

A warm dark matter cosmogony may yield more low-mass galaxy detections in 21-cm surveys than a cold dark matter one

Kyle A. Oman ^{1,2,3★}, Carlos S. Frenk^{1,3}, Robert A. Crain ⁴, Mark R. Lovell ^{1,3,5} and Joel Pfeffer ⁶

¹*Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK*

²*Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, UK*

³*Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*

⁴*Astrophysics Research Institute, Liverpool John Moores University, Liverpool L3 5RF, UK*

⁵*Centre for Astrophysics and Cosmology, Science Institute, University of Iceland, Dunhaga 5, 107 Reykjavík, Iceland*

⁶*Centre for Astrophysics & Supercomputing, Swinburne University, Hawthorn, VIC 3122, Australia*

Accepted 2024 July 19. Received 2024 June 7; in original form 2024 January 10

ABSTRACT

The 21-cm spectral line widths, w_{50} , of galaxies are an approximate tracer of their dynamical masses, such that the dark matter halo mass function is imprinted in the number density of galaxies as a function of w_{50} . Correcting observed number counts for survey incompleteness at the level of accuracy needed to place competitive constraints on warm dark matter (WDM) cosmological models is very challenging, but forward-modelling the results of cosmological hydrodynamical galaxy formation simulations into observational data space is more straightforward. We take this approach to make predictions for an ALFALFA-like survey from simulations using the EAGLE galaxy formation model in both cold (CDM) and WDM cosmogonies. We find that for WDM cosmogonies more galaxies are detected at the low- w_{50} end of the 21-cm velocity width function than in the CDM cosmogony, contrary to what might naïvely be expected from the suppression of power on small scales in such models. This is because low-mass galaxies form later and retain more gas in WDM cosmogonies (with EAGLE). While some shortcomings in the treatment of cold gas in the EAGLE model preclude placing definitive constraints on WDM scenarios, our analysis illustrates that near-future simulations with more accurate modelling of cold gas will likely make strong constraints possible, especially in conjunction with new 21-cm surveys such as WALLABY.

Key words: galaxies: abundances – galaxies: luminosity function, mass function – dark matter – radio lines: galaxies.

1 INTRODUCTION

The number density of dark haloes as a function of their mass, or of their (closely related) maximum circular velocity (v_{\max}), is a fundamental prediction of cosmological models with dark matter. The Λ CDM cosmogony predicts an approximately power-law subhalo v_{\max} function with a slope of about -3 (e.g. Klypin et al. 2015; Hellwing et al. 2016). Testing this prediction is particularly interesting at the low-velocity end ($v_{\max} \lesssim 80 \text{ km s}^{-1}$) where warm dark matter (WDM; Bond & Szalay 1983; Bardeen et al. 1986; Bode, Ostriker & Turok 2001) cosmogonies predict a flattening (followed by a cut-off at even lower v_{\max}) due to a truncation of the matter power spectrum leading to suppressed and delayed structure formation on small scales.

The subhalo v_{\max} function is not directly observable, so we must turn to observable proxies. The most direct approach is to choose a measurement that can be interpreted as a dynamical mass tracer. One such measurement is the width of the 21-cm emission line of neutral hydrogen (HI). Under optimistic assumptions, the spectral line width (for example the full width at half-maximum, w_{50}) is

related to v_{\max} by a simple geometric correction: $w_{50} \approx 2v_{\max} \sin i$ where i is the inclination angle. The first measurements of the HI velocity width function (hereafter, HI WF, Zavala et al. 2009; Zwaan, Meyer & Staveley-Smith 2010; Papastergis et al. 2011, see also Trujillo-Gomez et al. 2011; Klypin et al. 2015) found a flattening of the slope at the low-width end that could be interpreted as a consequence of WDM (Zavala et al. 2009; Papastergis et al. 2011; Schneider et al. 2017, but see Klypin et al. 2015), but the current consensus is that the flattening is mostly due to the fact that the HI WF is preferentially a tracer of the subhalo v_{\max} function of gas-rich galaxies (Obreschkow et al. 2013; Sawala et al. 2013; Brooks et al. 2017; Dutton, Obreja & Macciò 2019), rather than that of the global galaxy (or dark matter halo) population. However, as we will see below, provided that the bias towards gas-rich galaxies can be modelled in sufficient detail, the measured HI WF remains sensitive to a cutoff in the matter power spectrum like that due to a thermal relic with a mass of about 10 keV or less.

The conventional approach (e.g. Zavala et al. 2009; Zwaan et al. 2010; Papastergis et al. 2015) to constrain the v_{\max} function (often termed the ‘velocity function’) observationally is to assign v_{\max} values to surveyed galaxies based on their HI line widths, such as by assuming $w_{50} = 2v_{\max} \sin i$ or a more sophisticated but conceptually similar relationship. Once this is done the number density of surveyed

* E-mail: kyle.a.oman@durham.ac.uk

galaxies is corrected for the limited sensitivity of the survey to estimate the intrinsic number density of galaxies as a function of v_{\max} . Brooks et al. (2017) and Dutton et al. (2019, see also Brook & Shankar 2016; Sardone et al. 2024) have discussed some of the challenges of this approach, including: that the limited extent of HI discs in low-mass galaxies often fails to reach the radii where the maximum of the circular velocity curve is reached; that many low-mass galaxies have HI reservoirs too faint to be detected by current surveys; that HI is increasingly pressure-supported (rather than rotation-supported) in lower-mass galaxies. We explore some further challenges in this work.

Rather than attempting to transform the HI WF into an inferred measurement of a v_{\max} function suitable for comparison with ‘raw’ theoretical predictions, theoretical results can be recast into direct predictions for the HI WF for comparison with ‘raw’ observations, as advocated by Obreschkow et al. (2013) and Oman (2022). One way to approach this is to predict line widths from v_{\max} using essentially the reverse of the procedure described above: a variation on the theme of $w_{50} = 2v_{\max} \sin i$ (e.g. Obreschkow et al. 2009; Papastergis et al. 2011; Klypin et al. 2015; Chauhan et al. 2019; Brooks, Oman & Frenk 2023). This makes some of the challenges outlined above easier to deal with – for example, the extent of the HI disc can be folded into the construction of a spectrum from a circular velocity curve. However, such efforts often result in a population of model galaxies with gas kinematics and therefore spectra that are unrealistically regular and symmetric.

A promising alternative is to measure the HI spectrum of galaxies in cosmological hydrodynamical simulations as this should more accurately capture the diverse morphological and kinematic irregularities that shape the spectra of real galaxies – at least as far as these are accurately captured by the simulation in question. Some initial efforts along these lines have been made by Macciò et al. (2016), Brooks et al. (2017), and Dutton et al. (2019) and show that such a simulated survey can reproduce the observed HI WF quite accurately – and certainly more accurately than predicting line widths from v_{\max} using an analytic prescription for the same samples of simulated galaxies.

In this work, we develop this methodology further by carrying out the first direct survey of a hydrodynamical simulation volume equivalent to that of the ALFALFA 21-cm survey (with the caveat that some tiling of periodic simulation volumes was required). We use a suite of CDM and WDM simulations to show that the effects of WDM can sometimes be counter-intuitive in this context, in the sense that they can lead to increased source counts for low- w_{50} galaxies where we would naively expect fewer due to the suppression of the formation of low-mass haloes in WDM cosmologies, and explain the reasons for this. We discuss the prospects of using the HI WF as a constraint on WDM cosmologies – these are good, provided that a sufficiently faithful model of the galaxy population contributing to the measurement is realized in the simulations.

This paper is structured as follows. In Section 2, we describe the simulation data sets and survey catalogues used in our analysis. In Section 3, we describe our mock survey methodology. Our mock-observed HI WFs in CDM and WDM cosmologies are presented and analysed in Section 4. We discuss advantages of mock surveys based on hydrodynamical simulations over methodologies requiring assumptions loosely of the form $w_{50} \approx v_{\max} \sin i$ in Section 5. Finally, we summarize our findings and likely future advances as new simulations and 21-cm surveys become available in Section 6.

2 DATA

2.1 Simulations – EAGLE and variations

We make use of four cosmological hydrodynamical galaxy formation simulations. The simulations were carried out with the EAGLE galaxy formation model, which uses the pressure–entropy formulation of smoothed-particle hydrodynamics (Hopkins 2013) with improved time-stepping criteria (Durier & Dalla Vecchia 2012) and switches for artificial viscosity (Cullen & Dehnen 2010) and artificial conduction (Price 2008). It implements radiative cooling (Wiersma, Schaye & Smith 2009a), star formation (Schaye 2004; Schaye & Dalla Vecchia 2008), stellar evolution and chemical enrichment (Wiersma et al. 2009b), black hole growth (Springel 2005; Rosas-Guevara et al. 2015), energetic stellar (Dalla Vecchia & Schaye 2012) and black hole (Booth & Schaye 2009) feedback, and cosmic reionization (Haardt & Madau 2001; Wiersma et al. 2009b).

All four simulations were performed with mass resolution $8\times$ better than the flagship Ref-L100N1504 EAGLE simulation, yielding a baryon particle mass of $m_g = 2.26 \times 10^5 M_\odot$, and hence adopt the ‘Recal’ model (see Schaye et al. 2015, for further discussion). The first, which we consider to be our fiducial point of reference, is a $(25 \text{ Mpc})^3$ volume from the EAGLE suite run with a CDM cosmogony (that advocated by the Planck Collaboration XVI 2014), with identifier Recal-L0025N0752. We will refer to this as the ‘EAGLE25 cold dark matter’ (or ‘EAGLE25-CDM’) simulation. The second simulation that we use, introduced by Bastian et al. (2020) and also examined in the recent study of Mason et al. (2023), is a $(25 h^{-1} \text{ Mpc})^3 \approx (34 \text{ Mpc})^3$ volume followed at the same resolution, adopting the same cosmogony and the same Recal parameters for the EAGLE galaxy formation model. (This simulation also includes the E-MOSAICS globular cluster formation and evolution model of Pfeffer et al. 2018; Kruijssen et al. 2019, but this is immaterial as it does not influence any of the galaxy properties.) We label this simulation ‘EAGLE34 cold dark matter’ (or ‘EAGLE34-CDM’). The matched resolutions, cosmogonies, and galaxy formation models of these simulations offer a useful opportunity to control for cosmic variance effects using independent volumes.

The third and fourth simulations that we use are WDM counterparts to the first two. In each case, exactly the same galaxy formation model as in the CDM simulation is used, and the initial conditions are also created to have the same phases (Jenkins 2013), but with their power spectra modified to model the effect of WDM. The counterpart of the EAGLE25-CDM simulation uses the transfer function approximation of Bode et al. (2001) to represent the effect of a 1.5 keV thermal relic dark matter candidate. We choose this model as a relatively extreme case (e.g. it is incompatible with structure formation constraints: Enzi et al. 2021; Nadler et al. 2021) that is expected to emphasize any differences with respect to the CDM simulations. The model’s half-mode wavenumber is $k_{\text{hm}} = 9.88 \text{ Mpc}^{-1}$, and its half-mode mass is $M_{\text{hm}} = 5.3 \times 10^9 M_\odot$ (Bose et al. 2016). We label this ‘EAGLE25 1.5 keV thermal relic’ or simply ‘EAGLE25-WDM-1.5 keV’.

The counterpart of the EAGLE34-CDM simulation has a power spectrum modified to mimic the effect of a 7.1 keV sterile neutrino dark matter candidate with a mixing angle $\sin(2\theta) = 2 \times 10^{-11}$ (Meshveliani et al. 2024). This is the ‘warmest’ sterile neutrino model that is compatible with the X-ray emission line observed in galaxies and clusters and tentatively associated with dark matter decay radiation (Boyersky et al. 2014; Bulbul et al. 2014). The sterile neutrino distribution function was computed using the implementations of Laine & Shaposhnikov (2008) and Lovell et al. (2016), and the matter




Figure 1. Galaxy stellar mass functions of simulations used in this work, for all galaxies with $M_* > 10^7 M_\odot$ (approximately 50 stellar particles). Indicative counting uncertainties are illustrated with the shaded band for the EAGLE25-CDM simulation (these are slightly smaller in the larger EAGLE34 simulation volumes). The stellar mass functions agree closely, except at low-stellar masses ($M_* \lesssim 10^8 M_\odot$) where galaxy formation is suppressed in WDM cosmogonies.

power spectrum was calculated with a modified version of the CAMB Boltzmann-solver (Lewis, Challinor & Lasenby 2000). We label this simulation ‘EAGLE34 7 keV sterile neutrino $\sin(2\theta) = 2 \times 10^{-11}$ ’, or simply ‘EAGLE34-WDM-7 keV’.

Halo is identified in all cases with a friends-of-friends algorithm with a linking length equal to 0.2 times the mean inter-particle spacing (Davis et al. 1985), with gravitationally bound subhaloes subsequently identified using the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009). In WDM cosmogonies, the truncation of the power spectrum causes artificial fragmentation of filaments in N -body simulations that could be spuriously identified as subhaloes. We remove these from halo catalogues following the approach of Lovell et al. (2014). Such spurious subhaloes only outnumber ‘real’ subhaloes at maximum circular velocities $v_{\max} \lesssim 15 \text{ km s}^{-1}$. In this regime, subhaloes containing atomic hydrogen – which are our main focus in this work – are extremely rare.

As a check that the ‘Recal’ calibration of the EAGLE model remains applicable in the WDM cosmogonies that we consider, we show the galaxy stellar mass function in the 4 simulations listed above in Fig. 1. Outside of the regime where warm dark matter suppresses the formation of structure ($M_* \lesssim 10^8 M_\odot$ for the relevant cosmogonies), the stellar mass functions agree very closely.

The EAGLE model does not directly compute the HI fractions of simulation particles. Following Bahé et al. (2016, see also Crain et al. 2017), we partition neutral and ionized hydrogen following the prescription for self-shielding from the meta-galactic ionizing background radiation of Rahmati et al. (2013), then subdivide the neutral hydrogen into atomic (HI) and molecular (H_2) constituents using the empirically calibrated pressure-dependent correction defined by Blitz & Rosolowsky (2006).

2.2 Observations – The ALFALFA survey

The ALFALFA survey (Giovanelli et al. 2005) mapped about 7000 deg^2 at 21-cm wavelengths at declinations $0 \leq \delta/\text{deg} \leq 36$ (avoiding the Galactic plane) using the Arecibo radio telescope. We use the α -100 (full-survey; Haynes et al. 2018, see also Haynes

et al. 2011) release¹ of the ALFALFA extragalactic source catalogue. We apply quality cuts as are conventional in studies of the HIMF and HI WF (e.g. Papastergis et al. 2015; Jones et al. 2018; Oman 2022): we retain ‘Code 1’ ($S/N > 6.5$) sources with HI masses $10^6 \leq M_{\text{HI}}/M_\odot \leq 10^{11}$, with line widths² $1.2 \leq \log_{10}(w_{50}/\text{km s}^{-1}) < 2.95$, with recession velocities $v_{\text{CMB}} \leq 15000 \text{ km s}^{-1}$ and distances $D < 200 \text{ Mpc}$, that fall within the nominal survey footprint (defined in tables D1 & D2 of Jones et al. 2018). We use the updated distances of sources AGC 749 235 and AGC 220 210 reported by Oman (2022), and the TRGB distances for galaxies overlapping with the SHIELD sample reported by McQuinn et al. (2021). These cuts result in a final sample of 21 388 galaxies.

3 METHODS

3.1 Mock ALFALFA-like surveys

Whether a given galaxy is detected in the ALFALFA survey depends on its HI flux, S_{21} , and line width, w_{50} : a higher flux is of course easier to detect, and a narrower line width squeezes a given flux into a smaller number of channels in the detector, increasing the S/N. We therefore need to determine a flux and line width for each galaxy in a mock survey.

We begin by defining a mock survey volume with the same volume as our selection from the ALFALFA survey (see Section 2.2). Since our simulations are of random cosmological volumes, we simplify the geometry of the survey to cover a circular aperture on the sky out to a maximum distance of 200 Mpc. The ALFALFA survey footprint (Jones et al. 2018, tables D1 & D2) subtends 6517.7 deg^2 , corresponding to an opening angle of 46.8 deg from the axis for our circular survey footprint. Since the survey volume is 339 or 134 times larger than the $(25 \text{ Mpc})^3$ or $(34 \text{ Mpc})^3$ simulation volumes, we replicate the periodic simulation volumes such that the survey volume is completely covered. We note that given the relative volumes and the opening angle, there is no advantage in terms of minimizing aliasing effects of choosing any particular survey axis orientation with respect to the simulation volume Cartesian axes. Given the potentially large number of copies of a given simulated galaxy that could appear in our mock surveys, we have carefully checked that aliasing effects are not unduly driving or influencing any of our conclusions. The main effect of tiling a limited volume is that the rarest (high-HI mass, high w_{50}) galaxies that appear in the real survey are not sampled in the mocks. Since our focus is on the low-mass and low-velocity width regimes, this is of no concern. This focus also means that the galaxies that we are most interested in appear copied only up to about 10 times in our mock catalogues, because they are faint and therefore undetected at distances beyond about 60 Mpc. Finally, we note that no galaxies appear multiple times as exact copies since each instance of a galaxy has a different distance and is seen from a different viewing angle (and therefore has different S_{21} and w_{50} , as explained below) than its aliases.

We next catalogue every subhalo in the mock survey volume that has $M_{\text{HI}} > 10^6 M_\odot$ and $M_* > 0$, and measure S_{21} and w_{50} for each. The flux is simply measured as:

$$\frac{S_{21}}{\text{Jy km s}^{-1}} = (2.36 \times 10^5)^{-1} \frac{M_{\text{HI}}}{M_\odot} \left(\frac{D}{\text{Mpc}} \right)^{-2}, \quad (1)$$

¹<http://egg.astro.cornell.edu/alfalfa/data/>

²Throughout this work, we denote the natural logarithm $\log(\cdot)$, and the base-10 logarithm $\log_{10}(\cdot)$.




Figure 2. Sample 21-cm spectra from the mock survey of the EAGLE25-CDM simulation. *First panel:* a massive ($v_{\max} \approx 350 \text{ km s}^{-1}$) galaxy. The heavy line shows the spectrum of the nearest instance of the galaxy in the survey. Because the periodic simulation volume is replicated to cover the survey volume, the same galaxy appears more than once, but is seen at different distances and from different viewing angles each time – spectra of other realizations of this galaxy are shown with thinner lines, illustrating that each instance has a different total flux and spectral line width. *Second panel:* a spectrum of a lower-mass galaxy ($v_{\max} \approx 135 \text{ km s}^{-1}$) with a clear ‘double-horn’ shape. The peak flux densities on the approaching and receding sides, $S_{21,\text{peak}}^{\text{left}}$ and $S_{21,\text{peak}}^{\text{right}}$, and the locations where half of their amplitudes are crossed, are marked with dotted lines. The full-width at half-maximum w_{50} is marked with the thin solid line. *Third panel:* a spectrum of a yet lower mass galaxy ($v_{\max} \approx 70 \text{ km s}^{-1}$) with an approximately flat-topped spectral shape. *Fourth panel:* a spectrum of an even lower mass galaxy ($v_{\max} \approx 35 \text{ km s}^{-1}$) with an approximately Gaussian spectral shape.

where M_{HI} is the sum of the H I masses of gas particles identified by SUBFIND as bound to the subhalo and D is the distance from the mock observer.

The calculation of the H I line width is more involved. We first create a mock spectrum of each galaxy using the mock-observing tool MARTINI³ (Oman 2019, 2024; Oman et al. 2019). Each gas particle contributes a Gaussian spectral line profile to the spectrum with a centre determined by its velocity along the line of sight and a width:

$$\sigma = \sqrt{\frac{k_{\text{B}}T}{m_{\text{p}}}}, \quad (2)$$

where k_{B} is Boltzmann’s constant, T is the particle temperature and m_{p} is the proton mass. The typical gas temperature is $T \approx 8 \times 10^3 \text{ K}$, corresponding to $\sigma \approx 8 \text{ km s}^{-1}$. The initial spectral channels have a 4 km s^{-1} width, and we use 512 channels such that every spectrum is sampled in its entirety. As in ALFALFA, we then Hanning smooth the spectra for a final spectral resolution⁴ of 8 km s^{-1} . Since ALFALFA is a low-spatial resolution drift-scan survey with nearly uniform sensitivity across the survey area we do not take the morphology of the sources into account (i.e. we produce only a spectrum, not a data cube), nor do we model the primary beam. Since we process each subhalo separately, we implicitly neglect any source confusion. We also do not add any simulated noise to the measurements since the selection criteria include $S/N > 6.5$ – for our purposes, noise has a negligible influence on full-width at half maximum line widths in this regime. Fig. 2 shows some example spectra of galaxies of different masses. The first panel illustrates that replicas of a given simulated galaxy generally have significantly different spectra, both in terms of total flux and shape.

With spectra of all galaxies in the survey volume in hand, we measure line widths as the full width at half maximum of

each spectrum (w_{50}), loosely mimicking the approach used in the ALFALFA survey. Where a spectrum has local maxima on either side of the systemic velocity ($\Delta v = 0$ in Fig. 2), we identify the maximum on the left and on the right side separately. If these two maxima have comparable amplitudes (within a factor of 2), we treat them separately; in any other case, we use the single global maximum of the spectrum. We then linearly interpolate between channels to locate the velocity where 50 percent of the peak is crossed (using the separate peak flux density values, if applicable). This process is illustrated in the second panel of Fig. 2. In rare cases where the half-maximum flux level is crossed more than twice (keeping in mind that we allow the halo finder to separate the particles belonging to galaxies from their satellites or nearby companions), we use the crossings with the minimum and maximum Δv to measure the width.

Our process to measure w_{50} differs slightly from that used in constructing the ALFALFA catalogue. In particular, the decision of whether to use one or two peaks is informed by visual inspection of individual spectra, and polynomial (usually linear) fitting of the edge regions of the spectra is used to locate the bounds defining w_{50} rather than the direct interpolation between channels in our approach. Our decision to simplify the process somewhat (in particular to avoid the visual inspection step) is justified – we find that even relatively drastic changes to the method, such as using no interpolation but simply taking the nearest channel, do not lead to qualitative changes in our analysis or conclusions.

With S_{21} and w_{50} determined for each galaxy in the mock survey volume, we are in a position to evaluate which galaxies would actually be detected by an ALFALFA-like survey. We model ALFALFA’s bivariate selection function with the empirically determined form given in Oman (2022, equation A5, based on the approach of Haynes et al. 2011). We linearly interpolate between the 25, 50, and 90 percent completeness levels at fixed S_{21} , continuing the slope of the interpolation to extrapolate to 0 and 100 percent completeness levels. We decide whether a given simulated galaxy is detected by drawing a random number from a uniform distribution between 0 and 1 and comparing it to the survey completeness at its S_{21} and w_{50} – if the random number is less than the completeness, the galaxy is accepted as a detection.

³<https://github.com/kyleaoman/martini>, version 2.0.2.

⁴The channel spacing of the ‘raw’ ALFALFA spectra is about 5.5 km s^{-1} , or about 11 km s^{-1} after Hanning smoothing. However, we find no qualitative differences in our analysis or conclusions even if we skip the smoothing step and proceed with a spectral resolution of 4 km s^{-1} . We therefore surmise that the small difference in spectral resolution can be safely ignored.

In addition to S_{21} and w_{50} , we also catalogue the position within the survey cone (nominally, a right ascension, a declination, and a distance, although the former two carry little meaning since our simulations are of random cosmological volumes), and the maximum of the circular velocity curve $v_{\text{circ}}(r) = \sqrt{GM(<r)}/r$ determined by SUBFIND, where $M(<r)$ is the mass enclosed within a spherical aperture of radius r and G is Newton’s constant.

3.2 Matching galaxies across simulation volumes

The pairs of simulations at fixed side length (EAGLE25-CDM & EAGLE25-WDM-1.5 keV, and EAGLE34-CDM & EAGLE34-WDM-7 keV) share the same random phases defining their initial conditions, so the same large-scale structures form in the two simulations of each pair. Identifying a galaxy from each as the ‘same’ object is therefore well defined. In our analysis we will make some comparisons using such one-to-one matched galaxies. We identify the ‘best match’ of each galaxy in a CDM simulation as the galaxy in the corresponding WDM simulation that contains the largest number of dark matter particles that belong to the CDM realization of the galaxy.⁵ We then repeat the process to identify the ‘best match’ of each WDM galaxy in the corresponding CDM simulation. Candidate pairs are accepted as a match only if they are each others’ mutual best matches (i.e. the match is ‘bijective’). This inevitably leaves a large fraction of galaxies – especially satellites and the low-mass galaxies whose formation WDM physics suppresses – unmatched, but this does not pose any problem for the comparisons that we will make below.

4 RESULTS

4.1 The subhalo v_{max} function

The effect of WDM that we are aiming to constrain through observations is the suppression of the formation of low-mass galaxies. We therefore begin with a straightforward quantification of this suppression within our mock survey volumes – the subhalo v_{max} function. In the upper panel of Fig. 3, we show the number⁶ of subhaloes (including ‘central’ and ‘satellite’ subhaloes) in the mock survey volume as a function of their maximum circular velocity for each of our four simulations. Reassuringly, the v_{max} functions of the two CDM simulations (light- and dark-purple dashed lines) agree very closely with differences of more than 10 per cent only occurring at $v_{\text{max}} \gtrsim 100 \text{ km s}^{-1}$ (dashed dark purple line in centre panel) where sampling shot noise due to the limited volume of the simulations begins to become important.⁷

The two mock surveys of the WDM simulations have systematically fewer subhaloes at the low- v_{max} end. The ‘warmer’ DM model (EAGLE25-WDM-1.5 keV) has a factor of about 40 fewer haloes than the CDM simulations at $v_{\text{max}} \approx 10 \text{ km s}^{-1}$, while the EAGLE34-WDM-7 keV simulation has a factor of ≈ 10 fewer haloes

⁵Particles are labelled with an identifier that is determined before particle positions are perturbed according to the power spectrum in the initial conditions, so the ‘same’ particles in the CDM and WDM simulations share an identifier.

⁶Since we always work with the same fixed volume of $5.29 \times 10^6 \text{ Mpc}^3$, we choose to show number counts rather than number densities.

⁷For instance, in the EAGLE25-CDM mock survey, a subhalo would be replicated about 339 times, so at number counts less than $339 \times 100 = 3.4 \times 10^4$ we would expect shot noise at the level of $\approx (100)^{-1/2} = 0.1$, i.e. 10 per cent. This is the abundance reached at about 100 km s^{-1} .




Figure 3. *Upper panel:* the halo v_{max} function within a conical region with a volume equal to the ALFALFA survey for the four simulations used in this work (EAGLE25-CDM – light purple; EAGLE34-CDM – dark purple; EAGLE25-WDM-1.5 keV – red; EAGLE34-WDM-7 keV – orange). Spurious haloes in the WDM case have been removed (see Section 2.1). The bins in v_{max} have a width of 0.2 dex. The dashed lines include all haloes within the conical volume, while the solid lines include only haloes with $M_{\text{HI}} > 10^6 M_{\odot}$. The suppression of galaxy formation on small scales in WDM cosmogonies is evident below $v_{\text{max}} \approx 70 \text{ km s}^{-1}$. *Middle panel:* ratio of the curves from the upper panel to the corresponding EAGLE25-CDM curves. *Lower panel:* fraction of haloes that have $M_{\text{HI}} > 10^6 M_{\odot}$ as a function of v_{max} . More low-mass ($v_{\text{max}} \approx 40 \text{ km s}^{-1}$) haloes exceed this threshold in warmer dark matter cosmogonies.

than CDM at the same v_{max} . The suppression of low- v_{max} subhaloes continues up to about 80–100 km s^{-1} where differences are of about 10–20 per cent and begin to be difficult to distinguish from the shot noise in the subhalo v_{max} function due to the limited simulation volumes. Bose et al. (2016) note that the suppression in the halo number density due to WDM is clear up to a scale of about $10M_{\text{hm}}$ – for the EAGLE25-WDM-1.5 keV simulation this is a subhalo mass scale of about $5 \times 10^{10} M_{\odot}$, corresponding to a v_{max} scale of about 65 km s^{-1} , in good agreement with the curves in Fig. 3. We stress that at $v_{\text{max}} \lesssim 80 \text{ km s}^{-1}$ the counting uncertainties (including accounting for tiling the simulations volumes) are much smaller than the amplitude of the suppression of number counts relative to CDM.

To take one step closer to the ALFALFA observations that we wish to compare against, we also show the subhalo v_{max} function of galaxies with H I masses of at least $M_{\text{HI}} > 10^6 M_{\odot}$ in the upper panel of Fig. 3 (solid lines). At v_{max} of less than about 35 km s^{-1} , the number counts drop precipitously as most of these galaxies contain

very little or no H I gas. However, between $35 \lesssim v_{\max}/\text{km s}^{-1} \lesssim 80$, the reduction in the number counts driven by WDM is still clearly evident (and the two CDM simulations still agree closely with each other). Interestingly, the fraction of galaxies with $M_{\text{HI}} > 10^6 M_{\odot}$ at fixed v_{\max} in this range in v_{\max} becomes higher with increasingly warm DM (lower panel of Fig. 3). This turns out to strongly affect the relative number counts of galaxies that are detected in our mock surveys; we will explore this in detail below.

4.2 The H I MF and H I WF

Since the detection of a galaxy in ALFALFA depends primarily on its flux and line width, number counts as a bivariate function of S_{21} and w_{50} are, in a sense, the fundamental measurement obtained by the survey for population studies. Converting fluxes to masses (equation 1) helps to facilitate physical interpretation. Integrating the counts along the line-width axis yields the H I mass function, while the orthogonal integral along the mass axis gives the H I velocity width function. In a flux-limited survey number counts need to be corrected for the selection function of the survey in order to estimate number densities in a fixed volume (e.g. Zwaan et al. 2003, 2005). Such estimates are subject to substantial statistical and systematic uncertainties (e.g. Jones et al. 2018; Oman 2022; Brooks et al. 2023). Going in the other direction – drawing a flux-limited sample from a complete volume-limited catalogue – is much more straightforward: it depends only on the assumed selection function in a simple way. Since volume-limited catalogues are trivial to construct from our mock survey volumes, we adopt an approach where we sample from these (Section 3.1) and make comparisons with the ALFALFA survey in terms of observed number counts.

In Fig. 4, we show the H I MF (left panels) and H I WF (right panels) of our four mock survey volumes. The solid lines show the number counts of all galaxies in the volumes with $M_{\text{HI}} > 10^6 M_{\odot}$ and $w_{50} > 10 \text{ km s}^{-1}$, regardless of whether or not they are detected. Considering first the H I MF, the two CDM simulations agree closely (with differences of about the amplitude expected from Poisson noise) across almost the entire range in M_{HI} shown. At the highest masses ($M_{\text{HI}} > 2 \times 10^{10} M_{\odot}$), the EAGLE34-CDM curve turns up relative to the EAGLE25-CDM curve because the larger simulation volume of the former better samples the formation of these rare objects.

The two WDM simulations have very similar H I MFs to the CDM simulations across most of the range in M_{HI} , with no clear systematic trend with the ‘warmth’ of the DM in the range $2 \times 10^7 \lesssim M_{\text{HI}}/M_{\odot} \lesssim 5 \times 10^9$. In the most massive galaxies, the EAGLE25-WDM-1.5 keV model seems to cool substantially more gas than the other models and forms galaxies with very high central stellar densities (a point to which we will return below). At the low H I mass end, number counts at fixed M_{HI} decrease for increasingly warm DM, reflecting the lower abundance of dark matter haloes in WDM cosmogonies – the median v_{\max} of a halo hosting a galaxy with $M_{\text{HI}} = 10^7 M_{\odot}$ is about 35 km s^{-1} (in all 4 simulations). While this is a potentially interesting difference between the predictions of CDM and WDM models, it occurs at an H I mass scale below that where resolution-dependent numerical effects begin to occur in the EAGLE galaxy formation model ($M_{\text{HI}} \approx 5 \times 10^7 M_{\odot}$; see Crain et al. 2017, section 4.2.1 for a detailed discussion – we note that our selection excludes the ‘dark’ haloes discussed in that work). Observed number counts in the ALFALFA survey (green crosses; see Section 2.2) are also very low in this regime, with fewer than 30 galaxies detected with $M_{\text{HI}} < 10^7 M_{\odot}$; this will be mitigated in the

future by wider surveys of comparable sensitivity. Encouragingly, the observed number counts actually agree rather well with those in all simulations (coloured points) for H I masses $\gtrsim 3 \times 10^8 M_{\odot}$ – that is, the regime where resolution-dependent effects are not known to occur – with the exception of the most massive galaxies, which are not reproduced by our relatively small simulation volumes.

We turn next to the H I WF (right panels of Fig. 4) which, through its closer connection to the dynamical masses of galaxies, is perhaps a more promising means to discriminate reliably between CDM and WDM cosmogonies. The most prominent feature in the number count distribution of all galaxies (solid curves) is a peak at w_{50} of just over 20 km s^{-1} . The sharp decline towards lower line widths is due to the minimum line width imposed by thermal broadening of the H I spectra – almost all gas identified as H I in the simulations has a temperature of about $8 \times 10^3 \text{ K}$, corresponding to w_{50} of about $2\sqrt{2 \log(2)}(8 \text{ km s}^{-1}) \approx 19 \text{ km s}^{-1}$. The sharp increase in the number counts from w_{50} of about 40 towards 20 km s^{-1} is related to the bump in the H I MF at $M_{\text{HI}} \approx 5 \times 10^6\text{--}10^7 M_{\odot}$ – the median M_{HI} at a w_{50} of 20 km s^{-1} is $3\text{--}6 \times 10^6 M_{\odot}$ (depending on the particular simulation), with a 16th–84th percentile scatter of about 0.6 dex. In this extreme low- w_{50} regime, the suppression of number counts due to WDM physics is manifest (lower-right panel of Fig. 4), but surprisingly it only appears at $w_{50} \lesssim 35 \text{ km s}^{-1}$. In Fig. 3, we showed that number counts are suppressed in our WDM simulations at $v_{\max} \lesssim 80 \text{ km s}^{-1}$, and haloes with a given v_{\max} are expected to contribute to the H I WF at $w_{50} \lesssim 2v_{\max}$, so we might expect to see differences even up to $w_{50} \approx 160 \text{ km s}^{-1}$.

The result is even more surprising when we consider only those galaxies that are detected in our mock surveys (coloured points in right panels of Fig. 4). In this case, differences between CDM and WDM (mainly the more extreme EAGLE34-WDM-7 keV model) do persist to higher $w_{50} \gtrsim 100 \text{ km s}^{-1}$, but have the opposite sign to that which would naïvely be predicted – more galaxies are detected in the mock survey of the WDM simulation than of the CDM simulations! We will detail the reasons for this in the next subsection, but first, we briefly comment on the comparison of the H I WFs from our mock surveys with that observed by ALFALFA. These agree rather closely – within about 20 per cent or better in the case of the EAGLE25-CDM, EAGLE34-CDM, and EAGLE34-WDM-7 keV simulations – across most of the range in w_{50} . At $w_{50} \lesssim 35 \text{ km s}^{-1}$, resolution-dependent effects in the simulations begin to become important, as mentioned above, so the departure here likely reflects a failure of the simulations to adequately resolve H I gas physics in the lowest-mass galaxies.⁸ At the high- w_{50} end, the culprit is once again most plausibly the relatively small volume of our simulations precluding the formation of the rarest massive galaxies. This reasonably good agreement between the observed and mock-observed H I WFs is far from trivial – see for instance Brooks et al. (2023, fig. 4) for an example of a galaxy formation model that fails rather spectacularly in its prediction for the H I WF. We will discuss the origin of this agreement further in Section 5.1.

⁸The agreement at $w_{50} < 35 \text{ km s}^{-1}$, and also at $M_{\text{HI}} < 10^8 M_{\odot}$, is somewhat improved if simulated galaxies with $M_{\star} = 0$ are included in the selection, but we note that, all ALFALFA sources included in the sample have optical counterparts. See Crain et al. (2017, sec. 4.2.1) for further discussion of the (largely unphysical) origin of these ‘dark’ galaxies in the simulations.




Figure 4. *Upper left panel:* the HIMF in the four simulations used in this work, and as observed by ALFALFA. The solid curves show the counts of all galaxies in the mock survey cones, while points include only those that would be detected by an ALFALFA-like survey. The colour coding (legend in upper right panel) is as in Fig. 3. The ‘bump’ feature (solid curves) at low H I mass ($M_{\text{HI}} \lesssim 10^7 M_{\odot}$) has a numerical origin (see Section 4.2). The green crosses show the number counts from the ALFALFA extragalactic source catalogue. The bins in M_{HI} have a width of 0.1 dex. *Lower left panel:* ratio of the curves and points from the upper-left panel to the corresponding EAGLE (CDM) curve or points. *Upper right panel:* as upper left panel, but for the H I WF. The bins in w_{50} have a width of 0.1 dex. The suppression of number counts seen at low v_{max} in WDM cosmogonies is not evident at low w_{50} – in fact, the survey of the simulation with the ‘warmest’ DM (EAGLE with 1.5 keV thermal relic) has a factor of ≈ 2 more detections at $w_{50} \approx 70 \text{ km s}^{-1}$ than the surveys of CDM simulations. *Lower right panel:* as lower right panel, but with ratios relative to EAGLE (CDM) corresponding to points and curves in the upper right panel.

4.3 Why more galaxies are detected in WDM models

There is already a hint of the reason why more galaxies with line widths of about $40\text{--}100 \text{ km s}^{-1}$ are detected in our mock surveys of simulations with ‘warmer’ DM in the lower panel of Fig. 3: at fixed v_{max} between about $20\text{--}70 \text{ km s}^{-1}$ the fraction of haloes with H I masses greater than $10^6 M_{\odot}$ systematically increases for increasingly warm DM. We explore this further in Fig. 5. In the upper panel, we plot the median H I mass of galaxies as a function v_{max} . The solid lines show the trends when all galaxies are included. Below v_{max} of about 100 km s^{-1} , the median H I mass at fixed v_{max} systematically increases for increasingly warm DM. Considering only those galaxies that are detected in the mock surveys, we find that these are biased to higher median H I masses than the global galaxy population. This is intuitive – the surveys are biased towards the brightest sources.

The lower panel of Fig. 5 is similar to the upper panel, but shows the median w_{50} at fixed v_{max} . There is a trend for galaxies with $40 < v_{\text{max}}/\text{km s}^{-1} < 80$ to have slightly higher median w_{50} with increasingly warm DM; this is because their higher H I masses (as shown in the upper panel) cause their gas discs to be somewhat more extended and therefore trace the rotation curves to slightly larger radii. As should be expected from the very tight correlation between H I mass and size (see Stevens et al. 2019, for a detailed discussion of its origin), gas discs have the same sizes at fixed H I mass in all of the simulations (within 15 per cent or less), so we focus our attention on systematic offsets in H I mass rather than size. Considering only

those galaxies detected in our mock surveys, we find that at the low- v_{max} end (about $30 < v_{\text{max}}/\text{km s}^{-1} < 70$), these are biased to wider line widths than the overall galaxy population. Unlike the detection bias in H I mass, this runs contrary to intuition: all else being equal, a wider line has a lower signal-to-noise ratio than a narrow line and should therefore be more difficult to detect. This shows that the bias towards detecting galaxies with higher M_{HI} dominates over that towards detecting galaxies with lower w_{50} at any given v_{max} in our mock surveys.

The offset between the curves in the upper panel of Fig. 5 could be interpreted either as a systematic difference in M_{HI} or in v_{max} . That is, we could imagine either that a galaxy in a WDM simulation ends up with more H I gas than its counterpart in the corresponding CDM simulation, or instead that it ends up with the same amount of H I gas but a lower v_{max} (or a combination of both effects). Around v_{max} of 30 km s^{-1} , where the curves are nearly horizontal, it is clear that the driving effect must be a difference in M_{HI} , but around v_{max} of 70 km s^{-1} where the curves are steeper the situation is ambiguous. We resolve this ambiguity by using our individually matched pairs of haloes in WDM and corresponding CDM simulations (see Section 3.2 for details of our matching process).

In Fig. 6, we show the median ratio of v_{max} of haloes in the two WDM simulations and the v_{max} of the matched CDM halo in the upper panels, and similarly the median ratio of M_{HI} in the lower panel. This reveals that both quantities are systematically different




Figure 5. *Upper panel:* median H I mass of all galaxies with $M_{\text{HI}} > 10^6 M_{\odot}$ within the mock survey cones (solid lines) and of only those galaxies detected in the mock surveys (points), as a function of maximum circular velocity, v_{max} . The upturn in the curves at the low- v_{max} end is an edge effect due to the cut in H I mass. Colour coding (see legend) is as in Fig. 3. The interquartile scatter of the distributions is shown for the EAGLE (CDM) simulation as the shaded band and error bars; the scatter in the other simulations is very similar. *Lower panel:* as the upper panel, but showing the median H I line width w_{50} as a function of v_{max} .

in the two simulations. v_{max} is systematically slightly lower in WDM (by 10–30 per cent) at the low- v_{max} end, with little scatter. This is a familiar effect due to the lower concentrations of WDM haloes at fixed halo mass (see for example, Lovell & Zavala 2023, fig. 4). M_{HI} , on the other hand, is systematically higher by up to a factor of 2, with large scatter. Most of the offset between the curves in the upper panel of Fig. 5 therefore comes from systematic differences in H I mass, although systematic differences in v_{max} do play a small role.

The origin of the enhanced H I mass of low-mass galaxies in the WDM simulations seems to be related to their assembly histories. Galaxies with $30 < v_{\text{max}}/\text{km s}^{-1} < 50$, where the effect is most pronounced, form their stellar mass very late – after $z = 4$ (in EAGLE34-WDM-7 keV) or $z = 3$ (in EAGLE25-WDM-1.5 keV), as opposed to from about $z = 5.5$ in the CDM simulations. Yet, galaxies of a given v_{max} assemble almost exactly the same stellar mass in the CDM and corresponding WDM models. This lag in formation time persists up to about $v_{\text{max}} = 100 \text{ km s}^{-1}$, becoming smaller with increasing v_{max} ; this is also the value of v_{max} where the CDM and corresponding WDM simulations have about the same H I mass per galaxy at fixed v_{max} (the difference in H I mass grows again at larger v_{max}). The delayed formation of structure at the low-mass




Figure 6. *Upper panel:* the median maximum circular velocity v_{max} (solid line) and interquartile scatter (shaded band) of WDM haloes relative to the v_{max} of the matching halo from the corresponding CDM simulation. See Section 3.2 for details of the matching procedure. The v_{max} of low- v_{max} haloes is reduced by 10–30 per cent with little scatter in the WDM cosmogonies considered. The upturn in the EAGLE34-WDM-7 keV curve at $v_{\text{max}} \gtrsim 200 \text{ km s}^{-1}$ is due to high central baryon densities causing a central peak in the circular velocity curves of massive galaxies in the WDM simulation. *Lower panel:* as upper panel, but comparing the H I masses of WDM galaxies with those of their CDM counterparts. Low-mass galaxies have systematically higher gas masses by up to a factor of 2, with considerable scatter. This is similar to the trend shown in the upper panel of Fig. 5, but for matched haloes such that the systematic shift in v_{max} (upper panel of this figure) is removed.

end in WDM models, and the related delayed star formation in low-mass galaxies, likely allows galaxies to retain more gas over time. That galaxies near the cutoff scale in WDM form by direct collapse and accretion rather than mergers of smaller objects, which in CDM would likely have lost all of their gas through supernova feedback.⁹ A more definitive explanation would require a full merger tree analysis of the simulations which we defer to future work.

Although we are focused primarily on galaxies with $v_{\text{max}} \gtrsim 100 \text{ km s}^{-1}$ in this work, we feel compelled to comment on the conspicuous upturn of the EAGLE25-WDM-1.5 keV curve at $v_{\text{max}} \gtrsim 200 \text{ km s}^{-1}$ in the upper panel of Fig. 6. This comes from dense accumulations of stars and gas in the centres of massive galaxies in this simulation, causing a central peak in the circular velocity curve that much exceeds the circular velocity reached at larger radii. We have checked explicitly that no equivalent feature exists in a plot of $M_{200}/M_{200(\text{CDM})}$ versus M_{200} ; v_{max} is simply no longer a good tracer

⁹Similar effects are present in the ETHOS model of self-interacting dark matter (SIDM), which exhibits a matter power spectrum cutoff qualitatively similar to WDM. An analysis of SIDM-CDM matched pairs at $z > 6$ revealed systematically higher gas masses for galaxies in the SIDM model (Lovell, Zavala & Vogelsberger 2019).

of total dynamical mass in these galaxies. Since these objects are not our focus in this work, we will not dwell further on this detail.

5 DISCUSSION

In this section, we compare our mock survey approach with a simplified approach to highlight some of its strengths and comment on its potential to produce constraints on WDM cosmological models.

5.1 A simplified mock survey

As mentioned in Section 4.2, that a galaxy formation model will reproduce the H I WF even very approximately is far from a foregone conclusion. We explore some of the reasons for this in this section with the help of a simplified mock survey, using the EAGLE25-CDM simulation as an illustrative example.

We keep exactly the same set-up described in Section 3.1 except for the calculation of the line widths, w_{50} . We replace the entire calculation with a simple – if often used – approximation:

$$w_{50} = \sqrt{(2v_{\max} \sin i)^2 + (10 \text{ km s}^{-1})^2}. \quad (3)$$

This assumes an inclined (by i) rotating disc that reaches an amplitude v_{\max} . The second term in the sum in quadrature reflects the fact that H I gas is not perfectly dynamically cold but has an intrinsic velocity dispersion of about 10 km s^{-1} , which will dominate the line width once the rotation speed (or the inclination) becomes low enough.

We plot the median w_{50} obtained in this way as a function of v_{\max} in the left panel of Fig. 7, which displays the expected straight line with a slight upward inflection at the low- v_{\max} end where the second term begins to contribute. For comparison, we also show the curve using the mock-observed w_{50} values (reproduced from Fig. 5), which has a more complex shape.

We show the H I WF for each of these cases in the right panel of Fig. 7. There are several, often competing effects driving the many differences in the figure:

(i) In low-mass galaxies, the gas disc does not sample the rotation curve all the way to v_{\max} , but only the central part where it is still rising. This causes galaxies that had w_{50} of less than about 60 km s^{-1} in the fiducial mock survey to have larger w_{50} in the simplified mock survey – this is the origin of the peak in the solid line in the right panel of Fig. 7 at that line width.

(ii) The overestimated line widths in the simplified mock survey from (i) mean that more intrinsically faint galaxies have higher line widths, suppressing the fraction of galaxies that are detected at some line widths. For instance, even though there are a factor of several more galaxies with w_{50} of about 60 km s^{-1} in the simplified mock survey (solid line compared to dashed line), a factor of a few less are detected (solid points compared to open points) at the same line width.

(iii) In the simplified survey there is a long low- w_{50} tail in the distribution at fixed v_{\max} due to face-on galaxies, extending all the way to the velocity dispersion floor that we impose at 10 km s^{-1} . Galaxies in this tail have unrealistically narrow H I lines (because the discs are assumed to be unrealistically symmetric and without warps or other irregular features common in observed galaxies), enabling them to satisfy our detection criteria at unrealistically large distances. This is part of the reason why between w_{50} of 20 and 70 km s^{-1} , for instance, the relative galaxy counts in the fiducial and simplified mock surveys change by almost two orders of magnitude (compare

solid and dashed lines in right panel of Fig. 7), but the counts of detected galaxies varies by only about a factor of 5 over the same interval in w_{50} . Even though the tail of the w_{50} distribution has a low amplitude (e.g. the 0.1 percentile of the distribution shown in the left panel of Fig. 7), the fraction of galaxies detected at a given w_{50} is less than 1 per cent over much of the range in w_{50} , such that even a tiny fraction of affected galaxies can be important.

5.2 Implications for use of the H I WF to constrain dark matter

The above considerations make modelling the w_{50} function to infer a v_{\max} function an extremely challenging prospect. The first point was already identified by Brooks et al. (2017, see also Dutton et al. 2019), but the other two have not, to our knowledge, been previously discussed in the literature. The mock survey based on a semi-analytic model of Brooks et al. (2023) is similar to our simplified mock survey in some respects but accounts for the H I sizes of galaxies and how these sample their rotation curves. The H I WF that it predicts is more similar to that of our simplified mock survey than to our fiducial mock survey (although the agreement with ALFALFA observations is worse than either!), suggesting that rotation curves that do not reach v_{\max} are not the only effect at play.

Fundamentally, the issue is that given an observation of the H I line width of a source, there is ambiguity in what value of v_{\max} should be assigned to it. A common assumption (e.g. Zwaan et al. 2010; Papastergis et al. 2015) is precisely that of our simplified mock survey (equation 3). The left panel of Fig. 7 illustrates that a realistic mapping $v_{\max}(w_{50})$ is likely not single-valued (see also Brooks et al. 2017, fig. 3, and Dutton et al. 2019, fig. 4). Fig. 7 as a whole shows that (erroneously) assuming a relation similar to equation (3) is likely to have severe consequences for constraints on the galaxy population and, by extension, on the halo v_{\max} function.

The forward-modelling approach that we adopt in our fiducial mock surveys has the significant advantage of guaranteeing internal self-consistency: for a known halo v_{\max} function (that realized in a given simulation), the corresponding w_{50} function is unambiguously predicted and, given knowledge of the survey selection function, predictions of source counts are also unambiguous. This therefore seems like a much more robust way to test the consistency of cosmological models with observation.

This is not to say that forward-modelling the H I WF from a simulation is without caveats. The prediction for the line widths is limited by how faithfully the simulation models the gas ionization state and kinematics in galaxies which depends on physical processes that are notoriously difficult to model, such as supernova feedback. It can be hoped, however, that whether a given simulation is faithful enough in its reproduction of the detailed internal structure of galaxies can be assessed on the basis of comparison with similarly detailed observations of galaxies.

Another hurdle to using this forward-modelling approach for constraining cosmology is that it is challenging to show that a prediction for the H I WF is unique to a particular cosmological model and cannot be reproduced in a different cosmogony. The prediction is unlikely to be completely degenerate – it seems unlikely that if two cosmological models resulted in identical predictions for the H I WF that they would also predict otherwise indistinguishable populations of galaxies – so again there seems to be hope for progress through folding in additional observational constraints.

6 CONCLUSIONS

Our analysis suggests that if dark matter is warm (with an equivalent thermal relic mass of about 10 keV or less) then it is likely to

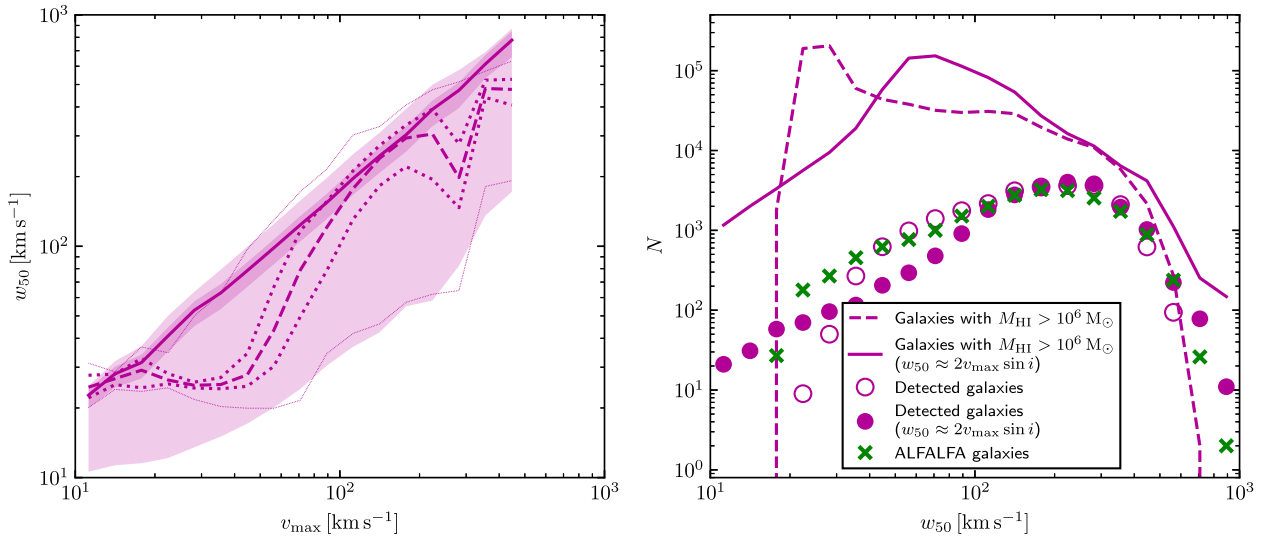


Figure 7. Comparison of a mock survey using our fiducial ‘observed’ H I line widths, w_{50} , with approximate line widths estimated as $w_{50} \approx 2v_{\max} \sin i$ (see Section 5.1 for details), for the EAGLE25- Λ CDM simulation. *Left panel:* median w_{50} as a function of v_{\max} using ‘observed’ (dashed line, repeated from Fig. 5) and approximate (solid line) values of w_{50} . The interquartile scatter is shown with the heavier dotted lines and darker shaded region, respectively, and the thinner dotted lines and lighter shaded region show the 0.5th–99.5th percentile range. *Right panel:* H I WF using ‘observed’ and approximate values of w_{50} . Lines are as in the left panel. The open symbols show the number of galaxies detected in the mock survey when the ‘observed’ values of w_{50} are used (repeated from the right panel of Fig. 4) while the filled symbols show the same when approximate w_{50} values are used. Number counts from the ALFALFA catalogue are also shown for comparison. There are several, often competing effects (detailed in Section 5.1) that shape the trends in the right panel.

leave a signature in the number counts as a function of spectral line width in 21-cm surveys. The (marginalized) H I WF is likely to be degenerate, or nearly degenerate, for combinations of different dark matter models and assumptions about galaxy formation physics, but this degeneracy can almost certainly be broken by considering the joint distribution of the spectral line width and other observables. For example, two combinations of different dark matter models and galaxy formation models that predict the same H I WF are unlikely to also predict identical galaxy stellar mass functions, and even less likely to predict the same joint distribution of stellar mass and H I line width. This opens up an avenue to attempt to place quantitative constraints on WDM cosmological models in a region of parameter space competitive with other methods, such as those obtained from Ly α forest measurements (e.g. Hansen et al. 2002; Viel et al. 2005; Seljak et al. 2006; Viel et al. 2006; Murgia, Iršič & Viel 2018; Palanque-Desabrouille et al. 2020; Garzilli et al. 2021).

As we have discussed above, the EAGLE galaxy formation model is rather limited in its treatment of H I by, amongst other shortcomings, the pressure floor imposed on the gas in the model (see Ploekinger et al. 2024, for a detailed discussion of the rationale for this, and of recent advances that obviate the need for it), the lack of local ionizing sources (Kannan et al. 2020; Li et al. 2020), and the mismatch in energetics between feedback events in nature and those modelled by the simulation (Crain et al. 2015, especially sec. 3.1). Our view is that this choice in modelling wreaks sufficient havoc at the low-H I mass and low- w_{50} end of the galaxy population as to make quantitative constraints on WDM models based on the H I WF using these simulations uncertain. We are optimistic that very recent (e.g. Feldmann et al. 2023) and future simulations, such as those with the forthcoming COLIBRE model (Schaye et al. in preparation), that explicitly follow gas to much lower temperatures and higher densities (while still recreating a realistic and statistically useful galaxy population) will offer opportunities for progress with the approach outlined in this work.

The ALFALFA survey is the current gold standard for work on the H I WF, but will very soon be superseded by new surveys including WALLABY (Koribalski et al. 2020) on ASKAP¹⁰ and CRAFTS (Zhang et al. 2019) on FAST,¹¹ and eventually the Square Kilometre Array (see Power et al. 2015, for some discussion around constraining dark matter models including WDM using SKA measurements of the H I WF). Notably, the WALLABY survey has somewhat better sensitivity than ALFALFA and when completed will cover triple the area on the sky, with projections of up to a factor of 25 more detections expected. CRAFTS is projected to detect up to 6×10^5 galaxies in a 20 000 deg² field out to $z < 0.35$. Of the WALLABY detections, about 20 000 (the total number of ‘Code 1’ sources detected in ALFALFA) will be at least marginally spatially resolved. Amongst many advantages, this offers the opportunity to identify galaxies where the gas does not sample the flat part of the rotation curve – an important systematic effect, as we have seen above.

The statistical (i.e. counting) uncertainties on the H I WF as determined from ALFALFA are already small (see Oman 2022) and are expected to shrink by a factor of 5 with WALLABY. This places us firmly in a systematic uncertainty-dominated regime; to make full use of the available data, the burden is on theoretical predictions to control for these at a level that allows for sufficient confidence in constraints on WDM models.

SOFTWARE

This work has made use of the following software packages: NUMPY (Harris et al. 2020), SCIPY (Virtanen et al. 2020), ASTROPY (Astropy Collaboration et al. 2013, 2022), MATPLOTLIB (Hunter 2007; Caswell

¹⁰Australian Square Kilometer Array Pathfinder.

¹¹Five-hundred-metre Aperture Spherical radio Telescope.

et al. 2023), MARTINI (Oman 2019, 2024; Oman et al. 2019), and PANDAS (McKinney 2010; The Pandas Development Team 2023).

ACKNOWLEDGEMENTS

We thank Adrian Jenkins for assistance with preparing initial conditions for the simulations used in this work, and the EAGLE and E-MOSAICS collaborations for making their simulations available to us.

KAO acknowledges support by the Royal Society through Dorothy Hodgkin Fellowship DHF/R1/231105 and by STFC through grant ST/T000244/1. KAO & CSF acknowledge support by the European Research Council (ERC) through an Advanced Investigator grant to C.S. Frenk, DMIDAS (GA 786910). JP is supported by the Australian government through the Australian Research Council's Discovery Projects funding scheme (DP220101863). This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. The study also made use of the Prospero HPC facility at Liverpool John Moores University. This research has made use of NASA's Astrophysics Data System.

DATA AVAILABILITY

The output of the simulations used in this work are available as follows:

- (i) EAGLE25-CDM: refer to public data release (McAlpine et al. 2016; The EAGLE team 2017).
- (ii) EAGLE34-CDM: available on reasonable request to RAC (R.A.Crain@ljmu.ac.uk).
- (iii) EAGLE25-WDM-1.5 keV: Available on reasonable request to KAO (kyle.a.oman@durham.ac.uk).
- (iv) EAGLE34-WDM-7 keV: available on reasonable request to RAC (R.A.Crain@ljmu.ac.uk).

Tabulated values for galaxies in our mock surveys as described in Section 3.1 are available from KAO on reasonable request.

REFERENCES

Astropy Collaboration et al., 2013, *A&A*, 558, A33
 Astropy Collaboration et al., 2022, *ApJ*, 935, 167
 Bahé Y. M. et al., 2016, *MNRAS*, 456, 1115
 Bardeen J. M., Bond J. R., Kaiser N., Szalay A. S., 1986, *ApJ*, 304, 15
 Bastian N., Pfeffer J., Kruijssen J. M. D., Crain R. A., Trujillo-Gomez S., Reina-Campos M., 2020, *MNRAS*, 498, 1050
 Blitz L., Rosolowsky E., 2006, *ApJ*, 650, 933
 Bode P., Ostriker J. P., Turok N., 2001, *ApJ*, 556, 93
 Bond J. R., Szalay A. S., 1983, *ApJ*, 274, 443
 Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
 Bose S., Hellwing W. A., Frenk C. S., Jenkins A., Lovell M. R., Helly J. C., Li B., 2016, *MNRAS*, 455, 318
 Boyarsky A., Ruchayskiy O., Iakubovskiy D., Franse J., 2014, *Phys. Rev. Lett.*, 113, 251301
 Brook C. B., Shankar F., 2016, *MNRAS*, 455, 3841
 Brooks A. M., Papastergis E., Christensen C. R., Governato F., Stilp A., Quinn T. R., Wadsley J., 2017, *ApJ*, 850, 97
 Brooks R. A. N., Oman K. A., Frenk C. S., 2023, *MNRAS*, 522, 4043
 Bulbul E., Markevitch M., Foster A., Smith R. K., Loewenstein M., Randall S. W., 2014, *ApJ*, 789, 13

Caswell T. A. et al., 2023, matplotlib/matplotlib: REL: v3.7.1. Zenodo
 Chauhan G., Lagos C. d. P., Obreschkow D., Power C., Oman K., Elahi P. J., 2019, *MNRAS*, 488, 5898
 Crain R. A. et al., 2015, *MNRAS*, 450, 1937
 Crain R. A. et al., 2017, *MNRAS*, 464, 4204
 Cullen L., Dehnen W., 2010, *MNRAS*, 408, 669
 Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
 Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
 Durier F., Dalla Vecchia C., 2012, *MNRAS*, 419, 465
 Dutton A. A., Obreja A., Macciò A. V., 2019, *MNRAS*, 482, 5606
 Enzi W. et al., 2021, *MNRAS*, 506, 5848
 Feldmann R. et al., 2023, *MNRAS*, 522, 3831
 Garzilli A., Magalich A., Ruchayskiy O., Boyarsky A., 2021, *MNRAS*, 502, 2356
 Giovanelli R. et al., 2005, *AJ*, 130, 2598
 Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, Clusters of Galaxies and the High Redshift Universe Observed in X-rays. p. 64, preprint (astro-ph/0106018), <https://inspirehep.net/literature/575192> and <https://moriond.in2p3.fr/previous-sessions.html>
 Hansen S. H., Lesgourgues J., Pastor S., Silk J., 2002, *MNRAS*, 333, 544
 Harris C. R. et al., 2020, *Nature*, 585, 357
 Haynes M. P. et al., 2011, *AJ*, 142, 170
 Haynes M. P. et al., 2018, *ApJ*, 861, 49
 Hellwing W. A., Frenk C. S., Cautun M., Bose S., Helly J., Jenkins A., Sawala T., Cytowski M., 2016, *MNRAS*, 457, 3492
 Hopkins P. F., 2013, *MNRAS*, 428, 2840
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Jenkins A., 2013, *MNRAS*, 434, 2094
 Jones M. G., Haynes M. P., Giovanelli R., Moorman C., 2018, *MNRAS*, 477, 2
 Kannan R., Marinacci F., Vogelsberger M., Sales L. V., Torrey P., Springel V., Hernquist L., 2020, *MNRAS*, 499, 5732
 Klypin A., Karachentsev I., Makarov D., Nasonova O., 2015, *MNRAS*, 454, 1798
 Koribalski B. S. et al., 2020, *Ap&SS*, 365, 118
 Kruijssen J. M. D., Pfeffer J. L., Crain R. A., Bastian N., 2019, *MNRAS*, 486, 3134
 Laine M., Shaposhnikov M., 2008, *J. Cosmol. Astropart. Phys.*, 2008, 031
 Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
 Li H., Vogelsberger M., Marinacci F., Sales L. V., Torrey P., 2020, *MNRAS*, 499, 5862
 Lovell M. R., Zavala J., 2023, *MNRAS*, 520, 1567
 Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, *MNRAS*, 439, 300
 Lovell M. R. et al., 2016, *MNRAS*, 461, 60
 Lovell M. R., Zavala J., Vogelsberger M., 2019, *MNRAS*, 485, 5474
 Macciò A. V., Udrescu S. M., Dutton A. A., Obreja A., Wang L., Stinson G. R., Kang X., 2016, *MNRAS*, 463, L69
 Mason A. C., Crain R. A., Schiavon R. P., Weinberg D. H., Pfeffer J., Schaye J., Schaller M., Theuns T., 2023, *MNRAS* preprint (arXiv:2311.00041)
 McAlpine S. et al., 2016, *Astron. Comput.*, 15, 72
 Meshveliani T., Lovell M. R., Crain R. A., Pfeffer J., 2024, *MNRAS*, 532, 1296
 McKinney W., 2010, in van der Walt S., Millman J., eds, Proceedings of the 9th Python in Science Conference. p. 56, <https://conference.scipy.org/proceedings/>
 McQuinn K. B. W. et al., 2021, *ApJ*, 918, 23
 Murgia R., Iršič V., Viel M., 2018, *Phys. Rev. D*, 98, 083540
 Nadler E. O., Birrer S., Gilman D., Wechsler R. H., Du X., Benson A., Nierenberg A. M., Treu T., 2021, *ApJ*, 917, 7
 Obreschkow D., Croton D., De Lucia G., Khochfar S., Rawlings S., 2009, *ApJ*, 698, 1467
 Obreschkow D., Ma X., Meyer M., Power C., Zwaan M., Staveley-Smith L., Drinkwater M. J., 2013, *ApJ*, 766, 137
 Oman K. A., 2019, in Allen A., DuPrie K., eds, MARTINI: Mock spatially resolved spectral line observations of simulated galaxies, Astrophysics

- Source Code Library. Published by ASCL at Michigan Technological University, Houghton, MI
- Oman K. A., 2022, *MNRAS*, 509, 3268
- Oman K. A., 2024, *J. Open Source Softw.*, 9, 6860
- Oman K. A., Marasco A., Navarro J. F., Frenk C. S., Schaye J., Benítez-Llambay A., 2019, *MNRAS*, 482, 821
- Palanque-Delabrouille N., Yèche C., Schöneberg N., Lesgourgues J., Walther M., Chabanier S., Armengaud E., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 038
- Papastergis E., Martin A. M., Giovanelli R., Haynes M. P., 2011, *ApJ*, 739, 38
- Papastergis E., Giovanelli R., Haynes M. P., Shankar F., 2015, *A&A*, 574, A113
- Pfeffer J., Kruijssen J. M. D., Crain R. A., Bastian N., 2018, *MNRAS*, 475, 4309
- Planck Collaboration XVI, 2014, *A&A*, 571, A16
- Ploekinger S., Nobels F. S. J., Schaller M., Schaye J., 2024, *MNRAS*, 528, 2930
- Power C. et al., 2015, in Proc. Advancing Astrophysics with the Square Kilometre Array (AASKA14). Giardini Naxos, Italy, Sissa Medialab srl Partita IVA: 01097780322, p. 133
- Price D. J., 2008, *J. Comput. Phys.*, 227, 10040
- Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, *MNRAS*, 430, 2427
- Rosas-Guevara Y. M. et al., 2015, *MNRAS*, 454, 1038
- Sardone A., Peter A. H. G., Brooks A. M., Kaczmarek J., 2024, *ApJ*, 964, 135
- Sawala T., Frenk C. S., Crain R. A., Jenkins A., Schaye J., Theuns T., Zavala J., 2013, *MNRAS*, 431, 1366
- Schaye J., 2004, *ApJ*, 609, 667
- Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schneider A., Trujillo-Gomez S., Papastergis E., Reed D. S., Lake G., 2017, *MNRAS*, 470, 1542
- Seljak U., Makarov A., McDonald P., Trac H., 2006, *Phys. Rev. Lett.*, 97, 191303
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Stevens A. R. H., Diemer B., Lagos C. d. P., Nelson D., Obreschkow D., Wang J., Marinacci F., 2019, *MNRAS*, 490, 96
- The EAGLE team, 2017, preprint ([arXiv:1706.09899](https://arxiv.org/abs/1706.09899))
- The Pandas Development Team, 2023, `pandas-dev/pandas`: Pandas. Zenodo
- Trujillo-Gomez S., Klypin A., Primack J., Romanowsky A. J., 2011, *ApJ*, 742, 16
- Viel M., Lesgourgues J., Haehnelt M. G., Matarrese S., Riotto A., 2005, *Phys. Rev. D*, 71, 063534
- Viel M., Lesgourgues J., Haehnelt M. G., Matarrese S., Riotto A., 2006, *Phys. Rev. Lett.*, 97, 071301
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- Wiersma R. P. C., Schaye J., Smith B. D., 2009a, *MNRAS*, 393, 99
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, *MNRAS*, 399, 574
- Zavala J., Jing Y. P., Faltenbacher A., Yepes G., Hoffman Y., Gottlöber S., Catinella B., 2009, *ApJ*, 700, 1779
- Zhang K. et al., 2019, *Sci. China Phys. Mech. Astron.*, 62, 959506
- Zwaan M. A. et al., 2003, *AJ*, 125, 2842
- Zwaan M. A., Meyer M. J., Staveley-Smith L., Webster R. L., 2005, *MNRAS*, 359, L30
- Zwaan M. A., Meyer M. J., Staveley-Smith L., 2010, *MNRAS*, 403, 1969

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.