UNIVERSITY OF AMSTERDAM

The Copernicus Complexio: Statistical properties of warm dark matter haloes

Bose, S.; Hellwing, W.A.; Frenk, C.S.; Jenkins, A.; Lovell, M.R.; Helly, J.C.; Li, B.

Link to publication

The Copernicus Complexio: statistical properties of warm dark matter haloes

Sownak Bose,[1][*] Wojciech A. Hellwing,[1,2] Carlos S. Frenk,[1] Adrian Jenkins,[1] Mark R. Lovell,[3,4] John C. Helly[1] and Baojiu Li[1]

[1]*Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK*
[2]*ICM, University of Warsaw, ul. Pawińskiego 5a, PL-02-106 Warsaw, Poland*
[3]*GRAPPA Institute, Universiteit van Amsterdam, Science Park 904, NL-1098 XH Amsterdam, the Netherlands*
[4]*Instituut-Lorentz for Theoretical Physics, Niels Bohrweg 2, NL-2333 CA Leiden, the Netherlands*

ABSTRACT

The recent detection of a 3.5 keV X-ray line from the centres of galaxies and clusters by Bulbul et al. and Boyarsky et al. has been interpreted as emission from the decay of 7 keV sterile neutrinos which could make up the (warm) dark matter (WDM). As part of the *Copernicus Complexio* (coco) programme, we investigate the properties of dark matter haloes formed in a high-resolution cosmological *N*-body simulation from initial conditions similar to those expected in a universe in which the dark matter consists of 7 keV sterile neutrinos. This simulation and its cold dark matter (CDM) counterpart have $\sim$13.4 bn particles, each of mass $\sim 10^5\,h^{-1}\,\mathrm{M_{\odot}}$, providing detailed information about halo structure and evolution down to dwarf galaxy mass scales. Non-linear structure formation on small scales ($M_{200} \lesssim 2 \times 10^9\,h^{-1}\,\mathrm{M_{\odot}}$) begins slightly later in coco-warm than in coco-cold. The halo mass function at the present day in the WDM model begins to drop below its CDM counterpart at a mass $\sim 2 \times 10^9\,h^{-1}\,\mathrm{M_{\odot}}$ and declines very rapidly towards lower masses so that there are five times fewer haloes of mass $M_{200} = 10^8\,h^{-1}\,\mathrm{M_{\odot}}$ in coco-warm than in coco-cold. Halo concentrations on dwarf galaxy scales are correspondingly smaller in coco-warm, and we provide a simple functional form that describes its evolution with redshift. The shapes of haloes are similar in the two cases, but the smallest haloes in coco-warm rotate slightly more slowly than their CDM counterparts.

Key words: methods: numerical – dark matter.

1 INTRODUCTION

The identity of dark matter, the dominant matter component of the Universe, has long been a subject of great interest in cosmology. In the last three decades, the model of non-relativistic dark matter consisting of heavy weakly interacting particles with negligible thermal velocities at early times, the cold dark matter (CDM) model, has become the cornerstone of the standard cosmological paradigm. The standard model with dark energy in the form of a cosmological constant, Λ (ΛCDM, henceforth just CDM) has been very successful in predicting and matching observational data on a wide range of scales, from the temperature fluctuations in the Cosmic Microwave Background (Planck Collaboration I 2014) to the statistics of galaxy clustering (Colless et al. 2001; Zehavi et al. 2002; Hawkins et al. 2003; Tegmark et al. 2004; Cole et al. 2005;

Eisenstein et al. 2005; for a comprehensive review on the subject, see Frenk & White 2012).

With the advent of the Large Hadron Collider (LHC) it was hoped that one of the best-motivated CDM candidates, the lightest supersymmetric particle (the neutralino) would be found. The lack of evidence for supersymmetry at the LHC and the absence of a convincing direct or indirect signal for CDM (but see Hooper & Goodenough 2011) has encouraged the exploration of viable alternatives. One of the most promising alternatives is the sterile neutrino (Dodelson & Widrow 1994; Asaka & Shaposhnikov 2005), which behaves as warm dark matter (WDM) due to the particles' non-negligible thermal velocities at early times. Being collisionless, this leads to free streaming and the damping of perturbations in the density field, creating a cutoff in the matter power spectrum on the scale of dwarf galaxies.

A simple extension of the Standard Model of particle physics, called the neutrino Minimal Standard Model (νMSM; Boyarsky, Ruchayskiy & Shaposhnikov 2009), consists of three right-handed sterile neutrinos in which, for a specific choice of parameters, one

[*]E-mail: sownak.bose@durham.ac.uk

of the sterile neutrinos behaves as a dark matter particle and the model explains neutrino flavour oscillations. Each one of this triplet of particles has its mass below the electroweak scale; one in the keV scale (denoted by M_1), and two in the GeV scale (denoted by M_2 and M_3). The former behaves as a relativistic particle at the time of neutrino decoupling and acts as WDM, and is then redshifted to non-relativistic energies during the radiation-dominated era. Unlike a thermal relic, the cutoff in the power spectrum introduced by a sterile neutrino of a fixed mass depends on a second parameter, the lepton asymmetry. As we explain later in the following section, it is possible to approximate the sterile neutrino power spectrum with a WDM thermal relic equivalent, particularly for very low and very high values of the lepton asymmetry.

The unidentified 3.53 keV X-ray line originally detected in the spectrum of a stack of galaxy clusters (Bulbul et al. 2014b) and in the spectra of M31 and the Perseus cluster (Boyarsky et al. 2014) could be a decay signal of sterile neutrino dark matter (or other kinds of decaying dark matter), with a particle mass of 7 keV. More recently, Boyarsky et al. (2015) have also identified a similar line in the centre of the Milky Way. While the excess at 3.5 keV has been seen in other studies (e.g. Urban et al. 2015), several groups have questioned the interpretation of this detection. For example, Riemer-Sorensen (2014) re-analysed the signal in *Chandra* observations of the Milky Way, attributing it to a Potassium (K VIII) line. Of course, the Galactic Centre is heavily contaminated by X-rays, which introduces uncertainties, a point made by Boyarsky et al. (2015).

Systematic effects can result from the atomic data used in modelling the plasma, as argued by Jeltema & Profumo (2014), who found no excess when re-analysing the Boyarsky et al. (2014) data and claimed that any signal at 3.5 keV could be explained by known K XVIII and Chlorine (Cl XVII) lines. Bulbul et al. (2014a) put this latter result down to the use of 'incorrect atomic data and inconsistent spectroscopic modelling' by Jeltema & Profumo (2015). A further non-detection was then reported in the stacked spectra of galaxies from *Chandra* and *XMM–Newton* (Anderson, Churazov & Bregman 2015), while most recently, Malyshev, Neronov & Eckert (2014) analysed the spectra of stacked dwarf galaxies from *XMM–Newton* and claimed to rule out the Andromeda signal detected by Bulbul et al. (2014b) at the 4.6σ confidence level. This has spurred other groups (see for example, Conlon & Day 2014) to associate the 3.53 keV signals to the conversion of a sterile neutrino into an axion, and its subsequent decay into photons. Such a mechanism requires a magnetic field, the presence and strength of which can vary from galaxy to galaxy, a scenario that could explain why this line is only seen in some objects.

Clearly, whether or not the 3.53 keV line really does correspond to a sterile neutrino decay remains an open question. It is, therefore, important to investigate the predictions for the formation of cosmic structures in a model in which the dark matter consists of particles that could decay producing such a line. Constraints on such models can be set from the observed clustering of the Lyman α forest at high redshift whose small-scale structure would be erased if the dark matter were warm. On these grounds, Viel et al. (2013) recently set a (current) lower limit of 3.3 keV for the mass of a dominant thermal WDM particle.

Coincidentally, the power spectrum of a 3.3 keV thermal WDM particle is well approximated by that of a 7 keV sterile neutrino for a lepton asymmetry of $L_6 = 8.66$. This corresponds to the smallest allowed value of the power spectrum cutoff length (i.e. to the 'coldest' power spectrum possible) for a sterile neutrino of mass 7 keV. This is the model that we will explore in this work.

Ruling out this model from astronomical data on small scales would rule out the entire family of 7 keV sterile neutrino candidates. To investigate the model we use high resolution *N*-body simulations whose results we compare with those of CDM simulations with the same phases in the initial conditions. We are interested exclusively in characterizing the properties of dark matter haloes of mass in the region of the power spectrum cutoff and, in this study, we ignore the effects of baryons. Such effects must be taken into account when comparing model predictions with observations. In the case of CDM, relevant baryon effects on the small scales of interest here have recently been quantified by Sawala et al. (2013, 2014, 2015) and Schaller et al. (2015).

The layout of this paper is as follows. In Section 2, we introduce the simulations used in this work, the modelling of the WDM component, and describe how we tackle the issue of spurious halo formation in our simulations. In Section 3, we present our main results from the comparison of WDM and CDM from our simulations, in terms of both the large-scale distribution of matter, and the internal structure of haloes. Finally, in Section 4, we summarize our findings and look into some future work that will be carried out with the same set of simulations.

2 THE SIMULATIONS

In this section, we provide an overview of the initial conditions and modelling of the WDM component in our simulations.

2.1 The simulation set-up

The *N*-body simulations presented in this paper are part of the *Copernicus Complexio* (COCO) simulation programme (Hellwing et al. 2015) being carried out by the *Virgo Consortium*. This is a set of cosmological 'zoom-in' simulations (Katz & White 1993; Frenk et al. 1996), as was done in the GIMIC simulations (Crain et al. 2009). The parent simulation, called the *Copernicus Complexio Low Resolution* (or COLOR) simulation, followed the evolution of 4.25 billion particles in a periodic box of size 70.4 h^{-1}Mpc. We extracted a roughly spherical region of radius $\sim$18 h^{-1} Mpc, and centred on the location $(42.2, 51.2, 8.8)h^{-1}$ Mpc in the COLOR volume. Both COLOR and COCO assume cosmological parameters derived from the 7-yr *Wilkinson Microwave Anisotropy Probe* (*WMAP* 7) data (Komatsu et al. 2011), with the parameters: $\Omega_m = 0.272$, $\Omega_\Lambda = 0.728$, $h = 0.704$, $n_s = 0.967$ and $\sigma_8 = 0.81$. Here, $\Omega_{\{m, \Lambda\}}$ represents the present-day fractional contribution of matter and the cosmological constant, respectively, in units of the critical density $\rho_c = 3H_0^2/8\pi G$, $h = H_0/100 \, \text{km s}^{-1} \, \text{Mpc}^{-1}$ is the dimensionless Hubble parameter, n_s is the spectral index of the primordial power spectrum and σ_8 is the linear rms density fluctuation in a sphere of radius 8 h^{-1} Mpc at $z = 0$.

Dark matter particles with three different masses are used in regions simulated at different resolutions within the parent simulation volume. Initially, the high-resolution region has a shape similar to an amoeba which approximates a sphere of radius $\sim$17.4 h^{-1} Mpc at the present time. It contains 12.9 billion particles of mass $1.135 \times 10^5 \, h^{-1} \, \text{M}_\odot$. The volume surrounding this region contains the medium- ($3.07 \times 10^6 \, h^{-1} \, \text{M}_\odot$) and low-resolution ($1.96 \times 10^8 h^{-1} \, \text{M}_\odot$) particles. We have taken care to minimize contamination of the high-resolution region by lower mass particles and all the haloes discussed in this study are entirely made up of the high-resolution particles. The gravitational softening was kept fixed at $\epsilon \sim 230 \, h^{-1}$ pc for the high-resolution particles, increasing by a factor of 10 each time for the medium- and low-resolution particles.

Table 1. Cosmological parameters used in the COCO simulations, and its parent volume, COLOR. Here, m_{WDM} is the mass of the thermal relic WDM particle, N_p is the total number of particles (of all types) used in the simulation, V_h is the approximate volume of the high-resolution region at $z = 0$, $m_{p,\,h}$ is the mass of an individual high-resolution dark matter particle, $N_{p,\,h}$ is the total number of particles of this species, whereas ϵ_h is the softening length applied to them. The cosmological parameters h, Ω_m, Ω_Λ and σ_8 are as described in the text. The phases for the parent COLOR simulation can be generated using the PANPHASIA phase descriptor provided in the last row. The blank fields in the COCO column mean that the parameter assumes the same value as in the parent simulation, COLOR.

Parameter	COLOR (Parent volume)	COCO (This paper)
Box size	70.4 h^{-1} Mpc	–
m_{WDM}	3.3 keV	–
N_p	4, 251, 528, 000	13, 384, 245, 248
V_h	$70.4^3\ h^{-3}$ Mpc^3	$\sim 2.2 \times 10^4\ h^{-3}$ Mpc^3
$m_{p,\,h}$	$6.196 \times 10^6\ h^{-1}$ M$_\odot$	$1.135 \times 10^5\ h^{-1}$ M$_\odot$
$N_{p,\,h}$	4, 251, 528, 000	12, 876, 807, 168
ϵ_h	1 h^{-1} kpc	230 h^{-1} pc
h	0.704	–
Ω_m	0.272	–
Ω_Λ	0.728	–
σ_8	0.81	–
Phase descriptor	[Panph1,L16,(31250,23438,39063),S12,CH582187950,COLOR]	–

The simulation ran from $z = 127$ to $z = 0$ using the GADGET3 code, which is an updated version of the publicly available GADGET2 code (Springel, Yoshida & White 2001a; Springel 2005). Phase information for the creation of the initial conditions for both COCO-WARM and COCO-COLD was obtained from the public Gaussian white noise field PANPHASIA (Jenkins 2013), and perturbations thereafter were calculated using the second-order Lagrangian Perturbation Theory algorithm presented in Jenkins (2010). The details of the simulation, along with the PANPHASIA phase descriptor, are summarized in Table 1.

The distinctive feature of WDM particles are non-negligible thermal velocities at early times, which result in free streaming that washes out perturbations in the matter distribution below the free streaming scale (Bond & Szalay 1983; Schneider et al. 2012; Benson et al. 2013). As a result, we expect the abundance, distribution and internal structure of WDM haloes to be different from those of CDM haloes. Indeed, thermal velocities introduce a limit to the fine-grained phase space density in dark matter haloes, creating cores in the density profile (Macciò et al. 2012; Shao et al. 2013). However, as shown in these papers, the cores produced by realistic thermal relics are only a few parsecs in size, and thus not astrophysically relevant. In our simulations we can neglect these thermal velocities, which at $z = 0$ are of the order of a few tens of metres per second (Lovell et al. 2012) so, over the course of the simulation, which starts at $z = 127$, the particles would travel only a few kiloparsecs, comparable to the mean interparticle spacing of the high-resolution particles.

The WDM power spectrum of density fluctuations is often modelled by the transfer function, $T(k)$, relative to the CDM case:

$$P_{\mathrm{WDM}}(k) = T^2(k) P_{\mathrm{CDM}}(k). \tag{1}$$

We approximate $T(k)$ using the fitting formula provided by Bode, Ostriker & Turok (2001):

$$T(k) = \left(1 + (\alpha k)^{2\nu}\right)^{-5/\nu}, \tag{2}$$

where α and ν are constants. As computed by Viel et al. (2005), for $k < 5\ h^{-1}$ Mpc, the value $\nu = 1.12$ provides the best-fitting transfer function. The value of α is dependent on the mass of the WDM particle (Viel et al. 2005):

$$\alpha = 0.049 \left[\frac{m_{\mathrm{WDM}}}{\mathrm{keV}}\right]^{-1.11} \left[\frac{\Omega_{\mathrm{WDM}}}{0.25}\right]^{0.11} \left[\frac{h}{0.7}\right]\ h^{-1}\ \mathrm{Mpc}, \tag{3}$$

and determines the scale of the cutoff due to free streaming in the WDM power spectrum relative to CDM. It should be noted that this transfer function is a fit to the full thermal relic power spectrum, obtained by solving the Boltzmann equation (Lovell et al., in preparation).

As we can see in equation (3), the 'warmer the dark matter particle (i.e. the lower its rest mass is), the larger the scale at which the cutoff in the power spectrum occurs. One way to define the characteristic scale in the power spectrum is through the 'half-mode' wavenumber, k_{hm}, where the transfer function in equation (2) drops by a factor of 2:

$$k_{\mathrm{hm}} = \frac{1}{\alpha} \left(2^{\nu/5} - 1\right)^{1/2\nu}. \tag{4}$$

The associated 'half-mode mass', M_{hm}, is the mean density enclosed within this half mode:

$$M_{\mathrm{hm}} = \frac{4\pi}{3} \bar{\rho} \left(\frac{\lambda_{\mathrm{hm}}}{2}\right)^3. \tag{5}$$

For the 3.3 keV model, this occurs at around $M_{\mathrm{hm}} \sim 2 \times 10^8\ h^{-1}$ M$_\odot$ (Colín, Valenzuela & Avila-Reese 2008; Angulo, Hahn & Abel 2013; Viel et al. 2013). We will show later that differences in the formation time of haloes in WDM and CDM begin to appear below $\sim 2 \times 10^9\ h^{-1}$ M$_\odot$, approximately an order of magnitude above the half-mode mass scale.

The form of the power spectrum for sterile neutrino dark matter is determined by two parameters – the mass of the particle, m_{ν_s}, and the lepton asymmetry, L_6:

$$L_6 \equiv 10^6 \left(\frac{n_{\nu_e} - n_{\bar{\nu}_e}}{s}\right), \tag{6}$$

where n_{ν_e} is the number density in electron neutrinos, $n_{\bar{\nu}_e}$ the number density in electron antineutrinos, and s is the entropy density of the Universe (Laine & Shaposhnikov 2008). The scales at which the power spectrum is suppressed for sterile neutrinos vary non-monotonically as a function of L_6. If L_6 is very small ($\ll 1$) the power spectrum exhibits a similar abrupt cutoff to that of a thermal

Figure 1. The (dimensionless) matter power spectrum for: a thermal 3.3 keV WDM (red), a sterile neutrino of mass $m_{\nu_s} = 7$ keV and lepton asymmetry $L_6 = 8.66$ (blue) and CDM (black). Both the WDM and sterile neutrino power spectra have significantly suppressed power at small scales, with the deviation from CDM case at almost identical scales: $\log(k) \gtrsim 1.0 \, h \, \mathrm{Mpc}^{-1}$. Also shown as thin coloured lines are power spectra for 7 keV sterile neutrinos with different values of L_6, as indicated in the legend.

relic. As L_6 is increased, the cutoff becomes gentler and k_{hm} shifts to larger values. At some value of L_6 (typically between 8 and 25 depending on the sterile neutrino mass), k_{hm} reaches a maximum; for still higher L_6, k_{hm} retreats to lower k and returns to its original shape and position (Shi & Fuller 1999; Abazajian 2014; Lovell et al., in preparation).

The power spectra used in the COCO simulations are shown as thick lines in Fig. 1: CDM in black, 3.3 keV WDM in red and 7 keV sterile neutrinos with $L_6 = 8.66$ in blue. All three power spectra agree on large scales. On small scales, the two WDM models differ from CDM. k_{hm} for the sterile neutrino case occurs at a very similar scale, and the cutoff has a similar shape to that for the thermal relic case. On smaller scales still, the sterile neutrino power spectrum has more power than its thermal counterpart, but the differences only become significant on scales where the amplitude is, at most, a few per cent of the peak amplitude. These differences are negligible and can be safely ignored in our simulations. The thin lines in the figure correspond to 7 keV sterile neutrino power spectra for different values of the lepton asymmetry, L_6. The $L_6 = 8.66$ model that we have simulated corresponds to the 'coldest' possible 7 keV sterile neutrino.

2.2 Halo identification and matching

Haloes were identified in our simulations using the friend-of-friend (FOF) algorithm (Davis et al. 1985) with a linking length of 0.2 times the mean interparticle separation, and a minimum of 20 particles. Gravitationally bound substructures within these groups were then identified using the SUBFIND algorithm (Springel et al. 2001b), although in this paper, we will be mostly concerned with the properties of the WDM FOF groups. We determine the halo centre using the 'shrinking sphere' method of Power et al. (2003).

Table 2. Number of groups and subhaloes identified by the FOF algorithm and SUBFIND in COLOR and COCO at $z = 0$.

Simulation	$N_{\mathrm{FOF}}(z=0)$	$N_{\mathrm{subs}}(z=0)$
COLOR-COLD	3, 961, 192	4, 770, 041
COLOR-WARM	2, 609, 122	3, 082, 275
COCO-COLD	8, 896, 811	10, 502, 187
COCO-WARM	2, 548, 743	2, 830, 514

In short, we recursively compute the centre of mass of all particles within a shrinking sphere, until a convergence criterion is met. In each iteration, the radius of the sphere is reduced by 5 per cent, and stopped when only 1000 particles or 1 per cent of the particles of the initial sphere (whichever is smaller) are left.

Comparing halo statistics between sets of simulations requires consistent definitions for the various properties of the haloes. In this work, we make use of two definitions of mass: M_{FOF}, which is the mass of all particles identified by the algorithm as belonging to the FOF group, and M_{200}, which is the mass contained within a sphere of radius r_{200} (centred on the 'shrinking sphere' centre defined above), within which the average density is 200 times the critical density of the Universe (ρ_c) at the specified redshift. Another common radius used to define a halo edge is the virial radius, r_{vir}, within which the density of the halo $\bar{\rho}(< r_{\mathrm{vir}}) = \Delta \rho_c$, where $\Delta \sim 178 \Omega_{\mathrm{m}}^{0.45}$ (motivated by the spherical collapse model, Eke, Cole & Frenk 1996). (Note that this definition is consistent with the virial overdensity relation in Bryan & Norman 1998.) Table 2 summarizes the total number of groups and self-bound substructures identified at $z = 0$ in our simulations.

Since both COCO-WARM and its COLD counterpart were simulated using the same initial phases, we are able to match many objects between the two simulations. This allows us to make also object-by-object comparisons in addition to comparing just statistical distributions of halo properties. In order to correctly match the haloes we do the following: first, we take the 50 most-bound particles from a COCO-WARM halo, and look for the COCO-COLD halo in which there are at least 25 (50 per cent) of these particles. We then confirm the match by repeating the same process, this time starting with the COCO-COLD haloes, in decreasing order of mass. This results in a bijective match between haloes in the two simulations. Using this method, we are able to match around 97 per cent of haloes with $M_{200} > 10^8 \, h^{-1} \, \mathrm{M}_\odot$.

2.3 Spurious haloes and their removal

Number counts of haloes and subhaloes are fundamental statistics of the halo population, so the correct identification of haloes is of primary importance. It has been known for some time (Wang & White 2007; Angulo et al. 2013; Lovell et al. 2014) that in simulations in which the initial power spectrum has a resolved cutoff, as is the case for COCO-WARM, small-scale structure is seeded in part by the discreteness of the particle set. In other words, a substructure finder will identify density peaks that have arisen not as a result of gravitational instabilities from a cosmological perturbation. These artificial fragments can often by identified 'by eye' as they tend to be regularly spaced along filaments of the mass distribution. They produce a power-law-like upturn at small masses in the WDM mass function. Since this is just a numerical (and resolution-dependent) artefact of our WDM simulations, care must be taken to identify these spurious haloes and, if appropriate, remove them from the halo catalogue. While it is, in principle, possible to eliminate these

structures by increasing the resolution of the simulation, this is computationally prohibitive: Wang & White (2007) have shown that the mass at which spurious structures dominate the mass function scales with the number of particles in the simulation, N, as $M \propto N^{-1/3}$.

Lovell et al. (2014) developed an algorithm to identify spurious clumps in WDM simulations. A large number of them can be removed by performing a mass cut below a resolution-dependent scale, as suggested by Wang & White (2007):

$$M_{\rm lim} = 10.1\bar{\rho}\, d\, k_{\rm peak}^{-2}, \qquad (7)$$

where d is the mean interparticle separation and $k_{\rm peak}$ is the spatial frequency at which the dimensionless power spectrum, $\Delta^2(k)$, has its maximum. Applying this condition on its own would also remove some genuine haloes that form below this scale. Lovell et al. (2014) refined this criterion by also making a cut on the basis of the shapes of the initial Lagrangian regions from which WDM haloes form. They find that the spurious candidates tend to have much more flattened configurations in their (unperturbed) initial positions than genuine haloes, as judged from a CDM simulation. Defining the sphericity, s, of haloes as the axis ratio, c/a, of the minor to major axes in the diagonalized moment of inertia tensor of the initial particle load, the sphericity cut is made such that 99 per cent of the CDM haloes at that redshift lie above the threshold.

Following exactly the methodology of Lovell et al. (2014), we clean the COCO-WARM catalogue as follows: (1) remove all (sub)haloes with $s_{\rm half-max} < 0.165$,[1] irrespective of mass; (2) for those that pass (1), remove (sub)haloes with $M_{\rm max} < 0.5 M_{\rm lim}$. Here, $M_{\rm max}$ is the maximum mass attained by a (sub)halo during its evolution, and $s_{\rm half-max}$ is the sphericity ($=c/a$) of the (sub)halo at the half-maximum mass snapshot. This is chosen so as to identify a (sub)halo at a time well before it falls into a larger host, when its particles are subject to tidal stripping. The factor of 0.5 in condition (2) is calibrated by matching between resolutions in the AQUARIUS simulations (see Lovell et al. 2014 for details). Having done so, we find that over 91 per cent of the (FOF) haloes formed in COCO-WARM are in fact spurious, and are rejected from the halo catalogue when computing properties like mass functions. The elements of this section are summarized in Fig. 2.

3 RESULTS

In both cold and warm models, dark matter haloes assemble in a hierarchical way, acquiring mass by merging with other haloes and by smoothly accreting ambient mass (e.g. Press & Schechter 1974; Frenk et al. 1985; Lacey & Cole 1993; Wechsler et al. 2002). In this section, we focus on global halo properties such as formation times, abundance and internal structure. We make a direct comparison between our cold and WDM models. On scales much larger than the WDM suppression scale in the initial power spectrum, we expect the properties of haloes to be very similar in the two cases, but differences should become increasingly important at $\sim 2 \times 10^9\, h^{-1}\,{\rm M}_\odot$ and below.

3.1 Redshift of formation

The absence of primordial perturbations below the cutoff scale in the WDM power spectrum induces differences in the formation

[1] The criterion $s_{\rm half-max} < 0.165$ is appropriate for haloes identified at $z = 0$; for higher redshifts, one needs to determine the 1 per cent sphericity cut at *that* redshift.

Figure 2. Number density of haloes in the sphericity versus maximum mass space in COCO-WARM at $z = 0$. The dashed black lines show the cuts on sphericity and mass that we use to clean the halo catalogue. Rejected (spurious) candidates are those that fail the cuts in the manner described in the text.

times of the smallest haloes. We can visualize these differences directly by examining the images displayed in Fig. 3. At early times, the projected density in COCO-WARM (right-hand panels) is visibly smoother than the equivalent projection in COCO-COLD (left-hand panels), which has a 'grainier' appearance owing to the very large number of haloes below $\sim 10^9\, h^{-1}\,{\rm M}_\odot$ that form in this case, well before the first objects have collapsed in COCO-WARM. Thus, the onset of the structure formation process in this simulation is delayed relative to its CDM counterpart.

In order to quantify the different halo formation epochs in COCO-WARM and COCO-COLD, we trace the evolution of each FOF group through its merger tree, and define the redshift of formation as the first time when the mass of the most massive progenitor exceeds half the final FOF mass: $M(z_{\rm form}) = M(z=0)/2$ (e.g. Harker et al. 2006; Neto et al. 2007). Other definitions of halo formation time also exist in the literature (e.g. Navarro, Frenk & White 1996, 1997), which should be borne in mind when making comparisons.

The result, for *all* haloes in COCO-WARM (including spurious objects) and COCO-COLD is shown in Fig. 4. The formation redshifts of haloes of mass $M_{200} \gtrsim 2 \times 10^9\, h^{-1}\,{\rm M}_\odot$, are very similar in COCO-WARM and COCO-COLD, as expected. The difference between the two begins to manifest below a mass of $M_{200} \sim 2 \times 10^9\, h^{-1}\,{\rm M}_\odot$, an order of magnitude above the half-mode mass scale for a 3.3 keV WDM particle (c.f. Section 2.1). For these smaller haloes, $z_{\rm form}$ is lower for WDM than CDM. The sudden upturn in the WDM $z_{\rm form}$ for $M_{200} < 10^8\, h^{-1}\,{\rm M}_\odot$ (shown in the open red circles) is a signature of the spurious haloes described in Section 2.3. From here on, we will exclude these spurious haloes and only show results from the cleaned COCO-WARM sample. The difference in formation times is a subject we will revisit when comparing the concentration–mass relations of WDM and CDM in Section 3.4. Note that in this figure, we include all haloes, and not necessarily matched between CDM

Figure 3. Redshift evolution of the projected dark matter density in COCO-COLD (left) and the 3.3 keV COCO-WARM universe (right). From top to bottom, the top three panels show snapshots at $z = 10$, $z = 6$, $z = 1$ of the projected mass density in cubes of side $2\,h^{-1}$ Mpc, centred on the most massive group at $z = 0$. The bottom panels show zooms of a $5 \times 10^{10}\,h^{-1}\,\mathrm{M_\odot}$ halo at $z = 0$ in a cube of side $150\,h^{-1}$ kpc. The emergence of small haloes at early times is apparent in the CDM case, when the WDM distribution is much smoother. The formation of large haloes occurs at roughly the same time in the two simulations and the subsequent growth of these haloes is similar in the two cases. In the zoom shown in the bottom panel, the lack of substructure in the WDM case compared to its CDM counterpart is stark.

Figure 4. The median redshift of formation of all FOF groups in coco-warm and coco-cold, as a function of the halo mass, M_{200}. The redshift z_{form} is defined in the text. The error bars/shaded region represent the bootstrapped errors on the median in each mass bin in coco-warm and coco-cold, respectively. As expected, there is good agreement at the high-mass end, whereas the differences between CDM and WDM become apparent below $\sim 2 \times 10^9\, h^{-1}\, M_\odot$. The thin red line is a fit to the coco-warm redshift of formation, using equation (8).

and WDM, which is why the medians at the largest mass bins are not exactly identical.

We find that the delay in the formation time of coco-warm haloes of a given mass, relative to coco-cold, is well described by the fitting function:

$$\frac{z_{form}^{WDM}}{z_{form}^{CDM}} = \left(1 + a\,\frac{M_{hm}}{M_{200}}\right)^{-b}, \tag{8}$$

where M_{hm} is the half-mode mass introduced in Section 2.1, $a = 1.23$ and $b = 0.56$. This fit is shown as the thin red line in Fig. 4.

3.2 Differential halo mass functions

Counting the number of dark matter haloes as a function of their mass is one of the simplest and most important population statistics that one can use to distinguish between WDM and CDM models, since fewer haloes will form in the former close to the half-mode mass.

In Fig. 5, we show the build-up of the halo population as a function of redshift in coco-cold (solid lines) and coco-warm (dashed lines). The shaded regions and error bars represent the Poisson uncertainty in both cases. Spurious haloes have been omitted from the WDM differential halo mass function (dHMF) at each redshift, using the methodology outlined in Section 2.3. The edge of the grey region marks the nominal resolution limit of our simulation which corresponds to a halo with at least 300 particles within r_{200} ($M_{200} \sim 3.4 \times 10^7\, h^{-1}\, M_\odot$). This 300-particle limit was derived by comparing the mass function of coco-cold with that of its lower resolution counterpart color-cold. Below this limit, the results of the simulations become increasingly unreliable. The results at high


Figure 5. The redshift evolution of the halo mass function in coco-cold and coco-warm. The solid lines show the CDM results, with the shaded regions representing the associated 1σ Poisson errors. The dashed lines with error bars represent the equivalent relation from coco-warm, with spurious haloes removed. The different colours show results for a selection of redshifts, as indicated in the legend. The grey shaded region corresponds to haloes with fewer than 300 particles within r_{200}.

masses are noisy because of the small number of high-mass haloes formed in the relatively small volume of our simulations.

The general trend across redshifts is similar: for haloes with $M_{200} > 2 \times 10^9\, h^{-1}\, M_\odot$, the dHMF in coco-warm and coco-cold are almost identical. The abundance of haloes below this mass scale is strongly suppressed in coco-warm, to the extent that, at $z = 10$, there are five times fewer $\sim 10^8\, h^{-1}\, M_\odot$ haloes than in coco-cold. The delayed non-linear structure formation below $\sim 2 \times 10^9\, h^{-1}\, M_\odot$ can also be seen from the fact that there are as many haloes with $M_{200} = 10^8\, h^{-1}\, M_\odot$ in coco-warm at $z = 10$, as there are haloes with $M_{200} = 6 \times 10^8\, h^{-1}\, M_\odot$ in coco-cold at that redshift.

Within the CDM paradigm, there are a number of analytic predictions for the dHMF, notably the Press–Schechter formula (Press & Schechter 1974; Bond et al. 1991; Lacey & Cole 1993), and the ellipsoidal collapse model (ST; Sheth & Tormen 1999, although this model is not fully analytic since it is tuned to numerical simulations). The dHMF is given by

$$\frac{dn}{d\log M} = \frac{\bar{\rho}}{M} f(\nu) \left|\frac{d\log \sigma^{-1}}{d\log M}\right|, \tag{9}$$

where $f(\nu)$ is the so-called *halo multiplicity function* and for hierarchical cosmologies has a universal form (see e.g. Jenkins et al. 2001; Reed et al. 2007; Tinker et al. 2008; Angulo et al. 2012). In the ST formalism, it is approximated by

$$f(\nu) = A \sqrt{\frac{2q\nu}{\pi}} \left[1 + (q\nu)^{-p}\right] e^{-q\nu/2}. \tag{10}$$

Here, $\nu \equiv \delta_c^2(z)/\sigma^2(M)$, $A = 0.3222$, $q = 0.707$ and $p = 0.3$. In linear theory, $\delta_c(z) \equiv 1.686/D(z)$, where $D(z)$ is the linear growth rate of perturbations. The value of δ_c is appropriate for the Einstein–de Sitter model, but differs slightly in ΛCDM due to a weak dependence on $\Omega_m(z)$. Finally, $\sigma^2(M)$ is the variance in the mass density

Figure 6. The fractional variance of the density field, $\sigma(M)$, calculated in equation (11) using a top-hat filter in real-space for CDM, and a sharp k-space filter for WDM. The flattening of the relation below $10^8 \, h^{-1} \mathrm{M_\odot}$ is due to the suppression of power below these scales in WDM, relative to CDM. The dashed line indicates the upper limit to the halo masses formed in our volume-limited simulations.

field on mass scale, M, given by

$$\sigma^2(M) = \int \frac{dk}{k} \Delta^2(k) \tilde{W}^2(k, M). \tag{11}$$

Here, $\tilde{W}(k, M)$ is the Fourier transform of a window function containing mass M, and $\Delta^2(k)$ is the dimensionless power spectrum as defined in Fig. 1.

In the Press–Schechter and Sheth–Tormen formalisms, the rms fluctuation amplitude, $\sigma^2(M)$, is assumed to be a monotonically increasing function of M. This is no longer true for the truncated power spectrum of WDM, so care must be taken when choosing an appropriate window function. In the CDM, $W(k, M)$ is usually chosen to be the real-space spherical top-hat function, a choice that results in an excellent match to the dHMF in cosmological N-body simulations. The same for WDM predicts an excess of low-mass haloes compared to simulations (Barkana, Haiman & Ostriker 2001; Menci, Fiore & Lamastra 2012; Schneider et al. 2012, but see also Schneider, Smith & Reed 2013). This problem was solved by Benson et al. (2013), who generalized the (extended) Press–Schechter formalism by using the correct solution for the excursion set barrier first-crossing distribution in WDM models. Rather than the top-hat real-space window function, they used a sharp k-space filter for WDM, so that the variance, $\sigma(M)$, remains flat up to the half-mode mass and then declines with increasing mass (see Fig. 6). In this formalism the smoothing scale, R, is defined as

$$R = \frac{a}{k_s}, \tag{12}$$

where $k_s = 2\pi\kappa/\alpha$, α as defined in equation (3), $\kappa = 0.361$ and $a = 2.5$. Benson et al. (2013) choose the free parameters such that the theoretical mass function turns over at the same scale as the halo mass function from simulations. This choice of parameters should


Figure 7. dHMFs from the COCO-WARM and COCO-COLD simulations, compared to the predictions of the ellipsoidal collapse formalism of Sheth & Tormen (1999). The solid lines show the predictions of the standard formalism applied to CDM; the dashed lines show the predictions of the modified, sharp k-space filter of Benson et al. (2013). The symbols represent results from our simulations as denoted in the legend: blue squares for COCO-COLD, green diamonds for all COCO-WARM FOF haloes, red circles for the genuine haloes and yellow stars for spurious haloes. The grey shaded region corresponds to haloes with fewer than 300 particles within r_{200}.

be applicable to all thermal WDM models, since the effect of the WDM suppression is captured in the value of α (equation 3).

In Fig. 7, we compare the $z = 0$ dHMF for COCO-COLD (blue squares), the full COCO-WARM (genuine and spurious objects; green diamonds), the spurious COCO-WARM objects only (yellow stars) and the genuine COCO-WARM haloes only (red circles).

The solid and dashed black lines in Fig. 7 show the ST predictions for the mass functions in CDM and WDM, respectively. For $M_{200} > 2 \times 10^9 \, h^{-1} \mathrm{M_\odot}$, the mass functions for CDM and WDM trace one another exactly, as expected. Below this mass, the WDM mass function begins to peel-off from the CDM case, reaching half the CDM amplitude at $M_{200} \sim 2 \times 10^8 \, h^{-1} \mathrm{M_\odot}$. This agrees with the half-mode mass scale, M_{hm}, introduced in Section 2.1. The raw WDM mass function (green diamonds) is entirely dominated by the spurious objects (yellow stars) below $\sim 4 \times 10^7 \, h^{-1} \mathrm{M_\odot}$, where the mass function shows an artificial upturn. On the other hand, the cleaned WDM sample, represented by the red circles, continues to fall off smoothly from the regime free of artificial haloes. The feature at $\sim 2 \times 10^7 \, h^{-1} \mathrm{M_\odot}$ could be related to the cut, $M_{max} = 3.2 \times 10^7 \, h^{-1} \mathrm{M_\odot}$, applied as part of the cleaning procedure (Section 2.3), but, in any case, this is very close to the resolution limit which is also the mass scale at which the spurious haloes begin to dominate the mass function.

The main conclusion to be drawn from Fig. 7 is that above the resolution limit, the modified ellipsoidal collapse model reproduces the WDM mass function remarkably accurately, over nearly six orders of magnitude in mass.
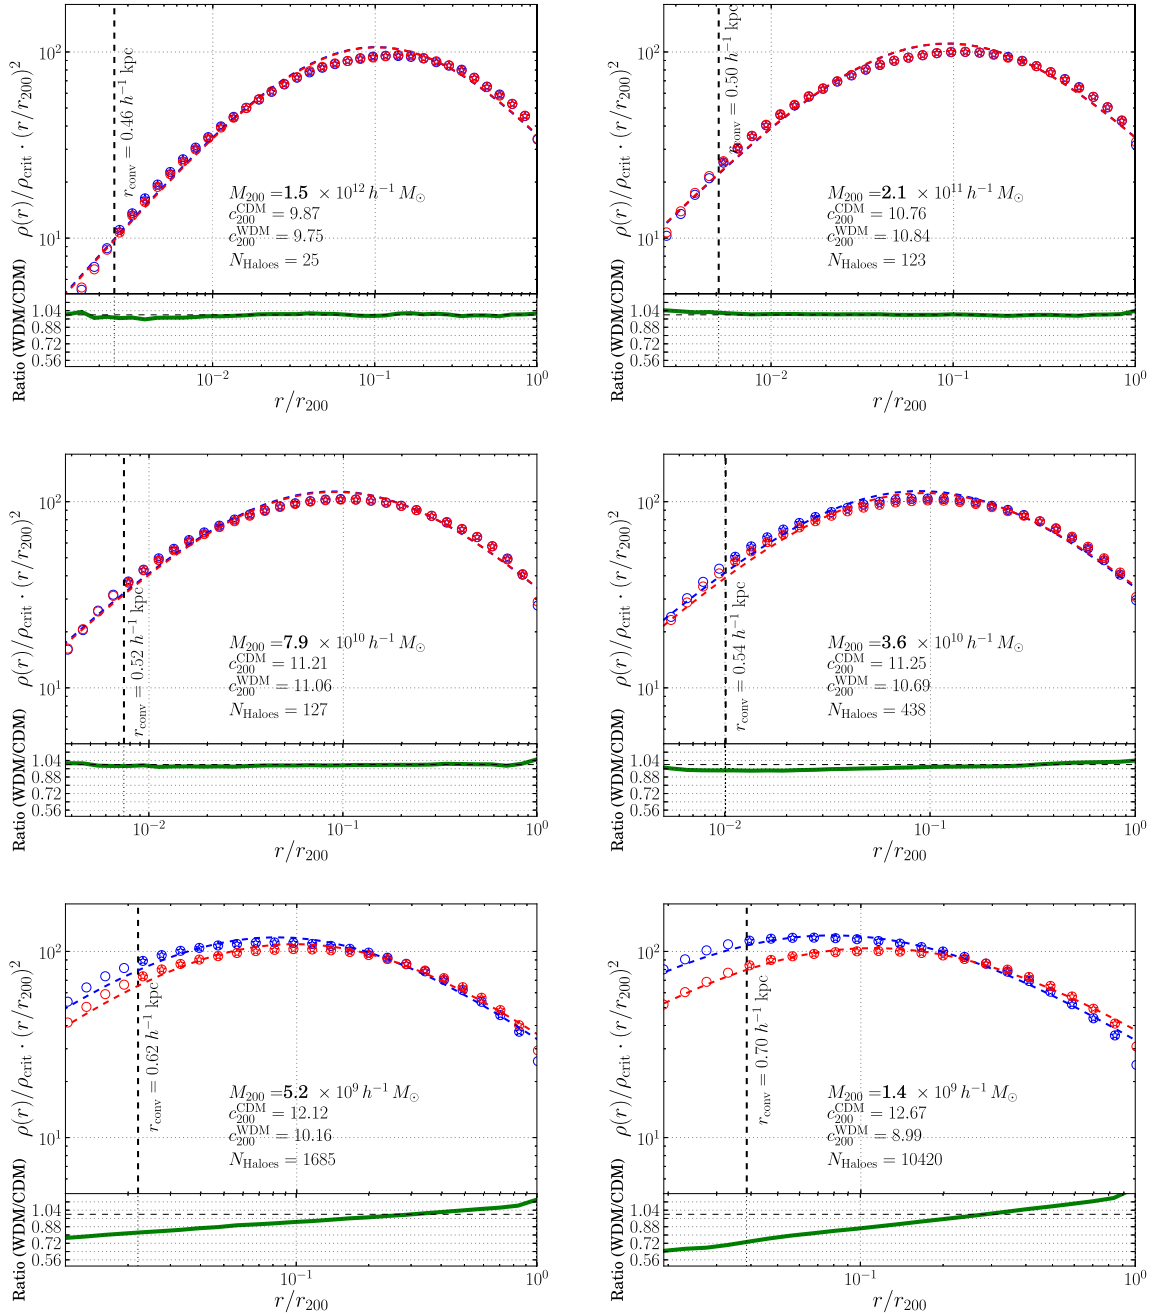
Figure 8. Stacked spherically averaged density profiles in COCO-WARM (red) and COCO-COLD (blue). For each mass bin we compare the profiles of only relaxed, matched haloes in the two simulations; the number in each bin is indicated in each subpanel. The vertical dashed line represents the convergence radius, $r_{\rm conv}$, and filled symbols indicate the range of the profile above this limit, whereas open symbols denote the radial range below it. The dashed red and blue lines are NFW fits to the WDM and CDM profiles, respectively. Note that the density profiles have been scaled by $(r/r_{200})^2$ so as to reduce the dynamic range on the vertical axis. The bottom panels show the ratio of the WDM and CDM densities in each bin.

3.3 Halo density and mass profiles

Spherically averaged radial density profiles provide the simplest and most direct descriptor of halo structure. We calculate profiles in radial shells equally spaced in $\log(r/r_{200})$. As we discussed in Section 2.2, haloes of mass above $10^8 h^{-1} M_\odot$ can be bijectively matched in COCO-WARM and COCO-COLD. To compare density profiles in the two models, we stack the individual profiles of matched and dynamically relaxed haloes in narrow bins of halo mass of width $\Delta \log(M_{200}) = 0.3$. To determine whether or not a halo is relaxed, we make use of the criteria for dynamical equilibrium set

out by Neto et al. (2007): (1) the displacement of the centre of mass from the potential centre should be less than $0.07 r_{\rm vir}$ and (2) less than 10 per cent of the mass within $r_{\rm vir}$ should be in the form of substructure.

The stacked differential density profiles are shown in Fig. 8 for a variety of mass bins, with the ratio of the densities shown in the bottom panels. For masses sufficiently larger than $\sim 2 \times 10^9 h^{-1} M_\odot$, we expect negligible differences in the properties of CDM and WDM haloes: this is apparent in mass bins with $M_{200} > 10^{11} h^{-1} M_\odot$. Systematic differences in the density profiles begin

to appear at around $M_{200} \sim 5 \times 10^{10}\,h^{-1}\,\mathrm{M}_{\odot}$: the WDM haloes have slightly but systematically lower central densities than their CDM counterparts. This halo mass is two orders of magnitude higher than the half-mode mass, and an order of magnitude higher than the scale at which the mass functions begin to differ (Fig. 7). The difference in central density grows as the mass decreases and reaches ~ 30 per cent at the smallest mass bin shown, $M_{200} \sim 1.4 \times 10^{9}\,h^{-1}\,\mathrm{M}_{\odot}$. We discuss the physical reason for this in the next section.

It is now well established that the density profiles of dark matter haloes in general are well described by the NFW profile (Navarro et al. 1996, 1997):

$$\frac{\rho(r)}{\rho_c} = \frac{\delta_c}{(r/r_s)(1+r/r_s)^2}, \tag{13}$$

where δ_c is a characteristic overdensity and r_s is a scale radius. These two parameters are strongly correlated and depend only on halo mass (Navarro et al. 1997). The NFW form is a nearly universal profile in the sense that it approximately fits the profiles of relaxed haloes of any mass formed by gravitational instability from all the initial conditions and cosmological parameters that have been tested so far. The universality of the NFW profile is intimately related to the way in which haloes are assembled (Ludlow et al. 2013).

We fit NFW profiles to the stacked density profiles of COCO-WARM and COCO-COLD in Fig. 8, between the radial range defined by the Power et al. (2003) convergence radius, r_{conv} (defined as the radius within which the relaxation time is of the order of the age of the Universe), and r_{200}, minimizing the following quantity:

$$\sigma_{\mathrm{fit}}^2 = \frac{1}{N_{\mathrm{bins}}-1} \sum_{i=1}^{N_{\mathrm{bins}}} \left[\ln \rho_i - \ln \rho_{\mathrm{NFW}}(\delta_c; r_s)\right]^2. \tag{14}$$

We obtain the best-fitting values of the scale radius, r_s, which defines the halo concentration, $c_{200} = r_{200}/r_s$. This parameter provides a unique characterization of the NFW density profile; the values of c_{200} for the stacked profiles are quoted in Fig. 8. There is a clear trend in that for large halo masses, where the density profiles in COCO-WARM and COCO-COLD are similar, the concentrations are nearly identical but, for masses below $\sim 5 \times 10^{10}\,h^{-1}\,\mathrm{M}_{\odot}$, the concentrations of WDM haloes are systematically lower than those of CDM haloes.

In many cases, even better fits to the density profile are provided by a formula first used by Einasto (1965) to describe star counts in the Milky Way. This formula, which has an additional free parameter, was dubbed the 'Einasto profile' by Navarro et al. (2004), who showed that it provides a very good fit to CDM haloes:

$$\ln\left(\frac{\rho}{\rho_{-2}}\right) = -\frac{2}{\alpha}\left[\left(\frac{r}{r_{-2}}\right)^{\alpha} - 1\right], \tag{15}$$

where ρ_{-2} is the density at $r = r_{-2}$, the radius at which the logarithmic slope of the profile is -2 (or where $r^2\rho$ has its maximum). The parameter r_{-2} in the Einasto profile is analogous to the scale radius, r_s, of the NFW profile. This allows an equivalent definition of halo concentration, $c_{200} = r_{200}/r_{-2}$. The parameter α (not to be confused with the one in equation 3) is a shape parameter that controls the curvature of the profile in the inner regions. A value of $\alpha \simeq 0.17$ results in a good match to CDM haloes over a wide range of masses (Navarro et al. 2004; Gao et al. 2008).

This is demonstrated in Fig. 9, which is similar to Fig. 8, but with Einasto profiles fitted instead of NFW profiles. It is apparent that the shape parameter, α, allows a better fit to the halo density profiles in both COCO-WARM and COCO-COLD, especially in the inner parts. It is also interesting to note that the concentrations inferred from the

Einasto profile fits tend to be slightly lower than those inferred from the NFW profile fits especially at higher masses.

In Fig. 10, we compare the ratio of M_{200} values for individually matched haloes in COCO-WARM and COCO-COLD at the present day. We consider only haloes with $M_{200} > 10^8\,h^{-1}\,\mathrm{M}_{\odot}$ for which we have almost complete matching (~ 97 per cent) between the two simulations, and plot the ratio, $M_{200}^{\mathrm{WDM}}/M_{200}^{\mathrm{CDM}}$ as a function of M_{200}^{CDM}. The solid red line shows the median ratio, whereas the dashed red lines represent the 16th and 84th percentiles. The masses are very similar for objects $>5 \times 10^{10}\,h^{-1}\,\mathrm{M}_{\odot}$, where the ratios agree to within 1 per cent. For masses lower than this, WDM haloes are systematically less massive than their CDM counterparts, with the deficit in WDM halo mass reaching ~ 30 per cent at $M_{200}^{\mathrm{CDM}} = 10^9\,h^{-1}\,\mathrm{M}_{\odot}$. Haloes of these masses in WDM form later than their CDM counterparts and thus have less time to grow.

In Fig. 11, we show the cumulative radial distribution of mass in haloes in COCO-WARM (red lines) and COCO-COLD (blue squares). The ratios are shown in the lower panels. From Fig. 10, we expect the cumulative profiles to be very similar at $r/r_{200} = 1$ except in the lowest mass bin, where WDM haloes are slightly (~ 10 per cent) less massive than their CDM matches. The same trend seen in the density profiles is apparent here: for $M_{200} < 5 \times 10^{10}\,h^{-1}\,\mathrm{M}_{\odot}$, the profiles are less concentrated in the central regions in COCO-WARM than in COCO-COLD. The reason for this difference is discussed in the next section.

3.4 The concentration–mass relation

As mentioned in the previous section, the density profile of a dark matter halo is characterized by its concentration. As a result of their hierarchical formation process, the inner parts of haloes in CDM and WDM are essentially in place even before the bulk of the halo mass is assembled (Wang et al. 2011). The concentration reflects the mean density of the Universe at the epoch when these inner regions are in place and the earlier a halo forms, the higher its concentration is (Navarro et al. 1997).

In Section 3.3, we found that the Einasto profile provides a slightly better fit to the density profiles of WDM and CDM haloes than does the conventional NFW profile. Furthermore, Einasto fits are less sensitive to the radial fitting range (Gao et al. 2008, but see also Ludlow et al. 2013). For these reasons, we proceed to derive the concentration–mass relation in our simulations using fits of the Einasto profile to the density profiles of *individual* haloes (not the stacks). Again, fitting is performed between the convergence radius, r_{conv}, and r_{200}, while minimizing the rms of the fit:

$$\sigma_{\mathrm{fit}}^2 = \frac{1}{N_{\mathrm{bins}}-1} \sum_{i=1}^{N_{\mathrm{bins}}} \left[\ln \rho_i - \ln \rho_{\mathrm{Ein}}(\rho_{-2}; r_{-2}; \alpha)\right]^2. \tag{16}$$

To obtain the halo M_{200}–c_{200} relation, we first split the haloes into bins equally spaced in logarithmic mass. We then fit an Einasto profile to each halo individually, removing all unrelaxed haloes according to the Neto et al. (2007) criteria. We then find the median halo concentration in each mass bin and estimate its uncertainty using bootstrap resampling.

Fig. 12 shows the (median) concentration–mass relations for COCO-COLD (dotted lines and shaded regions) and COCO-WARM (points with error bars) at redshifts $z = 0, 0.5, 1, 2, 3$ and 4 (different colours as indicated in the legend). These relations display the same qualitative behaviour seen in the density profiles in Fig. 11. For haloes with mass $M_{200} > 10^{11}\,h^{-1}\,\mathrm{M}_{\odot}$, the concentrations of CDM and WDM haloes agree well over all redshifts. For masses below this value,
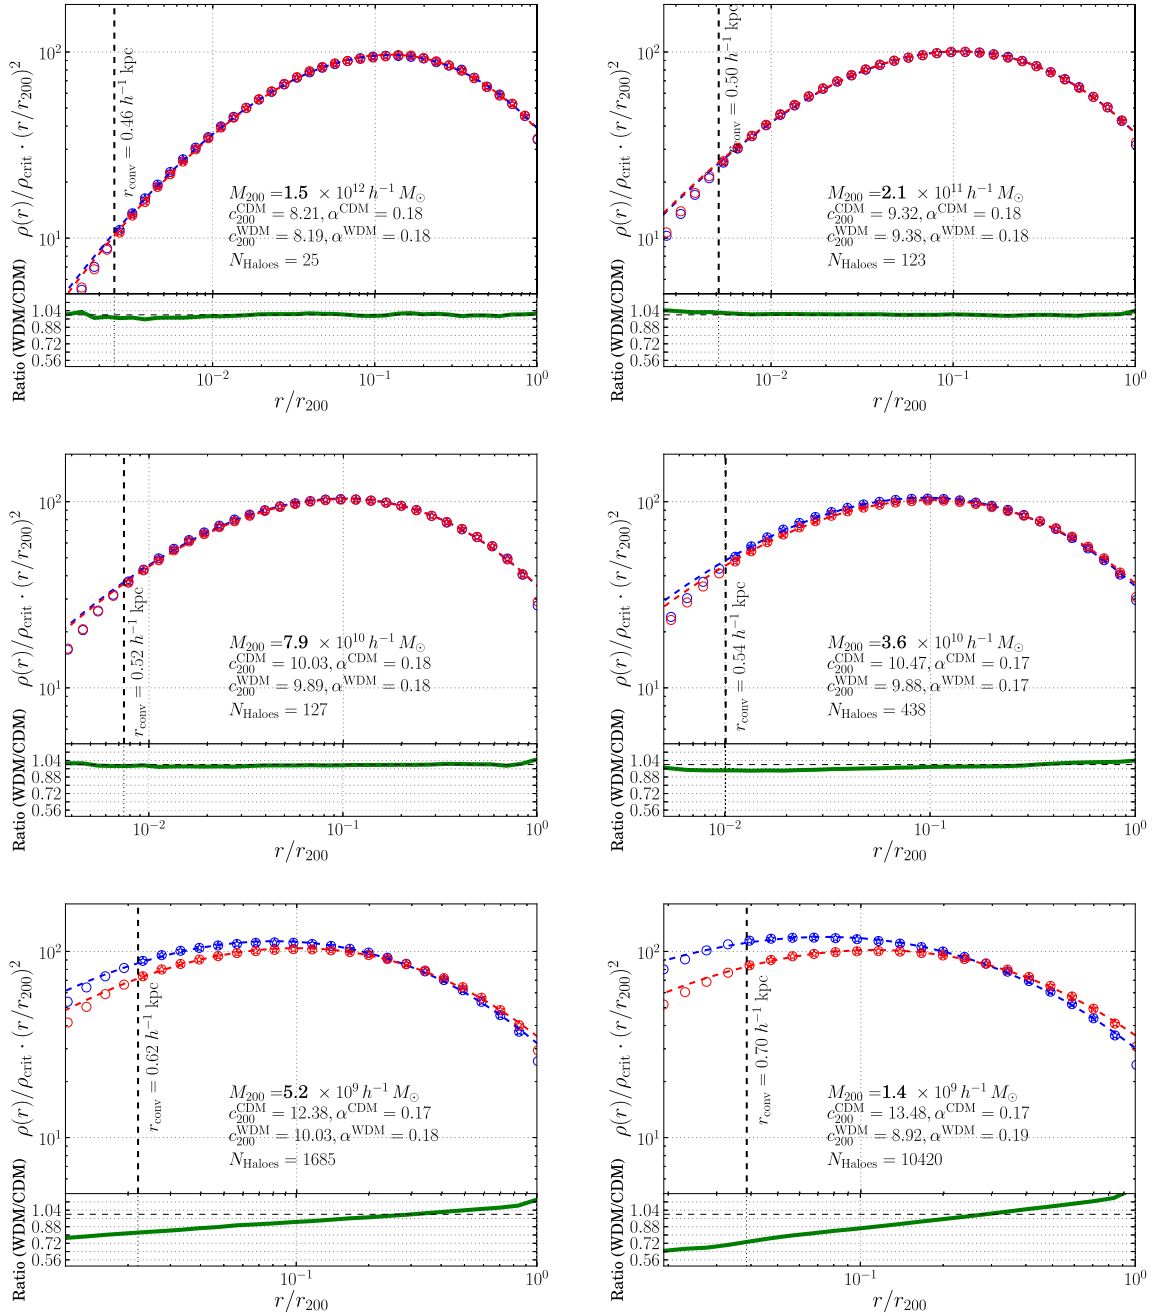
Figure 9. Same as Fig. 8, but with Einasto fits to the COCO-WARM and COCO-COLD density profiles.

WDM haloes are less concentrated than their CDM counterparts at all redshifts. This is a direct result of the later formation epoch of haloes of a given mass in WDM, and reflects the fact that the mass within r_{-2} in WDM haloes is assembled when the background density of the universe is lower than in the CDM case.

Whereas the CDM halo concentrations continue to increase as power laws towards lower masses, reflecting hierarchical growth, the WDM halo concentrations turn over at $M_{200} < 5 \times 10^{10} h^{-1} M_\odot$ and eventually begin to decrease (see also Schneider et al. 2012; Macciò et al. 2013). This echoes the finding in Fig. 11 that the mass in the central regions of WDM haloes begins to fall below that in the CDM case roughly below this mass. This mass is an order of magnitude larger than the mass scale at which the mass functions begin to differ ($\sim 2 \times 10^9 h^{-1} M_\odot$, see Figs 4 and 7). This result

is not entirely surprising: the concentration is sensitive to the inner parts of the profile and it is this inner mass (which we can roughly identify with the matter contained within r_{-2}) which is assembled later in WDM than in CDM, while most of the mass actually lies in the outer parts of the halo.

The lower panel of Fig. 12 shows the ratio of the concentrations in COCO-WARM and COCO-COLD, $c_{200}^{WDM}/c_{200}^{CDM}$. There are two interesting features of note: first, for all redshifts, the downturn in the WDM halo concentrations occurs at roughly the same halo mass, $M_{200} \sim 5 \times 10^{10} h^{-1} M_\odot$; and secondly, at fixed mass, the ratio decreases with decreasing redshift. The fact that the mass at which WDM halo concentrations begin to peel-off from the CDM relation is almost independent of redshift reflects the narrow redshift range in which the inner parts of WDM haloes form.
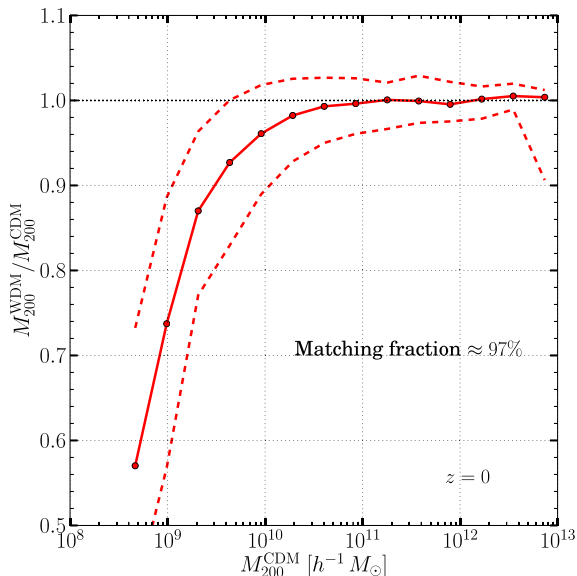
Figure 10. Ratio of halo mass (M_{200}) for all (relaxed and unrelaxed) matched haloes above $M_{200} > 10^8\,h^{-1}\,\mathrm{M}_\odot$ in COCO-WARM and COCO-COLD, as function of M_{200}^{CDM}. The solid red line shows the median relation in bins of M_{200}^{CDM}, whereas the dashed red lines indicate the 16th and 84th percentiles.

In COCO-WARM we also find that the evolution of the mass–concentration relation over redshift can be approximated using a simple functional form motivated by equation (8) (see Schneider et al. 2012), with an extra redshift-dependent component:

$$\frac{c_{200}^{\mathrm{WDM}}}{c_{200}^{\mathrm{CDM}}} = \left(1 + \gamma_1 \frac{M_{\mathrm{hm}}}{M_{200}}\right)^{-\gamma_2} \times (1+z)^{\beta(z)}. \tag{17}$$

Here, M_{hm} is the half-mode mass, z is the redshift of interest, $\gamma_1 = 60$, $\gamma_2 = 0.17$ and $\beta(z) = 0.026z - 0.04$. The predictions of our model are shown in the upper panel of Fig. 12 using the thin colour lines. While the model does not fully capture the relatively flat relationship at $z = 3$ and 4 in COCO-WARM, it generally reproduces the trends in the simulation and provides a good fit up to $z = 2$, over nearly five orders of magnitude in halo mass.

3.5 The shapes and spins of haloes

In this section, we examine the shapes and spins of WDM haloes. The shapes are most commonly quantified by the *triaxiality*, defined through the halo inertia tensor:

$$I_{ij} = m_p \sum_{n=1}^{N_{200}} x_{n,i} x_{n,j}, \tag{18}$$

where N_{200} is the number of particles within r_{200}, m_p is the mass of the simulation particle and $x_{n,i}$ is the ith coordinate of the nth particle relative to the halo centre. The eigenvalues of the inertia tensor define the axial lengths of an equivalent uniform density ellipsoid, $a \geq b \geq c$, which can be related to those of the halo itself (Bett et al. 2007). The sphericity is defined as c/a (as in Section 2.3): the higher its value, the less aspherical the ellipsoid's projection. The triaxiality is defined as $T = (a^2 - b^2)/(a^2 - c^2)$: large values correspond to prolate ellipsoids, small values to oblate ellipsoids.

The results for our simulations are shown in Fig. 13, where blue represents CDM and red WDM, with the top panel comparing the median triaxiality, and the lower the median sphericity. Errors on the median quantities were obtained by bootstrap resampling. Previous

N-body simulations of CDM haloes have shown that triaxiality correlates with halo mass, with triaxility decreasing with decreasing halo mass (Frenk et al. 1988; Allgood et al. 2006; Muñoz-Cuartas et al. 2011; Macciò et al. 2013). This trend reflects, in part, the younger dynamical age of more massive haloes. Fig. 13 shows that the same trend is present for WDM haloes but below $M_{200} \sim 10^{10}\,h^{-1}\,\mathrm{M}_\odot$, WDM haloes are slightly less triaxial than their CDM counterparts.

A more significant trend is revealed when comparing the spin of haloes in the two simulations. The spin is best characterized by the parameter, λ, defined as

$$\lambda = \frac{J\sqrt{|E|}}{GM^{5/2}} \tag{19}$$

(Peebles 1969), where J is the magnitude of the angular momentum of the halo, E is its total energy and M is the mass (which we take to be M_{200}). Haloes acquire a net angular momentum through tidal torques during growth in the linear regime which can be subsequently modified and rearranged by mergers (Peebles 1969; Doroshkevich 1970; White 1984). Since the merger histories are different for CDM and WDM haloes, we might expect some differences in their final angular momentum configurations. In particular, given that tidal forces associated with mergers tend to redistribute angular momentum from the central parts of haloes to the rest of the halo, the smaller frequency of mergers in WDM might facilitate the formation of extended spinning galactic discs (Frenk et al. 1988; Navarro & Benz 1991; Navarro & White 1994).[2]

The spin parameters in our two simulations are compared in the top panel of Fig. 14. Previous cosmological CDM simulations showed a very weak correlation between spin and halo mass, with a median value of $\lambda \approx 0.033$, across a wide range of halo masses (Davis et al. 1985; Barnes & Efstathiou 1987; Warren et al. 1992; Steinmetz & Bartelmann 1995; Cole & Lacey 1996; Mo, Mao & White 1998; Bett et al. 2007). Our COCO-COLD simulation reproduces this trend and extends it to lower masses, $M_{200} = 10^8\,h^{-1}\,\mathrm{M}_\odot$.

For $M_{200} > 5 \times 10^{10}\,h^{-1}\,\mathrm{M}_\odot$, the λ values for WDM haloes are almost identical to those of their CDM counterparts. However, for smaller halo masses λ decreases systematically with decreasing mass and is lower than the CDM value by almost 30 per cent at $M_{200} \sim 10^8\,h^{-1}\,\mathrm{M}_\odot$. This is consistent with the results of Bullock, Kravtsov & Colín (2002), who found that three out of four haloes below the WDM cutoff in their simulation had lower values of λ than the equivalent CDM matches. Note that in the top panel of Fig. 14 we include all haloes, not necessarily matches, which explains why in some of the largest mass bins, the median spins are not exactly the same in WDM and CDM. In addition, we only include haloes with more than 1000 particles within r_{200} since particle shot noise dominates the estimates of angular momentum for low particle numbers (Frenk et al. 1988; Bett et al. 2007, although we use a more conservative lower limit than the latter's choice of 300 particles).

To investigate why the spins of dwarf galaxy haloes are lower in WDM than in CDM we consider the relative contributions of energy, angular momentum and M_{200} to λ, illustrated in the bottom panel of Fig. 14, this time for bijectively matched haloes. The ratio of the median spin parameters is shown by the black squares and the ratio

---

[2] We note that the inability of many early simulations to form extended discs in the CDM model – the so-called angular momentum problem – is readily solved when appropriate prescriptions for supernovae feedback are included in the simulations (see e.g. Okamoto et al. 2005; Scannapieco et al. 2011).

Figure 11. Stacked cumulative mass profiles of relaxed, matched haloes in different mass bins for WDM (solid red lines) and CDM (blue squares). The lower panels show the ratio of the WDM mass to the CDM mass as a function of radius from the centre of the halo (in units of r_{200}). For haloes with $M_{200} > 10^{11} \, h^{-1} \, \mathrm{M_\odot}$, the mass profiles are nearly identical, but below $M_{200} \leq 5 \times 10^{10} \, h^{-1} \, \mathrm{M_\odot}$ they differ noticeably.

of the geometric means of the quantities that enter into equation (19) are shown by the other colour lines [magenta for $J_{\mathrm{CDM}}/J_{\mathrm{WDM}}$, cyan for $|E_{\mathrm{CDM}}/E_{\mathrm{WDM}}|^{1/2}$ and yellow for $(M_{200, \, \mathrm{CDM}}/M_{200, \, \mathrm{WDM}})^{5/2}]$. The combination of these ratios in equation (19) should reproduce the ratio of spin parameters, and this is shown in the thick green line. Part of the reason for lower WDM spins below $\sim 10^{10} \, h^{-1} \, \mathrm{M_\odot}$ is their slightly lower total energy which results from their lower concentration. The dominant factor, however, is their lower angular momentum relative to CDM haloes, ~ 25 per cent at $10^8 \, h^{-1} \, \mathrm{M_\odot}$. The cause of this could be related to the differing merger histories in WDM and CDM and the likely more quiescent mass accretion

of WDM haloes which can result in smaller spins (Bullock et al. 2002; Vitvitska et al. 2002; Hetznecker & Burkert 2006).

4 CONCLUSIONS

We have presented results from the COCO project, a set of cosmological 'zoom' simulations in which the dark matter is assumed to be either CDM (COCO-COLD) or a thermal 3.3 keV WDM particle (COCO-WARM). The combination of mass resolution and volume of our simulations provides a rich statistical sample of haloes over seven decades in mass. This WDM model is particularly interesting

Figure 12. The median concentration–mass relation and its redshift evolution for haloes in COCO-COLD and COCO-WARM. The colour dotted lines show the median relation over redshift for CDM haloes, as indicated in the legend. The shaded regions represent the errors in the median, as estimated by bootstrap resampling. The points with the error bars show the corresponding redshift relation in WDM. Only relaxed haloes are included. The thin colour lines show the results of the fitting formula introduced in equation (17).

Figure 13. Median halo triaxiality (top panel) and halo sphericity (lower panel) in COCO-WARM (red points) and COCO-COLD (blue lines). The errors on the median were obtained by bootstrapping 100 different samples in each case, and is represented by the red error bars for WDM and the blue shaded region for CDM. Only particles within r_{200} were used to compute these properties from the inertia tensor.

Figure 14. Top panel: the median halo spin–mass relation at $z = 0$ for COCO-WARM (red points) and COCO-COLD (blue line). Errors on the median for the WDM (shown by error bars) and for CDM (shown as the shaded region) haloes were calculated by bootstrap resampling. Bottom panel: the relative contributions of energy, angular momentum and halo mass to the spin of the halo. The black squares show the ratio (CDM to WDM) of the median spin parameters (from the top panel). The magenta, cyan and yellow lines show the ratios of the geometric means of the angular momentum, energy and M_{200}, respectively, which when multiplied together appropriately yield the thick green line, which show the ratio of the geometric means of λ_{CDM} and λ_{WDM}. As expected, the squares trace out the ratio of the geometric means. Note that ratios of all quantities are taken between the bijectively matched COCO-WARM and COCO-COLD haloes.

because it corresponds to the 'warmest' particle allowed by current Lyman α constraints (Viel et al. 2013) and has a linear power spectrum cutoff similar to that for the 'coldest' 7 keV sterile neutrino, evidence for which has recently been claimed to be found in galaxies and clusters (Boyarsky et al. 2014; Bulbul et al. 2014b). This cutoff – manifest in haloes of $M_{200} \leq 2 \times 10^9 \, h^{-1} \, \mathrm{M_\odot}$ for our assumed particle mass – is reflected both in the population statistics and the structure of individual haloes.

The formation of structure begins significantly later in COCO-WARM than in COCO-COLD. Across all redshifts, differences in the halo mass function between COCO-WARM and COCO-COLD begin to appear at a mass roughly one order of magnitude larger than the nominal half-mode mass. Below $\sim 2 \times 10^9 \, h^{-1} \, \mathrm{M_\odot}$, the WDM mass function declines rapidly but there are still some small haloes present at surprisingly large redshifts: at $z = 10$, for example, there are almost five times as many haloes with $M_{200} \sim 10^8 \, h^{-1} \, \mathrm{M_\odot}$ in COCO-COLD than in COCO-WARM. We find that the $z = 0$ halo mass functions in both COCO-WARM and COCO-COLD are well described by previous analytic fits to the CDM halo mass function (e.g. Sheth & Tormen 1999) down to our resolution limit, $M_{200} \sim 3 \times 10^7 \, h^{-1} \, \mathrm{M_\odot}$, provided that the window function used to compute the mass variance, $\sigma^2(M)$, in the WDM case is calculated using a sharp k-space filter, as described by Benson et al. (2013).

Just as for COCO-COLD, the spherically averaged density profiles of haloes in COCO-WARM, down to dwarf galaxy scales, are well described by NFW or Einasto profiles. The concentration–mass relation, M_{200}–c_{200} (where we have defined concentration using the

Einasto profile), in COCO-WARM begins to peel-off from the corresponding relation in COCO-COLD at a mass of $\sim 5 \times 10^{10} \, h^{-1} \, \mathrm{M}_\odot$, reflecting the later formation epoch of haloes of a given mass in WDM compared to CDM. This mass is larger than the scale below which the WDM mass function is suppressed because halo concentration is determined by the epoch when the *inner* regions of a halo form. The mass at which the concentration begins to differ in the two simulations is almost independent of redshift out to $z \simeq 4$. At the present day, the typical concentration of a halo of mass $10^9 \, h^{-1} \, \mathrm{M}_\odot$ in COCO-WARM is $c_{200} \simeq 8$ compared to $c_{200} \simeq 12.7$ in COCO-COLD. The trends and evolution of the concentration–mass relation can be approximated by the fitting formula provided in equation (17).

The generally triaxial shapes of haloes in COCO-WARM and COCO-COLD are very similar. However, we find that, for masses below $\sim 5 \times 10^{10} \, h^{-1} \, \mathrm{M}_\odot$, WDM haloes have slightly lower values of the spin parameter, λ, (up to 30 per cent) than their CDM counterparts.

In principle, gravitational lensing is one of the most promising techniques for distinguishing between WDM and CDM, as it directly probes the halo mass function (see for example Vegetti & Koopmans 2009). In the parent volumes of the COCO simulations, the non-linear power spectrum, $P(k)$, for COLOR-WARM is suppressed by ~ 3 per cent relative to COLOR-COLD on scales $k \leq 5h \, \mathrm{Mpc}^{-1}$ (consistent with the simulations of Viel et al. 2012, which bracket the 3.3 keV model). While the weak lensing signal on these scales should be measurable by surveys such as Dark Energy Spectroscopic Instrument (DESI) and Euclid, this difference is smaller than the differences introduced by baryon effects on the dark matter-only $P(k)$, which is of the order of 5–10 per cent, as seen in hydrodynamic simulations (van Daalen et al. 2014; van Daalen & Schaye 2015, Hellwing et al., in preparation). It is therefore necessary to use hydrodynamic simulations to check for any residual signal of the nature of the dark matter species, both in the power spectrum and in other observable properties of the galaxy population.

REFERENCES

Abazajian K. N., 2014, Phys. Rev. Lett., 112, 161303
Allgood B., Flores R. A., Primack J. R., Kravtsov A. V., Wechsler R. H., Faltenbacher A., Bullock J. S., 2006, MNRAS, 367, 1781
Anderson M. E., Churazov E., Bregman J. N., 2015, MNRAS, 452, 3905
Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, MNRAS, 426, 2046
Angulo R. E., Hahn O., Abel T., 2013, MNRAS, 434, 3337
Asaka T., Shaposhnikov M., 2005, Phys. Lett. B, 620, 17
Barkana R., Haiman Z., Ostriker J. P., 2001, ApJ, 558, 482
Barnes J., Efstathiou G., 1987, ApJ, 319, 575
Benson A. J. et al., 2013, MNRAS, 428, 1774
Bett P., Eke V., Frenk C. S., Jenkins A., Helly J., Navarro J., 2007, MNRAS, 376, 215
Bode P., Ostriker J. P., Turok N., 2001, ApJ, 556, 93
Bond J. R., Szalay A. S., 1983, ApJ, 274, 443
Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, ApJ, 379, 440
Boyarsky A., Ruchayskiy O., Shaposhnikov M., 2009, Annu. Rev. Nucl. Part. Sci., 59, 191
Boyarsky A., Franse J., Iakubovskyi D., Ruchayskiy O., 2015, Phys. Rev. Lett., 115, 161301
Boyarsky A., Ruchayskiy O., Iakubovskyi D., Franse J., 2014, Phys. Rev. Lett., 113, 251301
Bryan G. L., Norman M. L., 1998, ApJ, 495, 80
Bulbul E., Markevitch M., Foster A. R., Smith R. K., Loewenstein M., Randall S. W., 2014a, preprint (arXiv:1409.4143)
Bulbul E., Markevitch M., Foster A., Smith R. K., Loewenstein M., Randall S. W., 2014b, ApJ, 789, 13
Bullock J. S., Kravtsov a. A. V., Colín P., 2002, ApJ, 564, L1
Cole S., Lacey C., 1996, MNRAS, 281, 716
Cole S. et al., 2005, MNRAS, 362, 505
Colín P., Valenzuela O., Avila-Reese V., 2008, ApJ, 673, 203
Colless M. et al., 2001, MNRAS, 328, 1039
Conlon J. P., Day F. V., 2014, J. Cosmol. Astropart. Phys., 11, 33
Crain R. A. et al., 2009, MNRAS, 399, 1773
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
Dodelson S., Widrow L. M., 1994, Phys. Rev. Lett., 72, 17
Doroshkevich A. G., 1970, Astrophysics, 6, 320
Einasto J., 1965, Tr. Astrofizicheskogo Inst. Alma-Ata, 5, 87
Eisenstein D. J. et al., 2005, ApJ, 633, 560
Eke V. R., Cole S., Frenk C. S., 1996, MNRAS, 282, 263
Frenk C. S., White S. D. M., 2012, Ann. Phys., Lpz., 524, 507
Frenk C. S., White S. D. M., Efstathiou G., Davis M., 1985, Nature, 317, 595
Frenk C. S., White S. D. M., Davis M., Efstathiou G., 1988, ApJ, 327, 507
Frenk C. S., Evrard A. E., White S. D. M., Summers F. J., 1996, ApJ, 472, 460
Gao L., Navarro J. F., Cole S., Frenk C. S., White S. D. M., Springel V., Jenkins A., Neto A. F., 2008, MNRAS, 387, 536
Harker G., Cole S., Helly J., Frenk C., Jenkins A., 2006, MNRAS, 367, 1039
Hawkins E. et al., 2003, MNRAS, 346, 78
Hellwing W. A., Frenk C. S., Cautun M., Bose S., Helly J., Jenkins A., Sawala T., Cytowski M., 2015, preprint (arXiv:1505.06436)
Hetznecker H., Burkert A., 2006, MNRAS, 370, 1905
Hooper D., Goodenough L., 2011, Phys. Lett. B, 697, 412
Jeltema T. E., Profumo S., 2015, MNRAS, 450, 2143
Jenkins A., 2010, MNRAS, 403, 1859
Jenkins A., 2013, MNRAS, 434, 2094
Jenkins A., Frenk C. S., White S. D. M., Colberg J. M., Cole S., Evrard A. E., Couchman H. M. P., Yoshida N., 2001, MNRAS, 321, 372
Katz N., White S. D. M., 1993, ApJ, 412, 455
Komatsu E. et al., 2011, ApJS, 192, 18

Lacey C., Cole S., 1993, MNRAS, 262, 627

Laine M., Shaposhnikov M., 2008, J. Cosmol. Astropart. Phys., 6, 31

Lovell M. R. et al., 2012, MNRAS, 420, 2318

Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, MNRAS, 439, 300

Ludlow A. D. et al., 2013, MNRAS, 432, 1103

Macciò A. V., Paduroiu S., Anderhalden D., Schneider A., Moore B., 2012, MNRAS, 424, 1105

Macciò A. V., Ruchayskiy O., Boyarsky A., Muñoz-Cuartas J. C., 2013, MNRAS, 428, 882

Malyshev D., Neronov A., Eckert D., 2014, Phys. Rev. D, 90, 103506

Menci N., Fiore F., Lamastra A., 2012, MNRAS, 421, 2384

Mo H. J., Mao S., White S. D. M., 1998, MNRAS, 295, 319

Muñoz-Cuartas J. C., Macciò A. V., Gottlöber S., Dutton A. A., 2011, MNRAS, 411, 584

Navarro J. F., Benz W., 1991, ApJ, 380, 320

Navarro J. F., White S. D. M., 1994, MNRAS, 267, 401

Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563

Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493

Navarro J. F. et al., 2004, MNRAS, 349, 1039

Neto A. F. et al., 2007, MNRAS, 381, 1450

Okamoto T., Eke V. R., Frenk C. S., Jenkins A., 2005, MNRAS, 363, 1299

Peebles P. J. E., 1969, ApJ, 155, 393

Planck Collaboration I, 2014, A&A, 571, A1

Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, MNRAS, 338, 14

Press W. H., Schechter P., 1974, ApJ, 187, 425

Reed D. S., Bower R., Frenk C. S., Jenkins A., Theuns T., 2007, MNRAS, 374, 2

Riemer-Sorensen S., 2014, preprint (arXiv:1405.7943)

Sawala T., Frenk C. S., Crain R. A., Jenkins A., Schaye J., Theuns T., Zavala J., 2013, MNRAS, 431, 1366

Sawala T. et al., 2014, preprint (arXiv:1406.6362)

Sawala T. et al., 2015, MNRAS, 448, 2941

Scannapieco C., White S. D. M., Springel V., Tissera P. B., 2011, MNRAS, 417, 154

Schaller M. et al., 2015, MNRAS, 451, 1247

Schneider A., Smith R. E., Macciò A. V., Moore B., 2012, MNRAS, 424, 684

Schneider A., Smith R. E., Reed D., 2013, MNRAS, 433, 1573

Shao S., Gao L., Theuns T., Frenk C. S., 2013, MNRAS, 430, 2346

Sheth R. K., Tormen G., 1999, MNRAS, 308, 119

Shi X., Fuller G. M., 1999, Phys. Rev. Lett., 82, 2832

Springel V., 2005, MNRAS, 364, 1105

Springel V., Yoshida N., White S. D. M., 2001a, Nature, 6, 79

Springel V., White S. D. M., Tormen G., Kauffmann G., 2001b, MNRAS, 328, 726

Steinmetz M., Bartelmann M., 1995, MNRAS, 272, 570

Tegmark M. et al., 2004, Phys. Rev. D, 69, 103501

Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, ApJ, 688, 709

Urban O., Werner N., Allen S. W., Simionescu A., Kaastra J. S., Strigari L. E., 2015, MNRAS, 451, 2447

van Daalen M. P., Schaye J., 2015, MNRAS, 452, 2247

van Daalen M. P., Schaye J., McCarthy I. G., Booth C. M., Dalla Vecchia C., 2014, MNRAS, 440, 2997

Vegetti S., Koopmans L. V. E., 2009, MNRAS, 400, 1583

Viel M., Lesgourgues J., Haehnelt M. G., Matarrese S., Riotto A., 2005, Phys. Rev. D, 71, 063534

Viel M., Marković K., Baldi M., Weller J., 2012, MNRAS, 421, 50

Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, Phys. Rev. D, 88, 043502

Vitvitska M., Klypin A. A., Kravtsov A. V., Wechsler R. H., Primack J. R., Bullock J. S., 2002, ApJ, 581, 799

Wang J., White S. D. M., 2007, MNRAS, 380, 93

Wang J. et al., 2011, MNRAS, 413, 1373

Warren M. S., Quinn P. J., Salmon J. K., Zurek W. H., 1992, ApJ, 399, 405

Wechsler R. H., Bullock J. S., Primack J. R., Kravtsov A. V., Dekel A., 2002, ApJ, 568, 52

White S. D. M., 1984, ApJ, 286, 38

Zehavi I. et al., 2002, ApJ, 571, 172

This paper has been typeset from a TEX/LATEX file prepared by the author.