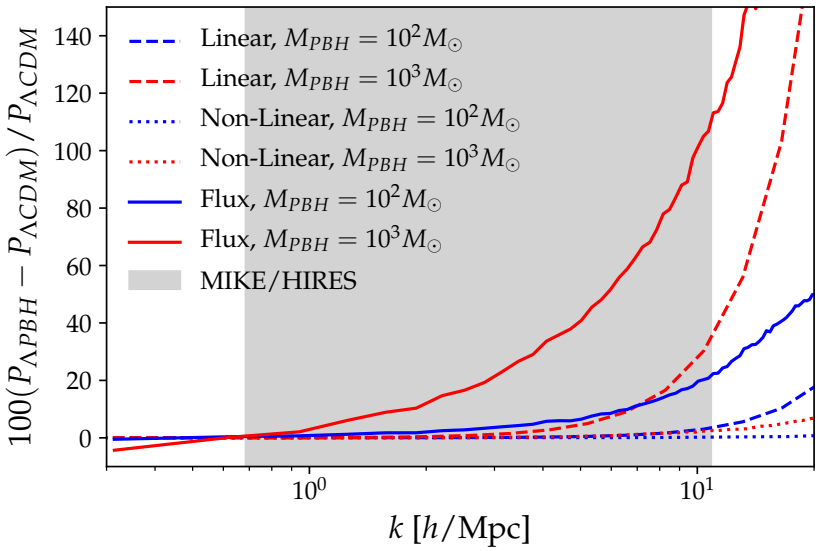

Figure 5.1: Relative difference, at redshift $z = 5$, between ΛCDM and ΛPBH scenarios for the linear matter (dashed), non-linear matter (dotted) and 1D flux (solid) power spectra. Blue/red colours correspond to $M_{\rm PBH} = \{10^2, 10^3\}\ M_\odot$, respectively, for a monochromatic scenario with $f_{\rm PBH} = 1$. The grey shaded area refers to the scales covered by Lyman-α forest data.

adiabatic power spectrum evolves as k^{-3} at large k's, the isocurvature contribution is expected to become important only at the scales probed by the Lyman-α forest. The parameter $A_{\rm iso}$ sets the amplitude of the isocurvature modes, depending on the PBH mass of the considered model. Indeed, by combining the equations reported in this Section, we can eventually express the isocurvature-to-adiabatic amplitude ratio as:

$$f_{\rm iso} = \sqrt{\frac{A_{\rm iso}}{A_s}} = \sqrt{\frac{k_*^3 f_{\rm PBH}^2}{2\pi^2 n_{\rm PBH}} \frac{1}{A_{\rm s}}} = \sqrt{\frac{k_*^3 M_{\rm PBH} f_{\rm PBH}}{2\pi^2 \Omega_{\rm CDM} \rho_{\rm cr}} \frac{1}{A_{\rm s}}}\,, \tag{5.5}$$

where the last equality holds only for MMDs. In this case it is immediate to see a degeneracy between $f_{\rm PBH}$ and $M_{\rm PBH}$: different combinations of PBH mass and abundance correspond to the same isocurvature-to-adiabatic amplitude ratio if the quantity $f_{\rm PBH} M_{\rm PBH}$ is the same.

In our framework, the effect on the linear matter power spectrum due to the presence of isocurvature modes consists of a power enhancement with respect to the standard ΛCDM spectrum, in the form of a small-scale plateau. As a straightforward consequence of Equation (5.1), the larger is the PBH mass, the stronger is the small-scale power enhancement.

In Figure 5.1 we provide the relative differences with respect to a pure ΛCDM scenario for the 3D linear and non-linear matter power spectra, at redshift $z = 5$, for ΛPBH models with $M_{\rm PBH} = \{10^2, 10^3\} M_\odot$, assuming $f_{\rm PBH} = 1$. We also show the 1D flux power spectra, which are the Lyman-α forest observables, associated to

the same ΛPBH models. The grey shaded area refers to the scales covered by our Lyman-α data set, obtained from MIKE/HIRES spectrographs (see Section 5.4 for details). The non-linear power spectra have been extracted from the snapshots of cosmological simulations, thus they include both the linear contribution (encoded in the initial conditions) and, on top of that, the effects of the non-linear evolution computed by the numerical simulation itself. The PBH contribution is thereby included in the initial conditions, whereas during the non-linear evolution both the isocurvature and adiabatic DM modes are treated as cold and collisionless (see Section 5.4). It can be easily seen how non-linearities in the 3D matter power spectrum wash out the differences induced by the presence of PBHs. Conversely, as we already stressed in the previous Chapters, the 1D flux spectra are a much more effective observable to probe the small-scale power, being a projection of the 3D ones.

5.3 Extended Mass Distributions

The PBH formation is, in the most standard case, a consequence of large perturbations in the primordial power spectrum; while the exact details of the peak required to form PBH and how this is linked to the real-space overdensities are still unclear, it seems not unlikely that the most realistic scenario involves the presence of a PBH population with an extended mass function. Moreover, a non-monochromatic mass distribution would be created by different merger and accretion history of each PBH. Focusing on EMDs is also intriguing due to the possibility that PBHs in the high-mass tail of the distribution may account for the seeds of supermassive BHs (Bernal et al. 2018; Carr & Silk 2018). A general method to convert MMD constraints to limits on Extended Mass Distributions (EMDs) has been developed by Carr et al. (2017); Bellomo et al. (2018).

The extension to EMDs of the observable considered for our analyses arises naturally from the second equality in Equation (5.5), by directly considering the PBH number density corresponding to a given EMD. Let us then consider EMDs in the form:

$$\frac{dn_{\rm PBH}}{d\ln M_{\rm PBH}} = f_{\rm PBH}\rho_{\rm DM}\frac{d\Phi_{\rm PBH}}{dM_{\rm PBH}}\,, \tag{5.6}$$

where the function $d\Phi_{\rm PBH}/dM_{\rm PBH}$ describes the shape of the EMD, and $\rho_{\rm DM} = \Omega_{\rm DM}\rho_{\rm cr}$. Given a certain EMD, one can define the so-called Equivalent Mass $M_{\rm eq}$, which is the mass of a MMD providing the same observational effect (i.e., in our case, the same $A_{\rm iso} \propto f^2_{\rm PBH}/n_{\rm PBH}$) (Bellomo et al. 2018).
Therefore, the conversion from MMD($M_{\rm eq}$) to EMD($M_{\rm PBH}$) is given by

$$f^2_{\rm PBH}\left[\frac{\Omega_{\rm DM}\rho_{\rm cr}f_{\rm PBH}}{M_{\rm eq}}\right]^{-1} = \frac{f^2_{\rm PBH}}{n_{\rm PBH}} = f^2_{\rm PBH}\left[\int\frac{dn_{\rm PBH}}{dM_{\rm PBH}}dM_{\rm PBH}\right]^{-1} \tag{5.7}$$

where we assume that the PBH abundances are the same for both the MMD and EMD cases. We finally have:

$$M_{\rm eq} = \left[\int \frac{1}{M_{\rm PBH}} \frac{d\Phi}{dM_{\rm PBH}} dM_{\rm PBH} \right]^{-1} . \tag{5.8}$$

Let us now estimate $M_{\rm eq}$ for two popular EMDs: *Lognormal* and *Powerlaw*.

The *Lognormal* EMD (Dolgov & Silk 1993) is defined by

$$\frac{d\Phi_{\rm PBH}}{dM_{\rm PBH}} = \frac{\exp\left\{-\frac{\ln^2(M_{\rm PBH}/\mu)}{2\sigma^2}\right\}}{\sqrt{2\pi}\sigma M_{\rm PBH}} , \tag{5.9}$$

where σ and μ are the standard deviation and the mean of the PBH mass, respectively. Such function is representative of a wide family of EMDs, since it describes well the scenario of PBHs forming from a smooth symmetric peak in the inflationary power spectrum (Green 2016; Kannike et al. 2017).

The *Powerlaw* EMD, corresponding to PBHs formed from collapsing cosmic strings or scale-invariant density fluctuations (Carr 1975), is given by

$$\frac{d\Phi_{\rm PBH}}{dM_{\rm PBH}} = \frac{\mathcal{N}_{\rm PL}}{M_{\rm PBH}^{1-\tilde{\gamma}}} \Theta(M_{\rm PBH} - M_{\rm min}) \Theta(M_{\rm max} - M_{\rm PBH}) , \tag{5.10}$$

characterised by an exponent $\tilde{\gamma} \in (-1, +1)$, a mass interval $\{M_{\rm min}, M_{\rm max}\}$, and a normalisation factor $\mathcal{N}_{\rm PL}$; Θ is the Heaviside step function.

5.4 Data set and methods

In order to extract limits on the properties of PBHs as DM candidates from the Lyman-α forest, we adapted the method proposed in Murgia et al. (2018) and discussed in Chapter 2 in the context of suppressed linear matter power spectra. To do so, we have built a new grid of hydrodynamic simulations in terms of the properties of PBHs, corresponding to initial linear matter power spectra featuring a small-scale plateau. Beside that, our analyses rely on a pre-computed multidimensional grid of hydrodynamic simulations, associated to several values of the astrophysical and cosmological parameters affecting the Lyman-α flux power spectrum, sampling all the viable volume of the corresponding parameter space. The simulations have been performed with `GADGET-3`, a modified version of the publicly available numerical code `GADGET-2` (Springel et al. 2001b; Springel 2005). The initial conditions have been produced with the numerical code `2LPTic` (Crocce et al. 2006), at redshift $z = 199$, with input linear matter power spectra for the ΛPBH models obtained by turning on the isocurvature mode in the numerical Boltzmann solver `CLASS` (Blas et al. 2011).

As in Chapters 2 and 4, our reference model simulation follows Murgia et al. (2018); Iršič et al. (2017b); it has a box length of $20/h$ comoving Mpc with 2×768^3 gas and CDM particles in a flat ΛCDM universe with cosmological parameters (Ade et al. 2016a): $\Omega_m = 0.301$, $\Omega_b = 0.0457$, $n_s = 0.961$, $H_0 = 70.2$ km s^{-1} Mpc^{-1}, $\sigma_8 = 0.829$, and $z_{\rm reio} = 9$.

For the cosmological parameters to be varied, we sample different values of σ_8, i.e., the normalisation of the linear matter power spectrum, and $n_{\rm eff} = d\ln P_m(k)/d\ln k|_{k_\alpha}$, namely the slope of the matter power spectrum evaluated at the scale probed by the Lyman-α forest ($k_\alpha = 0.009$ s/km) (Seljak et al. 2006; McDonald et al. 2006; Arinyo-i Prats et al. 2015). Hence, we have included five different simulations for both σ_8 (in the range $[0.754, 0.904]$) and $n_{\rm eff}$ (in the interval $[-2.3474, -2.2674]$). Additionally, we included simulations corresponding to different values for the instantaneous reionisation redshift, i.e., $z_{\rm reio} = \{7, 9, 15\}$.

Regarding the astrophysical parameters, we modelled the thermal history of the IGM with amplitude T_0 and slope γ of its temperature-density relation, parametrised as $T = T_0(1+\delta_{\rm IGM})^{\gamma-1}$, with $\delta_{\rm IGM}$ being the IGM overdensity (Hui & Gnedin 1997). We ran simulations with temperatures at mean density $T_0(z = 4.2) = \{6000, 9200, 12600\}$ K, evolving with redshift, as well as a set of three values for the slope of the temperature-density relation, $\gamma(z = 4.2) = \{0.88, 1.24, 1.47\}$. The redshift evolution of both T_0 and γ are parametrised as power laws, such that $T_0(z) = T_0^A[(1+z)/(1+z_p)]^{T_0^S}$ and $\gamma(z) = \gamma^A[(1+z)/(1+z_p)]^{\gamma^S}$, where the pivot redshift z_p is the redshift at which most of the Lyman-α forest pixels are coming from ($z_p = 4.5$ for MIKE/HIRES). The reference thermal history is defined by $T_0(z = 4.2) = 9200$ and $\gamma(z = 4.2) = 1.47$, since such values provide a good fit to observations (Bolton et al. 2017). Furthermore, we have considered the effect of ultraviolet (UV) fluctuations of the ionising background, the impact of which is controlled by the parameter $f_{\rm UV}$. Its template is built from a set of three simulations with $f_{\rm UV} = \{0, 0.5, 1\}$, where $f_{\rm UV} = 0$ corresponds to a spatially uniform UV background (Iršič et al. 2017b). We have also included 9 grid points obtained by re-scaling the mean Lyman-α forest flux $\langle F(z)\rangle$, namely $\{0.6, 0.7, 0.8, 0.9, 1.0, 1.1, 1.2, 1.3, 1.4\} \times \langle F_{\rm REF}\rangle$, with the reference values being the ones of the SDSS-III/BOSS measurements (Palanque-Delabrouille et al. 2013). With the goal to have a more refined grid in terms of mean fluxes, we have also included 8 additional values, obtained by re-scaling the optical depth $\tau = -\ln\langle F\rangle$, i.e. $\{0.6, 0.7, 0.8, 0.9, 1.1, 1.2, 1.3, 1.4\} \times \tau_{\rm REF}$.

Concerning the PBH properties, we have extracted the flux power spectra from 12 hydrodynamic simulations (512^3 particles; 20 comoving Mpc/h box length) corresponding to the following PBH mass and fraction products: $\log(M_{\rm PBH} f_{\rm PBH}) = \{1.0, 1.5, 2.0, 2.2, 2.3, 2.4, 2.5, 2.6, 2.7, 3.0, 3.5, 4.0\}$[1], where $M_{\rm PBH}$ is expressed in units

[1]Hereafter $\log \equiv \log_{10}$.

	Flat prior on $z_{\rm reio}$		Gaussian prior on $z_{\rm reio}$	
Parameter	**(2 σ)**	**Best Fit**	**(2 σ)**	**Best Fit**
$\langle F(z=4.2)\rangle$	[0.35, 0.41]	0.37	[0.35, 0.41]	0.37
$\langle F(z=4.6)\rangle$	[0.26, 0.34]	0.28	[0.27, 0.34]	0.28
$\langle F(z=5.0)\rangle$	[0.15, 0.25]	0.20	[0.15, 0.23]	0.16
$\langle F(z=5.4)\rangle$	[0.03, 0.12]	0.08	[0.04, 0.11]	0.05
T_0^A [10^4 K]	[0.44, 1.36]	0.72	[0.46, 1.44]	0.84
T_0^S	[-5.00, 3.34]	-4.47	[-5.00, 3.35]	-4.53
γ^A	[1.21, 1.60]	1.51	[1.19, 1.61]	1.44
γ^S	[-2.43, 1.30]	-1.76	[-2.25, 1.51]	0.46
σ_8	[0.72, 0.91]	0.79	[0.72, 0.91]	0.81
$z_{\rm reio}$	[7.00, 15.00]	14.19	[7.12, 10.25]	9.07
$n_{\rm eff}$	[-2.40, -2.22]	-2.30	[-2.41, -2.22]	-2.33
$f_{\rm UV}$	[0.00, 1.00]	0.02	[0.00, 1.00]	0.03
$\log(f_{\rm PBH}M_{\rm PBH})$	< 2.24	1.96	< 1.78	0.34
χ^2/d.o.f.		32/42		33/43

Table 5.1: 2 σ limits and best fit values for all the parameters of our analyses, for the two different prior choices on $z_{\rm reio}$ that we adopted. The values for $M_{\rm PBH}$ are expressed in units of $M_\odot$.

of $M_\odot$. For this set of simulations, all the astrophysical and cosmological parameters have been fixed to their reference values, and the equivalent ΛCDM flux power spectrum has also been determined.

As in Chapters 2 and 4, our analysis is based on the *Ordinary Kriging* interpolation method (Webster & Oliver 2007), and the interpolation is done in terms of ratios between the flux power spectra of the ΛPBH models and the reference ΛCDM one. Along the lines of Murgia et al. (2018), we first interpolate in the astrophysical and cosmological parameter space for the standard ΛCDM case; we then correct all the $\{M_{\rm PBH}f_{\rm PBH}\}$-grid points accordingly, and we finally interpolate in the $\{M_{\rm PBH}f_{\rm PBH}\}$-space.

Our data set is constituted by the MIKE/HIRES samples of quasar spectra, obtained with the MIKE/Magellan spectrographs and the HIRES/KECK, at redshift $z = \{4.2, 4.6, 5.0, 5.4\}$, in 10 k-bins in the range $[0.001 - 0.08]$ s/km, with spectral resolution of 13.6 and 6.7 s/km, respectively (Viel et al. 2013). As in Chapters 2 and 4, we applied a cut on the flux power spectra, by using only the measurements at $k > 0.005$ s/km, to avoid large-scale systematic uncertainties due to continuum fitting. Moreover, we did not include in our analyses the highest redshift bin for MIKE data, for which the errors on the flux power spectra are very large (Viel et al. 2013). We have thus used a total of 49 (k, z) data points.

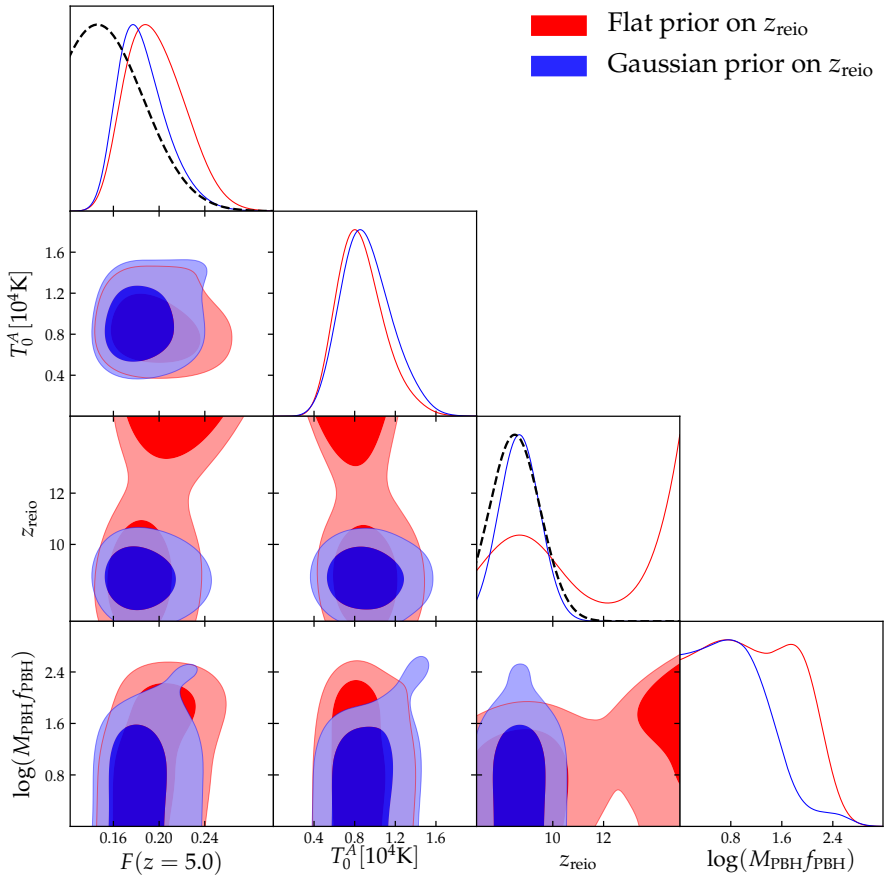

Figure 5.2: 1 and 2 σ contour plots for some of the parameters of our analyses, for the two different prior choices on $z_{\rm reio}$ that we adopted. The values for $M_{\rm PBH}$ are expressed in units of $M_\odot$. The dashed lines represent the Gaussian priors imposed on $\bar{F}(z=5)$ and $z_{\rm reio}$, with the latter referring to the blue plots.

5.5 Results and discussion

We have determined the constraints on both astrophysical, cosmological, and PBH properties, by maximising a Gaussian likelihood, using the publicly available affine-invariant MCMC sampler `emcee` (Foreman-Mackey et al. 2013). We adopted Gaussian priors on the mean fluxes $\langle F(z)\rangle$, centred on their reference values, with standard deviation $\sigma = 0.04$ (Iršič et al. 2017b), and on σ_8 and $n_{\rm eff}$, centred on their Planck values (Ade et al. 2016a), with standard deviation $\sigma = 0.05$, given that the latter two parameters, whereas very well constrained by CMB data, are poorly constrained by Lyman-α forest data alone (Murgia et al. 2018). Concerning the IGM thermal history, we adopt flat priors on both T_0^A and T_0^S, in the ranges $[0,2]\cdot 10^4$ K and $[-5,5]$, respectively. When the corresponding $T_0(z)$ are determined, they can assume values not enclosed by our template of simulations. When this occurs, the corresponding values of the flux power spectra are linearly extrapolated. Regarding γ^S and γ^A, we impose flat priors on the

corresponding $\gamma(z)$ (in the interval $[1, 1.7]$). The priors on $z_{\rm reio}$ and $f_{\rm UV}$ are flat within the boundaries defined by our grid of simulations. We adopt a logarithmic prior on $f_{\rm PBH}M_{\rm PBH}$. However, let us stress that our results are not affected by this choice. By re-running our analyses assuming a flat prior on it, the constraints are indeed unmodified, provided that a conservative Gaussian prior on the IGM temperature is also imposed, namely $T_0^A = 7500 \pm 1500$ (1σ) (motivated by the most up-to-date IGM studies by Boera et al. (2018)). All the other priors are flat within the boundaries defined by our grid of simulations.

Let us firstly focus on the simple case of PBHs featuring a MMD. In Table 5.1 we report the marginalised 2 σ constraints and the best fit values for all the parameters considered in our analyses. The first two columns refer to the case in which a flat prior is applied to the reionisation redshift. The limit on the PBH abundance under this assumption corresponds to

$$f_{\rm PBH}M_{\rm PBH} \lesssim 170\ M_\odot\ (2\ \sigma), \tag{5.11}$$

However, both Planck CMB results and measurements by Boera et al. (2018) clearly favour $z_{\rm reio}$ values around 8.5. For this reason, we have repeated our analysis by imposing on the reionisation redshift a Gaussian prior centred around $z_{\rm reio} = 8.5$, with standard deviation $\sigma = 1.0$. The results obtained under such assumption are shown in the last two columns of Table 5.1, and in this case we have

$$f_{\rm PBH}M_{\rm PBH} \lesssim 60\ M_\odot\ (2\ \sigma). \tag{5.12}$$

In the simplest case where all DM is made by PBHs ($f_{\rm PBH} = 1$), these constraints can be interpreted as absolute limits on the PBH mass. On the other hand, such absolute upper bounds on $M_{\rm PBH}$ weaken linearly when a smaller PBH abundance is assumed $(0.05 < f_{\rm PBH} < 1)$[2].

The degeneracy between $z_{\rm reio}$ and the PBH mass can be understood as follows: a higher reionisation redshift, having fixed the thermal history at late times, corresponds to a more effective (i.e., larger) filtering scale, and thus to a suppression of power that is compensated by a larger value of the PBH mass. The fact that the degeneracies are much more prominent for this parameter, unlike the nCDM case in which the cut-off is more degenerate with the thermal cut-off (see Chapters 2 and 3), is telling us that the increase of power at small scales is a distinctive feature whose effect is more likely to be degenerate with a different gas filtering scale.

In Figure 5.2 we show the 1 and 2 σ contour plots for some of the parameters considered in our analyses, for both the aforementioned prior choices on $z_{\rm reio}$. The

[2]The lower limit on $f_{\rm PBH}$ is given by the fact that, for the monochromatic case, at $z = 199$, i.e. the redshift of the initial conditions of our simulations, if $f_{\rm PBH}$ is smaller, the Poisson effect is subdominant with respect to the *seed effect*, the treatment of which goes beyond our purposes (see Carr & Silk (2018) for further details).

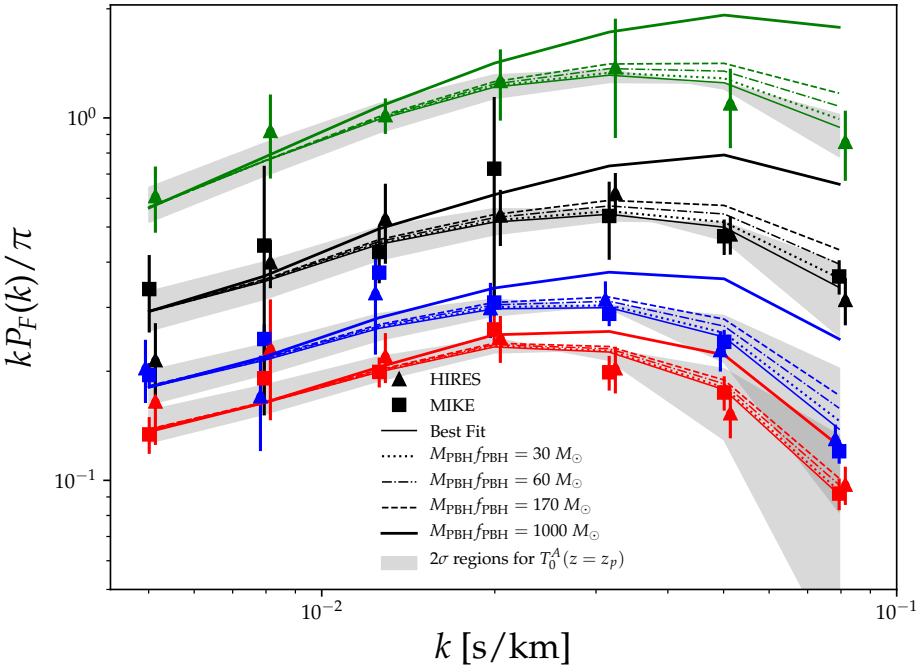

Figure 5.3: 1D flux spectra for ΛCDM and ΛPBH, for different PBH masses. Symbols are data from MIKE/HIRES, lines are obtained by interpolating in the $\{M_{\rm PBH}f_{\rm PBH}\}$-space defined by our simulations; while the best fit is technically for $M_{\rm PBH} \neq 0$, it is indistinguishable from the ΛCDM case. Red, blue, black and green indicate $z = 4.2, 4.6, 5.0, 5.4$, respectively. The grey shaded areas represent regions sampled by flux power spectra corresponding to values for T_0^A spanning its marginalised 2 σ interval.

mild degeneracy between the amplitude of the IGM temperature $T_0^A(z = z_p)$ and the PBH mass derives from the opposite effects on the flux power spectra due to the increase of the two parameters. A hotter IGM implies a small-scale power suppression which can be balanced by increasing $M_{\rm PBH}f_{\rm PBH}$. Slightly larger values for the mean fluxes $\langle F(z)\rangle$ are also required for accommodating the power enhancement induced by relatively large values of the PBH mass. The dashed lines represent the Gaussian priors imposed on $\bar{F}(z = 5)$ and $z_{\rm reio}$, with the latter referring to the blue plots. Note that our MCMC analyses favour higher values for $\bar{F}(z = 5)$ (still in agreement with its prior distribution), allowing in turn a larger power enhancement due to PBHs. This is a further hint of the conservativity of our constraints.

In order to test the stability of our results, we also performed an analysis with flat priors both on σ_8 and $n_{\rm eff}$. Under these assumptions, the constraint (using a Gaussian prior for $z_{\rm reio}$) on $f_{\rm PBH}M_{\rm PBH}$ is mildly weakened, up to 100 $M_\odot$. However, the largest values for the PBH mass are allowed only in combination with extremely low values for $n_{\rm eff}$, allowed in turn by our data set due to its poor constraining power on such parameter. As we have already stated, this is the main reason to impose a (still conservative) Gaussian prior motivated by CMB measurements on $n_{\rm eff}$.

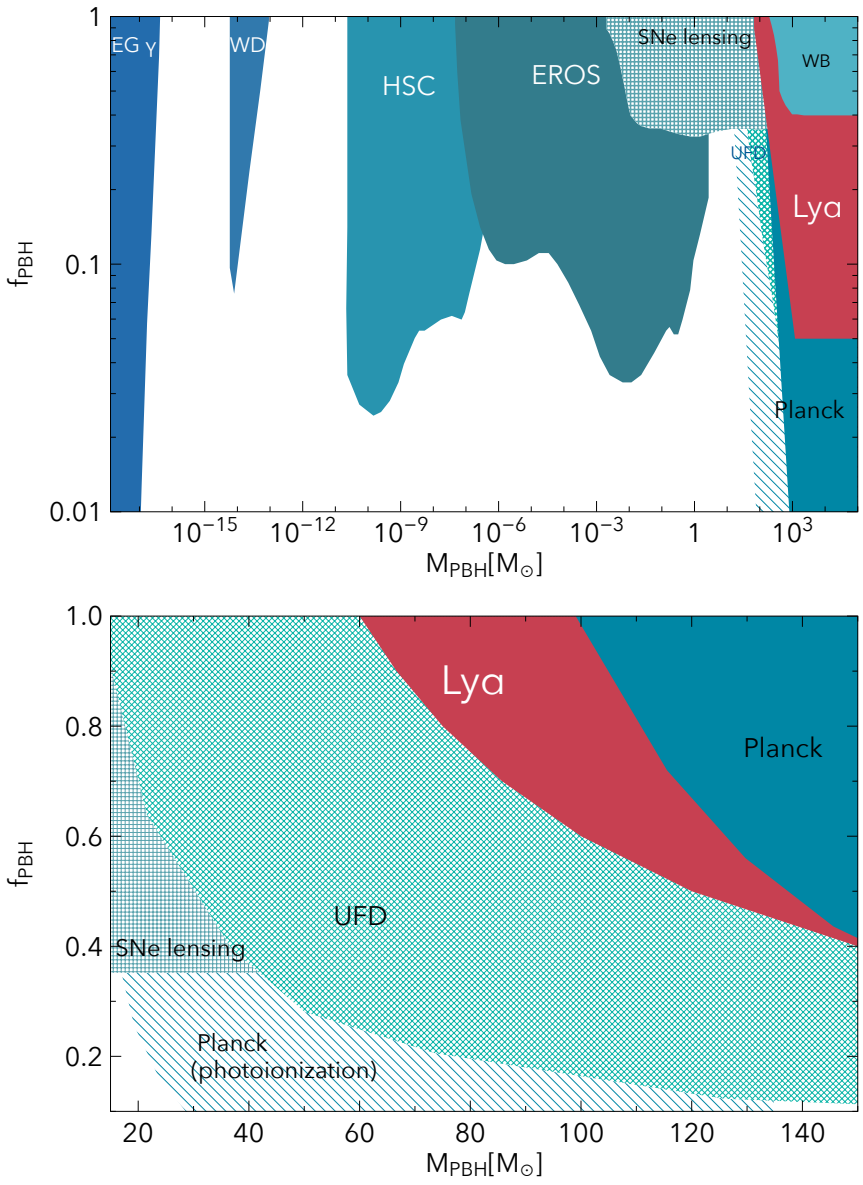

Figure 5.4: Overview on the present experimental constraints on the abundance of PBH for a monochromatic mass distribution (from Bartolo et al. (2019); Zumalacarregui & Seljak (2018); Ali-Haïmoud & Kamionkowski (2017)), in shades of blue, and in red are the limits that we have obtained. Patterned areas show limits that are the most dependent on astrophysical assumptions (see, e.g. Carr et al. (2016)). The bottom panel is zoomed in the LIGO mass range.

In Figure 5.3 we show the dimensionless 1D flux power spectra for our best fit model (last column in Table 5.1), together with the spectra corresponding to different values of $M_{\rm PBH}$. The latter ones are obtained by keeping fixed to their best fit values all other cosmological and astrophysical parameters. Symbols refer to MIKE/HIRES data. To exhibit the variations in the flux power induced by different IGM thermal histories, we also show, as grey shaded areas, the impact of different IGM temperature evolutions.

In Figure 5.4 we report the updated plot with the constraints on the fraction of DM in PBHs, as functions of their masses, in the monochromatic case. The "LIGO window" of PBHs between $\sim 20 - 80 M_\odot$ initially suggested by Bird et al. (2016) has been probed and tentatively closed by constraints from Ultra-Faint Dwarf galaxies (Brandt 2016) and lensing of Type Ia Supernovæ (Zumalacarregui & Seljak 2018); these constraints have been questioned because of astrophysics uncertainties (see, e.g., Li et al. (2017); Carr et al. (2016)[3]): we show them in a patterned area. With our results we robustly close the higher mass part of that remaining window. There remains however, an interesting possibility in the very low-mass range, $\lesssim 10^{-10} M_\odot$ (see, e.g., Pi et al. (2018); Bartolo et al. (2019)).

By defining an equivalent mass $M_{\rm eq}$ one can easily convert the aforementioned marginalised 2 σ limits on $f_{\rm PBH} M_{\rm PBH}$ for the MMD case, to bounds on the parameters of a given EMD. In Figure 5.5 we provide such bounds, similarly to what was shown in Figure 3 from Carr et al. (2017) for other observational constraints, by plotting the parameters of the two EMDs discussed in this Chapter, as functions of $M_{\rm eq}$. In other words, each of the panels maps the limits on MMDs peaked at $M_{\rm eq}$ to constraints on EMDs. The left panel shows the *Powerlaw* EMD, with $\tilde{\gamma} = 0$, focusing on the following mass range: $M_{\rm min}$, $M_{\rm max} \in [10^{-2}, 10^{7}]\, M_\odot$. In the right panel we focus on the *Lognormal* EMD, scanning the parameter space defined by $\mu \in [10^{-2}, 10^{7}]\, M_\odot$, and $\sigma \in [0, 5]$. The two black lines correspond to the marginalised 2 σ constraints quoted above, i.e. $M_{\rm eq} = 60\, M_\odot$ (*solid*), and $M_{\rm eq} = 170\, M_\odot$ (*dashed*). The blue regions are therefore admitted by our analyses, while the red areas are ruled out.

All our analyses are based on the straightforward assumption that the PBH number density $n_{\rm PBH}$ is fixed during the cosmic time investigated by our simulations, i.e. from $z = 199$ to $z = 4.2$, which indeed corresponds to epochs when PBHs do not form anymore. However, the possibility that PBHs can fully evaporate during such time interval would alter our conclusions, since in that case $n_{\rm PBH}$ would vary with time. Nevertheless, from $z = 199$ to now, only PBHs with masses smaller than $\mathcal{O}(10^{-18})$ $M_\odot$ might completely evaporate (Hawking (1974)). For this reason, such value has to lie below the PBH mass ranges investigated in this

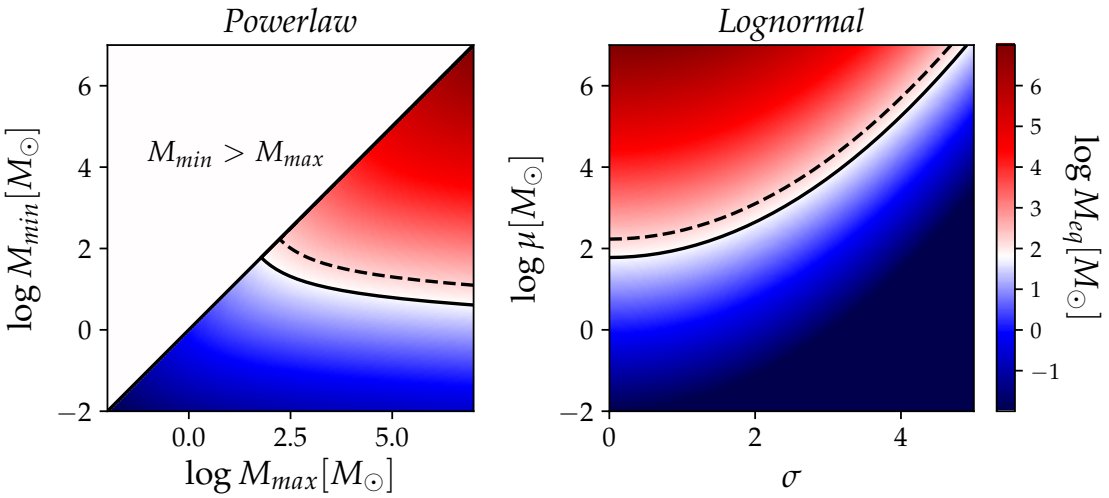

Figure 5.5: Equivalent Mass $M_{\rm eq}$ for EMDs. *Left: Powerlaw* with $M_{\rm min}$, $M_{\rm max} \in [10^{-2}, 10^{7}]\ M_\odot$, and $\gamma = 0$. *Right: Lognormal* with $\mu \in [10^{-2}, 10^{7}]\ M_\odot$, $\sigma \in [0, 5]$. The solid line refers to the case of $M_{\rm eq} = 60\ M_\odot$, while the dashed line refers to $M_{\rm eq} = 170\ M_\odot$, i.e. they correspond to the marginalised 2 σ limits. The red regions are thereby ruled out by our analyses, whereas the blue ones are admitted.

work. We have thus set $10^{-18}\ M_\odot$ as lower limit for the integral in Equation (5.8), and $10^{7}\ M_\odot$ as upper limit.

In the near future, it is expected that a larger number of high-redshift, high-resolution and high-signal-to-noise quasar spectra collected with the ESPRESSO spectrograph (Pepe et al. 2013) or at the E-ELT could further allow to use the small-scale flux power spectrum to achieve tighter constraints. Another relevant aspect would be an accurate modelling of the heating and ionisation due to accretion effects around the PBHs, to quantify how and if they could impact on the (much larger) scales of the Lyman-α forest.

Whereas PBHs with mass $\mathcal{O}(10)$ can potentially solve some tensions in the cosmic infrared background (Kashlinsky 2016, 2005; Kashlinsky et al. 2005), the accumulation of limits on the PBHs as DM model in the mass range probed by LIGO seems to suggest that the hypothesis of $30 M_\odot$ PBHs being the DM is less and less likely to be true. It has however become clear that these studies brought a plethora of astrophysical information, and even the exclusion of certain PBH mass ranges will bring information on some of the processes happening in the very early universe.

[3]For a discussion on some of these points, see Primordial versus Astrophysical Origin of Black Holes – CERN workshop

"The task is not so much to see what no one has yet seen;
but to think what nobody has yet thought,
about that which everybody sees."

Erwin Schrödinger

"Rubare, in teatro, è cosa buona,
*copiare è da co***oni."*

Dario Fo

6 Conclusions

6.1 Summary

The standard ΛCDM model appears to be in excellent agreement with both CMB and LSS observations. Some tensions exist, though, on the cosmological and local determination of the Hubble parameter H_0, and on the amplitude of the linear power spectrum σ_8. Furthermore, the standard CDM paradigm shows some discrepancies with structure formation data at sub-galactic scales (the CDM *small-scale crisis*). At such relatively small scales, observations tend indeed to favour cosmological models with suppressed matter power spectra with respect to the standard CDM scenario. Moreover, despite huge experimental efforts, DM in the form of WIMPs has so far eluded detection in direct and indirect searches, as well as at colliders. For all these reasons, the interest in models beyond the standard CDM paradigm has recently reinvigorated.

Combining together various independent observations is the key for shedding light on the dark sector, by building bridges to consistently link different studies, and exploring the interfaces between complementary lines of research, given that different scale and redshift coverage generally lead to different constraints and degeneracies. The main goal of all the works included in this thesis was to develop, step by step, a simple and versatile framework providing a direct link between particle physics model building and structure formation observations, addressed to both the cosmology and particle physics communities, enabling to systematically test a large variety of alternative cosmological models. Chapters 2 and 4 are entirely and explicitly dedicated to this subject. Whereas Chapters 3 and 5 are focused on studying specific non-standard DM models (i.e., Fuzzy DM and Primordial Black Holes, respectively), also those analyses might be seen in light of our efforts

towards devising a unified framework to model and constrain any departure from the standard CDM paradigm. Besides CMB and LSS observations, which accurately probe the linear matter power spectrum at large scales, high-redshift Lyman-α data make nowadays possible to study the detailed slope of the power spectrum at large k's, rather than simply the position of the small-scale cut-off. Such investigation is thereby compelling *per se*, independently of whether *non-cold* DM (nCDM) models provide a better match to the data, and it is worth to be performed with a simple and fully general approach.

Chapter 1 is a general introduction into basic and well known concepts in cosmology, intended to give a clearer understanding of the following Chapters. We briefly described the standard ΛCDM scenario, highlighting its remarkable success, its puzzles and possible alternatives. We then introduced the basics of cosmic structure formation and evolution, primarily focusing on the IGM and its main manifestation, i.e., the Lyman-α forest, which is one of the most powerful probes to investigate possible departures from the CDM paradigm.

In Chapter 2 we presented a new analytic fitting formula for the linear matter power spectrum, simple yet flexible enough to reproduce the clustering signal of large classes of non-thermal nCDM models, which are not at all adequately described by the oversimplified notion of thermal WDM. We showed that the formula is able to fully cover the parameter space of sterile neutrinos (whether resonantly produced or from particle decay), mixed cold and warm models, fuzzy DM, as well as other models suggested by effective theory of structure formation (ETHOS). We performed a large suite of N-body simulations and we extracted important non-linear statistics, such as the matter power spectrum and the halo mass function. We first presented preliminary astrophysical constraints, based on linear perturbation theory, from both the number of MW satellites and the Lyman-α forest. We then presented fully accurate limits from the Lyman-α forest, determined through an extensive analysis of the high-resolution, high-redshift data obtained by the MIKE/HIRES spectrographs. By using a large set of hydrodynamic simulations, we provided constraints on both astrophysical, cosmological, and nCDM parameters by performing a comprehensive MCMC analysis. We obtained a marginalised upper limit on the largest possible scale at which a power suppression induced by nearly any nCDM scenario can occur, i.e., $\alpha < 0.03$ Mpc$/h$ (2σ C.L.). We explicitly described how to test several of the most viable nCDM scenarios without the need to run any specific numerical simulations, due to the novel parametrisation proposed, and due to a new scheme that interpolates between the cosmological models explored. We showed that the shape of the linear matter power spectrum for standard thermal WDM models is in mild tension ($\sim 2\sigma$ C.L.) with the data, compared to non-thermal scenarios. We also showed that a DM fluid composed by both a warm (thermal) and a cold component is in tension with the Lyman-α forest, at least for large α-values.

In Chapter 3 we investigated constraints on ultra-light scalar DM (*a.k.a.* Fuzzy DM (FDM)) by analysing the most updated high-resolution Lyman-α forest data, and by studying the cosmological consequences at high and low redshift. For FDM that constitutes more than 30% of the total DM density, we obtained a lower limit $m \gtrsim 10^{-21}\,\mathrm{eV}$ for the FDM mass. This implies an upper limit on the initial field displacement (or the decay constant for an axion-like field) of $\phi \lesssim 10^{16}\,\mathrm{GeV}$. We also derived limits on the energy scale of cosmic inflation and established an upper bound on the tensor-to-scalar ratio of $r < 10^{-3}$, in the presence of FDM. Furthermore, we showed that there is very little room for FDM to solve the small-scale crisis of CDM without spoiling the Lyman-α forest results. We then focused on quantifying the impact of the Quantum Potential (QP) during the non-linear dynamical evolution of the FDM, by going beyond the universally adopted approximation of comprising the non-standard nature of the DM candidate in the suppressed linear power spectrum used to produce the initial conditions for the simulations. With the use of the new N-body cosmological hydrodynamic code `AX-GADGET` (Nori & Baldi 2018), we performed and analysed several hydrodynamic simulations in order to constrain the FDM mass by quantifying its impact on Lyman-α forest observations, as obtained for the first time in a N-body set-up without approximating the FDM dynamics. These simulations allowed us to perform a fully consistent comparison of mock Lyman-α observations with available data, and to update the existing constraints on the allowed FDM mass range. As the new constraints are not significantly different from the previous ones, this represents the first direct validation of the approximations adopted in all previous works. Furthermore, our large halo sample allowed us to perform an extensive characterisation of the properties of DM halos in the context of FDM scenarios, highlighting the typical mass scale below which QP effects start to appear.

In Chapter 4 we focused on the ETHOS parametrisation (Cyr-Racine et al. 2016) of interacting DM models, which includes a DR-DM scattering rate scaling like a power-law of the temperature, T^n. Scenarios with $n = 0$, 2, or 4 can easily be realised in concrete dark sector set-ups. We updated constraints on these three scenarios using recent CMB, BAO, and high-resolution Lyman-α data. Based on the general approach discussed in Chapter 2, we introduced a new Lyman-α likelihood, which is applicable to a wide range of cosmological models with a suppression of the matter power spectrum on small scales. For $n = 2$ and 4, we found that Lyman-α data strengthen the CMB+BAO bounds on the DM-DR interaction rate by many orders of magnitude. However, models offering a possible solution to the missing satellite problem are still compatible with our new bounds. For $n = 0$, high-resolution Lyman-α data bring no stronger constraints on the interaction rate than CMB+BAO data, except for extremely small values of the DR density.

Using CMB+BAO data and a theory-motivated prior on the minimal density of DR, we found that the $n = 0$ model can reduce the Hubble tension from 4.1σ to 2.7σ, while simultaneously accommodating smaller values of the σ_8 and S_8 parameters hinted by cosmic shear data.

In Chapter 5 we presented new limits on the Primordial Black Hole (PBH) mass and abundance, from a comprehensive analysis of high-resolution, high-redshift Lyman-α forest data. Using a grid of hydrodynamic simulations exploring different values of the astrophysical and cosmological parameters, we obtained a marginalised upper limit on the PBH mass of $f_{\rm PBH} M_{\rm PBH} \sim 60\ M_\odot$ at 2σ C.L., when a Gaussian prior on the reionisation redshift is imposed, preventing its posterior distribution to peak on very high values, which are disfavoured by the most recent estimates obtained both through CMB and IGM observations. Such bound weakens to $f_{\rm PBH} M_{\rm PBH} \sim 170\ M_\odot$, when a conservative flat prior is instead assumed. Both limits significantly improves previous constraints from the same physical observable. We also extended our predictions to non-monochromatic PBH mass distributions, ruling out large regions of the parameter space for some of the most viable PBH extended mass functions.

6.2 Future perspectives

The numerical tools that we have presented and used in this thesis are suitable to constrain a large variety of theoretical models inducing a small-scale departure in the linear matter power spectrum with respect to ΛCDM. Therefore, they may also be applied, e.g., to investigate the properties of massive neutrinos, by combining Lyman-α forest data with cluster abundance and/or cosmic shear measurements. Furthermore, they can easily be adapted to test non-standard cosmological scenarios with many other astrophysical probes. As we have already stressed, combining different data sets with complementary redshift and scale coverages, is paramount to break degeneracies, improve the present constraints, and make forecasts for upcoming surveys. Some of the possible directions for pursuing these tasks are the following:

- As we briefly discussed in Chapter 4 (Section 4.4), we are already working on a refinement of the method presented in Chapter 2, which will enable us to test alternative DM models with arbitrary fractions of interacting DM. One of the main goals is to apply such newer parametrisation to constrain DM-baryon interaction scenarios, since the interest towards such class of models has revived after the recent measurements by the *Experiment to Detect the Global Epoch of Reionization Signature* (EDGES) (Bowman et al. 2018; Barkana 2018).

- More accurate measurements of the IGM thermal history will provide strong priors for future analyses, allowing to better constrain the shape of the linear power spectrum at small scales (see, e.g., Boera et al. (2018)). Such improvements, in addition with new, higher-resolution Lyman-α forest spectra – as the ones collected with the ESPRESSO spectrograph (Pepe et al. 2013) or at the ELT[1] – are in principle expected to tighten all the limits reported in this thesis, both the ones on nCDM particle candidates (Chapters 2, 3 and 4) and the ones on PBHs (Chapter 5).

- Forthcoming instruments like DESI[2], LSST (Bechtol et al. 2019), or the Euclid[3] satellite might help to constrain the nCDM parameters by exploiting the weak lensing signal and the clustering of galaxies at small scales. These surveys will clarify the connection between luminous galaxies and the DM cosmic web, enabling high-precision tests applicable to all of the non-standard cosmological models discussed throughout this thesis.

- The study of the sub-galactic power spectrum through strong gravitational lensing is a very promising probe for testing the small-scale implications of nCDM scenarios, at cosmological distances from the MW. Investigating the gravitational imprints of DM structures on the lensed images in galaxy-scale strong lens systems allows indeed to put observational constraints on the substructure power spectrum, to be confronted against the results of very high-resolution hydrodynamic simulations (see, e.g., Vegetti et al. (2012, 2014); Bayer et al. (2018); Diaz Rivero et al. (2018); Díaz Rivero et al. (2018)). This will make it possible to disentangle more efficiently the consequences of different nCDM models. For instance, as we showed in Chapter 3, it is expected that higher-resolution Fuzzy DM simulations will soon allow us to observe the formation of solitonic cores.

- Another very intriguing possibility for testing alternative DM models is to make use of the global intensity mapping signal of the 21 cm transition produced by neutral hydrogen, which has the advantage of being at very high redshifts, where the non-linear evolution is less prominent (see, e.g., Sitwell et al. (2014); Carucci et al. (2015); Carucci et al. (2017a,b)). The cosmological 21 cm signal, studied both in absorption (Bowman et al. 2018) and in emission (Villaescusa-Navarro et al. 2018; Bacon et al. 2018), offers indeed a new probe, complementary to the existing ones, that could open a fresh window on the early universe to further test the imprints of alternative DM models (see, e.g., Schneider (2018); Kovetz et al. (2018); Lopez-Honorez et al. (2019)).

[1] `https://www.eso.org/public/teles-instr/elt/`
[2] `https://www.desi.lbl.gov/`
[3] `https://www.euclid-ec.org/`

To sum up, structure formation at small cosmological scales represents a pivotal frontier for DM research. Astrophysical observations are extremely complementary to high-energy physics programs and other experimental searches, given that it is basically impossible to separate the microscopic DM physics from its macroscopic distribution. Concurrently with ongoing and future experimental efforts, further theoretical work is needed to interpret observations in terms of fundamental models, to combine outcomes from different observational methods, and to develop novel DM probes.

Appendices

A

Quasi-degeneracy between α and γ

By considering the high-k limit of Equation (2.4), we can easily notice the quasi-degeneracy between the two parameters α and γ. For large k's (i.e., small scales), such that $(\alpha k)^\beta \gg 1$, we have indeed:

$$T(k)|_{k\alpha\gg 1} \simeq \alpha^{\beta\gamma} k^{\beta\gamma}. \tag{A.1}$$

If we now change α to a new value $\alpha \to \tilde{\alpha} = x \cdot \alpha$, where x is some real number, we can absorb this changes into a k-dependent change of γ:

$$\alpha^{\beta\gamma} k^{\beta\gamma} \to \tilde{\alpha}^{\beta\gamma} k^{\beta\gamma} = \alpha^{\beta\gamma} x^{\beta\gamma} k^{\beta\gamma} = \alpha^{\beta\gamma} \left(k^{\ln x/\ln k}\right)^{\beta\gamma} k^{\beta\gamma} = \alpha^{\beta\gamma} k^{\beta\gamma(1+\ln x/\ln k)}, \tag{A.2}$$

where we have used the obvious identity $x = k^{\ln x/\ln k}$. Thus, in Equation (A.1), we can trade a change in α for a k-dependent change in γ:

$$\gamma \to \tilde{\gamma}(k) = \gamma(1 + \ln x/\ln k), \tag{A.3}$$

which reproduces the effect of the change from α to $\tilde{\alpha}$. Now, this does not seem like a real degeneracy – and mathematically it is not, due to the k-dependence of $\tilde{\gamma}$. However, this dependence is only very weak (logarithmic), so that one can, in spite of the k-dependence, view $\tilde{\gamma}$ as approximately constant:

$$\tilde{\gamma}(k) \simeq \text{const.} \tag{A.4}$$

This approximation will be even better if $|\log_k x| = |\ln x/\ln k| \ll 1$, due to this quantity only appearing in the sum in Equation (A.3). Thus, indeed, to a very good approximation there is a degeneracy between α and γ: a change in one of them can be traded for a change in the other. Hence, by covering several different values of γ in Figure 2.1, we have information about many different α's at the same time.

B

Comparing simulated mass functions with theoretical predictions

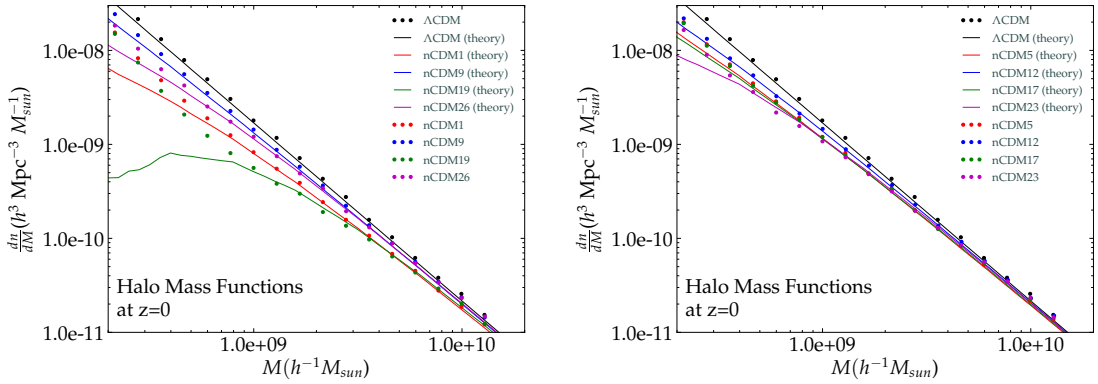

Figure B.1: For some of the models that we tested, we compare the theoretical DM halo mass function predicted by Equation (B.1) with the DM halo mass function extracted from the corresponding N-body simulation. The former is plotted as a solid line, the latter as a dotted one. Different colours refer to different models. The good agreement between them ceases to hold below masses of the order of $10^9\ M_\odot/h$, where artificial clumping strongly affects the results of the simulations. That mass indeed corresponds to the upturn highlighted in Figure 2.9.

We already highlighted in Section 2.4.1 the low-mass upturn in the halo mass functions, due to artificial clumping. To account for this effect and subtract the corresponding numerical artefacts, we estimated the number of subhalos predicted by our models by integrating over Equation (2.13), which is the theoretical subhalo mass function derived by an extended Press-Schechter approach (Press & Schechter (1974); Sheth & Tormen (1999)) based on a sharp-k window (Bertschinger (2006);

Benson et al. (2013)). This leads to the following expression (Schneider et al. (2013); Schneider (2015)):

$$\frac{\mathrm{d}n}{\mathrm{d}M} = \frac{1}{12\pi^2}\frac{\overline{\rho}}{M^2}\nu f(\nu)\frac{P(1/R)}{\delta_c^2 R^3}, \tag{B.1}$$

where $\overline{\rho}$ is the average density of the universe, $\nu = \delta_{c,0}^2/S(R)$ is the peak-height of the perturbations (at $z = 0$), and $f(\nu)$ is obtained by the excursion-set approach (Bond et al. (1991)). The relation between the sharp-k filter scale and mass and the variance $S(R)$ are defined in Equation (2.14).

In Figure B.1, we compare example theoretical DM halo mass functions predicted through Equation (B.1) to the DM halo mass functions extracted from the corresponding N-body simulations. The former is plotted as a solid line, the latter as a dotted line. Consistently with the trend outlined in Figure 2.9, below masses of the order of $10^9\ M_\odot/h$, artificial clumping induces the upturn in the simulated halo mass functions, and hence the discrepancies with the theory which are manifest in Figure B.1, below masses of the order of $10^9\ M_\odot/h$.

C

Reproducing the thermal Warm Dark Matter limits

This Appendix is dedicated to a comparison between the predictions that we obtained, when limiting our Lyman-α analysis to the thermal WDM case, and the most updated published results obtained with the standard approach (Iršič et al. (2017b)). The goal is to check the accuracy of the novel interpolation scheme and sampling method, introduced in Chapter 2, in order to safely extend our analyses to the full $\{\alpha, \beta, \gamma\}$-space, as it is done in both Chapters 2 and 4.

In Figure C.1a we report a comparison between the 1D and 2D posterior distributions for the main parameters of the analysis with freely floating IGM temperature. The overall agreement between the previous results (green contours) and ours (blue and orange contours) is evident. Both have been obtained by using the same data set, i.e. MIKE/HIRES data, as described in Section 2.6.1. Firstly, we notice the expected degeneracy between $T_0(z)$ and α. Due to this degeneracy, different prior choices on the IGM thermal history may sensibly affect the limits on the WDM mass. That is why, by looking at the 1D distribution for $m_{\rm WDM}$ and $T_0(z = 5.4)$, one can note a discrepancy between our prediction (blue) and the one by Iršič et al. (2017b) (green). In the previous study, in fact, flat priors on $T_0(z)$ were assumed, only forbidding unphysical jumps $\Delta T_0 > 5000$ K between adjacent redshift bins, whereas the blue contours are obtained by imposing broad Gaussian priors on $T_0(z)$ centred around their reference values, with standard deviation $\sigma = 3000$ K. The latter choice still constitutes a conservative assumption, given that it only prevents the temperatures to peak at $T_0 = 0$ at high redshifts, without precluding them from assuming reasonably cold values. Even though the blue 1D distribution is peaked at higher temperatures, only extremely cold values are excluded from

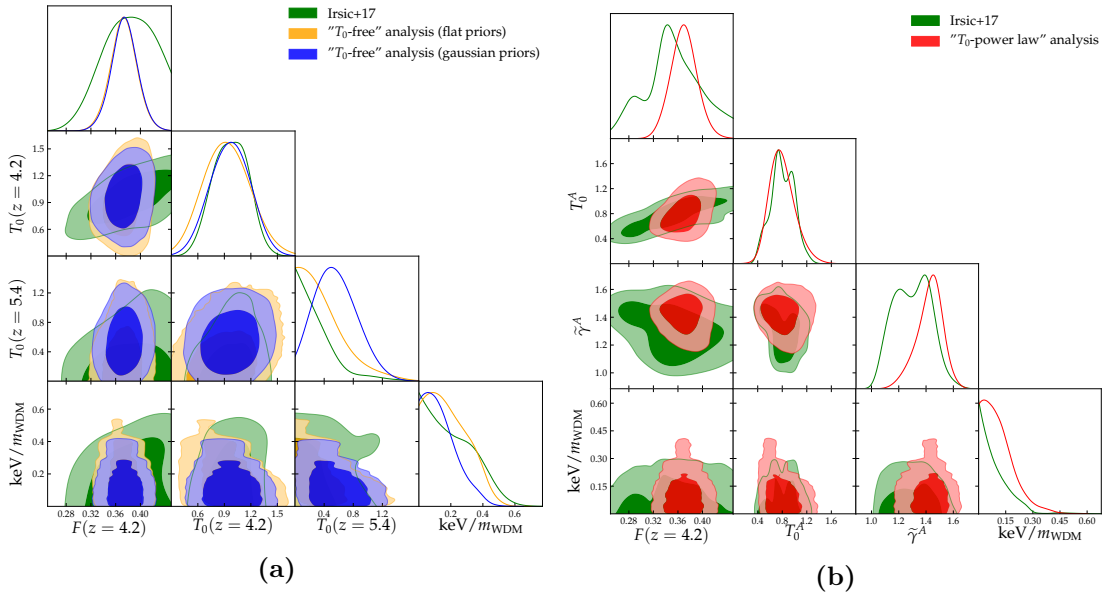

Figure C.1: Here we compare the 1 σ and 2σ C.L. exclusion plots obtained with our thermal WDM analyses against the results by Iršič et al. (2017b), which are displayed as green contours. The values of the temperatures are expressed in 10^4 K units. (a) In the left panel we focus on the freely floating temperature analysis (blue and orange contours). (b) In the right panel we focus on the power law analysis (red contours).

the corresponding 2 σ C.L. region. Quantitatively, the previous analysis led indeed to a 2 σ lower limit on the thermal WDM mass $m_{\rm WDM} > 2.1$ keV, while we have obtained $m_{\rm WDM} > 2.7$ keV. For a more direct comparison with the previous results, we also performed an analysis assuming flat priors on $T_0(z)$, shown by the orange contours in Figure C.1a. Under this assumption, which exactly corresponds to the prior choice adopted by Iršič et al. (2017b) for the freely floating temperature case, we obtained $m_{\rm WDM} > 2.2$ keV (2σ C.L.), in excellent agreement with the previously published result. Such agreement is also manifest when looking at the 1D posterior distribution for both $m_{\rm WDM}$ and $T_0(5.4)$. The orange curves are in very good agreement with the previous results, probing that the discrepancies between blue and green contours are driven by the different thermal history choice, rather than by some of the approximations characterising the new interpolation scheme.

In Figure C.1b we compare the 1D and 2D distributions for the main free parameters of the power law analysis, chosen to be the reference case by Iršič et al. (2017b) as well as by us. Analogously to the freely floating temperature plots, the green contours refer to the results by Iršič et al. (2017b); the red contours represent our results. Most of the considerations that we did above apply to this case too. The main difference between the results of the two analyses consists in a significant tightening of the upper bound on the thermal WDM mass, due

to the less conservative prior choice, as it has been explained in Section 2.6.1. The previously published 2σ C.L. limit is $m_{\rm WDM} > 4.1$ keV, whereas we obtained a slightly weaker constraint, namely $m_{\rm WDM} > 3.6$ keV. In the aforementioned freely floating T_0 case, our study yielded to a more aggressive limit with respect to the one by Iršič et al. (2017b). Conversely, in our reference analysis, the lack of cross-simulations, which is the main difference characterising our work (see Section 2.6.1), resulted in a weaker upper limit on the thermal WDM mass. From this point of view, the power law analysis that we adopted as our reference MCMC analysis can be considered conservative.

D

Exact solution of Klein–Gordon equation

In this Appendix we provide the exact solution of the homogeneous Klein–Gordon equation (Equation (3.29)) in a radiation dominated universe whose Hubble rate redshifts as

$$H \propto a^{-2}. \tag{D.1}$$

Here we should note that the redshifting of the Hubble rate can depart from $\propto a^{-2}$ when the effective number of relativistic degrees of freedom $g_{(s)*}$ changes in time; however as long as $g_{(s)*}$ stays constant while the scalar makes the transition from vacuum energy-like (non-oscillatory) to matter-like (oscillatory), then the solution under Equation (D.1) can be used to accurately compute the scalar density in the asymptotic future.

The solution of Equation (3.29) with Equation (D.1) is given in terms of a Bessel function of the first kind as

$$\phi = \phi_\star \, \Gamma\left(\frac{5}{4}\right) \left(\frac{4H}{m}\right)^{1/4} J_{1/4}\left(\frac{m}{2H}\right), \tag{D.2}$$

where we have chosen the initial condition for ϕ such that it approaches a constant value $\phi \to \phi_\star$ as $m/H \to 0$. Hence, after the scalar starts to oscillate, its density, given by Equation (3.30), asymptotes to

$$\lim_{\frac{m}{H}\to\infty} \rho_\phi = \frac{4}{\pi}\left[\Gamma\left(\frac{5}{4}\right)\right]^2 m^{1/2}\phi_\star^2 \, H^{3/2}. \tag{D.3}$$

Equating this with Equation (3.32):

$$\lim_{\frac{m}{H}\to\infty} \rho_\phi = \frac{1}{2}m^2\phi_\star^2\left(\frac{a_{\rm osc}}{a}\right)^3 = \frac{1}{2}m^2\phi_\star^2\left(\frac{H}{H_{\rm osc}}\right)^{3/2}, \tag{D.4}$$

yields the ratio between the scalar mass and $H_{\rm osc}$ as

$$\frac{m^2}{H_{\rm osc}^2} = \left(\frac{8}{\pi}\right)^{4/3} \left[\Gamma\left(\frac{5}{4}\right)\right]^{8/3} \approx 2.68. \tag{D.5}$$

With this $H_{\rm osc}$, one can compute the scalar density not only during radiation domination, but also in the subsequent epochs as demonstrated in Section 3.4.1.

List of Figures

List of Tables

Bibliography

Abazajian, K., Fuller, G. M., & Patel, M. 2001, Phys. Rev., D64, 023501

Abbott, B. P. et al. 2016a, Phys. Rev. Lett., 116, 061102

Abbott, B. P. et al. 2016b, Phys. Rev. Lett., 116, 241102

Abbott, T. M. C. et al. 2018, Phys. Rev., D98, 043526

Ade, P. et al. 2014, Astron.Astrophys., 571, A16

Ade, P. A. R. et al. 2016a, Astron. Astrophys., 594, A13

Ade, P. A. R. et al. 2016b, Astron. Astrophys., 594, A15

Ade, P. A. R. et al. 2016c, Astron. Astrophys., 594, A20

Adhikari, R. et al. 2017, JCAP, 1701, 025

Afshordi, N., McDonald, P., & Spergel, D. N. 2003, ApJ, 594, L71

Aghanim, N. et al. 2016, Astron. Astrophys., 594, A11

Aghanim, N. et al. 2018, 1807.06209

Akrami, Y. et al. 2018, 1807.06205

Alam, S. et al. 2017, Mon. Not. Roy. Astron. Soc., 470, 2617

Alcock, C. et al. 2001, ApJ Letters, 550, L169

Ali-Haïmoud, Y. & Kamionkowski, M. 2017, Phys. Rev. D, 95, 043534

Ali-Haïmoud, Y., Kovetz, E. D., & Kamionkowski, M. 2017, Phys. Rev. D, 96, 123523

Aloni, D., Blum, K., & Flauger, R. 2017, JCAP, 2017, 017

Alpher, R., Bethe, H., & Gamow, G. 1948, Phys.Rev., 73, 803

Angulo, R. E., Hahn, O., & Abel, T. 2013, MNRAS, 434, 3337

Archidiacono, M., Bohr, S., Hannestad, S., JÃžrgensen, J. H., & Lesgourgues, J. 2017, JCAP, 1711, 010

Archidiacono, M. & Hannestad, S. 2014, JCAP, 1407, 046

Archidiacono, M., Hooper, D. C., Murgia, R., et al. 2019, 1907.01496

Arinyo-i Prats, A., Miralda-Escude, J., Viel, M., & Cen, R. 2015, JCAP, 1512, 017

Armengaud, E., Palanque-Delabrouille, N., Yèche, C., Marsh, D. J. E., & Baur, J. 2017, Mon. Not. Roy. Astron. Soc., 471, 4606

Arvanitaki, A., Dimopoulos, S., Dubovsky, S., Kaloper, N., & March-Russell, J. 2010, Phys. Rev., D81, 123530

Audren, B., Lesgourgues, J., Bird, S., Haehnelt, M. G., & Viel, M. 2013, JCAP, 1301, 026

Bacon, D. J. et al. 2018, 1811.02743

Bae, K. J., Huh, J.-H., & Kim, J. E. 2008, JCAP, 0809, 005

Baldi, M., Pettorino, V., Robbers, G., & Springel, V. 2010, MNRAS, 403, 1684

Bardeen, J. M., Steinhardt, P. J., & Turner, M. S. 1983, Phys. Rev. D, 28, 679

Bardeen, J. M., Szalay, A., Kaiser, N., & Bond, J. 1985, ApJ, 304, 15

Barkana, R. 2018, Nature, 555, 71

Barkana, R., Haiman, Z., & Ostriker, J. P. 2001, Astrophys. J., 558, 482

Barnacka, A., Glicenstein, J.-F., & Moderski, R. 2012, Phys. Rev. D, 86, 043001

Bartolo, N., De Luca, V., Franciolini, G., et al. 2019, Phys. Rev., D99, 103521

Baur, J., Palanque-Delabrouille, N., Yeche, C., Magneville, C., & Viel, M. 2016, JCAP, 1608, 012

Bayer, D., Chatterjee, S., Koopmans, L. V. E., et al. 2018, 1803.05952

Bechtol, K. et al. 2019, 1903.04425

Bellomo, N., Bernal, J. L., Raccanelli, A., & Verde, L. 2018, JCAP, 2018, 004

Beltran, M., Garcia-Bellido, J., Lesgourgues, J., & Viel, M. 2005, Phys. Rev., D72, 103515

Bennett, C. L. et al. 2013, Astrophys. J. S., 208, 20

Benson, A. J., Farahi, A., Cole, S., et al. 2013, MNRAS, 428, 1774

Berezin, V., Kuzmin, V., & Tkachev, I. 1983, Physics Letters B, 120, 91

Bernal, J. L., Bellomo, N., Raccanelli, A., & Verde, L. 2017, JCAP, 2017, 052

Bernal, J. L., Raccanelli, A., Verde, L., & Silk, J. 2018, JCAP, 1805, 017

Bertone, G. & Tait, Tim, M. P. 2018, Nature, 562, 51

Bertschinger, E. 2006, Phys. Rev., D74, 063509

Beutler, F., Blake, C., Colless, M., et al. 2011, Mon. Not. Roy. Astron. Soc., 416, 3017

Binder, T., Covi, L., Kamada, A., et al. 2016, JCAP, 1611, 043

Bird, S., Cholis, I., Muñoz, J. B., et al. 2016, Phys. Rev. Lett., 116, 201301

Blas, D., Floerchinger, S., Garny, M., Tetradis, N., & Wiedemann, U. A. 2015, JCAP, 1511, 049

Blas, D., Lesgourgues, J., & Tram, T. 2011, JCAP, 1107, 034

Boddy, K. K., Feng, J. L., Kaplinghat, M., & Tait, T. M. P. 2014, Phys. Rev., D89, 115017

Bode, P., Ostriker, J. P., & Turok, N. 2001, ApJ, 556, 93

Boehm, C., Schewtschenko, J. A., Wilkinson, R. J., Baugh, C. M., & Pascoli, S. 2014, MNRAS, 445, L31

Boera, E., Becker, G. D., Bolton, J. S., & Nasir, F. 2018, 1809.06980

Bolton, J. S., Puchwein, E., Sijacki, D., et al. 2017, MNRAS, 464, 897

Bond, J. R., Cole, S., Efstathiou, G., & Kaiser, N. 1991, Astrophys. J., 379, 440

Bose, S., Vogelsberger, M., Zavala, J., et al. 2019, Mon. Not. Roy. Astron. Soc., 487, 522

Bowman, J. D., Rogers, A. E. E., Monsalve, R. A., Mozdzen, T. J., & Mahesh, N. 2018, Nature, 555, 67

Boyarsky, A., Lesgourgues, J., Ruchayskiy, O., & Viel, M. 2009, JCAP, 0905, 012

Boylan-Kolchin, M., Bullock, J. S., & Kaplinghat, M. 2011, Mon. Not. Roy. Astron. Soc., 415, L40

Brandt, T. D. 2016, ApJ Letters, 824, L31

Brinckmann, T. & Lesgourgues, J. 2018, 1804.07261

Bringmann, T., Ihle, H. T., Kersten, J., & Walia, P. 2016, Phys. Rev., D94, 103529

Brookshaw, L. 1985, in Proceedings of the Astronomical Society of Australia, Vol. 6, 207–210

Brust, C., Cui, Y., & Sigurdson, K. 2017, JCAP, 1708, 020

Buckley, M. R. & Fox, P. J. 2010, Phys. Rev., D81, 083522

Buckley, M. R., Zavala, J., Cyr-Racine, F.-Y., Sigurdson, K., & Vogelsberger, M. 2014, Phys. Rev., D90, 043524

Buen-Abad, M. A., Marques-Tavares, G., & Schmaltz, M. 2015, Phys. Rev., D92, 023531

Buen-Abad, M. A., Schmaltz, M., Lesgourgues, J., & Brinckmann, T. 2018, JCAP, 1801, 008

Calabrese, E. & Spergel, D. N. 2016, MNRAS, 460, 4397

Calchi Novati, S., Mirzoyan, S., Jetzer, P., & Scarpetta, G. 2013, MNRAS, 435, 1582

Canetti, L., Drewes, M., Frossard, T., & Shaposhnikov, M. 2013, Phys. Rev., D87, 093006

Capela, F., Pshirkov, M., & Tinyakov, P. 2013, Phys. Rev. D, 87, 123524

Carr, B., Kühnel, F., & Sandstad, M. 2016, Phys. Rev. D, 94, 083504

Carr, B., Raidal, M., Tenkanen, T., Vaskonen, V., & Veermäe, H. 2017, Phys. Rev., D96, 023514

Carr, B. & Silk, J. 2018, MNRAS, 478, 3756

Carr, B. J. 1975, ApJ, 201, 1

Carucci, I. P., Corasaniti, P.-S., & Viel, M. 2017a, JCAP, 1712, 018

Carucci, I. P., Villaescusa-Navarro, F., & Viel, M. 2017b, JCAP, 1704, 001

Carucci, I. P., Villaescusa-Navarro, F., Viel, M., & Lapi, A. 2015, JCAP, 7, 047

Cautun, M., Frenk, C. S., van de Weygaert, R., Hellwing, W. A., & Jones, B. J. T. 2014, Mon. Not. Roy. Astron. Soc., 445, 2049

Cen, R., McDonald, P., Trac, H., & Loeb, A. 2009, ApJ, 706, L164

Cherry, J. F. & Horiuchi, S. 2017, Phys. Rev., D95, 083015

Cholis, I., Kovetz, E. D., Ali-Haïmoud, Y., et al. 2016, Phys. Rev., D94, 084013

Chun, E. J. 2014, Phys. Lett., B735, 164

Cleary, P. W. & Monaghan, J. J. 1999, Journal of Computational Physics, 148, 227

Clowe, D., Bradac, M., Gonzalez, A. H., et al. 2006, Astrophys.J., 648, L109

Colin, F., Egli, R., & Lin, F. Y. 2006, J. Comput. Phys., 217, 680

Colombi, S., Dodelson, S., & Widrow, L. M. 1996, ApJ, 458, 1

Corbelli, E. & Salucci, P. 2000, Mon.Not.Roy.Astron.Soc., 311, 441

Crawford, M. & Schramm, D. N. 1982, Nature, 298, 538

Crocce, M., Pueblas, S., & Scoccimarro, R. 2006, MNRAS, 373, 369

Croft, R. A. C., Weinberg, D. H., Bolte, M., et al. 2002, Astrophysical Journal, 581, 20

Cyr-Racine, F.-Y., de Putter, R., Raccanelli, A., & Sigurdson, K. 2014, Phys. Rev., D89, 063517

Cyr-Racine, F.-Y., Sigurdson, K., Zavala, J., et al. 2016, Phys. Rev., D93, 123527

D'Aloisio, A., McQuinn, M., & Trac, H. 2015, ApJ, 813, L38

Das, S., Mondal, R., Rentala, V., & Suresh, S. 2018, Journal of Cosmology and Astro-Particle Physics, 2018, 045

Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, ApJ, 292, 371

de Vega, H. J., Salucci, P., & Sanchez, N. G. 2012, New Astron., 17, 653

Destri, C., de Vega, H. J., & Sanchez, N. G. 2013, Phys. Rev., D88, 083512

Diamanti, R., Ando, S., Gariazzo, S., Mena, O., & Weniger, C. 2017, JCAP, 1706, 008

Diaz Rivero, A., Cyr-Racine, F.-Y., & Dvorkin, C. 2018, Phys. Rev., D97, 023001

Dicke, R., Peebles, P., Roll, P., & Wilkinson, D. 1965, Astrophys.J., 142, 414

Dodelson, S. & Widrow, L. 1994, Phys. Rev. Lett., 72, 17

Dolgov, A. & Silk, J. 1993, Phys. Rev. D, 47, 4244

Donato, F., Gentile, G., Salucci, P., et al. 2009, Mon. Not. Roy. Astron. Soc., 397, 1169

Dvorkin, C., Blum, K., & Kamionkowski, M. 2014, Phys. Rev., D89, 023519

Díaz Rivero, A., Dvorkin, C., Cyr-Racine, F.-Y., Zavala, J., & Vogelsberger, M. 2018, Phys. Rev., D98, 103517

Enqvist, K., Kainulainen, K., & Maalampi, J. 1990, Phys. Lett., B249, 531

Fairbairn, M., Hogan, R., & Marsh, D. J. E. 2015, Phys. Rev., D91, 023509

Fluri, J., Kacprzak, T., Lucchi, A., et al. 2019, 1906.03156

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306

Frigerio, M. & Yaguna, C. E. 2015, Eur. Phys. J., C75, 31

Gaggero, D., Bertone, G., Calore, F., et al. 2017, Phys. Rev. Lett., 118, 241101

Gamow, G. et al. 1948, Nature, 162, 680

García-Bellido, J., Linde, A., & Wands, D. 1996, Phys. Rev. D, 54, 6040

Gariazzo, S., Escudero, M., Diamanti, R., & Mena, O. 2017, Phys. Rev., D96, 043501

Garny, M., Konstandin, T., Sagunski, L., & Tulin, S. 2018, JCAP, 1809, 011

Garrison-Kimmel, S. et al. 2017, Mon. Not. Roy. Astron. Soc., 471, 1709

Garzilli, A., Boyarsky, A., & Ruchayskiy, O. 2017, Phys. Lett., B773, 258

Garzilli, A., Magalich, A., Theuns, T., et al. 2018, 1809.06585

Ghiglieri, J. & Laine, M. 2015, JHEP, 11, 171

Gluscevic, V. et al. 2019, 1903.05140

Gnedin, N. Y. & Hamilton, A. J. S. 2002, MNRAS, 334, 107

Gong, J.-O. & Kitajima, N. 2017, JCAP, 1708, 017

Gong, J.-O. & Kitajima, N. 2018, Journal of Cosmology and Astroparticle Physics, 2018, 041

Governato, F., Zolotov, A., Pontzen, A., et al. 2012, MNRAS, 422, 1231

Graham, P. W., Rajendran, S., & Varela, J. 2015, Phys. Rev. D, 92, 063007

Green, A. M. 2016, Phys. Rev. D, 94, 063530

Griest, K., Cieplak, A. M., & Lehner, M. J. 2014, ApJ, 786, 158

Guth, A. H. & Pi, S. Y. 1982, Phys. Rev. Lett., 49, 1110

Hawking, S. 1971, MNRAS, 152, 75

Hawking, S. 1989, Physics Letters B, 231, 237

Hawking, S. W. 1974, Nature, 248, 30

Higaki, T., Jeong, K. S., & Takahashi, F. 2014, Phys. Lett., B734, 21

Hildebrandt, H. et al. 2018, 1812.06076

Hlozek, R., Grin, D., Marsh, D. J. E., & Ferreira, P. G. 2015, Phys. Rev. D, 91, 103512

Hou, Z., Keisler, R., Knox, L., Millea, M., & Reichardt, C. 2013, Phys. Rev., D87, 083008

Hu, W., Barkana, R., & Gruzinov, A. 2000, Phys. Rev. Lett., 85, 1158

Hubble, E. 1929, Proc.Nat.Acad.Sci., 15, 168

Hui, L. & Gnedin, N. Y. 1997, MNRAS, 292, 27

Hui, L., Ostriker, J. P., Tremaine, S., & Witten, E. 2017, Phys. Rev., D95, 043541

Hwang, J.-c. 1997, Phys. Lett., B401, 241

Ipser, J. & Sikivie, P. 1984, Phys. Rev. D, 30, 712

Iršič, V., Viel, M., Haehnelt, M. G., Bolton, J. S., & Becker, G. D. 2017a, Phys. Rev. Lett., 119, 031302

Iršič, V. et al. 2017b, Phys. Rev. D, 96, 023522

Iršič, V. et al. 2017c, MNRAS, 466, 4332

Ivanov, P. 1998, Phys. Rev. D, 57, 7145

Ivanov, P., Naselsky, P., & Novikov, I. 1994, Phys. Rev. D, 50, 7173

Jethwa, P., Erkal, D., & Belokurov, V. 2018, Mon. Not. Roy. Astron. Soc., 473, 2060

Joudaki, S. et al. 2019, 1906.09262

Jungman, G., Kamionkowski, M., & Griest, K. 1996, Phys. Rept., 267, 195

Kannike, K., Marzola, L., Raidal, M., & Veermäe, H. 2017, JCAP, 2017, 020

Kashlinsky, A. 2005, Phys. Rept., 409, 361

Kashlinsky, A. 2016, Astrophys. J., 823, L25

Kashlinsky, A., Arendt, R. G., Mather, J. C., & Moseley, S. H. 2005, Nature, 438, 45

Kashlinsky, A. et al. 2019, 1903.04424

Katz, A., Kopp, J., Sibiryakov, S., & Xue, W. 2018, JCAP, 1812, 005

Kennedy, R., Frenk, C., Cole, S., & Benson, A. 2014, Mon. Not. Roy. Astron. Soc., 442, 2487

Khlopov, M., Malomed, B. A., & Zeldovich, I. B. 1985, Mon. Not. Roy. Astron. Soc., 215, 575

Klypin, A. A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, Astrophys. J., 522, 82

Kobayashi, T., Kurematsu, R., & Takahashi, F. 2013, JCAP, 1309, 032

Kobayashi, T., Murgia, R., De Simone, A., Iršič, V., & Viel, M. 2017, Phys. Rev. D, 96, 123514

Kobayashi, T. & Takahashi, F. 2016, JCAP, 1608, 056

König, J., Merle, A., & Totzauer, M. 2016, JCAP, 1611, 038

Kovetz, E. D. 2017, Phys. Rev. Lett., 119, 131301

Kovetz, E. D., Cholis, I., Breysse, P. C., & Kamionkowski, M. 2017, Phys. Rev., D95, 103010

Kovetz, E. D., Poulin, V., Gluscevic, V., et al. 2018, Phys. Rev., D98, 103529

Krall, R., Cyr-Racine, F.-Y., & Dvorkin, C. 2017, JCAP, 1709, 003

Kusenko, A. 2006, Phys. Rev. Lett., 97, 241301

La, D. & Steinhardt, P. J. 1989, Physics Letters B, 220, 375

Langacker, P. 1989, UPR-0401T

Lapi, A. & Danese, L. 2015, JCAP, 1509, 003

Lello, L. & Boyanovsky, D. 2015, Phys. Rev., D91, 063502

Lesgourgues, J. 2011, 1104.2932

Lesgourgues, J., Marques-Tavares, G., & Schmaltz, M. 2016, JCAP, 1602, 037

Lesgourgues, J. & Tram, T. 2011, JCAP, 1109, 032

Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473

Li, T. S. et al. 2017, Astrophys. J., 838, 8

Liddle, A. R. 1999, 260

Linde, A. D. 1991, Phys. Lett., B259, 38

Linde, A. D. & Lyth, D. H. 1990, Phys. Lett., B246, 353

López, S., D'Odorico, V., Ellison, S. L., et al. 2016, A&A, 594, A91

Lopez-Honorez, L., Mena, O., & Villanueva-Domingo, P. 2019, Phys. Rev., D99, 023522

Lovell, M. R., Bose, S., Boyarsky, A., et al. 2016, MNRAS, 461, 60

Lovell, M. R., Frenk, C. S., Eke, V. R., et al. 2014, MNRAS, 439, 300

Lovell, M. R., Gonzalez-Perez, V., Bose, S., et al. 2017, MNRAS, 468, 2836

Lyth, D. H. 1992, Phys. Rev., D45, 3394

Lyth, D. H. & Riotto, A. 1999, Phys. Reports, 314, 1

Magee, R., Deutsch, A.-S., McClincy, P., et al. 2018, Phys. Rev. D, 98, 103024

Marsh, D. J. E. 2016a, Phys. Rept., 643, 1

Marsh, D. J. E. 2016b, 1605.05973

Marsh, D. J. E. & Ferreira, P. G. 2010, Phys. Rev., D82, 103528

Marsh, D. J. E. & Pop, A.-R. 2015, MNRAS, 451, 2479

Marsh, D. J. E. & Silk, J. 2014, MNRAS, 437, 2652

McDonald, P. et al. 2005, Astrophys. J., 635, 761

McDonald, P. et al. 2006, Astrophys. J. Suppl., 163, 80

McQuinn, M. 2016, Ann. Rev. Astron. Astrophys., 54, 313

Mediavilla, E., Munoz, J. A., Falco, E., et al. 2009, ApJ, 706, 1451

Meiksin, A. & White, M. J. 2004, Mon. Not. Roy. Astron. Soc., 350, 1107

Meiksin, A. A. 2009, Rev. Mod. Phys., 81, 1405

Melott, A. L. & Shandarin, S. F. 1989, ApJ, 343, 26

Menci, N., Merle, A., Totzauer, M., et al. 2017, Astrophys. J., 836, 61

Merle, A. 2017

Merle, A., Niro, V., & Schmidt, D. 2014, JCAP, 1403, 028

Merle, A., Schneider, A., & Totzauer, M. 2016, JCAP, 1604, 003

Merle, A. & Totzauer, M. 2015, JCAP, 1506, 011

Meszaros, P. 1975, Astronomy and Astrophysics, 38, 5

Miller, C., Erickcek, A., & Murgia, R. 2019, 1908.10369

Mocz, P. & Succi, S. 2015, Phys. Rev., E91, 053304

Moore, B., Ghigna, S., Governato, F., et al. 1999, Astrophys. J., 524, L19

Muñoz, J. B., Kovetz, E. D., Dai, L., & Kamionkowski, M. 2016, Phys. Rev. Lett., 117, 091301

Murgia, R. 2016, 1612.02282

Murgia, R. 2018, J. Phys. Conf. Ser., 956, 012005

Murgia, R., Gariazzo, S., & Fornengo, N. 2016, JCAP, 1604, 014

Murgia, R., Iršič, V., & Viel, M. 2018, Phys. Rev. D, 98, 083540

Murgia, R., Merle, A., Viel, M., Totzauer, M., & Schneider, A. 2017, JCAP, 1711, 046

Murgia, R., Scelfo, G., Viel, M., & Raccanelli, A. 2019, Phys. Rev. Lett., 123, 071102

Nakama, T., Silk, J., & Kamionkowski, M. 2017, Phys. Rev. D, 95, 043511

Nambu, Y. & Sasaki, M. 1990, Phys. Rev., D42, 3918

Niikura, H., Takada, M., Yasuda, N., et al. 2017, 1701.02151

Nori, M. & Baldi, M. 2018, Mon. Not. Roy. Astron. Soc., 478, 3935

Nori, M., Murgia, R., Iršič, V., Baldi, M., & Viel, M. 2019, Mon. Not. Roy. Astron. Soc., 482, 3227

Okamoto, T., Gao, L., & Theuns, T. 2008, MNRAS, 390, 920

Oman, K. A. et al. 2015, Mon. Not. Roy. Astron. Soc., 452, 3650

Palanque-Delabrouille, N. et al. 2013, A&A, 559, A85

Pan, Z., Kaplinghat, M., & Knox, L. 2018, Phys. Rev., D97, 103531

Pawlowski, M. S., Famaey, B., Merritt, D., & Kroupa, P. 2015, Astrophys. J., 815, 19

Peebles, P. J. E. 2017, 1705.10683

Penzias, A. & Wilson, R. 1965, Astrophys. J., 142, 419

Pepe, F. et al. 2013, The Messenger, 153, 6

Perlmutter, S. et al. 1999, Astrophys. J., 517, 565

Petraki, K. & Kusenko, A. 2008, Phys. Rev., D77, 065014

Pi, S., Zhang, Y.-l., Huang, Q.-G., & Sasaki, M. 2018, JCAP, 1805, 042

Polisensky, E. & Ricotti, M. 2011, Phys. Rev., D83, 043506

Polnarev, A. & Zembowicz, R. 1991, Phys. Rev. D, 43, 1106

Poulin, V., Lesgourgues, J., & Serpico, P. D. 2017a, JCAP, 1703, 043

Poulin, V., Serpico, P. D., Calore, F., Clesse, S., & Kohri, K. 2017b, Phys. Rev. D, 96, 083524

Poulin, V., Serpico, P. D., & Lesgourgues, J. 2016, JCAP, 1608, 036

Press, W. H. & Schechter, P. 1974, Astrophys. J., 187, 425

Quinn, D. P., Wilkinson, M. I., Irwin, M. J., et al. 2009, MNRAS Letters, 396, L11

Raccanelli, A., Kovetz, E. D., Bird, S., Cholis, I., & Muñoz, J. B. 2016, Phys. Rev. D, 94, 023516

Ratra, B. 1991, Phys. Rev., D44, 352

Ricotti, M., Ostriker, J. P., & Mack, K. J. 2008, ApJ, 680, 829

Riess, A. G., Casertano, S., Yuan, W., Macri, L. M., & Scolnic, D. 2019, Astrophys. J., 876, 85

Riess, A. G. et al. 1998, Astron.J., 116, 1009

Rindler-Daller, T. & Shapiro, P. R. 2014, Mod. Phys. Lett., A29, 1430002

Ringwald, A. 2012, Phys. Dark Univ., 1, 116

Ross, A. J., Samushia, L., Howlett, C., et al. 2015, Mon. Not. Roy. Astron. Soc., 449, 835

Rossi, G., Yéche, C., Palanque-Delabrouille, N., & Lesgourgues, J. 2015, Phys. Rev., D92, 063505

Sasaki, M., Suyama, T., Tanaka, T., & Yokoyama, S. 2018, Classical and Quantum Gravity, 35, 063001

Sawala, T. et al. 2016, Mon. Not. Roy. Astron. Soc., 457, 1931

Scelfo, G., Bellomo, N., Raccanelli, A., Matarrese, S., & Verde, L. 2018, JCAP, 2018, 039

Schewtschenko, J. A., Baugh, C. M., Wilkinson, R. J., et al. 2016, Mon. Not. Roy. Astron. Soc., 461, 2282

Schive, H.-Y., Chiueh, T., & Broadhurst, T. 2014, Nature Phys., 10, 496

Schive, H.-Y., Chiueh, T., Broadhurst, T., & Huang, K.-W. 2016, Astrophys. J., 818, 89

Schive, H.-Y., Tsai, Y.-C., & Chiueh, T. 2010, Astrophys. J. Suppl., 186, 457

Schive, H.-Y., ZuHone, J. A., Goldbaum, N. J., et al. 2018, Mon. Not. Roy. Astron. Soc., 481, 4815

Schneider, A. 2015, MNRAS, 451, 3117

Schneider, A. 2016, JCAP, 1604, 059

Schneider, A. 2018, Phys. Rev., D98, 063021

Schneider, A., Smith, R. E., Maccio, A. V., & Moore, B. 2012, Mon. Not. Roy. Astron. Soc., 424, 684

Schneider, A., Smith, R. E., & Reed, D. 2013, MNRAS, 433, 1573

Schneider, A., Trujillo-Gomez, S., Papastergis, E., Reed, D. S., & Lake, G. 2017, MNRAS, 470, 1542

Seljak, U., Slosar, A., & McDonald, P. 2006, JCAP, 0610, 014

Shakya, B. 2016, Mod. Phys. Lett., A31, 1630005

Shandera, S., Jeong, D., & Gebhardt, H. S. G. 2018, Phys. Rev. Lett., 120, 241102

Sheth, R. K. & Tormen, G. 1999, MNRAS, 308, 119

Shi, X.-D. & Fuller, G. 1999, Phys. Rev. Lett., 82, 2832

Sikivie, P. 1982, Phys. Rev. Lett., 48, 1156

Sitwell, M., Mesinger, A., Ma, Y.-Z., & Sigurdson, K. 2014, Mon. Not. Roy. Astron. Soc., 438, 2664

Smoot, G., Bennet, C., et al. 1992, Astrophys. J., 396, L1

Springel, V. 2005, MNRAS, 364, 1105

Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001a, MNRAS, 328, 726

Springel, V., Yoshida, N., & White, S. D. M. 2001b, New Astron., 6, 79

Strobl, K. & Weiler, T. J. 1994, Phys. Rev., D50, 7690

Tisserand, P. et al. 2007, A&A, 469, 387

Tulin, S. & Yu, H.-B. 2018, Phys. Rept., 730, 1

Turner, M. S. 1986, Phys. Rev., D33, 889

Vegetti, S. & Koopmans, L. V. E. 2009, MNRAS, 400, 1583

Vegetti, S., Koopmans, L. V. E., Auger, M. W., Treu, T., & Bolton, A. S. 2014, MNRAS, 442, 2017

Vegetti, S., Lagattuta, D. J., McKean, J. P., et al. 2012, Nature, 481, 341

Veltmaat, J., Niemeyer, J. C., & Schwabe, B. 2018, Phys. Rev., D98, 043509

Venumadhav, T., Cyr-Racine, F.-Y., Abazajian, K. N., & Hirata, C. M. 2016, Phys. Rev., D94, 043515

Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013, Phys. Rev. D, 88, 043502

Viel, M. & Haehnelt, M. G. 2006, MNRAS, 365, 231

Viel, M., Haehnelt, M. G., & Springel, V. 2004, Mon. Not. Roy. Astron. Soc., 354, 684

Viel, M., Lesgourgues, J., Haehnelt, M. G., Matarrese, S., & Riotto, A. 2005, Phys. Rev. D, 71, 1

Viel, M., Markovic, K., Baldi, M., & Weller, J. 2012, MNRAS, 421, 50

Vilenkin, A. & Everett, A. E. 1982, Phys. Rev. Lett., 48, 1867

Villaescusa-Navarro, F. et al. 2018, Astrophys. J., 866, 135

Visinelli, L. 2017, Phys. Rev., D96, 023013

Vogelsberger, M., Zavala, J., Cyr-Racine, F.-Y., et al. 2016, MNRAS, 460, 1399

Voruz, L., Lesgourgues, J., & Tram, T. 2014, JCAP, 1403, 004

Wang, B., Abdalla, E., Atrio-Barandela, F., & Pavon, D. 2016, Rept. Prog. Phys., 79, 096901

Wang, J. & White, S. D. M. 2007, Mon. Not. Roy. Astron. Soc., 380, 93

Wang, W., Han, J., Cooper, A. P., et al. 2015, Mon. Not. Roy. Astron. Soc., 453, 377

Webster, R. & Oliver, M. A. 2007, Geostatistics for environmental scientists (John Wiley & Sons)

Weinberg, S. 2002, Astrophys. J., 581, 810

Wichoski, U. F., MacGibbon, J. H., & Brandenberger, R. H. 1998, Physics Reports, 307, 191

Widrow, L. M. & Kaiser, N. 1993, Astrophys. J., 416, L71

Wilkinson, P. N., Henstock, D. R., Browne, I. W. A., et al. 2001, Phys. Rev. Lett., 86, 584

Woo, T.-P. & Chiueh, T. 2009, The Astrophysical Journal, 697, 850

Xu, W. L., Dvorkin, C., & Chael, A. 2018, Phys. Rev., D97, 103530

Yèche, C., Palanque-Delabrouille, N., Baur, J., & du Mas des Bourboux, H. 2017, JCAP, 1706, 047

Zumalacarregui, M. & Seljak, U. 2018, Phys. Rev. Lett., 121, 141101