MNRAS 000, 1–22 (0000) Preprint 2 February 2022 Compiled using MNRAS LATEX style file v3.0

The distribution of HI velocity profiles in a ΛCDM universe

Aseem Paranjape[1]⋆, R. Srianand[1]†, Tirthankar Roy Choudhury[2]‡ & Ravi K. Sheth[3,4]§

[1] *Inter-University Centre for Astronomy & Astrophysics, Ganeshkhind, Post Bag 4, Pune 411007, India*
[2] *National Centre for Radio Astrophysics, TIFR, Post Bag 3, Ganeshkhind, Pune 411007, India*
[3] *Center for Particle Cosmology, University of Pennsylvania, 209 S. 33rd St., Philadelphia, PA 19104, USA*
[4] *The Abdus Salam International Center for Theoretical Physics, Strada Costiera, 11, Trieste 34151, Italy*

2 February 2022

ABSTRACT

We model the distribution of the observed profiles of 21 cm line emission from neutral hydrogen (HI) in central galaxies selected from a statistically representative mock catalog of the local Universe in the Lambda-cold dark matter framework. The distribution of these HI velocity profiles (specifically, their widths W_{50}) has been observationally constrained, but has not been systematically studied theoretically. Our model profiles derive from rotation curves of realistically baryonified haloes in an N-body simulation, including the quasi-adiabatic relaxation of the dark matter profile of each halo in response to its baryons. We study the predicted W_{50} distribution using a realistic pipeline applied to noisy profiles extracted from our luminosity-complete mock catalog with an ALFALFA-like survey geometry and redshift selection. Our default mock is in good agreement with observed ALFALFA results for $W_{50} \gtrsim 700\,\mathrm{km\,s^{-1}}$, being incomplete at lower widths due to the intrinsic threshold of $M_r \leq -19$. Variations around the default model show that the velocity width function at $W_{50} \gtrsim 300\,\mathrm{km\,s^{-1}}$ is most sensitive to a possible correlation between galaxy inclination and host concentration, followed by the physics of quasi-adiabatic relaxation. We also study the *excess kurtosis* of noiseless velocity profiles, obtaining a distribution which tightly correlates with W_{50}, with a shape and scatter that depend on the properties of the turbulent HI disk. Our results open the door towards using the shapes of HI velocity profiles as a novel statistical probe of the baryon-dark matter connection.

Key words: galaxies: formation - cosmology: theory, dark matter, large-scale structure of Universe - methods: numerical

1 INTRODUCTION

The distribution of baryons in the Universe, particularly those locked up inside galaxies, is of fundamental interest for theories of structure formation. In the context of the Lambda-cold dark matter (ΛCDM) paradigm, a key goal is to robustly establish and theoretically interpret the details of the observed galaxy-dark matter connection. A host of observational probes is typically employed in this exercise, ranging from the distribution of masses and baryonic content of galaxy clusters (Vikhlinin et al. 2009a,b), to the clustering of galaxies in the local (Zehavi et al. 2011; eBOSS Collaboration et al. 2020) and high-redshift Universe (de la Torre et al. 2011; Marulli et al. 2013; Laurent et al. 2017), to the effects of gravitational lensing on galaxy shapes (Vikram et al. 2015; Heymans et al. 2021), all the way down to spatially resolved spectroscopy yielding information on the stellar content and inter-stellar medium of individual galaxies (Bundy et al. 2015) and (for spiral galaxies) their rotation curves (Persic et al. 1996; McGaugh et al. 2001).

Galaxy rotation curves in particular have a long history as probes of not only galactic structure and content (e.g., Athanassoula et al. 1987; Sofue & Rubin 2001; Gentile et al. 2004) but also the nature of gravity itself (Begeman et al. 1991; Blais-Ouellette et al. 2001; de Almeida et al. 2016; Lelli et al. 2016; McGaugh et al. 2016). For relatively nearby (distance $\lesssim 100h^{-1}$Mpc) rotationally supported galaxies, rotation curves can be measured using either optical observations of the stellar content or radio-frequency observations of the cold gas content by exploiting the 21 cm line transition of neutral hydrogen (HI) (Begeman 1989; Blais-Ouellette et al. 2001, 2004; Lelli et al. 2016). At larger distances ($z \sim 0.1$), spatially resolved spectroscopy at radio frequencies becomes increasingly challenging due to the decreasing projected sizes of galaxies. Nevertheless, due to the *velocity* resolution of $\lesssim$ few $\mathrm{km\,s^{-1}}$ achieved by current radio telescopes, the rotation curves of HI-bearing galaxies can still be indirectly probed by observing the spatially integrated HI *velocity profiles* – i.e., the redshifted 21 cm flux as a function of observed frequency – of individual objects. This quantity forms the key observable in large-volume surveys of HI-selected galaxies such as the HI

⋆ E-mail: aseem@iucaa.in
† E-mail: anand@iucaa.in
‡ E-mail: tirth@ncra.tifr.res.in
§ E-mail: shethrk@physics.upenn.edu

Parkes All Sky Survey (HIPASS, Barnes et al. 2001; Meyer et al. 2004) or the Arecibo Legacy Fast ALFA (ALFALFA) survey (Giovanelli et al. 2005, 2007) and is the main focus of the present work. Ongoing and upcoming surveys of HI-bearing galaxies with the SKA precursors are expected to be wider and deeper than the present ones (e.g., WALLABY and DINGO using ASKAP, Duffy et al. 2012; Koribalski et al. 2020 and LADUMA using MeerKAT, Holwerda et al. 2012), which will extend the scope of studies like that presented in this paper.

There has been extensive work in the literature on the modelling of rotation curves in the ΛCDM framework, focused mainly on describing observed rotation curves by fitting them with static or dynamic mass models of the respective galaxy's baryonic and dark matter content (see, e.g., Athanassoula et al. 1987; Gentile et al. 2004; Begum & Chengalur 2004; Granados et al. 2017; Kurapati et al. 2020). In recent work, some of us have explored an alternate route, using synthetic rotation curves – produced as part of statistically representative mock galaxy catalogs – to *predict* the statistical properties of large samples of rotation curve data. The underlying mock catalogs are generated by populating gravity-only cosmological N-body simulations with galaxies, using an empirical halo occupation distribution (HOD) constrained by the observed galaxy abundances and luminosity-dependent clustering (see below). The present work continues along these lines, focusing on self-consistently predicting the observed distribution of velocity profiles of massive HI-bearing galaxies in large surveys. The main motivation behind this exercise is the realisation that HI velocity profiles are, in principle, sensitive to a number of baryonic physics details due to their connection with the underlying rotation curve and the nature of the HI disk, as described in detail below. To our knowledge, this aspect of HI velocity profiles has not been systematically explored or exploited in the literature previously. The only works we are aware of are by Papastergis et al. (2011) and Moorman et al. (2014) who presented measurements of the distribution of the velocity widths of HI-selected galaxies in the ALFALFA survey. As such, the distribution of shapes of HI velocity profiles is a hitherto unexplored probe of the baryon-dark matter connection at small scales.

With this in mind, in this work we explore the sensitivity of HI velocity profiles to various aspects of the baryon-dark matter physics, such as (i) scaling relations involving disk sizes, (ii) environmental effects, (iii) the physics of quasi-adiabatic relaxation of dark matter in the presence of baryons and (iv) the impact of baryonic physics involving the intrinsic dispersion of the HI 21cm line in a galactic disk. As mentioned above, we perform this analysis using a realistic mock catalog of low-redshift ($z \lesssim 0.1$) galaxies which is constrained to reproduce the abundances and clustering of optically selected galaxies in the Sloan Digital Sky Survey (SDSS, York et al. 2000),[1] and HI-selected galaxies in the ALFALFA survey. As part of our analysis, we perform an in-depth study of the extraction of velocity widths from our simulated velocity profiles in the presence of realistic noise, allowing us to compare with the published ALFALFA results from Papastergis et al. (2011) and Moorman et al. (2014). Additionally, we emphasize the utility of beyond-width statistics such as excess kurtosis as a novel probe of baryonic physics in HI disks.

The paper is organised as follows. In section 2 we describe our mock catalogs and the procedure to 'baryonify' the host halo of each HI-bearing central galaxy. In section 3, we show how the rotation curve of such a galaxy can be used to model the HI profile it would present to a distant observer, discussing in detail the sensitivity of the model to different parameters and assessing its potential as a mass-modelling tool. We further discuss the extraction of the velocity width from a velocity profile in the presence of realistic noise, along with the subsequent estimate of the distribution of widths of an HI-selected sample. In section 4, we present the results of applying this procedure for obtaining the velocity width function to our mock galaxy catalog, exploring a number of variations in sample selection and modelling choices around our default model, as mentioned above. In section 5, we move beyond the velocity width and propose the excess kurtosis of the velocity profile as a novel probe of the physics of turbulence in the HI disk. We summarise and conclude in section 6. The appendices present technical details related to some aspects of the analysis. Throughout, we assume a spatially flat ΛCDM background cosmology, with parameters $\{\Omega_{\rm m}, \Omega_{\rm b}, h, n_{\rm s}, \sigma_8\}$ given by $\{0.276, 0.045, 0.7, 0.961, 0.811\}$, compatible with the 7-year results of the *Wilkinson Microwave Anisotropy Probe* experiment (WMAP7, Komatsu et al. 2011). We denote the base-10 (natural) logarithm as log (ln).

2 MOCK GALAXY CATALOG

The mock galaxy catalog on which we build our analysis is constructed using the algorithm described by Paranjape et al. (2021, hereafter, PCS21) and summarised below.

2.1 Simulation and mock algorithm

In this work, we rely on one realisation of the L300_N1024 simulation box discussed by PCS21. The (gravity-only) simulation evolved 1024^3 particles in a $(300h^{-1}{\rm Mpc})^3$ cubic box with the code GADGET-2 (Springel 2005)[2]. Dark haloes were identified using the code ROCKSTAR (Behroozi et al. 2013a)[3] and relaxed objects were retained, discarding substructure. Further details of the simulation can be found in Paranjape & Alam (2020). In the following, $m_{\rm vir}$ and $R_{\rm vir}$ will refer to the total halo mass and virial radius. We define $R_{\rm vir} \equiv R_{\rm 200c}$, the radius at which the enclosed halo-centric density becomes 200 times the critical density $\rho_{\rm crit}$ of the Universe, so that $m_{\rm vir} = (4\pi/3)R_{\rm vir}^3 \times 200\rho_{\rm crit}$.

Mock central and satellite galaxies were populated in these host haloes using the PCS21 algorithm to produce a luminosity-complete sample of galaxies with an r-band absolute magnitude threshold $M_r \leq -19$. This algorithm is based on the halo occupation distribution (HOD) model and optical-HI scaling relation calibrated by Paul et al. (2018) and Paul et al. (2019), and additionally assigns each mock galaxy with realistic values of $g-r$ and $u-r$ colours and

[1] www.sdss.org

[2] http://www.mpa-garching.mpg.de/gadget/

[3] https://bitbucket.org/gfcstanford/rockstar

stellar mass m_*. Most importantly for the present work, approximately 60% of these galaxies are also assigned values of neutral hydrogen (HI) mass $m_{\rm HI}$ sampled from the optical-HI scaling relation. The HOD models underlying the algorithm are constrained by the observed abundances and clustering of optically selected galaxies in the SDSS and of HI-selected galaxies in the ALFALFA survey. The luminosity threshold of $M_r \leq -19$ leads to completeness limits of $10^{9.85}h^{-2}M_\odot$ and $10^{9.7}h^{-2}M_\odot$ in m_* and $m_{\rm HI}$, respectively. We refer the reader to PCS21 for various tests and predictions of the algorithm.

2.2 Baryonification and rotation curves

The host haloes of the central galaxies thus produced are 'baryonified' by the PCS21 algorithm according to a modified version of the prescription of Schneider & Teyssier (2015, hereafter, ST15) which we discuss next, focusing on galaxies containing HI. The host halo of each HI-bearing central galaxy is assigned spatial distributions of the following baryonic components:

- A 2-dimensional axisymmetric HI disk ('HI') with scale length $h_{\rm HI}$ (surface density $\Sigma_{\rm HI}(r_\perp) \propto {\rm e}^{-r_\perp/h_{\rm HI}}$ in the disk plane), for centrals with an assigned $m_{\rm HI}$ value. The scale length $h_{\rm HI}$ is assumed to follow the empirical scaling $h_{\rm HI} \propto m_{\rm HI}^{0.5}$ (Wang et al. 2016, see equation 8 of PCS21). The corresponding mass fraction is $f_{\rm HI} = 1.33\, m_{\rm HI}/m_{\rm vir}$, with the prefactor accounting for Helium correction.
- A spherical distribution of stars in the central galaxy ('cgal') with half-light radius $R_{\rm hl}$ constrained by observations (Kravtsov 2013) and a mass fraction $f_{\rm cgal} = m_*/m_{\rm vir}$. The model currently does not include a separate stellar disk, which remains an interesting future extension.
- Spherical distributions of gravitationally bound, hot ionized gas ('bgas') in hydrostatic equilibrium, and expelled gas ('egas') or the circum-galactic medium affected by feedback processes. The mass fraction $f_{\rm bgas}$ is extrapolated to low $m_{\rm vir}$ from the relation calibrated by ST15 from X-ray cluster observations; $f_{\rm bgas} \lesssim 0.01$ for typical HI-bearing centrals. The mass fraction $f_{\rm egas}$ is set by baryonic mass conservation (see PCS21 for details).

Finally, the presence of these baryonic components is assumed to backreact on the dark matter profile according to the prescription of ST15 (see appendix A of PCS21), leading to a quasi-adiabatic relaxation, approximately conserving angular momentum, which tends to contract the dark matter in the inner halo and slightly expand it the halo outskirts, on average (see, e.g., fig. 1 of Paranjape & Sheth 2021). The physics of this relaxation is parametrised by a quantity $q_{\rm rdm}$ (e.g., equation A1 of PCS21), such that $q_{\rm rdm} = 0$ corresponds to no baryonic backreaction and $q_{\rm rdm} = 1$ to perfect conservation of angular momentum. The default value adopted in the PCS21 mocks and used below is $q_{\rm rdm} = 0.68$, which was suggested by ST15 based on the hydrodynamical CDM simulation results of Teyssier et al. (2011).

We refer the reader to section 3.2 of PCS21 for details of the numerical implementation of this scheme, as well as all the underlying scalings of baryonic mass fractions and galaxy sizes with halo properties. Baryonification schemes of this type have been shown to successfully reproduce the small-scale spatial correlation statistics of cosmological hydrodynamical simulations (e.g., Chisari et al. 2018; Aricò et al. 2020).

The spatial distributions of baryons and dark matter produced by the scheme above allow for a calculation of the rotation curve of each mock central galaxy. For HI-bearing galaxies, we focus on the mid-plane of the thin exponential HI disk, which gives a circular velocity contribution $v_{\rm HI}(r)$ satisfying

$$v_{\rm HI}^2(r) = \frac{2 f_{\rm HI} V_{\rm vir}^2}{(h_{\rm HI}/R_{\rm vir})}\, y^2 \left[I_0(y)K_0(y) - I_1(y)K_1(y)\right] , \qquad (1)$$

where $y \equiv r/(2h_{\rm HI})$, $V_{\rm vir} = \sqrt{Gm_{\rm vir}/R_{\rm vir}}$ is the virial velocity and $I_n(y)$ and $K_n(y)$ are modified Bessel functions of the first and second kind, respectively. The rotation curve $v_{\rm rot}(r)$ for each mock galaxy is calculated using equation (11) of PCS21, which can be rewritten as

$$v_{\rm rot}^2(r) = v_{\rm HI}^2(r) + \sum_\alpha \frac{Gm_\alpha(<r)}{r} + \frac{Gm_{\rm rdm}(<r)}{r} , \qquad (2)$$

where the sum runs over $\alpha \in \{{\rm bgas}, {\rm cgal}, {\rm egas}\}$, $m_\alpha(<r)$ is the mass of component α enclosed in radius r and $m_{\rm rdm}(<r)$ is the corresponding mass of the quasi-adiabatically relaxed dark matter component. The rotation curves produced by the default baryonification model adopted by PCS21 have been shown to be in very good agreement with the median and scatter of the radial acceleration relation of low-redshift galaxies (Paranjape & Sheth 2021).

3 MODELLING HI VELOCITY PROFILES

A mock rotation curve, along with an assignment of an 'observed' redshift (see appendix A) and inclination angle to the galaxy, can be used to predict the observed velocity profile of the HI 21 cm emission line in a survey such as ALFALFA. In this section, we describe our methodology to predict the HI velocity profile $S_{\rm HI}(v)$ for each central galaxy, followed by an assessment of its potential as a mass-modelling tool, and a description of our technique for extracting the velocity width W_{50} in realistic observational samples.

3.1 From rotation curves to velocity profiles

The rotation curve of each HI-bearing galaxy can be converted into the observable $S_{\rm HI}(v)$ essentially using geometrical considerations and accounting for the Doppler-shifting of line emission from a differentially rotating system (Gordon 1971; Roberts 1978). We consider a thin HI disk as described in section 2.2, inclined at an angle i relative to the observer's line of sight (such that $i = 0^\circ$ for a face-on disk). We assume the optically thin regime, which is a good approximation for all but nearly edge-on disks. Finally, we assume that the observed HI 21 cm line has an intrinsic Gaussian velocity distribution $p(v)$ (Schulman et al. 1994) with width $\sigma_v \lesssim 10\ {\rm km\,s^{-1}}$ arising from turbulent motions in the disk (Sellwood & Balbus 1999).

The observed flux density $S_{\rm HI}(v)$ in a velocity channel $(v, v+$

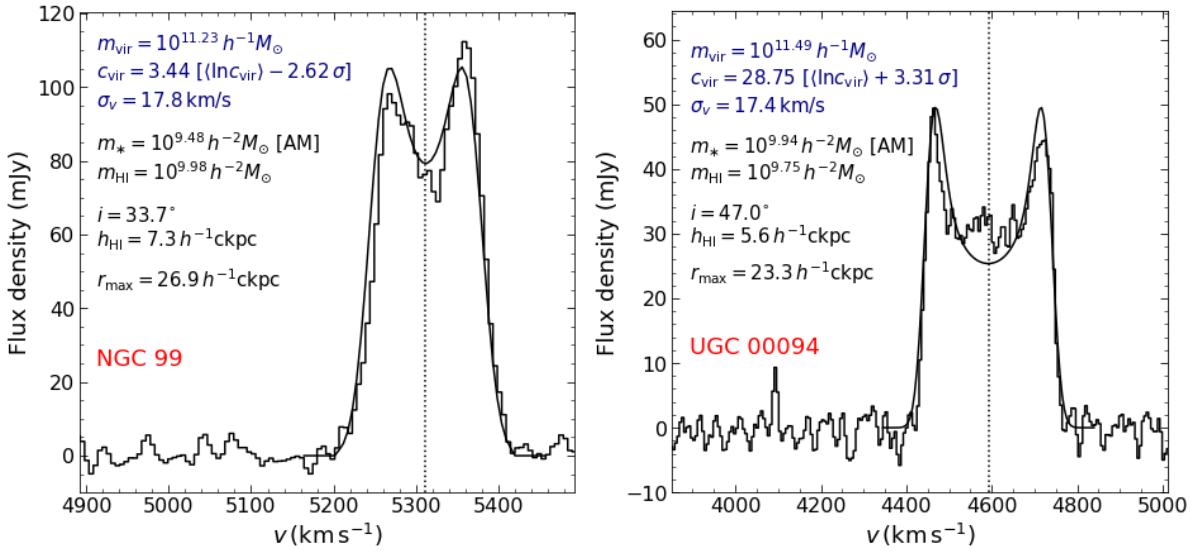

Figure 1. Velocity profile model for two observed galaxies. The stepped black line in the *left (right) panel* shows the observed profile of NGC 99 (UGC 00094) from the ALFALFA source catalog (Haynes et al. 2018), while the smooth black curve shows our best fit model. For each galaxy, the values of inclination i, $m_{\rm HI}$, $h_{\rm HI}$ and $r_{\rm max}$ were fixed as described in the text. The remaining parameters $m_{\rm vir}$, $c_{\rm vir}$ and σ_v were varied in a least squares calculation to obtain the best-fit values (marked in blue), with the stellar mass m_* set by abundance matching (AM) as described in the text. The profiles are centered at the systemic velocities reported by Haynes et al. (2018): $cz = 5312\,(4592)\ {\rm km\,s^{-1}}$ for NGC 99 (UGC 00094) shown as the vertical dotted line in each panel.

$\mathrm{d}v$) then satisfies (Gordon 1971; Schulman et al. 1994)

$$S_{\rm HI}(v) \propto \int \mathrm{d}v'\, p(v-v') \int_0^\pi \mathrm{d}\theta \int \mathrm{d}r_\perp\, 2\, r_\perp \Sigma_{\rm HI}(r_\perp)\, B(r_\perp,\theta) \times \delta_{\rm D}\left(v' - v_{\rm rot}(r_\perp)\sin i\cos\theta\right)$$

$$\approx \int_0^{r_{\rm max}} \mathrm{d}r_\perp\, 2\, r_\perp \Sigma_{\rm HI}(r_\perp) \times \int_0^\pi \mathrm{d}\theta\, p\left(v - v_{\rm rot}(r_\perp)\sin i\cos\theta\right)\,. \qquad (3)$$

Here $\delta_{\rm D}(x)$ is the Dirac delta distribution, $r_\perp$ and θ are the radial distance and azimuthal angle, respectively, in the disk plane and $B(r_\perp,\theta)$ is the telescope beam response converted to the disk reference frame. In the second line, we approximated the beam response as a simple tophat in $r_\perp$. Throughout, we will assume the relation

$$r_{\rm max} = (\theta_{\rm beam}/2) \times D_{\rm A}(z)\,, \qquad (4)$$

where $D_A(z)$ is the angular diameter distance to redshift z and $\theta_{\rm beam}$ is the instrument beam width in radians.

The normalisation of $S_{\rm HI}(v)$ is fixed by relating its integral to the HI mass $m_{\rm HI}$ and luminosity distance $D_{\rm L}$ of the galaxy (Roberts 1975; Giovanelli & Haynes 1988)

$$m_{\rm HI} = 2.356\times10^5\,h^{-2}M_\odot \left(\frac{D_{\rm L}}{h^{-1}{\rm Mpc}}\right)^2 \int \frac{\mathrm{d}v}{\rm km\,s^{-1}} \frac{S_{\rm HI}(v)}{\rm Jy}\,. \qquad (5)$$

We emphasize that $S_{\rm HI}(v)$ is sensitive to the entire matter content of the galaxy's host halo (stars, neutral gas, ionised gas, and their effect on the dark matter), not just the HI disk, through its dependence on the rotation curve $v_{\rm rot}(r)$ in equation (3). Our analysis self-consistently produces the velocity profile and rotation curve for a given baryonic composition of the host halo without, e.g., treating the rotation curve independently of the HI disk. In principle, the model can be made more complex by including the effects of (i) holes and warps in the HI surface density $\Sigma_{\rm HI}$, e.g., by separately modelling a stellar and gas disk, (ii) high velocity clouds (HVCs) modelled by changing the intrinsic velocity distribution $p(v)$ (Schulman et al. 1994) or (iii) a more realistic beam profile $B(r_\perp,\theta)$ (Gordon 1971). We will ignore the first two for simplicity, while the third is unlikely to be relevant for large beams which do not resolve individual galaxies.

3.2 Examples: parameter inference and sensitivity

In this section, we compare the results of numerically integrating the double integral in equation (3) with two example HI velocity profiles of real galaxies, by adjusting some of the model parameters. This allows us to assess the potential of our model as a mass-modelling parameter inference tool, and also explore its sensitivity to various parameters. Although not our primary aim in this work, this exercise will inform our subsequent exploration of the statistical distributions of velocity profile properties.

3.2.1 Modelling NGC 99 and UGC 00094

We consider two galaxies, NGC 99 and UGC 00094, whose velocity profiles we obtain from the ALFALFA source catalog presented by Haynes et al. (2018). (NGC 99 was also modelled using early Arecibo observations by Schulman et al. 1994, see their fig. 2). In each case, we fix the values of $m_{\rm HI}$ and $D_{\rm L}$ using, respectively, the integrated flux from the observed profile and the systemic velocity reported by Haynes et al. (2018). The value of the disk scale length $h_{\rm HI}$ is then fixed using the empirical scaling relation mentioned in section 2.2. We use the inclination reported by Sánchez et al. (2012, their table 1) and Di Teodoro & Fraternali (2014, their table A.1) for NGC

99 and UGC 00094, respectively, and use the Arecibo beam width of $\theta_{\rm beam} \simeq 3.5'$ to set $r_{\rm max}$ using equation (4). We then vary the values of the remaining parameters, namely halo mass $m_{\rm vir}$, halo concentration $c_{\rm vir}$ and intrinsic dispersion σ_v, using $m_{\rm vir}$ to fix the stellar mass m_* using the abundance matching (AM) prescription of Behroozi et al. (2013b), with recalibrated parameter values from Kravtsov et al. (2018). The values of $m_{\rm vir}$ and $c_{\rm vir}$ also fix other baryonification variables such as the stellar bulge size and the mass fractions and profiles of ionised and expelled gas (see section 2.2). We hold the value of the relaxation parameter fixed at the default $q_{\rm rdm} = 0.68$ in this exercise. The *left (right) panel* of Fig. 1 shows the observed and best-fit profile of NGC 99 (UGC 00094), along with the values of various parameters.

For NGC 99 (UGC 00094) the best-fit $m_{\rm vir}$ leads to a cold gas fraction $f_{\rm HI} \simeq 0.106\,(0.035)$ and a gas-to-stellar mass ratio $f_{\rm HI}/f_{\rm cgal} \simeq 4.3\,(0.9)$, implying that NGC 99 is a relatively gas-rich system compared to UGC 00094. The log-concentration for NGC 99 (UGC 00094) is $\sim 2.6\sigma$ lower ($\sim 3.3\sigma$ higher) than the median value for each halo mass (calculated using the calibration of Diemer & Kravtsov 2015). Although these values represent statistically rare fluctuations relative to the ΛCDM expectation, we note that $c_{\rm vir}$ is strongly degenerate with $m_{\rm vir}$ when both are left free as in our case. Almost equally acceptable fits can also be achieved in each case with more reasonable $c_{\rm vir}$ values, by adjusting $m_{\rm vir}$. This degeneracy is difficult to break with HI velocity profiles alone.[4]

The best-fit value of the intrinsic dispersion σ_v of $\sim$ 17-18 $\rm km\,s^{-1}$ in each example is substantially higher than the typical values of 6-10 $\rm km\,s^{-1}$ reported for individual systems using spatially resolved spectroscopy (e.g., Sellwood & Balbus 1999; Stilp et al. 2013). This could be due to unmodelled HVCs along the line-of-sight which can broaden the spatially integrated profile, especially affecting its tails (see, e.g., the discusion in Schulman et al. 1994, who model NGC 99 and other galaxies including HVCs). We have found that σ_v is also degenerate with $m_{\rm vir}$ and $c_{\rm vir}$, so that improving the modelling of HVCs would also, in general, affect their inferred values. For example, fixing $\sigma_v = 14\,\rm km\,s^{-1}$ for NGC 99, as suggested by Schulman et al. (1994), leads to best-fit values of $m_{\rm vir} = 10^{11.03} h^{-1} M_\odot$ and a log-concentration 0.33σ *above* the median.

3.2.2 Sensitivity to parameter variations

To better understand some of the degeneracies discussed above, we next explore the effects on $S_{\rm HI}(v)$ of varying each parameter individually, using our best-fit models for NGC 99 and UGC 00094 as the defaults. We do not attempt to model the obvious asymmetry between the two horns of each observed profile in Fig. 1, commenting on this aspect at the end of the section. We display results for NGC 99 (UGC 00094) in Fig. 2 (Fig. 3) for variations of $m_{\rm vir}$, $c_{\rm vir}$, σ_v, inclination i, $h_{\rm HI}$, m_*, $r_{\rm max}$ and the relaxation parameter $q_{\rm rdm}$, with the default case repeated as the black curve in each panel. For the chosen default parameter values, the model is visibly most sensitive to $h_{\rm HI}$ and i, followed by $m_{\rm vir}$, $c_{\rm vir}$ and m_*, while being less sensitive to σ_v, $q_{\rm rdm}$ and $r_{\rm max}$ (the last is understandable due to the large width of the Arecibo beam in comparison to the sizes of our chosen galaxies).

Most of these trends can be understood by inspecting equation (3). For an exponential surface density $\Sigma_{\rm HI} \propto {\rm e}^{-r_\perp/h_{\rm HI}}$, the locations of the two horns of the velocity profile are determined roughly by the combination $v_{\rm rot}(h_{\rm HI}) \sin i$. For example, increasing (decreasing) the inclination will cause the two horns to go further apart (come closer), making the profile broader (narrower) while keeping its integral fixed (Gordon 1971; Schulman et al. 1994). This is exactly the trend seen in the *upper right-most panels* of Figs. 2 and 3 (see also Fig. B1). Since the effect of $v_{\rm rot}(h_{\rm HI})$ is identical to that of $\sin i$, any variation that increases or decreases $v_{\rm rot}(h_{\rm HI})$ can be understood in the same manner. This is clearly the case for $m_{\rm vir}$ at fixed halo scale radius $r_{\rm s} = R_{\rm vir}/c_{\rm vir}$ *(upper left-most panels)*: changing $m_{\rm vir}$ primarily scales the overall amplitude of $v_{\rm rot}$ by changing $V_{\rm vir} \propto m_{\rm vir}^{1/3}$, apart from other effects due to changes in the various baryonic fractions. Increasing (decreasing) $m_{\rm vir}$ thus has a qualitatively similar effect to increasing (decreasing) $\sin i$. Similar reasoning also explains the trend seen with halo concentration $c_{\rm vir}$ at fixed $m_{\rm vir}$: high-concentration haloes tend to have higher peak rotation curve values, and hence higher $v_{\rm rot}(h_{\rm HI})$, as compared to low-concentration haloes of the same mass (e.g., Navarro et al. 1996), so that variations in $c_{\rm vir}$ are also qualitatively similar to those in $\sin i$.

Variations in stellar mass m_* (*lower middle-right panels* of Figs. 2 and 3) behave very similarly to those in $m_{\rm vir}$ and $c_{\rm vir}$. Increasing (decreasing) m_* affects the rotation curve (2) in two ways: (i) it increases (decreases) the contribution of the stellar profile $m_{\rm cgal}(< r)$ in the inner halo and, consequently, (ii) it leads to a stronger (weaker) contraction of the dark matter profile. Both effects conspire to make the halo more (less) centrally concentrated, thus explaining the trend. (Similar results would be true if we simultaneously varied $m_{\rm HI}$ and $D_{\rm L}$ keeping $m_{\rm HI}/D_{\rm L}^2$ fixed.) And, as expected for the large Arecibo beam, the beam width variable $r_{\rm max}$ has a relatively minor effect, being more prominent for UGC 00094 which is the closer of the two systems.

The remaining three variables explored in Figs. 2 and 3, namely σ_v, $h_{\rm HI}$ and $q_{\rm rdm}$, behave somewhat differently than the others. The intrinsic dispersion σ_v affects the width of the distribution $p(v)$ in equation (3) without changing its mean, so that increasing (decreasing) σ_v makes each horn broader (narrower) without changing its position, as is clearly seen in the *upper middle-right panels* of each figure. We will return to this effect in section 4.2.3. The disk size $h_{\rm HI}$ affects not only the location at which the rotation curve is effectively sampled due to the exponential surface density (as implied by our writing the combination $v_{\rm rot}(h_{\rm HI}) \sin i$ above), but also the shape of the rotation curve itself. For the rotation curves of HI-bearing centrals in our luminosity-complete mock catalog (section 2), we find that increasing $h_{\rm HI}$ for each galaxy while keeping all its other variables fixed tends to *decrease* $v_{\rm rot}(h_{\rm HI})$ on average, and vice-versa. This is consistent with the behaviour seen in the *lower left-most panels*, where increasing (decreasing) $h_{\rm HI}$ has an effect similar to decreasing (increasing) $\sin i$.

The relaxation parameter $q_{\rm rdm}$ controls the amount of contraction or expansion of the dark matter profile due to the baryonic components. This is a novel aspect of our model which has been generally ignored in the mass-modelling litera-

[4] For completeness, we report that using a simple NFW density profile and rotation curve without any baryonic components leads to completely unrealistic solutions; e.g., the inferred $c_{\rm vir}$ is more than 6σ away from the median relation.

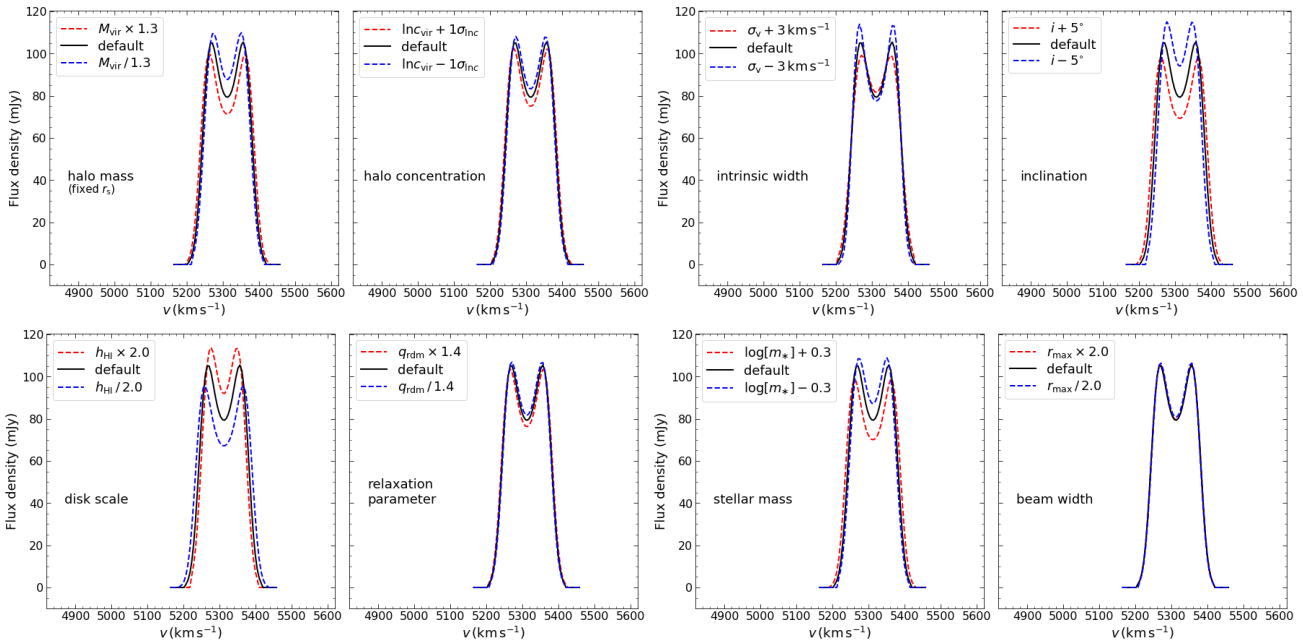

Figure 2. Parameter sensitivity of velocity profile. Each panel shows the result of varying one parameter at a time around the default model of NGC 99 from the left panel of Fig. 1, shown as the solid black curve in each panel. Upward (downward) variations of each parameter are shown as the red (blue) dashed curve in each panel.

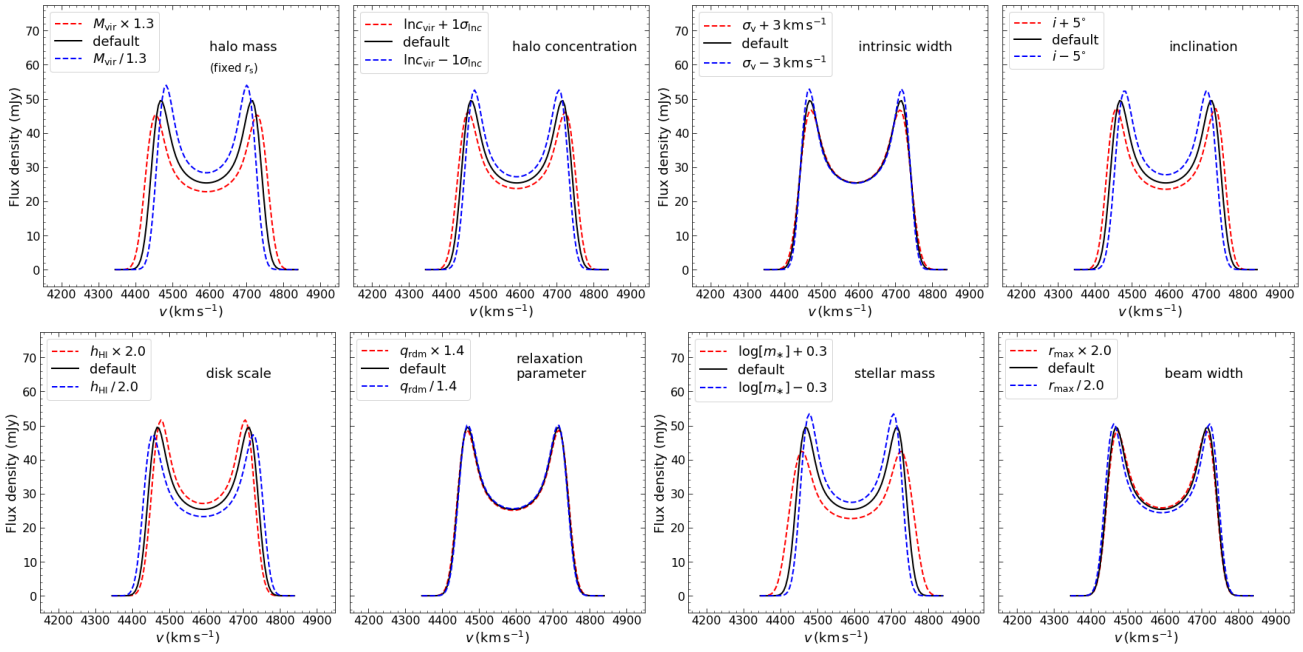

Figure 3. Same as Fig. 2 but using the best fit model for UGC 00094 from the right panel of Fig. 1 as the default.

ture. For the chosen default profiles, varying $q_{\rm rdm}$ has a weak effect *(lower middle-left panels)*, being more noticeable for NGC 99 in Fig. 2. The weakness of the effect follows from the fact that quasi-adiabatic relaxation largely affects the inner halo, while the double-horn structure of the HI profile is more sensitive to the peak or flat part of $v_{\rm rot}$. The trends seen are also sensible: a larger $q_{\rm rdm}$ leads to a stronger contraction of the dark matter profile of each halo, making it more centrally concentrated, so that the effect of $q_{\rm rdm}$ is qualitatively similar to that of $c_{\rm vir}$ (e.g., compare the *lower* and *upper middle-left panels* of Fig. 2).

The reason NGC 99 shows a more prominent effect than UGC 00094 is more subtle, however. The effects of quasi-adiabatic relaxation on the rotation curve at the mass scales of our interest depend on the combination of $m_{\rm vir}$, $c_{\rm vir}$, m_* and $m_{\rm HI}$, along with the spatial extents of the stars and cold gas (see, e.g., fig. 1 of Paranjape & Sheth 2021). To try and disentangle these effects, we varied the values of m_* and $m_{\rm HI}$

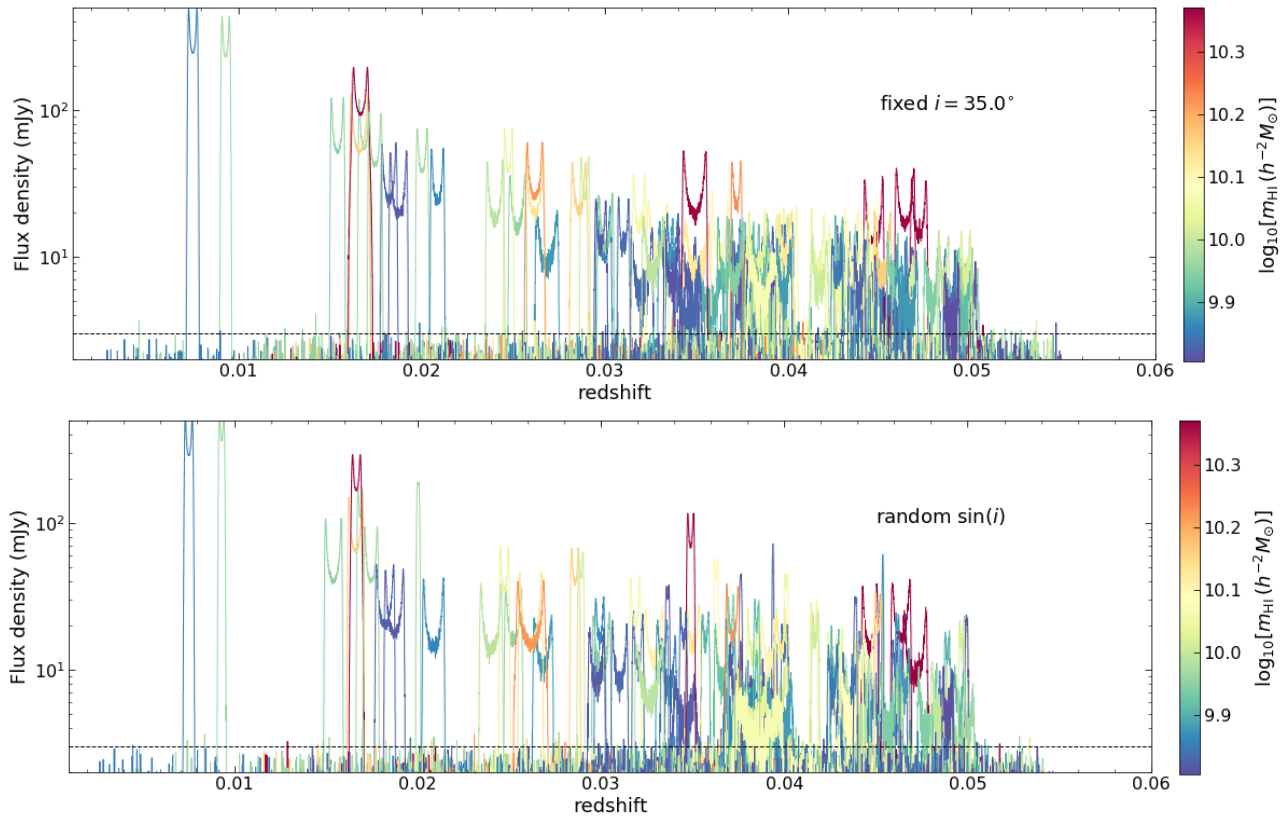

Figure 4. Mock line profiles of HI disks of 100 galaxies randomly chosen from a sample selected to have $M_r \leq -19$, $m_{\rm HI} \geq 10^{9.8}h^{-2}M_\odot$ and $z \leq 0.05$ in a $300h^{-1}$Mpc box with 2 km s^{-1} channel widths. The profiles were generated using our default model as described in section 3, with redshifts assigned as described in appendix A for arbitrary lines of sight and a Gaussian noise of 1 mJy per velocity channel added. The *upper panel* shows results for galaxies observed with a fixed inclination angle of $i = 35^\circ$, while the *lower panel* shows the same galaxies with randomised inclinations ($\sin i$ uniformly sampled in the range $[0, 1)$). Each curve is coloured by the value of HI mass $m_{\rm HI}$ as indicated by the colour bar. With fixed inclination, we clearly see the overall decrease in amplitude due to increasing distance, with only a few high-mass objects jutting out over the envelope. With randomised inclinations, low-mass objects at higher redshift can also have high amplitudes. The horizontal dashed line indicates the 3σ noise level.

independently for each of these examples, producing three sets of curves in addition to those shown in Figs. 2 and 3: one in which $m_{\rm HI}$ is increased by 0.3 dex, one in which m_* is increased by 0.3 dex and one in which both are increased by this amount while keeping their ratio fixed. In each case, we calculated the ratio of the profile widths obtained using the higher and lower values of $q_{\rm rdm}$. For both NGC 99 and UGC 00094, this ratio of widths is most sensitive to changes in m_* at fixed $m_{\rm HI}$, while the other two variations produced almost no effect. However, while the ratio rises from ~ 1.035 to ~ 1.06 for UGC 00094 when m_* is increased, it *decreases* from ~ 1.08 to ~ 1.035 for NGC 99. That is to say, while our model for UGC 00094 becomes more sensitive to $q_{\rm rdm}$ when m_* is increased, the opposite is true for NGC 99. This makes it interesting to ask how changes in $q_{\rm rdm}$ would affect the overall distribution of, say, profile widths for a statistically representative sample; we explore this later in section 4.2.2.[5]

[5] We have also checked that the effect of modifying the bound gas fraction scaling $f_{\rm bgas}(m_{\rm vir})$ is negligible, while only very large ($\gtrsim$ factor 2) variations in the stellar bulge size $R_{\rm hl}$ lead to appreciable changes in the velocity profiles of both NGC 99 and UGC 00094. We will therefore not discuss these two parameters further.

As mentioned previously, our model ignores the asymmetry of the observed profiles. This could easily bias the inferred values of $m_{\rm vir}$ and $c_{\rm vir}$ due to their degeneracy. Asymmetry in observed profiles could arise due to several reasons, from effects such as beam mis-centering for relatively nearby or large galaxies, to physical effects on the galaxy's morphology caused by interactions between the stellar and HI disk or with the environment, particularly in dense regions (see, e.g., Bok et al. 2019; Watts et al. 2020b). This would require making our disk model substantially more complex, with the inclusion of several new parameters. The lack of inherent asymmetry, and the fact that we do not model a stellar disk, also prevents us from testing the AM assumption by independently varying m_*: doing so leads to runaway behaviour, with extremely strong degeneracies appearing between m_* and $m_{\rm vir}$ as expected from Figs. 2 and 3. Similarly, opening up the inclination angle as a free variable also leads to runaway behavior, indicating that knowledge of $\sin i$ for the HI disk is a minimum requirement if our model is to be used for parameter inference. To conclude this discussion, we note that our model produces reasonably realistic descriptions of symmetric profiles, while the modelling of asymmetries is currently challenging.

This machinery can be used to generate 'observed' velocity profiles for our mock galaxies (in which all the parameters

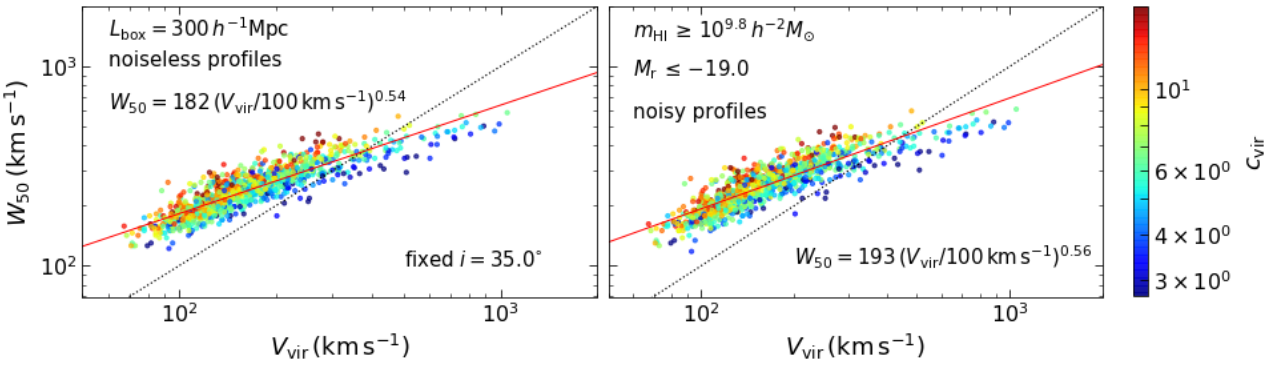

Figure 5. Velocity width (fixed inclination). W_{50} measured from mock HI line profiles of 1000 galaxies randomly chosen from a sample selected to have $M_r \leq -19$, $m_{HI} \geq 10^{9.8}h^{-2}M_\odot$ and $z \leq 0.05$ in a $300h^{-1}$Mpc box, observed with a *fixed* inclination angle $i = 35°$ (see also Fig. 4). Each galaxy is shown as a marker in the plane of W_{50} and halo virial velocity $V_{vir} = \sqrt{GM_{vir}/R_{vir}}$, coloured by the value of halo concentration c_{vir}. The *left (right) panel* shows results for noiseless (noisy) profiles, with noise corresponding to 1 mJy per 2 km s^{-1} channel. There is evidently a tight relation between $\log W_{50}$ and $\log V_{vir}$, quantified by the linear regression shown as the red solid line in each panel with parameters indicated in the labels. The dotted line in each panel shows the one-to-one relation for comparison. The effect of noise is clearly minimal, and the scatter around the mean relation is correlated with c_{vir}: high-concentration haloes at fixed V_{vir} have larger W_{50}.

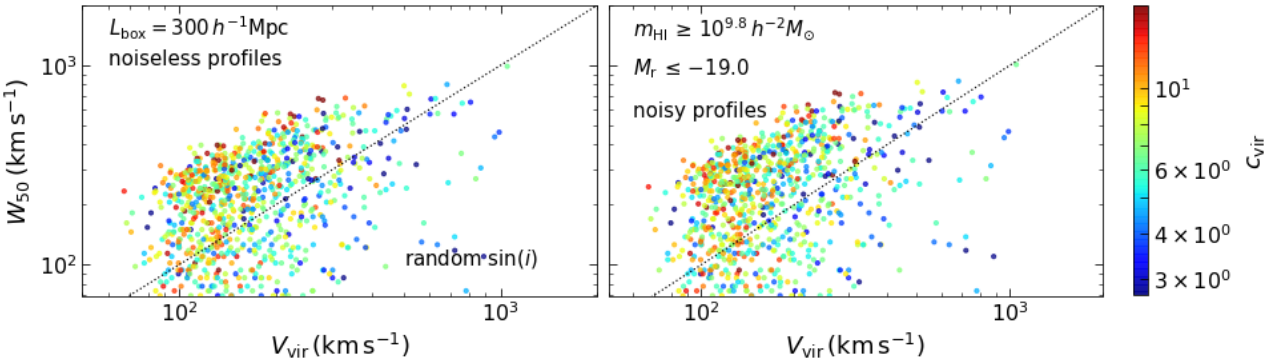

Figure 6. Velocity width (random inclination). Same as Fig. 5, for the same galaxies, now observed with randomised inclination angles. The scatter between W_{50} and V_{vir} is now much broader, and the correlation between W_{50} and halo concentration at fixed V_{vir} is also visibly weaker.

are known) by placing them in redshift space relative to an observer sitting at the center of one face of the simulation box. Appendix A describes our procedure to move galaxies into redshift space and assign them an observed redshift. Fig. 4 shows a sample of noisy velocity profiles of HI-selected galaxies in our default mock catalog, with the *upper panel* showing galaxies observed with a fixed inclination of $i = 35°$ and the *lower panel* showing the same galaxies observed with random inclination angles. For simplicity, we set $\sigma_v = 10$ km s^{-1} for all the objects. The 100 galaxies shown were randomly selected from a sample satisfying $M_r \leq -19$ and $m_{HI} \geq 10^{9.8}h^{-2}M_\odot$. The profiles are coloured by the value of m_{HI}, and we see in the upper panel that only high-mass objects jut out over the envelope of decreasing amplitude as a function of distance. With randomised inclinations, on the other hand, low-mass objects can also be detected with high significance depending on how close to face-on they are viewed. For this example, we used 2 km s^{-1} velocity channels with 1 mJy Gaussian noise added per channel, similar to the Arecibo observations used by Schulman et al. (1994). Below, we will discuss in detail the effects of noise in a realistic survey.

3.3 Velocity widths from velocity profiles

For each observed mock profile, we estimate the velocity width W_{50} using a modified version of the template-matching algorithm described by Saintonge (2007). This technique, which we describe in appendix B, will also be used later when discussing realistic surveys. For the present exercise, we do not smooth the data and also do not place any restriction on signal-to-noise when selecting galaxies. Fig. 5 shows the distribution of W_{50} versus virial velocity $V_{vir} \propto m_{vir}^{1/3}$, coloured by c_{vir}, for a sample of 1000 galaxies observed with a fixed inclination angle of $i = 35°$, with the *left (right) panel* showing results for noiseless (noisy) profiles. We see that W_{50} at fixed i is almost completely determined by V_{vir} and c_{vir}: there is a tight correlation between W_{50} and V_{vir}, with the scatter around the mean relation at fixed V_{vir} itself being quite tightly correlated with c_{vir} (we measure Spearman correlation coefficients between W_{50} and c_{vir} of $\gtrsim 0.2$ in bins of $V_{vir} \gtrsim 100$ km s^{-1}, rising to nearly ~ 0.8 at $V_{vir} \gtrsim 300$ km s^{-1}). The trends seen are also consistent with the m_{vir}-c_{vir} degeneracy discussed earlier in the context of mass-modelling. A comparison between the

two panels shows that the effect of the chosen level of noise is minimal (see also appendix B).

Fig. 6 is formatted identically to Fig. 5 and shows results for the same 1000 galaxies now oriented randomly (i.e., $\sin i$ uniformly sampled as in the bottom panel of Fig. 4). Randomising the inclinations clearly has a substantial effect, with a large scatter between W_{50} and $V_{\rm vir}$ and a correspondingly weaker correlation between W_{50} and $c_{\rm vir}$ at fixed $V_{\rm vir}$ (Spearman correlation coefficients now drop to $\lesssim 0.15$ over nearly the entire range of $V_{\rm vir}$). We return to a discussion of inclination effects in the context of the ability to constrain model variations in section 4.

3.4 Realistic samples

In order to be useful as a probe of small-scale ($\lesssim 10h^{-1}$kpc) physics, it is important that variations in the HI velocity width function be robust to observational systematics and errors. We therefore turn to constructing samples that mimic actual surveys such as ALFALFA (Giovanelli et al. 2005, 2007).

In the context of our mock profiles, this requires (i) setting the velocity channel width Δv and noise per channel $\sigma_{\Delta v}$ to values matching the required survey, (ii) processing the resulting noisy profile of each mock galaxy using a realistic template fitting procedure and (iii) calculating a signal-to-noise ratio S/N. The sample can then be constructed using a threshold on S/N. We use the following method to create an observed catalog of $S_{\rm HI}^{\rm (int)}$ and W_{50} values, with details provided in appendix C.

- For each mock galaxy in the chosen sample, we produce a noise-free HI velocity profile as detailed in section 3.1, with the channel width Δv set by equation (C5) evaluated at $z = 0$.
- We add independent Gaussian noise to each channel with width $\sigma_{\Delta v}$ from equation (C6).
- We smooth each profile using a 3-point Hann filter, which takes value 0.5 at the central channel and value 0.25 at each adjacent channel, being zero thereafter.
- We apply the template-matching procedure of appendix B and estimate W_{50} as the width at half the peak height of the (symmetric) best-fitting template for each noisy, smoothed profile.
- Knowing W_{50}, a spectral extent ΔW is set using equation (C4) and we estimate $S_{\rm HI}^{\rm (int)}$ for each object by integrating the smoothed profile over the range $v \in (-\Delta W/2, \Delta W/2)$ relative to the systemic velocity. We *do not* introduce errors in determining the systemic velocity, instead using the true value as produced by our mock algorithm.
- The S/N is then calculated using equation (C3).

Having generated a set of noisy measurements of $S_{\rm HI}^{\rm (int)}$ and W_{50} from each mock profile, we implement the 2-dimensional step-wise maximum likelihood (2DSWML) technique (e.g., Efstathiou et al. 1988) as described by Martin et al. (2010, see their appendix B) to infer the joint distribution $\phi_{\rm 2d}(\log[m_{\rm HI}], \log[W_{50}])$ of HI mass and velocity width. Briefly, the maximum likelihood solution for the shape of the 2-dimensional density of galaxies ϕ_{mw} in bins of log-mass (labelled by m) and log-width (labelled by w) takes the form

$$\phi_{mw} = N_{mw} \Big/ \sum_i \left(\frac{H_{imw}}{\sum_{m',w'} H_{im'w'} \phi_{m'w'}} \right) , \qquad (6)$$

Here N_{mw} is the observed galaxy count in the 2d bin, $\sum_i$ indicates a sum over all galaxies and H_{imw} is the 'completeness matrix' defined as

$$H_{imw} = \frac{1}{\Delta m \Delta w} \int_{w^-}^{w^+} {\rm d}\tilde{w} \int_{m^-}^{m^+} {\rm d}\tilde{m}\, \mathcal{C}_i(\tilde{m}, \tilde{w}) \, , \qquad (7)$$

where (w^-, w^+) and (m^-, m^+) indicate the bin edges, Δw and Δm are the corresponding bin widths and the completeness function $\mathcal{C}_i(m, w)$ for the redshift of the $i^{\rm th}$ galaxy is unity if the S/N returned by equation (C3) using this redshift and the mass-width pair (m, w) exceeds the chosen threshold $(S/N)_{\rm min}$, and is zero otherwise. In practice, due to our standardised choice of spectral extent for defining S/N, H_{imw} can be written in closed form as a function of $D_{\rm L}(z_i)$, m and w (parametrised by survey-dependent quantities such as channel width and noise r.m.s.). Equation (6) is then iterated to obtain a convergent solution for ϕ_{mw} (we have found that 10 iterations are more than sufficient).

Since equation (6) is insensitive to the normalisation of ϕ_{mw}, this is fixed as follows (appendix B1 of Martin et al. 2010). We first normalise ϕ_{mw} to unity, such that $\Delta m \Delta w \sum_{m,w} \phi_{mw} = 1$. We then estimate the number density of objects in the survey, accounting for survey incompleteness, using

$$n_{\rm sur} = V_{\rm sur}^{-1} \sum_i \frac{1}{\sum_{m,w} H_{imw} \phi_{mw}} \, , \qquad (8)$$

where $V_{\rm sur}$ is the survey volume. Finally, the required 2d number density $\phi_{\rm 2d}(\log[m_{\rm HI}], \log[W_{50}])$ is estimated as the product of $n_{\rm sur}$ and the unit-normalised ϕ_{mw}. Integrating $\phi_{\rm 2d}(\log[m_{\rm HI}], \log[W_{50}])$ over $\log[W_{50}]$ gives the HI mass function, while integrating over $\log[m_{\rm HI}]$ gives the HI velocity width function.

4 RESULTS

In this section, we present the results of our algorithm for our default model as well as a number of variations. In the following, we will use an ALFALFA-like survey configuration selected from the L300_N1024 box by placing the observer at the center of one box face (see appendix A) and selecting galaxies satisfying $z \leq 0.05$ and Dec $\geq 31^\circ$ which gives a survey area of $\sim 10,000\,{\rm deg}^2$ and a volume $\simeq (151h^{-1}{\rm Mpc})^3$. (For comparison, the complete ALFALFA survey covers $\sim 7000\,{\rm deg}^2$ with $z \lesssim 0.05$.) We select central galaxies having optical magnitude $M_r \leq -19$ (this is set by the resolution limit of the simulation box, see PCS21) and $m_{\rm HI} > 0$, which results in $\sim 27,400$ galaxies. As before, we assume a telescope beam width of $\theta_{\rm beam} = 3.5'$ matching the Arecibo value. Also, as in sections 3.1 and 3.3, we use $\sigma_v = 10\,{\rm km\,s^{-1}}$ for all galaxies in our default model.

4.1 Default model

The *left panel* of Fig. 7 shows the observed $S_{\rm HI}^{\rm (int)}$ and W_{50} obtained using the procedure outlined in section 3.4 on our default mock sample. Each marker shows the observation for an individual galaxy and is coloured by the galaxy's redshift. For reference, the black lines show various constant S/N values. We clearly see that low S/N objects preferentially occur at higher redshift, as expected, but otherwise span a wide range of velocity widths. The vertical streaks, particularly apparent

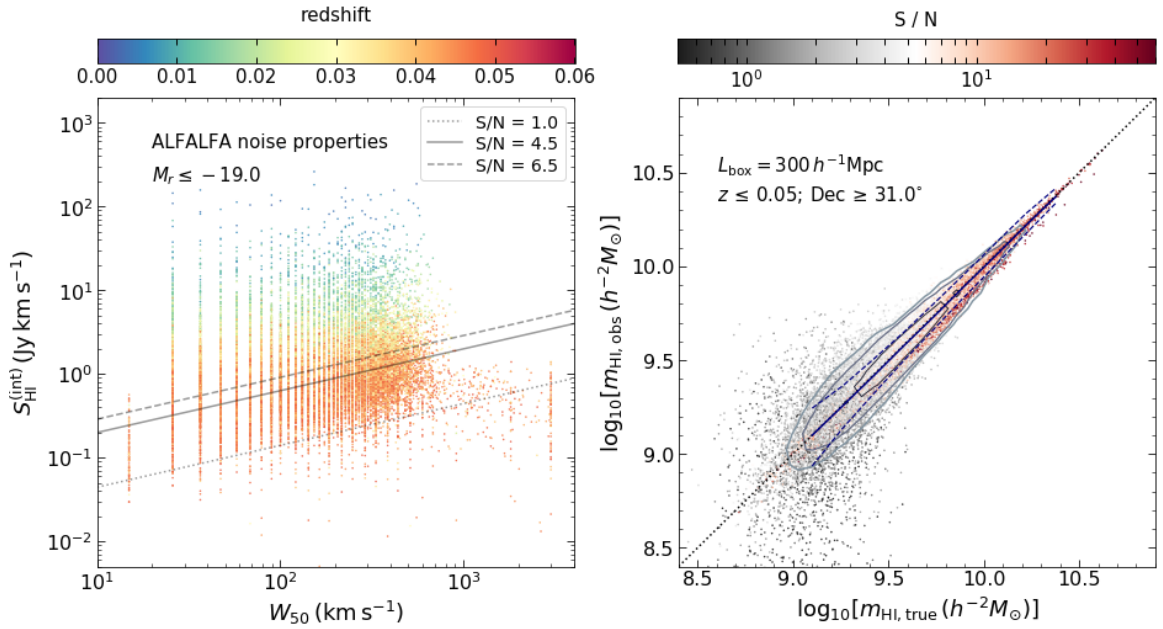

Figure 7. Mock HI survey. *(Left panel:)* Mock observations of HI disks for a central galaxy sample selected by optical luminosity $M_r \leq -19$ and having $m_{\rm HI} > 0$, for the WS survey configuration chosen from a $(300h^{-1}{\rm Mpc})^3$ simulation volume ($z \leq 0.05$ and Dec $\geq 31°$, survey area $\sim 10,000\,{\rm deg}^2$) containing $\sim 27,400$ galaxies. The disks are observed with randomised inclination angles assuming a 3.5′ beam width and velocity channel width and r.m.s. noise appropriate for the ALFALFA survey (see appendix C). Each marker shows the integrated flux density $S_{\rm HI}^{\rm (int)}$ and velocity width W_{50} of an individual galaxy, computed using the algorithm outlined in section 3.4 and coloured by the galaxy redshift (c.f., e.g., fig. 1 of Martin et al. 2010). Solid black line shows the threshold S/N = 4.5 used in the subsequent analysis, calculated using equation (C3). Dashed line shows the S/N threshold of 6.5 used in ALFALFA analyses, while the dotted line shows the value S/N = 1 for reference. *(Right panel:)* Quality of mass recovery as a function of S/N. Markers show the value of $m_{\rm HI}$ inferred from the observations in the left panel using equation (5) against the true $m_{\rm HI}$ value for each galaxy, coloured by the S/N of the observation. Gray lines show contours of equal number, with levels of (100, 200, 800) objects on a grid of pixel width 0.083 dex. Purple solid and dashed lines show the median and central 68% region of $m_{\rm HI,obs}$ in bins of $m_{\rm HI,true}$. Dotted black line shows the 1:1 relation.

at low W_{50}, reflect our choice of velocity channel width of $\Delta v \simeq 5\,{\rm km\,s^{-1}}$ (appendix C).[6] Below, we use the threshold S/N ≥ 4.5 when constructing samples for estimating the velocity width function (for comparison, ALFALFA analyses such as that of Martin et al. 2010, typically use a threshold of 6.5). We have checked that our results for HI abundances below are insensitive to small variations in this choice.

The *right panel* of Fig. 7 compares the HI mass $m_{\rm HI,obs}$ estimated from the observed $S_{\rm HI}^{\rm (int)}$ using equation (5) (replacing the integral on the right hand side with $S_{\rm HI}^{\rm (int)}$) with the true mass $m_{\rm HI,true}$ from the mock catalog. Each marker is coloured by the S/N. We see that large departures from the 1 : 1 relation (dotted black line) occur predominantly at low S/N. This is further quantified by the blue solid and dashed lines, which respectively show the median and central 68% region of $m_{\rm HI,obs}$ in bins of $m_{\rm HI,true}$: the solid line closely follows the 1 : 1 relation while the dashed lines enclose a narrow region at high mass, which broadens towards lower masses where the fraction of low S/N observations is higher.

Fig. 8 shows the 2-dimensional distribution $\phi_{\rm 2d}(\log[m_{\rm HI}], \log[W_{50}])$ estimated from these observations using the 2DSWML method, i.e., after correcting for the incompleteness caused by the S/N ≥ 4.5 threshold. There is a weak but distinct bimodality in the distribution along the W_{50} direction, with a prominent excess around $W_{50} \simeq 250\,{\rm km\,s^{-1}}$ and a somewhat smaller excess near $W_{50} \simeq 30\,{\rm km\,s^{-1}}$,

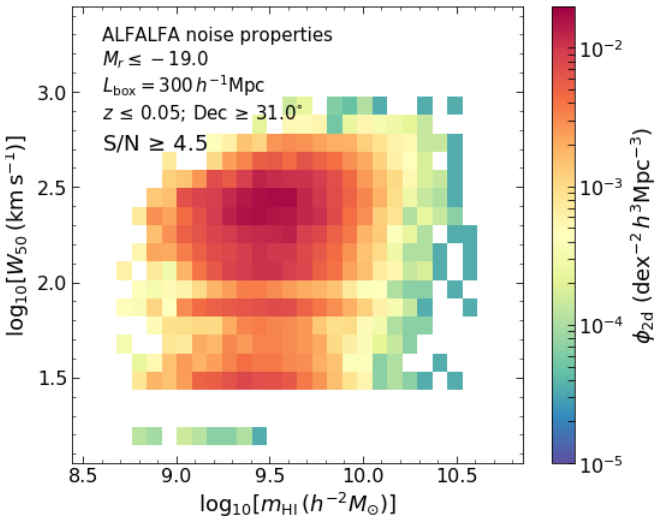

Figure 8. Galaxy abundances. 2-dimensional abundance $\phi_{\rm 2d}(\log[m_{\rm HI}], \log[W_{50}])$ inferred from the data in the left panel of Fig. 7 using the 2DSWML method.

[6] The small clumping of low S/N galaxies near $W_{50} \simeq 3000\,{\rm km\,s^{-1}}$ is due to a numerical choice in our analysis in which we only simulate HI profiles over the range $\pm 1500\,{\rm km\,s^{-1}}$ on either side of the object's systemic velocity.

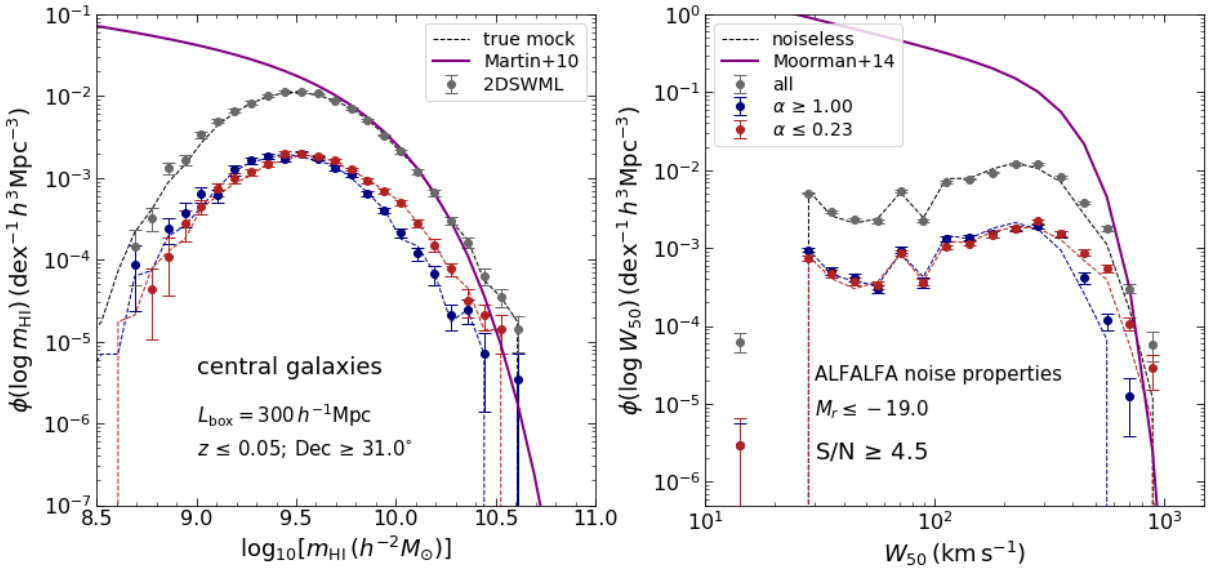

Figure 9. Galaxy abundances. HI mass function *(left panel)* and velocity width function *(right panel)* of central galaxies, calculated as the integral of the 2-d abundance $\phi_{2d}(\log[m_{HI}], \log[W_{50}])$ from Fig. 8 over $\log[W_{50}]$ and $\log[m_{HI}]$, respectively (gray symbols with error bars). Error bars were computed by applying the 2DSWML method separately to 50 bootstrap samples and taking the standard deviation of the resulting abundances. Dotted black lines show the underlying true distributions of m_{HI} and W_{50}, computed by histogramming the true m_{HI} values in the mock and estimates of W_{50} from the noiseless velocity profiles. Solid purple curves show the respective Schechter function fits from Martin et al. (2010) and Moorman et al. (2014) using the α0.4 ALFALFA sample. Blue (red) symbols with errors show the abundances for galaxies chosen to reside in anisotropic (isotropic) tidal environments defined by the tidal anisotropy variable α. Blue (red) dotted curves show the corresponding true distributions calculated similarly to the dotted black curves.

below which the distribution truncates sharply. This feature could be partly due to the incompleteness inherent in our base sample caused by the optical selection of $M_r \le -19$. This systematically misses HI-bearing galaxies progressively smaller than $m_{HI} \lesssim 10^{9.7}h^{-2}M_\odot$ (PCS21; see also below) and cannot be accounted for by the 2DSWML technique. So, e.g., it is possible that the missing galaxies would preferentially occupy widths $W_{50} \sim 10^{1.5}$-$10^2\,\mathrm{km\,s}^{-1}$, thus filling in the decrement between the two maxima. We see, however, that the bimodality in W_{50} persists even when focusing on galaxies with $m_{HI} \ge 10^{9.7}h^{-2}M_\odot$, and also in the absence of noise (not shown), indicating that this may be a genuine feature of the model.

This is explored further in Fig. 9 which shows the integrals over this 2-d distribution to yield the HI mass function *(left panel)* and velocity width function *(right panel)* as the gray points with errors, compared with the respective noiseless distributions in the mock shown as the dotted black lines. This comparison shows that the 2DSWML method accurately recovers the underlying distribution of m_{HI} and W_{50}, except perhaps at the largest W_{50} where the abundance is overestimated compared to the noiseless case, and the smallest W_{50} where some spurious counts are recorded. (The error bars were computed by applying the 2DSWML method to each of 50 bootstrap samples and taking the standard deviation of the resulting 1-d distributions.) To assess the level of incompleteness relative to actual ALFALFA observations, we show Schechter function fits to $\phi(m_{HI})$ and $\phi(W_{50})$ (solid purple curves) as calibrated by Martin et al. (2010) and Moorman et al. (2014), respectively. For $\phi(m_{HI})$, we reproduce the result alluded to above (see Paul et al. 2018, for a detailed discussion) that the HI mass function produced by the PCS21 algorithm is incomplete for $m_{HI} \lesssim 10^{9.7}h^{-2}M_\odot \simeq 1.12M_*$, where $M_* = 10^{9.65}h^{-2}M_\odot$ is the knee of the Martin et al. (2010) Schechter fit to $\phi(m_{HI})$. The distribution of W_{50}, on the other hand, clearly suffers more than that of m_{HI} from this inherent incompleteness of our mocks. We see that $\phi(W_{50})$ is only complete for $W_{50} \gtrsim 700\,\mathrm{km\,s}^{-1} \simeq 1.84W_*$, with $W_* = 380\,\mathrm{km\,s}^{-1}$ being the knee of the Moorman et al. (2014) Schechter fit to $\phi(W_{50})$. The bimodality in the W_{50} distribution mentioned above is apparent, although somewhat suppressed, in the right panel of Fig. 9 where $\phi(W_{50})$ traced out by the gray points shows a shallow minimum around $W_{50} \simeq 50\,\mathrm{km\,s}^{-1}$.

Of course, since our mocks are fundamentally limited by the resolution of the underlying HOD, an apples-to-apples comparison would require comparing them with optically selected subsamples of the ALFALFA survey. Alternatively, one could compare estimates of the conditional distribution $\phi(W_{50}|m_{HI} > 10^{9.7}h^{-2}M_\odot)$ for which our mocks are expected to produce complete results. Another option would be to explore AM techniques to access the low-m_{HI} regime. We leave such comparisons for future work.

4.2 Variations

Our primary motivation in studying HI velocity profiles was to investigate their potential in constraining the baryon-dark matter connection in the ΛCDM framework. The results of section 3.1 suggest that $\phi(W_{50})$ is likely to be sensitive to correlations involving inclination, disk size, halo mass, concentration and, to a lesser extent, the physics of quasi-adiabatic relaxation and the intrinsic width of the HI 21 cm line (see Figs. 2 and 3). In this section, we study the effect of such correlations on the shape of $\phi(W_{50})$.

4.2.1 Sensitivity to environment

All galaxy properties (except σ_v) in our default model are ultimately related to the mass of the host halo through the underlying HOD. Since halo mass correlates with environment, it is worth asking what the model predicts for the environment dependence of the HI observables.

The cosmic web environment of galaxies or their host haloes can be defined in a number of ways. While the large-scale overdensity of dark matter is perhaps the most commonly used discriminator of environment (e.g., Abbas & Sheth 2007; Goh et al. 2019), recent work has emphasized the importance of the local tidal anisotropy in explaining many environmental trends of dark matter haloes (Hahn et al. 2009; Borzyszkowski et al. 2017; Paranjape et al. 2018; Ramakrishnan et al. 2019). The red (blue) markers in Fig. 9 show abundances for galaxy samples selected by low (high) values of the halo-centric tidal anisotropy parameter α, which is inherited by each galaxy from its host halo and is defined at a scale $\sim 4\times$ the host radius R_{200b}. We refer the reader to Paranjape et al. (2018) for a detailed definition of α (see their equation 10) and a description of how it is measured in an N-body simulation, but only note here that values $\alpha \gtrsim 0.5$ correspond to haloes in filamentary environments while $\alpha \lesssim 0.2$ corresponds to node-like environments (which could occur for massive objects at the intersection of large filaments or low-mass, isolated objects in voids). The base sample from which these subsamples are created is the same S/N thresholded set of galaxies used for producing the gray markers in Fig. 9.

Our chosen thresholds $\alpha \leq 0.23$ and $\alpha \geq 1$ lead to subsamples of approximately equal number (~ 2500) before applying the S/N threshold. We see that there is a distinct difference between the two subsamples at both, large $m_{\rm HI}$ and large W_{50}, with the abundance of objects in filamentary environments being suppressed in each case. We can understand this as an effect of halo mass: filamentary haloes with high α tend to span a range of lower halo mass than node-like haloes which exist in all mass ranges, with massive haloes residing almost exclusively in low-α environments (see, e.g., fig. 7 of Paranjape et al. 2018). The suppression of abundances in filamentary environments is then a natural consequence of the correlation between velocity width and halo mass (see Fig. 6). At low $m_{\rm HI}$ and especially at low W_{50}, we see that the environmental cuts leave essentially no imprint on the abundances, apart from the obvious decrease due to reduced overall numbers.

We also repeated this exercise after splitting samples by the value of $\delta_{2h^{-1}{\rm Mpc}}$, the halo-centric dark matter density contrast, smoothed with a Gaussian filter of radius $2h^{-1}$Mpc. Upon choosing high and low thresholds $\delta_{2h^{-1}{\rm Mpc}}$ that give subsamples of approximately the same size as the α-split subsamples (i.e., ~ 2500 objects before applying the S/N threshold), we found that the resulting abundances of galaxies with high (low) $\delta_{2h^{-1}{\rm Mpc}}$ are *quantitatively* very similar to those of galaxies with low (high) α. To avoid clutter, we have not separately shown these results in Fig. 9. This similarity can be understood from the fact that, (i) there is a strong positive correlation between α and $\delta_{2h^{-1}{\rm Mpc}}$ (Spearman correlation of $\simeq 0.62$ for the sample shown by the gray markers in Fig. 9) and (ii) these environmental trends are ultimately derived from halo mass alone in our default model.

It will be very interesting to confront these predictions with corresponding observational results. Recently, Moorman et al. (2014) have reported results for the HI mass function and velocity width function in "void-like" and "wall-like" environments. This environmental classification was based on the void catalog constructed by Pan et al. (2012) which used the Void Finder algorithm of (El-Ad & Piran 1997; Hoyle & Vogeley 2002) in which wall galaxies are first identified based on a nearest neighbour criterion and voids are then constructed by growing empty spheres in the wall-galaxy sample. At $W_{50} \gtrsim 300\,{\rm km\,s^{-1}}$, void-like environments show a suppression in the velocity width function relative to wall-like environments (fig. 9 of Moorman et al. 2014), qualitatively in agreement with the difference between the blue and red points in Fig. 9 which correspond to low- and high-density environments, respectively. As mentioned above, an apples-to-apples comparison would require observational samples selected by optical properties, and also require using the same definitions of environment in both mocks and data, which we defer to future work.

4.2.2 Sensitivity to relaxation physics

As discussed in section 2.2, the quasi-adiabatic relaxation physics of dark matter in each host halo is parametrised by the quantity $q_{\rm rdm}$, whose default value is set to $q_{\rm rdm} = 0.68$. We also saw in section 3.2.2 that changing $q_{\rm rdm}$ has relatively small effects as compared to other variables, but that these effects arise from a complex combination of dark matter and baryonic variables. In this section, we study the predicted effects of these changes on $\phi(W_{50})$.[7]

We have repeated the procedure outlined in section 3.4 for two variations around the default model, setting $q_{\rm rdm} = 0.68 \times 1.4 \simeq 0.95$ in one and $q_{\rm rdm} = 0.68/1.4 \simeq 0.49$ in the other (the same as used in Figs. 2 and 3). The larger value thus represents near-perfect angular mometum conservation, while the lower value is observationally interesting for the radial acceleration relation in the high-acceleration regime (Paranjape & Sheth 2021). Fig. 10 shows the results for the velocity width function for the ALFALFA-like sample. We see that these variations lead to essentially no effect for $W_{50} \lesssim 300\,{\rm km\,s^{-1}}$, while larger widths show small but significant departures from the default model, with the difference between the upward and downward variation in $q_{\rm rdm}$ exceeding $\sim 20\%$ for $W_{50} \gtrsim 500\,{\rm km\,s^{-1}}$ *(bottom panel)*.

In the context of the discussion in section 3.2.2, these trends would be understandable if, at low $m_{\rm vir}$ (and hence low W_{50}), our mock galaxies had stellar masses m_* that were preferentially above the AM relation used in Figs. 2 and 3, while at high $m_{\rm vir}$ (high W_{50}) the mock m_* values were preferentially lower than the AM value. As we saw there, a low-$m_{\rm vir}$ halo with a larger-than-AM m_* would be much less sensitive to $q_{\rm rdm}$ than a high-$m_{\rm vir}$ halo with a lower-than-AM m_*. Indeed, the stellar mass incompleteness induced by our intrinsic luminosity threshold of $M_r \leq -19$ leads to exactly such an effect: fig. 12 of PCS21 shows that galaxies with $m_{\rm vir}$ lower (higher) than $\sim 10^{11.6} h^{-1} M_\odot$ have m_* values preferentially substantially above (slightly below) the AM relation. We conclude that the lack of sensitivity of the width function to $q_{\rm rdm}$ at

[7] The mass function $\phi(m_{\rm HI})$ is, by construction, totally insensitive to $q_{\rm rdm}$ in our model.

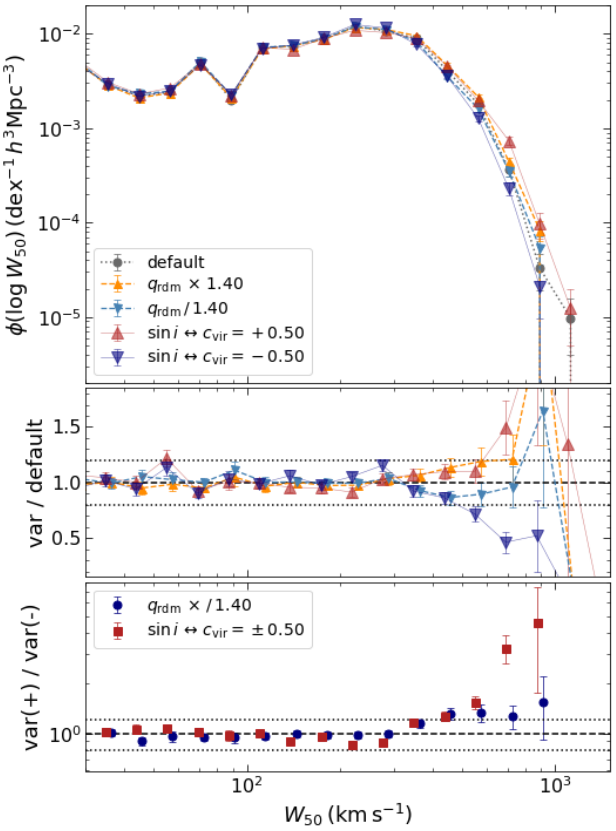

Figure 10. HI velocity width function in alternative models. *(Top panel:)* $\phi(\log[W_{50}])$ for the default model (gray circles joined with dotted line, repeated from the right panel of Fig. 9) compared with results when varying the relaxation parameter $q_{\rm rdm}$ (small triangles joined by dashed lines, described in section 4.2.2), or including a correlation between inclination i and halo concentration $c_{\rm vir}$ (large triangles joined by solid lines, described in section 4.2.5). For each alternative model, upward (downward) variations of the relevant parameter are shown using upward (downward) pointing triangles with warmer (cooler) colours, and used galaxy samples defined identically to the one used for the default model (see Fig. 9). *(Middle panel:)* Ratio of abundances in the alternative models to those in the default model, formatted identically to the *top panel*. *(Bottom panel:)* Ratio of abundances in the upward and downward parameter variations for each alternative model. Circles and squares respectively show the results when varying $q_{\rm rdm}$ and the sin $i \leftrightarrow c_{\rm vir}$ correlation. Horizontal dotted lines in the *middle* and *bottom panels* indicate ±20% deviations around unity (horizontal dashed line). Error bars in the *top panel* were estimated using 50 bootstrap samples for each case, as in Fig. 9, while those in the *middle* and *bottom* panels were estimated using error propagation. For $W_{50} \gtrsim 300\,\mathrm{km\,s^{-1}}$, we see $\sim 20\%$ effects when varying $q_{\rm rdm}$ and up to factor ~ 5 effects when varying the sin $i \leftrightarrow c_{\rm vir}$ correlation. For lower W_{50}, neither of the variations leads to any significant effect. The variations involving σ_v (section 4.2.3) and the correlation $h_{\rm HI} \leftrightarrow c_{\rm vir}$ (section 4.2.4) do not lead to any significant effect and are therefore not shown.

low W_{50} is likely due to the stellar mass incompleteness of our sample.

4.2.3 Correlation between gas surface mass density and intrinsic width

We saw in Figs. 2 and 3 (upper middle-right panels) that the shape of the HI velocity profile responds in a small but distinctive manner to the value of the intrinsic dispersion σ_v. Namely, increasing (decreasing) σ_v makes the individual horns broader (sharper). Since our default model used the constant $\sigma_v = 10\,\mathrm{km\,s^{-1}}$, it is interesting to ask whether variations in σ_v might leave an imprint in $\phi(W_{50})$ or related quantities. Observationally, while early work using small galaxy samples indicated that σ_v is remarkably insensitive to galaxy properties (Sellwood & Balbus 1999), later work has revealed strong correlations between σ_v and variables such as the surface density of HI mass ($\Sigma_{\rm HI}$), of stellar mass ($\Sigma_{\rm cgal}$) or of baryonic mass ($\Sigma_{\rm bary}$) (e.g., Stilp et al. 2013). Such correlations might be connected to the physics of supernova feedback, although this is not a settled question as yet (see, e.g., Utomo et al. 2019; Bacchini et al. 2020).

With this motivation, we have therefore explored the following variations around our default model: (a) setting $\sigma_v = 8\,\mathrm{km\,s^{-1}}$ and (b) setting σ_v as a Gaussian distributed variable with mean $8\,\mathrm{km\,s^{-1}}$ and standard deviation $2\,\mathrm{km\,s^{-1}}$, perfectly correlated or anti-correlated with the surface density $\Sigma_{\rm HI}$. In practice, for variation (b), we note that the HI disk scale $h_{\rm HI}$ in the default model has a lognormal scatter of 0.06 dex around a median value $\langle\, h_{\rm HI}|m_{\rm HI}\,\rangle \propto m_{\rm HI}^{0.5}$ at fixed $m_{\rm HI}$ given by equation 8 of PCS21. Due to this, the surface density $\Sigma_{\rm HI} \sim m_{\rm HI}/h_{\rm HI}^2$ in the default model has a lognormal scatter of 0.12 dex around a value independent of $m_{\rm HI}$, with the scatter in $\log(\Sigma_{\rm HI})$ being perfectly anti-correlated with that in $\log[h_{\rm HI}/\langle\, h_{\rm HI}|m_{\rm HI}\,\rangle]$. To construct variation (b), we therefore write $\sigma_v/(\,\mathrm{km\,s^{-1}}) = 8 \mp 2\epsilon$, where $\epsilon = \log[h_{\rm HI}/\langle\, h_{\rm HI}|m_{\rm HI}\,\rangle]/0.06$ is a standard normal deviate, with the minus (plus) sign leading to a perfect (anti-)correlation $\Sigma_{\rm HI} \leftrightarrow \sigma_v$. The variation (a) tests the model's sensitivity to the absolute value of σ_v, while the variation (b) further tests for the effect of a scatter in σ_v as well as any strong (anti-)correlation with $\Sigma_{\rm HI}$.

We found that $\phi(W_{50})$ for the ALFALFA-like sample shows essentially *no departure* (within errors) from the default model, for *any* of these variations. This is likely due to the fact that the changes we have explored in our σ_v model are comparable to or smaller than the velocity sampling width (equation C5) of an ALFALFA-like survey. To avoid clutter, we have omitted these results from Fig. 10. Thus, while the shapes of individual HI profiles are affected by the value of σ_v, there is no observable imprint on $\phi(W_{50})$. We will see later, however, that beyond-width statistics describing the profile shape are, in principle, sensitive to these variations.

4.2.4 Correlation between disk size and halo concentration

A potential correlation between disk size and halo concentration would be of great interest for galaxy formation models. As discussed by Paranjape & Sheth (2021), a correlation between stellar bulge size and halo concentration, motivated by the size-spin correlations typically predicted by semi-analytical models (Mo et al. 1998; Kravtsov 2013), leads to interesting features in the radial acceleration relation. We have therefore investigated whether a similar correlation between $h_{\rm HI}$ and $c_{\rm vir}$ leads to any effect in $\phi(W_{50})$. We follow Paranjape & Sheth (2021) and

assume that the entire scatter of 0.06 dex around the median $\langle h_{\rm HI}|m_{\rm HI}\rangle$ in the distribution of $h_{\rm HI}$ at fixed $m_{\rm HI}$ is caused by variations in $c_{\rm vir}$, which allows us to write a modified model of disk sizes: $h_{\rm HI} = \langle h_{\rm HI}|m_{\rm HI}\rangle \times (c_{\rm vir}/\langle c_{\rm vir}|m_{\rm vir}\rangle)^{\pm 0.375}$. Here $\langle c_{\rm vir}|m_{\rm vir}\rangle$ is the median concentration at fixed halo mass, and the value of the exponent is fixed by noting that halo concentrations in our model obey a Lognormal distribution with a scatter of 0.16 dex.

Interestingly, despite the strong effects of both $h_{\rm HI}$ and $c_{\rm vir}$ on individual profiles (see Figs. 2 and 3), we found *no significant effect* of this correlation on $\phi(W_{50})$, for either sign of the exponent, for the ALFALFA-like sample. We have checked that this absence of a signature in $\phi(W_{50})$ persists when binning galaxies by inclination (which could, in principle, be estimated from spatially resolved optical spectroscopy). To avoid clutter, we have not shown these results in Fig. 10. This lack of effect is likely due to the strong constraint of a small scatter in $h_{\rm HI}$ at fixed $m_{\rm HI}$, which our model treats as a purely observational input. An explanation of this small scatter in the ΛCDM framework would therefore be an interesting avenue of future research.

4.2.5 Correlation between inclination and halo concentration

The inclination angle of a galaxy relative to the observer is determined by the angular momentum vector of the rotating HI disk, which in turn is expected to correlate with the *halo* angular momentum vector, which further correlates with local environment. Although each correlation in this chain is expected to be weak, this 'intrinsic alignment' effect can, in principle, lead to an indirect correlation between inclination angles and halo properties such as concentration (since the latter also correlates with environment). We can ask whether the distribution of W_{50} is sensitive to the amplitude of such a correlation.

We therefore introduce a correlation between $\sin i$ and $c_{\rm vir}$ (whose distribution is Lognormal, see above) by first drawing a Gaussian random variable $y = a\ln(c_{\rm vir}/\langle c_{\rm vir}/m_{\rm vir}\rangle) + \epsilon$, where ϵ is a standard normal deviate uncorrelated with $c_{\rm vir}$. The values of $\sin i$ are then set by drawing uniform random numbers between zero and unity and rank ordering them according to the values of y. The constant a is fixed so that the Spearman rank correlation coefficient $\sin i \leftrightarrow c_{\rm vir}$ takes some desired value: in the following, we fix $\sin i \leftrightarrow c_{\rm vir} = \pm 0.5$. Although this is large in magnitude compared to what one might expect in reality, it allows us to cleanly study the resulting trends in $\phi(W_{50})$.

We see in Fig. 10 that this variation around the default model again leads to no effect in $\phi(W_{50})$ at $W_{50} \lesssim 300\,{\rm km\,s^{-1}}$, but shows large differences at higher widths. In particular, a positive (negative) correlation between inclination and halo concentration leads to larger (smaller) widths, with a corresponding increase (decrease) in the amplitude of $\phi(W_{50})$. The results in Figs. 2 and 3 show that these trends are sensible.

5 BEYOND-WIDTH STATISTICS: EXCESS KURTOSIS

The qualitative similarity between the effects of a $\sin i \leftrightarrow c_{\rm vir}$ correlation and changes in the relaxation parameter $q_{\rm rdm}$ on the velocity width function make it interesting to study other aspects of the shape of HI velocity profiles. To this end, in this section we study the predicted distribution of the next most interesting shape statistic for symmetric profiles beyond the profile width, namely the *excess kurtosis* κ.[8] We focus on noiseless profiles so as to understand the intrinsic prediction of our default model and the variations discussed above, and comment later on the requirements for measuring κ observationally.

For a noiseless, symmetric velocity profile $S(v)$ which is centered at its systemic velocity, κ can be written as

$$\kappa \equiv c_4/c_2^2 = \langle v^4\rangle / \langle v^2\rangle^2 - 3\,, \tag{9}$$

where $\langle v^n\rangle \equiv \int {\rm d}v\, S(v)\, v^n / \int {\rm d}v\, S(v)$ is the $n^{\rm th}$ moment of the profile and c_n is the $n^{\rm th}$ cumulant. A Gaussian-shaped profile would have $\kappa = 0$ due to the vanishing of all c_n with $n \geq 3$. More generally, the assumption of symmetry and centering mean that $c_1 = 0 = c_3$, so that $c_2 = \langle v^2\rangle$ and $c_4 = \langle v^4\rangle - 3c_2^2$, which leads to the second equality. The expression in equation (9) is equivalent to the usual definition of excess kurtosis as 'kurtosis minus 3', where the kurtosis is defined as the ratio of the fourth central moment to the square of the variance. In general, a non-vanishing κ is a measure of the relative importance of the tails of the profile as compared to a Gaussian shape (Westfall 2014), with $\kappa < 0$ ($\kappa > 0$) indicating that the tails of the distribution are lighter (heavier) than that of a Gaussian.

From equation (3), it is easy to show that the variance of $S_{\rm HI}(v)$ can be written as $\langle v^2\rangle = \sigma_v^2(1 + y^2/2) \approx (W_{50}/2)^2$, while κ takes the form

$$\kappa = -\frac{3}{8}\frac{y^4(2-Y)}{(1+y^2/2)^2}\,, \tag{10}$$

where we defined y and Y as

$$y^2 \equiv \langle v_{\rm rot}^2\rangle (\sin^2 i)/\sigma_v^2\,; \quad Y \equiv \langle v_{\rm rot}^4\rangle / \langle v_{\rm rot}^2\rangle^2\,, \tag{11}$$

with the averages appearing in y and Y being performed over the HI surface density, so that, e.g., $\langle v_{\rm rot}^n\rangle = \int_0^{r_{\rm max}} {\rm d}r_\perp\, r_\perp\, \Sigma_{\rm HI}(r_\perp)\, v_{\rm rot}^n(r_\perp) / \int_0^{r_{\rm max}} {\rm d}r_\perp\, r_\perp\, \Sigma_{\rm HI}(r_\perp)$.

We see that $\kappa < 0$ always, provided $Y < 2$. If the rotation curve $v_{\rm rot}(r_\perp)$ is in its flat part in the region where $r_\perp\, \Sigma_{\rm HI}(r_\perp)$ has its support (i.e., near $r_\perp \simeq h_{\rm HI}$), then $Y \simeq 1$ and κ becomes a function of (W_{50}/σ_v) alone. In general, since we expect $\langle v_{\rm rot}^2\rangle \gg \sigma_v^2$, we will have $y \gg 1$ except for nearly face-on galaxies. In this limit, which is where we expect most galaxies to be, $\kappa \to -(3/2)(2-Y) \simeq -3/2$, independent of inclination and nearly independent of W_{50}. For low-inclination galaxies such that $y \lesssim 1$, $\kappa \propto -y^4(2-Y)$, thus becoming a strong function of both inclination and the intrinsic width $\langle v_{\rm rot}^2\rangle$.

Fig. 11 shows the joint distributions of κ, W_{50} and $\sin i$ for the noiseless profiles in our default model, using the ALFALFA-like mock sample shown in Fig. 7.[9] We see all the trends discussed above. There is a tight and non-linear anti-correlation

[8] For intrinsically asymmetric profiles, the *skewness* derived from the third moment of the profile would also be interesting. Since the skewness vanishes for the symmetric profiles discussed in this work, we do not discuss it here.

[9] We remind the reader that our default model uses $\sigma_v = 10\,{\rm km\,s^{-1}}$ for all galaxies. Also, as in appendix B, W_{50} for each noiseless profile is directly estimated as the width at half its common peak height, without matching to any template.

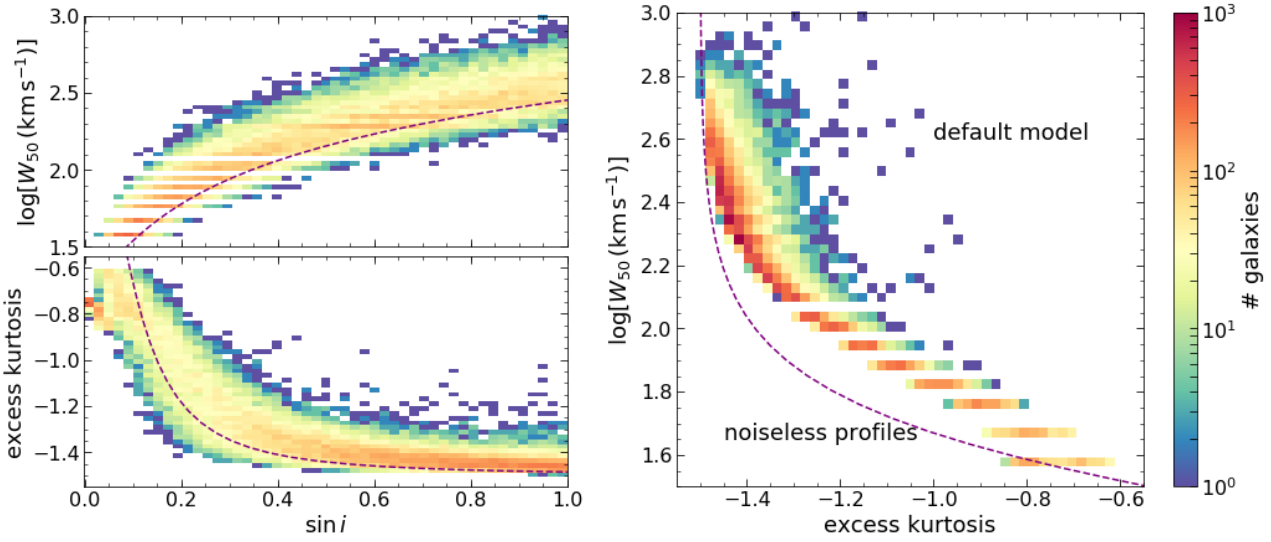

Figure 11. Excess kurtosis, width and inclination. Distributions of inclination $\sin i$ against velocity width W_{50} *(top left panel)* and excess kurtosis κ *(bottom left panel)* and joint distribution of W_{50} and κ *(right panel)* for the *noiseless profiles* generated using the default model for the same galaxies used in Fig. 12. The horizontal streaks along W_{50} are due to the discrete channel width used in sampling each velocity profile (see also appendix B). The blue dashed curves show analytical approximations setting $\langle\, v_{\rm rot}^2 \,\rangle^{1/2} = 200\,{\rm km\,s^{-1}}$ for the *left panels* and $Y = 1$ in equation (10) for the *right panel*. These qualitatively describe the main trends but are quantitatively different; the text discusses the implications of this comparison, especially for the tight relation seen in the *right panel*.

between κ and $\log[W_{50}]$ *(right panel)*, such that most galaxies are found near $\kappa \simeq -1.45$, with a smaller cluster near $\kappa \simeq -0.8$. The dashed purple line shows the approximation $Y \simeq 1$ discussed above. While this broadly traces the κ-W_{50} anti-correlation, it misses most of the distribution and has a different shape. This difference, as well as the scatter in the measured relation, can be attributed entirely to the fact that $Y \neq 1$ for all galaxies in equation (10). The shape and scatter of the measured κ-W_{50} relation, therefore, are potentially sensitive to the physics governing the distribution of Y.

The *left panels* of Fig. 11 show that, as expected, both κ and W_{50} correlate with inclination at low values of $\sin i$, with κ becoming nearly independent of inclination for $\sin i \gtrsim 0.5$. The dashed purple curves in each panel show the prediction if we set $\langle\, v_{\rm rot}^2 \,\rangle = (200\,{\rm km\,s^{-1}})^2$; this clearly provides a reasonable description of the qualitative trends. Since our mocks are incomplete at low W_{50} (see Fig. 9), the structure and position of the $\kappa \simeq -0.8$ cluster of galaxies (which also all occur at the lowest W_{50}) is quite possibly not representative of an HI mass-complete sample, but should rather only be interpreted for an optical luminosity-complete sample with $M_r \leq -19$.

We now ask how sensitive the 1-dimensional κ distribution is to variations around our default model, leaving a detailed study of the κ-W_{50} relation to future work. The *top panel* of Fig. 12 shows the noiseless distribution $p(\kappa)$ for our default model (thick dashed black curve; same as integrating over W_{50} in the right panel of Fig. 11) and variations (coloured lines). The bimodality mentioned above is now readily apparent. The variations around the default model we have explored mostly do not appear to affect $p(\kappa)$ substantially, as seen in the *middle* and *bottom panels* of Fig. 12, except for a clear dependence on the value of σ_v. In particular, the variation (a) from section 4.2.3 ($\sigma_v = 8\,{\rm km\,s^{-1}}$ instead of the default $10\,{\rm km\,s^{-1}}$) leads to a shift in $p(\kappa)$ to lower values, with both modes being somewhat narrower than in the default case. The variation (b) from section 4.2.3 ($\sigma_v = 8 \pm 2\,{\rm km\,s^{-1}}$ with a Gaussian distribution) leads to even more interesting results. In this case, $p(\kappa)$ is identical to that of variation (a) for $\kappa \lesssim -1$, but has a distinctly broader high-κ mode (compare the black dotted line with the red and blue dotted lines for $\kappa \gtrsim -0.9$). Finally, the *sign* of the correlation between σ_v and $\Sigma_{\rm HI}$ does not lead to any noticeable difference (red and blue dotted lines are nearly identical, see also the *bottom panel*). These effects of changing σ_v are all naturally explained by equation (10), keeping in mind that decreasing σ_v will increase y.

Finally, as regards observational estimates of κ, the presence of noise in realistic HI velocity profiles means that the integrals involved in measuring κ in real data must be performed carefully. One approach would be to directly integrate the best-fitting templates obtained using the method outlined in appendix C, provided the template shapes are flexible enough to capture the range of κ seen in the noiseless profiles. The examples shown in Fig. B1 indicate that this would require the inclusion of at least Ψ_4, in addition to Ψ_0 and Ψ_2, in the Hermite function basis set used for building templates. Consequently, the least squares exercise would involve at least one more free parameter. We will explore the feasibility of this exercise, including the minimal requirements on the template basis functions, in future work.

6 CONCLUSIONS

We have studied the distribution of HI velocity profiles as measured by an observer in a ΛCDM universe, which constitutes a hitherto unexplored statistical probe of the small-scale baryon-dark matter connection.

As is well known, the velocity profile of an HI disk as seen by a distant observer can be derived using the galaxy's rotation curve (modulated by its observed inclination angle) and

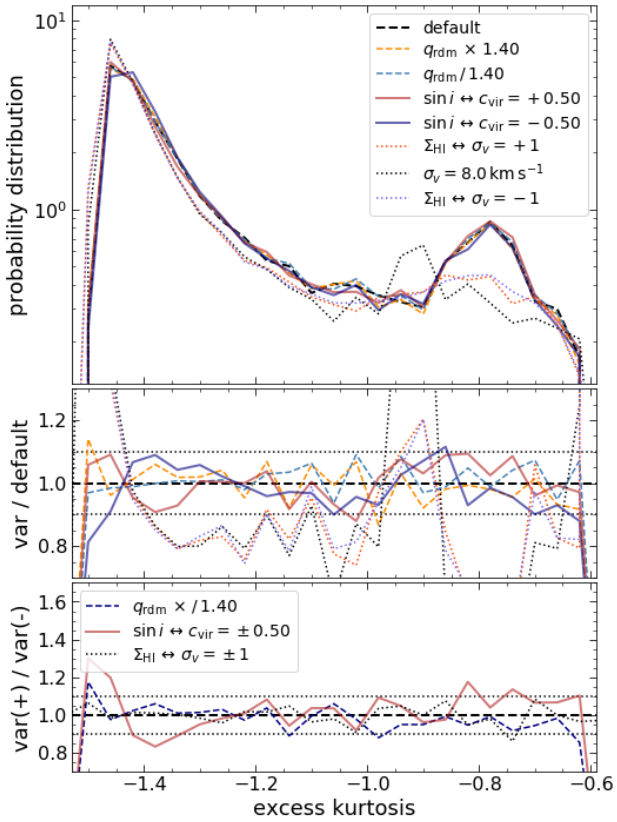

Figure 12. Excess kurtosis distribution. *(Top panel:)* Probability distribution of excess kurtosis κ (equation 9) for *noiseless* HI velocity profiles of all objects in the ALFALFA-like mock sample shown in Fig. 7. Thick dashed black line shows the default model. Solid and thin dashed lines show the variations involving $q_{\rm rdm}$ and a $\sin i \leftrightarrow c_{\rm vir}$ correlation discussed in sections 4.2.2 and 4.2.5, respectively (see also Fig. 10). Dotted lines show the variations involving σ_v discussed in section 4.2.3. These include setting $\sigma_v = 8\,{\rm km\,s^{-1}}$ (dotted black) or distributed as a Gaussian with mean $8\,{\rm km\,s^{-1}}$ and standard deviation $2\,{\rm km\,s^{-1}}$ (red and blue dotted lines), with the red (blue) lines showing the case for perfect (anti-)correlation between σ_v and the surface density of HI mass $\Sigma_{\rm HI}$. *(Middle panel:)* Ratio of the κ distribution for each variation with that in the default model. *(Bottom panel:)* For each variation involving $q_{\rm rdm}$, $\sin i \leftrightarrow c_{\rm vir}$ and $\Sigma_{\rm HI} \leftrightarrow \sigma_v$, the curves show the ratio of the distribution in the upward variation to that in the downward variation. Horizontal dotted lines in the *middle* and *bottom panels* indicate ±10% deviations around unity (horizontal dashed line). We see that only the variations involving σ_v leave any noticeable imprint on the κ distribution, which we discuss further in the main text. Similarly to Fig. 10, the variation involving the correlation $h_{\rm HI} \leftrightarrow c_{\rm vir}$ discussed in section 4.2.4 does not lead to any significant effect and is therefore not shown.

the mass distribution of HI in the disk (e.g., Schulman et al. 1994, see section 3.1). Our analysis applied this calculation to the rotation curves of HI-bearing central galaxies having optical magnitude $M_r \leq -19$ in a statistically realistic mock catalog of galaxies in a $(300h^{-1}{\rm Mpc})^3$ box (Paranjape et al. 2021, hereafter, PCS21) constructed using an optical+HI halo occupation distribution (HOD) model (Paul et al. 2018, 2019, see section 2). The HOD is constrained to reproduce the abundances and luminosity- and colour-dependent clustering of optically selected galaxies in SDSS, as well as the abundances and HI-dependent clustering of massive HI-selected galaxies in the ALFALFA survey. The rotation curves derived from the baryonified host haloes of these central galaxies have been shown to be in very good agreement with the median and scatter of the observed radial acceleration relation in the local Universe (Paranjape & Sheth 2021).

We showed in section 3.2 that, when constrained by observed HI profiles of nearby galaxies, along with knowledge of the disk inclination, our baryonification model produces realistic descriptions of their dark matter and baryonic content. Additionally, our model accounts for the quasi-adiabatic relaxation of dark matter in the presence of baryons in each halo. This suggests that our technique for generating HI disks could be a useful mass-modelling tool, particularly for objects with spatially resolved optical and radio spectra available. Our novel sample of HI velocity profiles, on the other hand (e.g., Fig. 4), and the resulting statistics derived from our mock catalog by 'observing' galaxies in redshift space (appendix A) constitute the first theoretical study of the statistical properties of velocity profiles in a ΛCDM universe.

In addition to our default model for generating rotation curves and velocity profiles, we have explored a number of variations which could, in principle, affect the shapes of HI velocity profiles. These include changing the quasi-adiabatic relaxation physics (section 4.2.2), a correlation between gas surface mass density $\Sigma_{\rm HI}$ and the HI intrinsic velocity dispersion σ_v (section 4.2.3), a correlation between HI disk size and halo concentration (section 4.2.4), and a correlation between galaxy inclination and halo concentration (section 4.2.5).

A commonly used statistic derived from an HI velocity profile is its width W_{50}, which is sensitive to not only the galaxy's inclination but also other physical properties such as host halo mass and concentration (section 3.3, Figs. 5 and 6), as well as baryonic properties such as the HI disk size and intrinsic velocity dispersion (Figs. 2 and 3). Along with the HI mass function $\phi(m_{\rm HI})$ (Zwaan et al. 2005; Martin et al. 2010), the HI velocity width function $\phi(W_{50})$ is a natural product of large-volume surveys of HI-selected galaxies (Papastergis et al. 2011; Moorman et al. 2014), although it is only $\phi(m_{\rm HI})$ which has been typically used for constraining models of galaxy evolution. In order to assess the constraining power of $\phi(W_{50})$, we therefore set up a realistic procedure for estimating W_{50} by template-matching noisy HI velocity profiles measured in an ALFALFA-like survey (appendix B) and consequently estimating $\phi(m_{\rm HI})$ and $\phi(W_{50})$ using the 2DSWML method (section 3.4 and appendix C). Our main results in this regard are as follows.

- For our default model, the 2DSWML method applied to ALFALFA-like noisy data accurately recovers the intrinsic $\phi(m_{\rm HI})$ and $\phi(W_{50})$, except at $W_{50} \gtrsim 400\,{\rm km\,s^{-1}}$ where it overestimates $\phi(W_{50})$ and in the lowest W_{50} bin where it returns a spurious count (Fig. 9; see also Fig. 8).
- Our default model for rotation curves, applied to a luminosity-complete mock catalog of central galaxies with $M_r \leq -19$, leads to an HI mass function that is complete for $m_{\rm HI} \gtrsim 10^{9.7}h^{-2}M_\odot$ (Paul et al. 2018; PCS21) but a velocity function that is complete only for $W_{50} \gtrsim 700\,{\rm km\,s^{-1}} \simeq 1.84W_*$, where W_* is the knee of the observed ALFALFA velocity width function (Fig. 9). As such, all our results should be interpreted for samples that are complete in optical luminosity rather than HI mass.

• The default model (which is based on a 'halo mass only' HOD) predicts distinct differences in $\phi(W_{50})$ for galaxies in tidally anisotropic (underdense) and isotropic (overdense) environments (blue and red symbols, respectively, in Fig. 9; see section 4.2.1 for a discussion).

• Among the variations around the default model mentioned above, the strongest imprints on $\phi(W_{50})$ are seen when introducing a correlation between galaxy inclination and halo concentration, followed by variations in the quasi-adiabatic relaxation physics (Fig. 10). The effects of these variations are, however, degenerate with each other. The remaining variations showed no discernable effects on the $\phi(W_{50})$ for an ALFALFA-like survey.

We have also performed a preliminary study of beyond-width statistics, focusing on the excess kurtosis κ (equation 9) of noiseless profiles of a luminosity-complete sample in an ALFALFA-like survey geometry, which led to the following conclusions.

• The analytical understanding of κ (equation 10) predicts that κ is negative and restricted to values $\gtrsim -1.5$, being a strong function of W_{50} at any inclination, and of $\sin i$ at low inclination. This is borne out by Fig. 12.

• The shape and scatter of the κ-W_{50} relation are predicted to be sensitive to the distribution of the ratio $Y = \langle\, v_{\rm rot}^4 \,\rangle \,/\, \langle\, v_{\rm rot}^2 \,\rangle^2$ of HI-mass-weighted averages of galaxy rotation curves (section 5).

• Among the variations around the default model, it is now the one involving changes in the intrinsic width σ_v which leads to strong effects in the 1-dimensional κ distribution at low inclinations, while the other variations lead to essentially *no* effect (Fig. 12). The response of the κ-W_{50} relation to such variations deserves further attention. The distribution of κ could thus be a sensitive probe of baryonic physics in the turbulent HI disk, provided κ can be robustly estimated from noisy profiles. Independent estimates of the inclination would make such analyses even more sensitive.

We end with a discussion of possible improvements and extensions of our model. Our analysis above was restricted to central galaxies, because it relies on the baryonification scheme described in section 2.2 which has not yet been developed for the (subhalo) hosts of satellite galaxies. Indeed, our mocks do not use subhalo information from the N-body simulation at all, relying instead on empirical models for the spatial distribution and properties of (point-like) satellites (PCS21). Observationally, the clustering of HI-selected galaxies with projected separations $\lesssim 300h^{-1}$kpc does require the inclusion of a small but significant number of HI-bearing satellite galaxies in groups (Guo et al. 2017; Paul et al. 2018). Such satellites are also likely to contain spatially disturbed distributions of HI due to tidal interactions with their dense environments and with other galaxies, possibly leading to preferentially asymmetric HI velocity profiles (Watts et al. 2020b). Tidal interactions would also strip away dark matter from a satellite's subhalo host, while interactions with the hot halo gas in massive groups can affect the star formation properties and gas content of the satellite itself (e.g., van den Bosch et al. 2008). All of these would affect the mass profile and hence rotation curve of the satellite, thus making it imperative to robustly model such effects using, e.g., subhalo demographics from high-resolution N-body experiments (e.g., van den Bosch et al. 2005; Jiang & van den Bosch 2016) along with empirical models for the stellar and HI spatial distribution. The modelling of satellite rotation curves, allowing for asymmetries such as warps in the HI distribution, is therefore a clear direction for future improvements in our model. The modelling of asymmetries in HI velocity profiles is, in general, an interesting avenue of research, although the statistical characterisation of asymmetry in observed samples, along with its connection to galaxy properties, is yet to be settled (see, e.g., Bok et al. 2019; Watts et al. 2020a; Deg et al. 2020; Watts et al. 2021).

Our analysis above also did not fully exploit the spatial distribution of the galaxies in the surrounding cosmic web. It will be interesting to study the predictions of our model for clustering statistics such as mark correlations (Sheth 2005; Skibba et al. 2013) using W_{50} and/or κ as marks. The presence of high-velocity clouds (HVCs) of HI due to substructure in the vicinity of an HI disk, which is currently not included in our model, could alter the shapes of individual HI velocity profiles, particularly in the tails (e.g., Schulman et al. 1994), and possibly also leave an imprint in clustering statistics. More generally, it would be interesting to develop compact summary statistics (e.g., using wavelet transformations) that can capture aspects of an individual HI velocity profile such as the shape of the individual horns, the height between each horn summit and the central trough, etc., which might be sensitive to the underlying baryonic and dark matter variables in different ways and therefore useful in breaking degeneracies.

Finally, weak gravitational lensing leaves a number of interesting signatures on the observed properties of rotating disks. In a spatially resolved galaxy spectrum, the axes along which the radial velocity is zero and maximum are perpendicular to one another if the object is not lensed. The amount by which this angle differs from 90° is a measure of the lensing signal (Blain 2002; Morales 2006). Lensing will also modify the axis lengths of the image (while preserving surface brightness), producing an offset from the Tully-Fisher relation – an effect known as Kinematic Lensing (Huff et al. 2013). These are subtle effects that can be detected with even higher signal-to-noise if other photometric parameters (e.g., colour) are known (Croft et al. 2017). Our mock catalogs contain all the required spectroscopic and photometric information that is required to make realistic estimates of the strength of the expected signal from massive galaxies, simplifying the process of forecasting the constraints that HI surveys may place on the lensing potential (Wittman & Self 2021). We will return to these ideas in future work.

ACKNOWLEDGMENTS

AP thanks Nishikanta Khandai for valuable discussions. The research of AP is supported by the Associateship Scheme of ICTP, Trieste and the Ramanujan Fellowship awarded by the Department of Science and Technology, Government of India. TRC acknowledges support of the Department of Atomic Energy, Government of India, under project no. 12-R&D-TFR-5.02-0700 and the Associateship Scheme of ICTP. This work made extensive use of the open source computing packages

NumPy (Van Der Walt et al. 2011),[10] SciPy (Virtanen et al. 2020),[11] Matplotlib (Hunter 2007),[12] and Jupyter Notebook.[13]

DATA AVAILABILITY

The mock catalogs generated by our algorithm will be shared upon reasonable request to the authors.

REFERENCES

Abbas U., Sheth R. K., 2007, MNRAS, 378, 641
Aricò G., Angulo R. E., Hernández-Monteagudo C., Contreras S., Zennaro M., 2020, arXiv e-prints, p. arXiv:2009.14225
Athanassoula E., Bosma A., Papaioannou S., 1987, A&A, 179, 23
Bacchini C., Fraternali F., Iorio G., Pezzulli G., Marasco A., Nipoti C., 2020, A&A, 641, A70
Barnes D. G., et al., 2001, MNRAS, 322, 486
Begeman K. G., 1989, A&A, 223, 47
Begeman K. G., Broeils A. H., Sanders R. H., 1991, MNRAS, 249, 523
Begum A., Chengalur J. N., 2004, A&A, 413, 525
Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013a, ApJ, 762, 109
Behroozi P. S., Wechsler R. H., Conroy C., 2013b, ApJ, 770, 57
Blain A. W., 2002, ApJ, 570, L51
Blais-Ouellette S., Amram P., Carignan C., 2001, AJ, 121, 1952
Blais-Ouellette S., Amram P., Carignan C., Swaters R., 2004, A&A, 420, 147
Bok J., Blyth S. L., Gilbank D. G., Elson E. C., 2019, MNRAS, 484, 582
Borzyszkowski M., Porciani C., Romano-Díaz E., Garaldi E., 2017, MNRAS, 469, 594
Bundy K., et al., 2015, ApJ, 798, 7
Chisari N. E., et al., 2018, MNRAS, 480, 3962
Croft R. A. C., Freeman P. E., Schuster T. S., Schafer C. M., 2017, MNRAS, 469, 4422
Deg N., Blyth S. L., Hank N., Kruger S., Carignan C., 2020, MNRAS, 495, 1984
Di Teodoro E. M., Fraternali F., 2014, A&A, 567, A68
Diemer B., Kravtsov A. V., 2015, ApJ, 799, 108
Duffy A. R., Meyer M. J., Staveley-Smith L., Bernyk M., Croton D. J., Koribalski B. S., Gerstmann D., Westerlund S., 2012, MNRAS, 426, 3385
Efstathiou G., Ellis R. S., Peterson B. A., 1988, MNRAS, 232, 431
El-Ad H., Piran T., 1997, ApJ, 491, 421
Gentile G., Salucci P., Klein U., Vergani D., Kalberla P., 2004, MNRAS, 351, 903
Giovanelli R., Haynes M. P., 1988, in Kellermann K. I., Verschuur G. L., eds, , Galactic and Extragalactic Radio Astronomy. Springer, pp 522–562
Giovanelli R., et al., 2005, AJ, 130, 2598
Giovanelli R., et al., 2007, AJ, 133, 2569
Goh T., et al., 2019, MNRAS, 483, 2101
Gordon K. J., 1971, ApJ, 169, 235
Granados A., Torres D., Castañeda L., Henao-O. J. L., Vanegas S., 2017, arXiv e-prints, p. arXiv:1705.01665
Guo H., Li C., Zheng Z., Mo H. J., Jing Y. P., Zu Y., Lim S. H., Xu H., 2017, ApJ, 846, 61
Hahn O., Porciani C., Dekel A., Carollo C. M., 2009, MNRAS, 398, 1742
Haynes M. P., et al., 2018, ApJ, 861, 49
Heymans C., et al., 2021, A&A, 646, A140
Holwerda B. W., Blyth S. L., Baker A. J., 2012, in Tuffs R. J., Popescu C. C., eds, Proc. IAU Symp. 284 Vol. 284, The Spectral Energy Distribution of Galaxies - SED 2011. pp 496–499 (arXiv:1109.5605), doi:10.1017/S1743921312009702
Hoyle F., Vogeley M. S., 2002, ApJ, 566, 641
Huff E. M., Krause E., Eifler T., Fang X., George M. R., Schlegel D., 2013, arXiv e-prints, p. arXiv:1311.1489
Hunter J. D., 2007, Computing In Science & Engineering, 9, 90
Jiang F., van den Bosch F. C., 2016, MNRAS, 458, 2848
Komatsu E., et al., 2011, ApJS, 192, 18
Koribalski B. S., et al., 2020, Ap&SS, 365, 118
Kravtsov A. V., 2013, ApJ, 764, L31
Kravtsov A. V., Vikhlinin A. A., Meshcheryakov A. V., 2018, Astronomy Letters, 44, 8
Kurapati S., Chengalur J. N., Kamphuis P., Pustilnik S., 2020, MNRAS, 491, 4993
Laurent P., et al., 2017, J. Cosmology Astropart. Phys., 2017, 017
Lelli F., McGaugh S. S., Schombert J. M., 2016, AJ, 152, 157
Martin A. M., Papastergis E., Giovanelli R., Haynes M. P., Springob C. M., Stierwalt S., 2010, ApJ, 723, 1359
Marulli F., et al., 2013, A&A, 557, A17
McGaugh S. S., Rubin V. C., de Blok W. J. G., 2001, AJ, 122, 2381
McGaugh S. S., Lelli F., Schombert J. M., 2016, Phys. Rev. Lett., 117, 201101
Meyer M. J., et al., 2004, MNRAS, 350, 1195
Mo H. J., Mao S., White S. D. M., 1998, MNRAS, 295, 319
Moorman C. M., Vogeley M. S., Hoyle F., Pan D. C., Haynes M. P., Giovanelli R., 2014, MNRAS, 444, 3559
Morales M. F., 2006, ApJ, 650, L21
Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563
Pan D. C., Vogeley M. S., Hoyle F., Choi Y.-Y., Park C., 2012, MNRAS, 421, 926
Papastergis E., Martin A. M., Giovanelli R., Haynes M. P., 2011, ApJ, 739, 38
Paranjape A., Alam S., 2020, MNRAS, 495, 3233
Paranjape A., Sheth R. K., 2021, arXiv e-prints, p. arXiv:2102.13116
Paranjape A., Hahn O., Sheth R. K., 2018, MNRAS, 476, 3631
Paranjape A., Choudhury T. R., Sheth R. K., 2021, MNRAS, 503, 4147
Paul N., Choudhury T. R., Paranjape A., 2018, MNRAS, 479, 1627
Paul N., Pahwa I., Paranjape A., 2019, MNRAS, 488, 1220
Persic M., Salucci P., Stel F., 1996, MNRAS, 281, 27
Ramakrishnan S., Paranjape A., Hahn O., Sheth R. K., 2019, MNRAS, 489, 2977
Roberts M. S., 1975, in Sandage A., Sandage M., Kristian J., eds, , Galaxies and the Universe. Univ. of Chicago Press, p. 309
Roberts M. S., 1978, AJ, 83, 1026
Saintonge A., 2007, AJ, 133, 2087
Sánchez S. F., et al., 2012, A&A, 546, A2
Schneider A., Teyssier R., 2015, J. Cosmology Astropart. Phys., 2015, 049
Schulman E., Bregman J. N., Roberts M. S., 1994, ApJ, 423, 180
Sellwood J. A., Balbus S. A., 1999, ApJ, 511, 660
Sheth R. K., 2005, MNRAS, 364, 796
Skibba R. A., Sheth R. K., Croton D. J., Muldrew S. I., Abbas U., Pearce F. R., Shattow G. M., 2013, MNRAS, 429, 458
Sofue Y., Rubin V., 2001, ARA&A, 39, 137
Springel V., 2005, MNRAS, 364, 1105
Stilp A. M., Dalcanton J. J., Skillman E., Warren S. R., Ott J., Koribalski B., 2013, ApJ, 773, 88
Teyssier R., Moore B., Martizzi D., Dubois Y., Mayer L., 2011, MNRAS, 414, 195
Utomo D., Blitz L., Falgarone E., 2019, ApJ, 871, 17
Van Der Walt S., Colbert S. C., Varoquaux G., 2011, preprint, (arXiv:1102.1523)
Vikhlinin A., et al., 2009a, ApJ, 692, 1060

[10] http://www.numpy.org
[11] http://www.scipy.org
[12] https://matplotlib.org/
[13] https://jupyter.org

Vikhlinin A., et al., 2009b, ApJ, 692, 1033
Vikram V., et al., 2015, Phys. Rev. D, 92, 022006
Virtanen P., et al., 2020, Nature Methods, 17, 261
Wang J., Koribalski B. S., Serra P., van der Hulst T., Roychowdhury S., Kamphuis P., Chengalur J. N., 2016, MNRAS, 460, 2143
Watts A. B., Catinella B., Cortese L., Power C., 2020a, MNRAS, 492, 3672
Watts A. B., Power C., Catinella B., Cortese L., Stevens A. R. H., 2020b, MNRAS, 499, 5205
Watts A. B., Catinella B., Cortese L., Power C., Ellison S. L., 2021, arXiv e-prints, p. arXiv:2104.05995
Westfall P. H., 2014, The American Statistician, 68, 191
Wittman D., Self M., 2021, ApJ, 908, 34
York D. G., et al., 2000, AJ, 120, 1579
Zehavi I., et al., 2011, ApJ, 736, 59
Zwaan M. A., Meyer M. J., Staveley-Smith L., Webster R. L., 2005, MNRAS, 359, L30
de Almeida Á. O. F., Piattella O. F., Rodrigues D. C., 2016, MNRAS, 462, 2706
de la Torre S., et al., 2011, MNRAS, 412, 825
eBOSS Collaboration et al., 2020, arXiv e-prints, p. arXiv:2007.08991
van den Bosch F. C., Tormen G., Giocoli C., 2005, MNRAS, 359, 1029
van den Bosch F. C., Aquino D., Yang X., Mo H. J., Pasquali A., McIntosh D. H., Weinmann S. M., Kang X., 2008, MNRAS, 387, 79

APPENDIX A: REDSHIFT SPACE

Here we collect some relations that are useful when moving objects into redshift space and for determining observed redshifts based on local positions and velocities. Throughout, we consider a flat FLRW cosmology and assume that peculiar velocities are locally non-relativistic. Below, z will generically denote redshift, (X, Y, Z) will denote comoving Cartesian coordinates centered at the observer and (v_X, v_Y, v_Z) will denote physical peculiar velocities relative to the Cartesian grid.

Consider a source at comoving distance R from the observer, emitting at cosmic time t corresponding to redshift $z = 1/a(t) - 1$ and observed at current epoch t_0. Let the source have a peculiar velocity $v_\parallel$ along the observer's line of sight. Then the light propagation integrals for two pulses separated by one wavelength λ at the source are

$$\begin{aligned} \text{pulse 1}: &\quad \int_0^R \mathrm{d}r = \int_t^{t_0} \frac{c\,\mathrm{d}t}{a(t)} \\ \text{pulse 2}: &\quad \int_0^{R+v_\parallel\,\delta t} \mathrm{d}r = \int_{t+\delta t}^{t_0+\delta t_0} \frac{c\,\mathrm{d}t}{a(t)} \end{aligned} \tag{A1}$$

where $\delta t = \lambda/c$ and $\delta t_0 = \lambda_{\rm obs}/c$, with $\lambda_{\rm obs}$ being the observed wavelength. Straightforward manipulation leads to the 'cosmic Doppler' formula

$$1 + z_{\rm obs} \equiv \frac{\lambda_{\rm obs}}{\lambda} = (1+z)\left(1 + \frac{v_\parallel}{c}\right) \tag{A2}$$

Consider now a cubic, periodic simulation box of comoving length $L_{\rm com}$ at cosmic time $t_{\rm sim}$ or redshift $z_{\rm sim} = 1/a(t_{\rm sim}) - 1$. We wish to assign an 'observed' redshift to a tracer (halo, galaxy, etc.) at a comoving position $\mathbf{r}_{\rm com} = (X, Y, Z)$ with peculiar velocity $\mathbf{v} = (v_X, v_Y, v_Z)$. Let us first do this using the so-called distant observer approximation and later generalise to arbitrary lines of sight.

A1 Distant observer approximation

Assume that the simulation box is sufficiently far from the observer along the Cartesian Z-direction, such that the comoving position vector $\mathbf{r}_{\rm com}$ of any tracer in the box relative to the observer satisfies $\mathbf{r}_{\rm com} = \hat{n}\, r_{\rm com} \approx \hat{Z}\, Z$. In other words, the line of sight $\hat{n}$ to *any* tracer is approximately $\hat{n} = \hat{Z}$.

Let us write $Z = \bar{Z} + \delta Z$, where

$$\bar{Z} \equiv L_{\rm com}/2 + r_{\rm com}(z_{\rm sim}) \equiv L_{\rm com}/2 + \int_0^{z_{\rm sim}} \frac{c\,\mathrm{d}z}{H(z)} \tag{A3}$$

is essentially the comoving distance to redshift $z_{\rm sim}$ in the FLRW geometry, and δZ is the actual comoving position of the tracer along the Z-direction in the simulation box, relative to the box center. We have chosen a convention in which the observer sits on one face of the box if $z_{\rm sim} = 0$. We can then convert δZ into a residual cosmic redshift δz (in the absence of peculiar motion) using

$$\begin{aligned} \delta Z &= \int_{z_{\rm sim}}^{z_{\rm sim}+\delta z} \frac{c\,\mathrm{d}z}{H(z)} - L_{\rm com}/2 \\ \implies \delta z &\approx \frac{(1+z_{\rm sim})\,(\delta Z + L_{\rm com}/2)}{\ell_{\rm H,com}(z_{\rm sim})}\,, \end{aligned} \tag{A4}$$

where the second line assumes that the box size $L_{\rm com}$ is much smaller than the comoving Hubble length

$$\ell_{\rm H,com}(z_{\rm sim}) = (1+z_{\rm sim})\,\ell_{\rm H}(z_{\rm sim}) \equiv \frac{c(1+z_{\rm sim})}{H(z_{\rm sim})}\,. \tag{A5}$$

Using this in the cosmic Doppler formula (A2) gives us the observed redshift of a tracer under the distant observer approximation (with the line of sight along the Z-direction)

$$\begin{aligned} \frac{1+z_{\rm ob}}{1+z_{\rm sim}} &= \left(1 + \frac{(\delta Z + L_{\rm com}/2)}{\ell_{\rm H,com}(z_{\rm sim})}\right)\left(1 + \frac{v_Z}{c}\right) \\ &\approx 1 + \frac{(\delta Z + L_{\rm com}/2)}{\ell_{\rm H,com}(z_{\rm sim})} + \frac{v_Z}{c}\,. \end{aligned} \tag{A6}$$

For a simulation snapshot at $z_{\rm sim} = 0$, this reduces to the familiar formula for the comoving redshift space position $\delta Z_{\rm S}$ along the line of sight: $\delta Z_{\rm S} = c\delta z/H_0 - L_{\rm com}/2 = \delta Z + v_Z/H_0$.

A2 Arbitrary line of sight

For a simulation box whose center is at $(0, 0, \bar{Z})$ relative to the observer, with $\bar{Z}$ (equation A3) not necessarily large, it is straightforward to show that the residual cosmic redshift δz for a tracer at location $(\delta X, \delta Y, \delta Z)$ relative to the box center can be obtained by solving

$$\begin{aligned} \delta R_{\rm S} &\equiv \left(\bar{Z}^2 + 2\bar{Z}\,\delta Z + \delta R^2\right)^{1/2} - \bar{Z} + L_{\rm com}/2 \\ &= \int_{z_{\rm sim}}^{z_{\rm sim}+\delta z} \frac{c\,\mathrm{d}z}{H(z)}\,, \end{aligned} \tag{A7}$$

where $\delta R^2 \equiv \delta X^2 + \delta Y^2 + \delta Z^2$ and the first line defines the redshift space comoving distance residual $\delta R_{\rm S}$. The cosmic Doppler formula then becomes

$$1 + z_{\rm obs} = \Big(1 + z_{\rm sim} + \delta z(\delta R_{\rm S})\Big)\left(1 + \frac{\mathbf{v}\cdot\hat{n}}{c}\right)\,, \tag{A8}$$

which assumes non-relativistic peculiar velocities but *does not* assume a small box. Here $\delta z(\delta R_{\rm S})$ must be obtained by inverting equation (A7) and the line of sight direction $\hat{n}$ is given by

$$\hat{n} = (\delta X, \delta Y, \bar{Z} + \delta Z)/(\bar{Z} - L_{\rm com}/2 + \delta R_{\rm S})\,. \tag{A9}$$

As a limiting case, we can recover the distant observer approximation by setting $\delta X = 0 = \delta Y$, so that $\delta R_\mathrm{S} = \delta Z + L_\mathrm{com}/2$, $\hat{n} = (0,0,1) = \hat{Z}$ and equation (A7) reduces to the first line in equation (A4). Further assuming a small box then leads to equation (A6).

A3 Periodicity

The above did not account for periodic boundary conditions imposed by typical cosmological simulations. For clustering studies which rely on relative distances between multiple tracers, we must also ensure that the periodicity of the simulation box is respected when moving objects into redshift space. This can be done for the general case as follows.

- First, use the value of z_obs from the cosmic Doppler formula (A8) to calculate the new box-centric comoving vector position $\mathbf{x}_\mathrm{S} = r_\mathrm{com}(z_\mathrm{obs})\hat{n} - \bar{Z}\,\hat{Z}$ of the tracer, where $\hat{n}$ is given by equation (A9) and $\bar{Z}$ by equation (A3).
- Replace $\mathbf{x}_\mathrm{S} \to (\mathbf{x}_\mathrm{S} + \mathbf{b})\%L_\mathrm{com} - \mathbf{b}$ where $\mathbf{b} \equiv (1,1,1) \times L_\mathrm{com}/2$, i.e., wrap each coordinate around L_com and maintain the centering around the box center.
- Re-calculate z_obs by inverting the relation $r_\mathrm{com}(z_\mathrm{obs}) = \left\|\mathbf{x}_\mathrm{S} + \bar{Z}\,\hat{Z}\right\|$.

As a consequence, no observed redshift will correspond to an object outside the comoving space of the box. The scheme above also ensures that no object will have a negative redshift. For example, in the distant observer limit with $z_\mathrm{sim} = 0$ and a small box, we have $\hat{n} = \hat{Z}$ and $r_\mathrm{com}(z_\mathrm{obs}) \simeq c\,\delta z/H_0 = \delta Z + v_Z/H_0 + L_\mathrm{com}/2$, which must be wrapped around the Z-axis of the box.

Caution: The scheme above will produce consistent redshift space positions which can be used in clustering studies, but the corresponding *redshifts* themselves do not account for the fact that no two tracers can be more than a comoving distance $\sqrt{3}L_\mathrm{com}/2$ apart in a periodic box. So the redshifts z_obs and redshift space positions $\mathbf{x}_\mathrm{S}$ should not be combined. The values of z_obs would typically be useful in combination with survey selection strategies to assess the impact of selection effects.

APPENDIX B: TEMPLATE-MATCHING

Here we describe a simple algorithm, based on the one presented by Saintonge (2007), for performing a robust least-squares estimate of the width W_{50} (or FWHM) of each observed HI line profile, which is then used for estimating the S/N of the profile, in addition to being an observable in its own right.

Since the detailed shape of the profile is less relevant at this stage, it is useful to build templates using simple functions with well-defined analytical properties. Following Saintonge (2007), we use the first two symmetric, orthogonal Hermite functions $\Psi_0(v;\sigma) \propto \mathrm{e}^{-v^2/(2\sigma^2)}$ and $\Psi_2(v;\sigma) \propto \partial_v^2\Psi_0$ (both analytically normalised such that $\int \mathrm{d}v\,|\Psi_n|^2 = 1$) to define a template

$$t(v;\sigma,\lambda) = \Psi_0(v;\sigma) + \lambda\Psi_2(v;\sigma)\,. \tag{B1}$$

The signal $s(v)$ is then modelled as $s(v) \sim A\,t(v;\sigma,\lambda)$, with the overall amplitude A, width σ and relative amplitude λ being free parameters.[14] Realistic signals require $0 \le \lambda \le \sqrt{2}$, with the lower limit corresponding to the face-on case of a single horn and the upper limit leading to an extreme double horn with zero flux density at $v = 0$.

As discussed by Saintonge (2007), a least-squares analysis of the signal $s(v)$ relative to the template $A\,t(v;\sigma,\lambda)$ shows that the best-fitting amplitude A satisfies $A = \sigma_s\,c/\sigma_t$, where σ_s^2 and σ_t^2 are the signal and template variance and c is their correlation coefficient. Using this, the χ^2 reduces to $\chi^2 \propto \sigma_s^2(1-c^2)$, so that minimising χ^2 is equivalent to maximising c. We therefore perform a 2-dimensional maximisation of $c(\sigma,\lambda)$ for a given signal. In detail, we first search for the location of the maximum on a 2-dimensional grid in $(\log\sigma,\lambda)$ and then refine this estimate using a 5-point interpolation assuming that $c(\log\sigma,\lambda)$ can be approximated by a bi-variate quadratic form in the vicinity of its maximum. W_{50} is then estimated as the full width of the best-fitting template $t(v;\sigma,\lambda)$ at half of its (common) peak height, with the template being evaluated on the array of velocity channels for the given survey.

We have checked that this technique accurately recovers the full shape of an injected signal (after adding Gaussian noise), relatively independently of the noise level, when the signal itself is chosen to be one of the templates. Turning to the recovery of more realistic signals, Fig. B1 shows the performance of this technique on four injected signals (thick solid black curves) derived from our mock catalog. The *top (bottom) row* used a galaxy with $\log[m_\mathrm{HI}/h^{-2}M_\odot] \simeq 9.7\,(10.8)$ placed at $z \simeq 0.04$ (luminosity distance $D_\mathrm{L} = 120h^{-1}$Mpc) and viewed at an inclination $i = 12^\circ$ *(left panels)* and $i = 65^\circ$ *(right panels)*. For each noiseless profile calculated using equations (3) and (5), we generate three noisy profiles by adding Gaussian noise using values of the per-pixel r.m.s. $\sigma_{\Delta v} = 1, 3, 10$ mJy, assuming a channel width $\Delta v = 5.3\,\mathrm{km\,s^{-1}}$ (which is appropriate for an ALFALFA-like survey, see appendix C). These examples therefore allow us to explore the effects of inclination as well as overall S/N on the recovery of W_{50}.

For most of these cases, it is visually apparent that the best-fitting templates do not exactly match the detailed shape of the input profile, which is not surprising since they are limited by the shapes of the two Hermite functions. Nevertheless, the recovered values of W_{50} differ from the true value $W_{50,\mathrm{true}}$ by $\lesssim 0.5\%$ at low inclination and high S/N. We estimate $W_{50,\mathrm{true}}$ as the full width at half of the (common) peak height of each noiseless profile evaluated on the same discrete velocity channels as the noisy profiles. Indeed, inclination plays a dominant role in causing a systematic difference between W_{50} and $W_{50,\mathrm{true}}$, with a $\sim 10\%$ overestimate at high inclinations (nearly edge-on galaxies). This is also not surprising, since edge-on disks have sharp peaks in their line profiles separated by a long, flat portion, which cannot be accurately captured by a linear combination of Ψ_0 and Ψ_2 alone. At large $\sigma_{\Delta v}$, the effect of noise becomes more apparent, especially when combined with a lower signal strength (low m_HI and/or large D_L). We now see larger variations in $W_{50}/W_{50,\mathrm{true}}$ for both high and low inclinations. Fig. B2 shows that, for the ALFALFA-like sample from Fig. 7, the recovery of W_{50} is

[14] Strictly speaking, one should include a fourth parameter δ to capture the unknown redshift of the galaxy, and model the signal as $s(v) \sim A\,t(v-\delta;\sigma,\lambda)$. For simplicity, we will assume perfect knowledge of each redshift and center all profiles at $\delta = 0$.

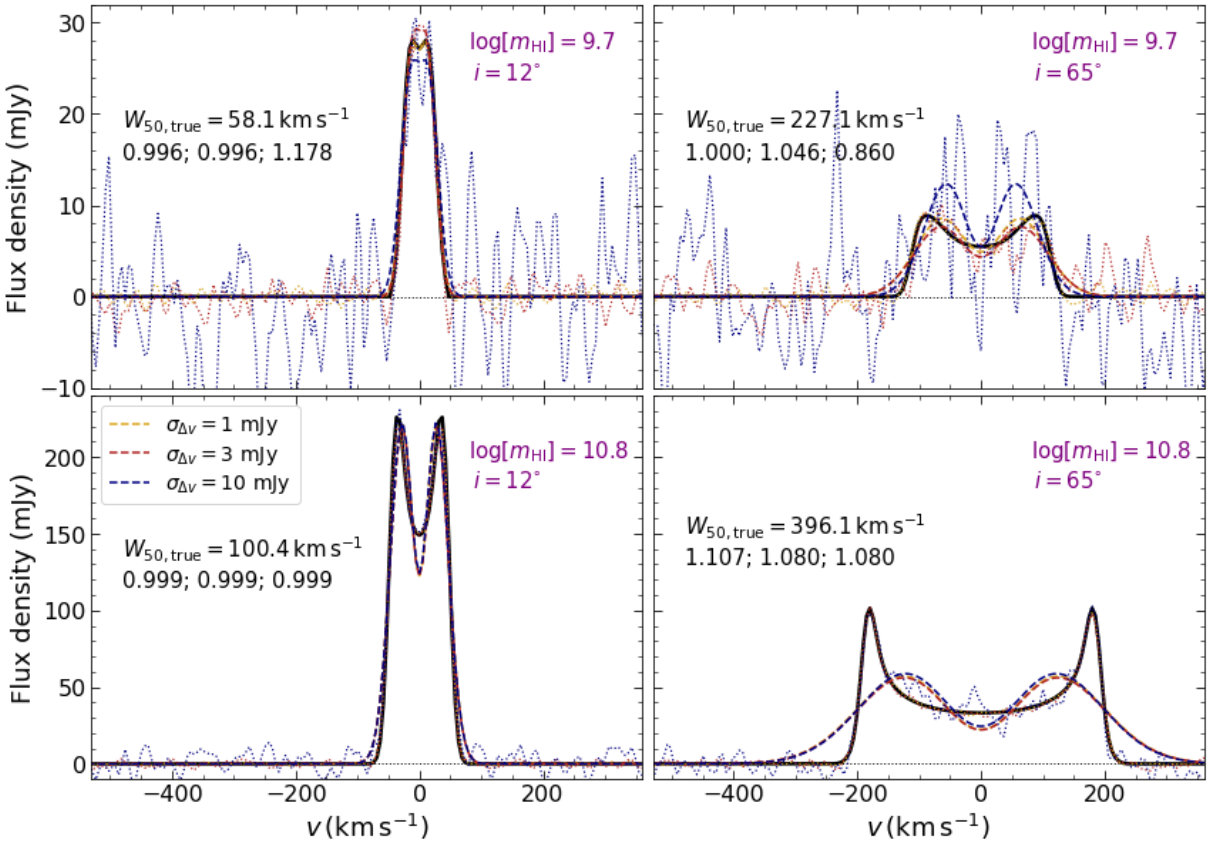

Figure B1. Template-matching performance. The thick solid black curves in the *top (bottom) row* show the noiseless HI velocity profile of a realistic mock galaxy with $\log[m_{\rm HI}/(h^{-2}M_\odot)] = 9.7\,(10.8)$, placed at a luminosity distance $D_{\rm L} = 100h^{-1}$Mpc from the observer and viewed at an inclination of $i = 12°$ *(left panels)* and $i = 65°$ *(right panels)*. The true FWHM $W_{50,\rm true}$ in each case is indicated as a label. Dotted coloured curves show the corresponding noisy profiles with $\sigma_{\Delta v}$ as indicated in the legend of the bottom left panel. Dashed coloured curves show the corresponding best-fit template. The ratios of the corresponding W_{50} with $W_{50,\rm true}$ are displayed as text labels, in order of increasing $\sigma_{\Delta v}$ from left to right. Overall, inclination plays a dominant role in causing a systematic offset in the recovered W_{50} ($\sim 10\%$ overestimate at low inclinations), with the effects of noise becoming important at low S/N (large $\sigma_{\Delta v}$ coupled with low mass and/or large $D_{\rm L}$; see blue dashed curve in the *top right panel*).

essentially perfect at $i \lesssim 23.5°$, while higher inclinations lead to the $\lesssim 10\%$ offset discussed above.

Overall, these examples show that the template-matching technique described above leads to a reasonably robust recovery ($\sim 10\%$ systematic error) of W_{50} for all but the lowest S/N objects. The main text quantifies this further, showing that the mass $m_{\rm HI}$ inferred from each profile using its estimated W_{50} deviates substantially from the true mass only at low S/N (see Fig. 7).

APPENDIX C: SIGNAL-TO-NOISE

As described by Giovanelli et al. (2007), the ALFALFA signal extraction pipeline detailed in Saintonge (2007) first uses a least-squares template-matching method to produce an initial catalog, with signal-to-noise (S/N) values for each candidate detection determined using the matched templates. A cut is imposed on these S/N values and each object surviving this cut is then visually inspected and processed further. Properties including the velocity width W_{50}, integrated flux density $S_{\rm HI}^{\rm (int)}$ and consequently a S/N ratio depending on

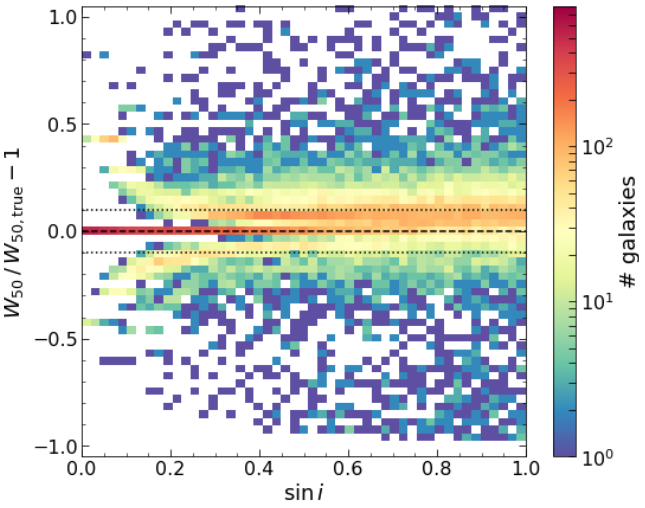

Figure B2. Quality of W_{50} recovery. Joint distribution of inclination $\sin i$ and relative difference between estimated and true W_{50} for the ALFALFA-like mock sample shown in Fig. 7. We see that low inclinations ($\sin i \lesssim 0.4$) lead to essentially perfect recovery while higher inclinations lead to a $\sim 10\%$ overestimate of W_{50}.

these (e.g., equation 16 of Saintonge 2007, see also below) are calculated.

In particular, the integrated flux density is extracted over the 'spectral extent' of the signal, which involves a subjective choice for each object (see section 5 of Giovanelli et al. 2007). The initial use of template-matching, as well as the subjective choice of integration range involved in estimating the integrated flux density, leads to a specific relation between $S_{\rm HI}^{\rm (int)}$ and W_{50} for objects near the threshold of detection, which changes behaviour for $W_{50} \gtrsim 400\ {\rm km\,s^{-1}}$ and is discussed in detail by Giovanelli et al. (2007) and Martin et al. (2010). To simplify our analysis while still keeping it realistic, we do use the template-matching technique described in appendix B, but choose to standardise the choice of integration range in estimating the integrated flux density. We also examine the effects of this standardisation on the statistics of our interest.

The calculation of S/N requires fixing an integration range of length ΔW (in ${\rm km\,s^{-1}}$) for the measured velocity profile $S_{\rm HI}(v)$, which we assume to be centered on the systemic velocity $c\,z$ of the galaxy. The integrated flux can then be approximated by

$$S_{\rm HI}^{\rm (int)} \simeq \Delta v \sum_{v=-\Delta W/2}^{\Delta W/2} S_{\rm HI}(v)\,, \tag{C1}$$

whose measurement error is

$$\sigma_S = \sqrt{\Delta v\,\Delta W}\,\sigma_{\Delta v}\,. \tag{C2}$$

Similarly to Saintonge (2007), we define the S/N as being based on one half of the signal, so that

$$\mathrm{S/N} \equiv \frac{S_{\rm HI}^{\rm (int)}}{\sqrt{2}\,\sigma_S} = \frac{\left(S_{\rm HI}^{\rm (int)}/\Delta W\right)}{\sigma_{\Delta v}} \left(\frac{\Delta W/2}{\Delta v}\right)^{1/2}. \tag{C3}$$

The second equality highlights that the S/N is the ratio of mean flux density over the signal extent to the r.m.s. noise per velocity channel, scaled up by the square-root of the number of independent channels available in half the signal width. In order to standardise the integration range ΔW and avoid subjective choices, in the following we will assume

$$\Delta W = 1.4 \times W_{50}\,, \tag{C4}$$

with the assumption that the profile will typically contribute only noise in channels with $|v| \gtrsim 1.4(W_{50}/2)$ relative to the central velocity. The value of the prefactor is a compromise between maximising S/N and minimising the bias in the recovery of $m_{\rm HI}$; small values of the prefactor will tend to systematically underestimate $m_{\rm HI}$, while large values will integrate over noise and degrade the S/N. We have checked that small variations of the prefactor (values between 1.3 to 1.6) do not affect our results. Larger variations (values of, say 1.0 or 1.8) lead to a biased inference of the $m_{\rm HI}$ and W_{50} abundances relative to the noise-free case, with the bias being relatively insensitive to the chosen S/N threshold. We therefore use equation (C4) as our default choice in the entire analysis.

The frequency resolution of the ALFALFA observations prior to spectral smoothing is $\Delta\nu = 25\,{\rm kHz}$ (Giovanelli et al. 2007). Using $\Delta\nu/\nu_0 = \Delta v/c$ with $\nu_0 = 1420\,{\rm MHz}\,(1+z)^{-1}$ gives us a channel width

$$\Delta v \simeq 5.3\ {\rm km\,s^{-1}}\,(1+z)\,. \tag{C5}$$

Spectra are smoothed with a 3-point Hann filter (Saintonge 2007). This effectively degrades the spectral resolution to $\simeq 10\ {\rm km\,s^{-1}}$ at $z \simeq 0$, but does not drastically affect equation (C3) for the S/N, so we will continue to use that relation in the following. The noise properties of the ALFALFA data cubes *after* Hann smoothing give an r.m.s. $\sigma_{\rm rms} \simeq 2.23\,{\rm mJy}$ (see fig. 4 of Saintonge 2007), which implies a pre-smoothing value of the per-pixel width $\sigma_{\Delta v}$ of

$$\sigma_{\Delta v} = \sqrt{8/3}\,\sigma_{\rm rms} \simeq 3.64\,{\rm mJy}\,, \tag{C6}$$

which we use in our analysis.