The neutral hydrogen content of galaxies in cosmological hydrodynamic simulations

Romeel Davé,[1,2,3,4]★ Neal Katz,[5] Benjamin D. Oppenheimer,[6] Juna A. Kollmeier[7] and David H. Weinberg[8]

[1]*University of the Western Cape, Bellville, Cape Town 7535, South Africa*
[2]*South African Astronomical Observatories, Observatory, Cape Town 7925, South Africa*
[3]*African Institute for Mathematical Sciences, Muizenberg, Cape Town 7945, South Africa*
[4]*Astronomy Department, University of Arizona, Tucson, AZ 85721, USA*
[5]*Astronomy Department, University of Massachusetts, Amherst, MA 01003, USA*
[6]*Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, the Netherlands*
[7]*Observatories of the Carnegie Institute of Washington, Pasadena, CA 91101, USA*
[8]*Astronomy Department and CCAPP, Ohio State University, Columbus, OH 43210, USA*

ABSTRACT

We examine the global H I properties of galaxies in quarter billion particle cosmological simulations using GADGET-2, focusing on how galactic outflows impact H I content. We consider four outflow models, including a new one (ezw) motivated by recent interstellar medium simulations in which the wind speed and mass loading factor scale as expected for momentum-driven outflows for larger galaxies and energy-driven outflows for dwarfs ($\sigma < 75\,\mathrm{km\,s^{-1}}$). To obtain predicted H I masses, we employ a simple but effective local correction for particle self-shielding and an observationally constrained transition from neutral to molecular hydrogen. Our ezw simulation produces an H I mass function whose faint-end slope of -1.3 agrees well with observations from the Arecibo Fast Legacy ALFA survey; other models agree less well. Satellite galaxies have a bimodal distribution in H I fraction versus halo mass, with smaller satellites and/or those in larger haloes more often being H I deficient. At a given stellar mass, H I content correlates with the star formation rate and inversely correlates with metallicity, as expected if driven by stochasticity in the accretion rate. To higher redshifts, massive H I galaxies disappear and the mass function steepens. The global cosmic H I density conspires to remain fairly constant from $z \sim 5 \to 0$, but the relative contribution from smaller galaxies increases with redshift.

Key words: ISM: atoms – galaxies: abundances – galaxies: evolution – galaxies: formation – galaxies: ISM.

1 INTRODUCTION

The current paradigm for galaxy evolution rests on the tenet that gas flows into and out of galaxies are primarily responsible for governing galaxy growth (e.g. Kereš et al. 2005; Dekel et al. 2009; Davé, Finlator & Oppenheimer 2012, and references therein). Such accretion and outflow processes are often difficult to detect directly owing to the tenuous and multiphase nature of the gas, but they are expected to leave clear imprints on the stellar and gaseous content of galaxies. A key region where such baryon cycling processes might be probed is in the outskirts of star-forming galaxies, where reservoirs of neutral hydrogen (H I) are expected to hold the raw material that ultimately fuels star formation (SF). This reservoir is thought to be continually replenished by gravitationally driven accretion from the highly ionized intergalactic medium (IGM), possibly augmented by the re-accretion of ejected gas.

Observations of the neutral hydrogen content of galaxies using 21 cm fine structure line emission are progressing rapidly, and promise to improve further. The H I Parkes All-Sky Survey (Meyer et al. 2004) provided a comprehensive deep census of the H I content of low-redshift galaxies. This was followed by the Arecibo Fast Legacy ALFA (ALFALFA) survey (Giovanelli et al. 2005), which surveyed 7000 deg^2 for H I 21 cm emitting objects. More recently, the *GALEX* Arecibo SDSS Survey (GASS; Catinella et al. 2010) has obtained high-fidelity H I data on optically selected galaxies, to provide a more H I-unbiased census. Upgrades to the Jansky Very Large Array (JVLA) promise further improvements. Construction is underway for the MeerKAT telescope, which will provide the ability to detect H I in galaxies out to unprecedented redshifts

★ E-mail: rad@as.arizona.edu

[e.g. the Looking At the Distant Universe with the MeerKAT Array (LADUMA) survey; Holwerda et al. 2011], and the Australian SKA Pathfinder (ASKAP), which will analogously perform the Wallaby Survey. Ultimately, the Square Kilometre Array (SKA) will enable studies of H I both nearby and out to higher redshifts that far exceed today's capabilities. Complementary information is provided by studies of damped Lyα absorbers (DLAs) that trace H I in and around galaxies in the spectra of background quasars (e.g. Battisti et al. 2012; Noterdaeme et al. 2012).

It is, therefore, timely and crucial to develop a theoretical framework for understanding the physics that governs the H I content of galaxies and its evolution. Because galactic H I reservoirs are thought to represent a transient phase of baryons (Prochaska & Wolfe 2009) as they pass from the diffuse IGM (e.g. Davé et al. 2010) to higher density molecular gas, which eventually converts into stars, it is essential that models for cosmic H I be dynamical in nature. Cosmological hydrodynamic simulations provide one such type of dynamical model for this purpose, allowing one to directly track the physical state of gas as it flows through the H I reservoir. Particularly, the inclusion of galactic outflows in simulations provides some interesting new twists on the accretion paradigm, as enriched outflows that return back into galaxies may provide a significant source of accretion ('wind recycling'; Oppenheimer et al. 2010). These inflow, outflow and recycling processes are expected to manifest themselves in the H I content of galaxies.

In this work, we examine the H I content of galaxies in cosmological hydrodynamic simulations with a variety of outflow models. This work follows on our earlier study in Popping et al. (2009), but it uses a significantly improved simulation suite and focuses on the physical processes that govern the H I content in addition to making predictions for H I observables and their evolution. This work is also comparable to that of Duffy et al. (2012), who used simulations to examine both the atomic and molecular content of simulated galaxies. Compared to that work, our model explores a different range of feedback parameters (including ones that match H I observations substantially better), probes down to significantly smaller systems (albeit in a smaller volume) and uses a slightly different approach to compute the H I content of galaxies. Nonetheless, for overlapping models and mass ranges, our results generally agree with those of Duffy et al. (2012).

The paper is organized as follows. In Section 2, we introduce our simulations, and describe our method for calculating the atomic and molecular gas content of our simulated galaxies. In Section 3, we examine the H I mass function (HIMF) and its evolution, and we examine the complementary constraint of H I richness versus stellar mass in Section 4. In Section 5, we study how the H I content is impacted by environment, and in Section 6 we show how it relates to galaxy metallicity and star formation rate (SFR). In Section 7, we make predictions for the evolution of $\Omega_{\rm H\,I}$ and compare to observations out to $z \sim 4$. In Section 8, we discuss the robustness of our H I predictions to variations in resolution, molecular gas prescription and numerical method. Finally, we summarize and discuss our results in Section 9.

2 METHODS

2.1 Simulations

Our simulations are evolved with an extended version of the GADGET-2 *N*-body + smoothed particle hydrodynamic (SPH) code (Springel 2005). We assume a Λ cold dark matter (ΛCDM) cosmology (Hinshaw et al. 2009): $\Omega_{\rm M} = 0.28$, $\Omega_{\Lambda} = 0.72$,

$h \equiv H_0/(100 \, {\rm km\,s^{-1}\,Mpc^{-1}}) = 0.7$, a primordial power spectrum index $n = 0.96$, an amplitude of the mass fluctuations scaled to $\sigma_8 = 0.82$ and $\Omega_{\rm b} = 0.046$. We call this cosmology our *r*-series, where our general naming convention is r[box size]n[particles/side][wind model]. Our primary simulations use a cubic volume of $32 \, h^{-1}$ Mpc on a side with 512^3 dark matter and 512^3 gas particles, and a softening length of $\epsilon = 1.25 \, h^{-1}$kpc (comoving, Plummer equivalent). The gas particle mass is $4.5 \times 10^6 \, {\rm M_\odot}$, and the dark matter particles masses are approximately five times larger. We can thus reliably resolve galaxies down to stellar masses of $M_{*,\,{\rm lim}} = 1.4 \times 10^8 \, {\rm M_\odot}$ (see the discussion below).

Our version of GADGET-2 includes cooling processes using the primordial abundances as described by Katz, Weinberg & Hernquist (1996), with additional cooling from metal lines assuming photoionization equilibrium from Wiersma et al. (2009). SF is modelled using a subgrid recipe introduced by Springel & Hernquist (2003a), where a gas particle above a density threshold of $n_{\rm H} = 0.13 \, {\rm cm^{-3}}$ is modelled as a fraction of cold clouds embedded in a warm ionized medium following McKee & Ostriker (1977). SF follows a Schmidt law (Schmidt 1959), where the SF rate is proportional to $n_{\rm H}^{1.5}$, with the SF time-scale scaled to match the $z = 0$ Kennicutt relation (Kennicutt 1998). We use a Chabrier (2003) initial mass function throughout. We account for metal enrichment from Type II supernovae (SNe), Type Ia SNe and asymptotic giant branch stars, and we track four elements (C, O, Si, Fe) individually, as described by Oppenheimer & Davé (2008). We note that the entropy-conserving (EC) SPH algorithm in GADGET-2 is known to have some deficiencies in properly modelling hydrodynamical instabilities; in Section 8 we discuss this further and argue that this should not have a large impact on our results.

Galactic outflows are implemented using a Monte Carlo approach analogous to SF. Outflows are directly tied to the SFR, using the relation $\dot{M}_{\rm wind} = \eta \times {\rm SFR}$, where η is the outflow mass loading factor. The probability for a gas particle to spawn a star particle is calculated from the subgrid model described above, and the probability to be launched in a wind is η times the SF probability. If the particle is selected to be launched, it is given a velocity boost of v_w in the direction of $\boldsymbol{v} \times \boldsymbol{a}$, where $\boldsymbol{v}$ and $\boldsymbol{a}$ are the particle's instantaneous velocity and acceleration, respectively. Once a gas particle is launched, its hydrodynamic (not gravitational) forces are turned off until either $1.95 \times 10^{10}/(v_{\rm w}({\rm km\,s^{-1}}))$ years have passed or, as more often occurs, the gas particle has reached a density that is 10 per cent of the SF critical density (i.e. $0.013 \, {\rm cm^{-3}}$). This attempts to mock up chimneys generated by outflows that would allow a relatively unfettered escape from the galactic interstellar medium (ISM), a process not properly captured by the spherically averaging SPH algorithm at ∼kpc resolution. It also yields results that are less sensitive to numerical resolution (Springel & Hernquist 2003b) than models that do not turn off hydrodynamic forces. For a further discussion of hydrodynamic decoupling, see Dalla Vecchia & Schaye (2008).

Choices of the parameters η and v_w define the 'wind model'. For this paper, we make use of the following four wind models:

(i) *No winds* (*nw*), where we do not include outflows (i.e. $\eta = 0$). This model fails to match a wide range of observables (e.g. Davé, Oppenheimer & Finlator 2011a; Davé, Finlator & Oppenheimer 2011b), but is included to establish a baseline for the overall impact of winds.

(ii) *Constant winds* (*cw*), where $\eta = 2$ and $v_{\rm w} = 680 \, {\rm km\,s^{-1}}$ for all galaxies. This model is similar to the wind model employed by Duffy et al. (2012).

(iii) *Momentum-conserving winds* (*vzw*), where the wind speed and the mass loading factor depend on the galaxy velocity dispersion σ (Murray, Quataert & Thompson 2005), using the relations (see Oppenheimer & Davé 2008)

$$v_{\mathrm{w}} = 3\sigma \sqrt{f_L - 1}, \tag{1}$$

$$\eta = \frac{\sigma_0}{\sigma}, \tag{2}$$

where $f_L = [1.05, 2]$ is the luminosity factor in units of the galactic Eddington luminosity (i.e. the critical luminosity necessary to expel gas from the galaxy potential) and $\sigma_0 = 150\,\mathrm{km\,s^{-1}}$ is the normalization of the mass loading factor. Choices for the former are taken from observations (Rupke, Veilleux & Sanders 2005), while the latter is broadly constrained to match high-redshift IGM enrichment (Oppenheimer & Davé 2008). Galaxies are identified 'on the fly' during the simulation using a friends-of-friends (FOF) algorithm applied to the gas, star and dark matter particles. We choose a smaller-than-usual linking length to pick out just the galaxies. This linking length evolves with redshift and in terms of the mean interparticle separation of all particles is

$$0.06 \left(\frac{H(z)}{H_0} \right)^{1/3}. \tag{3}$$

We estimate the velocity dispersions necessary for our wind scaling laws from the total FOF galaxy mass M_{gal}[1] using the relation

$$\sigma = 200 \left(\frac{M_{\mathrm{gal}}}{5 \times 10^{12}\, h^{-1}\,\mathrm{M_\odot}} \frac{H(z)}{H_0} \right)^{1/3} \mathrm{km\,s^{-1}}, \tag{4}$$

which, we have empirically determined, gives an accurate measure of the velocity dispersion in our simulations.

(iv) *Hybrid energy/momentum-driven winds* (*ezw*), which employs the vzw scalings for galaxies with $\sigma > 75\,\mathrm{km\,s^{-1}}$, and then switches over to a steeper dependence of $\eta \propto \sigma^{-2}$ in $\sigma < 75\,\mathrm{km\,s^{-1}}$ systems. The wind speed still scales proportionally to σ as in the vzw model. This model roughly captures the behaviour in recent analytic and hydrodynamic models of outflows from interstellar media by Murray, Quataert & Thompson (2010) and Hopkins, Quataert & Murray (2012), respectively. The basic idea is that in dwarf galaxies, the energy from SNe plays a dominant role in driving outflows, while in larger systems the momentum flux from young stars and/or SNe is the dominant driver. As a result, the outflow scalings switch from momentum driven at high masses to energy driven at low masses. We make this transition abruptly at $\sigma = 75\,\mathrm{km\,s^{-1}}$ guided by the analytic models of Murray et al. (2005, 2010) and the high-resolution ISM simulations of Hopkins et al. (2012), although it should perhaps be more gradual (note that η itself is continuous across this boundary). In any case, this model captures the gist of the most up-to-date small-scale outflow models. As we will see, the ezw outflow model fares somewhat better than vzw, which in turn fares much better than cw or nw, when compared with observations.

Additionally, in our fiducial ezw simulation, we employ a heuristic prescription to quench SF in massive galaxies. This is not a physical model, but simply a tuned parametrization to limit SF in massive systems to reproduce the observed exponential high-mass

cutoff in the stellar mass function and to make the simulation run faster. In this prescription, we quench SF in an *entire galaxy* according to a probability, P_Q, given by the equation

$$P_Q = 1 - \frac{1}{2}\mathrm{erfc}\left[\frac{\log(\sigma) - \log(\sigma_{\mathrm{Qmed}})}{\log(\sigma_{\mathrm{Qspr}})} \right], \tag{5}$$

where the median σ at which a galaxy has a 50 per cent chance of being quenched is $\sigma_{\mathrm{Qmed}} = 110\,\mathrm{km\,s^{-1}}$, corresponding to $M_{\mathrm{halo}} = 10^{12.1}\,\mathrm{M_\odot}$ at $z = 0$, and the spread in σ is $\sigma_{\mathrm{Qspr}} = 32\,\mathrm{km\,s^{-1}}$. We also require $\sigma > 75\,\mathrm{km\,s^{-1}}$ to fully suppress the already low probability of lower mass galaxies being quenched. We have found that these parameter choices nicely reproduce the high-mass end of the stellar mass function, as we will show in Section 2.3.

Every time we identify galaxies using our FOF group finder, which we do to calculate σ for the vzw and ezw wind models every 10 Myr, we compute this quenching probability for each identified galaxy. If the galaxy is 'quenched', then each time a gas particle would have formed a star over the time interval until the next time we identify galaxies using our FOF group finder; it is instead heated to 50 times the virial temperature, T_{vir}, where

$$T_{\mathrm{vir}} = 5 \times 10^6 \left(\frac{\sigma}{200\,\mathrm{km\,s^{-1}}} \right)^2 \mathrm{K}. \tag{6}$$

The motivation for heating the gas to such extreme temperatures is primarily to prevent it from re-accreting at later times and thus requiring multiple ejections; the total energy input is thus less for higher quenching temperatures. The median temperature to which the ISM is heated by quenching is $10^{8.1}\,\mathrm{K}$ and arises from a median halo mass of $10^{12.2}\,\mathrm{M_\odot}$. The energy input from quenching averages to $10^{40.5}\,\mathrm{erg\,s^{-1}\,Mpc^{-3}}$ between $z = 0.75$ and 2.5, corresponding to the peak of AGN activity, and the integrated energy input until $z = 0$ is $8 \times 10^{14}\,\mathrm{erg\,g^{-1}}$, which equals 9×10^{-7} of the rest mass energy of all baryons. Considering that 6.5 per cent of baryons are in stars in this simulation at $z = 0$, and using the assumption that supermassive black holes (SMBHs) have 10^{-3} of the mass in stars, the quenching energy corresponds to 1.4 per cent of the rest-mass energy of SMBHs, which is comparable to the energy imparted from AGN feedback in cosmological simulations that self-consistently track black hole growth and feedback (Di Matteo, Springel & Hernquist 2005; Booth & Schaye 2009).

In this paper, we are mostly concerned with galaxies around L^* and below, which are mostly unquenched. We emphasize that this quenching prescription is not intended to be a realistic physical model, and is actually in large part a computational convenience, as removing gas from the largest galaxies substantially speeds up the simulation at low redshifts and thus allows us to run our large simulations to $z = 0$ within practical time frames. We refer the reader to Duffy et al. (2012) for a discussion of H I in larger volume simulations that attempt to quench massive galaxies based on directly tracking black hole growth and feedback.

Finally, our no-wind simulation employs a volume of $16\,h^{-1}\mathrm{Mpc}$ with 2×256^3 particles. This results in the same mass and spatial resolution as the wind simulations, but in a volume that is eight times smaller. As we mostly use the no-wind simulation to qualitatively show the impact of winds at small masses, the reduced volume for the simulation will not significantly affect our conclusions.

---

[1] We have actually been using this algorithm, which slightly differs from Oppenheimer & Davé (2008), in all of our recent simulations starting with Oppenheimer et al. (2010).

2.2 Computing the H I content

We use Spline Kernel Interpolative Denmax (SKID)[2] to identify galaxies as bound groups of star-forming gas and stars (Kereš et al. 2005; Oppenheimer et al. 2010). Our galaxy stellar mass limit is set to be ≥ 64 star particles (Finlator et al. 2006), resulting in a minimum resolved mass of $M_{*,\,\mathrm{min}} = 1.4 \times 10^8\,\mathrm{M_\odot}$. We will only consider galaxies with stellar masses $M_* \geq M_{*,\,\mathrm{min}}$ in our analysis, regardless of their H I content. Our resolution convergence tests in Section 8 show that the H I properties are reasonably well converged even at this stellar mass threshold.

We identify dark matter haloes via a spherical overdensity algorithm (Kereš et al. 2005) out to a virial overdensity given in Davé et al. (2010, see their equation 1), centred on each galaxy. We separate galaxies into central and satellite galaxies by associating each galaxy with a halo; if a galaxy's centre lies within the virial radius of a larger (by stellar mass) galaxy, we consider it to be a satellite of that galaxy, and the haloes of those two galaxies are merged. Note that for galaxies near the edge of a larger halo, this can result in a halo that has 'bumps' along its (mostly spherical) surface.

To compute the H I content of galaxies, we need to isolate the gas that is in the (predominantly) neutral phase. This involves defining two 'boundaries': the division between gas that is exposed to the full metagalactic ionizing flux and gas that is self-shielded, and the division between atomic and molecular gas.

To calculate the self-shielding, we first compute the neutral hydrogen component of each gas particle under the assumption that it is *not* self-shielded. We follow Popping et al. (2009), who employed a simple hydrogen ionization balance calculation that yielded the following formula to determine the neutral fraction:

$$f_{\mathrm{H\,I}} = \frac{2C + 1 - \sqrt{(2C+1)^2 - 4C^2}}{2C} \tag{7}$$

with

$$C = \frac{n\beta(T)}{\Gamma_{\mathrm{H\,I}}}, \tag{8}$$

where n is the hydrogen number density, T is the gas temperature, $\Gamma_{\mathrm{H\,I}}$ is the H I photoionization rate and the recombination rate coefficient β is well fitted by the function (Verner & Ferland 1996)

$$\beta(T) = a\left[\sqrt{(T/T_0)}(1 + \sqrt{(T/T_0)})^{1-b}(1 + \sqrt{(T/T_1)})^{1+b}\right]^{-1}. \tag{9}$$

For H I, the fitting parameters are $a = 7.982 \times 10^{-11}\,\mathrm{cm^3\,s^{-1}}$, $b = 0.7480$, $T_0 = 3.148\,\mathrm{K}$ and $T_1 = 7.036 \times 10^5\,\mathrm{K}$. We take $\Gamma_{\mathrm{H\,I}}$ from the Haardt & Madau (2001) ionizing background, whose amplitude is adjusted to match the observed mean flux decrement in the Lyα forest in these simulations (see Davé et al. 2010); we describe this further below.

With the optically thin neutral fraction computed for each gas particle, we employ a simple particle-by-particle post-processing correction for self-shielding of the H I. We resort to this approach because it is computationally prohibitive to do the full radiative line transfer on these simulations; instead, we will calibrate our approach to radiative transfer simulations by Faucher-Giguere, Kereš & Ma (2010).

We begin by assuming that each particle has a density profile given by the SPH spline kernel $W(r)$ (see Springel 2005

for definition). We compute the radial column density profile as follows:

$$N_{\mathrm{H\,I}}(r) = \frac{0.76 f_{\mathrm{H\,I}} \rho_g}{m_p} \int_r^h W(r')\,\mathrm{d}r', \tag{10}$$

where ρ_g is the SPH density of the gas particle, m_p is the proton mass and h is the particle's smoothing length. We then determine the radius R where $N_{\mathrm{H\,I}}(R) = N_{\mathrm{H\,I,lim}}$, with $N_{\mathrm{H\,I,lim}} = 10^{17.2}\,\mathrm{cm^{-2}}$ where the particle becomes optically thick ($\tau = 1$) to continuum photons at the Lyman limit. If no such radius R exists, then the particle is optically thin. If R exists, then we compute the unshielded and shielded mass fraction using $4\pi \int W(r) r^2 \mathrm{d}r$ integrated from $h \to R$ and $R \to 0$, respectively. The neutral fraction of the particle is then the mass-weighted mean neutral fraction, assuming that the portion from $R \to 0$ is 90 per cent neutral (since some ionized gas is seen to exist even within dense regions). By this procedure, particles in dense regions 'autoshield' themselves from the ambient flux, though the effects of shielding from nearby gas are neglected.

This computation relies on knowing the metagalactic photoionizing flux, since this determines the neutral fraction in the optically thin regime. Our simulations assume a Haardt & Madau (2001) ionizing background, but more detailed constraints can be placed on the amplitude of the H I photoionization rate by using the measured mean flux decrement in the Lyα forest. In particular, at each redshift for each simulation, we determine a 'flux factor', which is the value by which we must multiply the strength of the Haardt & Madau (2001) background to achieve consistency with the observed flux decrement. For the observed flux decrement, we employ at $z < 2$ the determination by Kirkman et al. (2007), while for $z \geq 2$ we take the values from Becker et al. (2012). As described more fully in Davé et al. (2010), we extract 1000 spectra and iteratively adjust the flux factor until the mean flux is within 1 per cent of the observed value. The resulting flux factors for our ezw simulation at $z = 0, 1, 2, 3, 4$ and 5 are 1.72, 2.03, 1.66, 0.92, 1.13 and 1.70, respectively. For our vzw simulation, the $z = 0$ flux factor is 1.63 and $z = 3$ is 0.94, while for cw, it is 1.41 at $z = 0$ (we do not consider these models at other redshifts).

Next, we separate the molecular component from the atomic H I. To do so, we employ the observed ISM pressure relation from The H I Nearby Galaxy Survey (THINGS; Leroy et al. 2008); specifically, we use their 'combined spiral subsample' fit, which gives the ratio of molecular to atomic gas as

$$R_{\mathrm{mol}} = (P/P_0)^\alpha, \tag{11}$$

where $P_0 = 1.7 \times 10^4\,\mathrm{cm^{-3}\,K}$ and $\alpha = 0.8$. We compute the gas pressure based on the density and two-phase medium temperature of each star-forming gas particle (our prescription follows Springel & Hernquist 2003a). Note that we only compute molecular fractions for star-forming gas particles, which in our simulations is gas with $n_{\mathrm{H}} > 0.13\,\mathrm{cm^{-3}}$. Gas that is not star forming is assumed to have zero molecular content, and in any case the formula above would yield a very small molecular fraction. This prescription follows that employed by Duffy et al. (2012), but differs from Popping et al. (2009) who used a fixed pressure threshold of $810\,\mathrm{cm^{-3}\,K}$ to separate atomic from molecular gas.

Recent theoretical work by Krumholz & Gnedin (2011) has also provided a prescription for separating molecular from atomic gas, which includes a metallicity dependence. We describe this method and examine its impact in Section 8. To encapsulate the results, this theoretically based prescription has little effect at $z = 0$, but at $z = 3$ the lower metallicity causes somewhat less molecular gas to form, and hence the H I content is slightly higher.

Figure 1. Gas particle properties as a function of the number density n_H in our r48n384vzw simulation. The left-hand panels show $z = 0$ and the right-hand panels show $z = 3$. Top: H I fraction (blue) and H_2 fraction (red) based on the calculation described in the text. We show a random sampling of 1 per cent of all the particles. Middle: total mass density Ω per unit $\log n_H$ in H I (blue) and H_2 (red). Bottom: pressure–density relation. The green points show unshielded particles (either below the pressure threshold or above the density threshold), the blue points show H I-dominated gas and the red points show molecular-dominated gas.

Fig. 1 illustrates the resulting gas particle atomic and molecular fractions. The top-left panel shows the H I and H_2 mass fractions in particles from our momentum-driven wind scaling (vzw) simulations at $z = 0$. In the optically thin regime (mostly not depicted here; see Davé et al. 2010 for a full phase space diagram), $f_{H I}$ scales linearly with n_H. Then there is a sharp upturn at $n_H \sim 2 \times 10^{-3} \text{cm}^{-3}$, when autoshielding becomes important.

At high redshifts, the stronger ionizing background causes autoshielding to set in at a higher density. A simple scaling shows that $N_{H I,lim} \propto f_{H I} n_H l$, where the length $l \propto n_H^{-1/3}$ for a spherical cloud (or SPH particle) and $f_{H I} \propto n_H / \Gamma_{H I}$; hence, $N_{H I,lim} \propto n_H^{5/3} / \Gamma_{H I}$. Therefore, a factor of 10 increase in $\Gamma_{H I}$, which occurs between $z = 0 \rightarrow 3$, will lead to a factor of 4 increase in the n_H where self-shielding becomes important, because $n_H \propto \Gamma_{H I}^{3/5}$ for constant $N_{H I,lim}$. Fig. 1 (top-right panel) shows that autoshielding at $z = 3$ is effective above $n_H \sim 10^{-2} \text{cm}^{-3}$ as expected for the stronger $z = 3$ background UV field, which also matches the $z = 3$ radiative transfer simulations of Faucher-Giguere et al. (2010, see their fig. 3). Indeed the overall shape of $f_{H I}(n_H)$ is actually quite similar to theirs, although there is less scatter at a given n_H in our prescription owing to the fact that we do not consider shielding from neighbouring particles. This shows that our physically motivated choice of autoshielding with $N_{H I,lim}$ at the Lyman limit is a reasonable approximation to much more detailed radiative transfer models.

Moving to higher densities, eventually we reach the star-forming density threshold (the vertical dashed line), above which the molec-

ular fraction becomes non-zero. At $n_H \gtrsim 0.5 \text{ cm}^{-3}$, corresponding to a pressure close to P_0, the gas becomes molecular dominated. The pressure relation is shown in the bottom panels of Fig. 1, with the particles colour-coded by their dominant phase of hydrogen. The THINGS pressure threshold is much higher than the $P/k = 810 \text{ cm}^{-3} \text{ K}$ assumed in Popping et al. (2009) and, moreover, our prescription produces a more gradual transition between atomic and molecular-dominated gas. A change in slope at $n_H = 0.13 \text{ cm}^{-3}$ occurs because this is the density above which we allow stars to form and the gas particles become two-phase.

In this paper, we are mostly interested in the regime $10^{-2.5} \lesssim n_H \lesssim 0.5 \text{ cm}^{-3}$, where the majority of cosmic H I resides (at $z = 0$), above which gas becomes mostly molecular and below which it is mostly ionized. This is shown in the middle panels of Fig. 1, where we plot the total mass density Ω in H I (blue) and H_2 (red) per unit interval of $\log n_H$; the peak is around $n_H \sim 10^{-1} \text{ cm}^{-3}$. This H I-dominant density regime is reasonably well resolved in our simulations, and hence the predicted H I content is expected to be robust, despite it being a transitory phase. In contrast, the transition from molecular gas to stars typically occurs at densities well above what we can resolve directly, and hence the molecular content may not be quite as robustly predicted. Still, our SF prescription is consistent with the Kennicutt (1998) relation, which is well established on the $\sim$kpc-scale resolution of our simulations. Our predictions for the evolution from $z \sim 2 \rightarrow 0$ of the total star-forming gas in our momentum-driven wind (vzw) simulation are in good agreement with observations from the Plateau de Bure High-z Blue Sequence Survey (Tacconi et al. 2012, see their fig. 13); other wind models fare less well. Finally, we note that the stellar mass growth rate is fairly robust because it is not a transitory phase but an end state of accreted gas, and is typically limited by the gas supply rate modulated by outflows (e.g. Finlator & Davé 2008; Bouché et al. 2010; Davé et al. 2012).

We note that, by the above prescription to separate molecular from atomic gas, 55 per cent of the gas at $z = 0$ that is above our SF density threshold of $n_H = 0.13 \text{ cm}^{-3}$ is actually neutral (it is much lower at high redshift). This gas could in principle form stars in our simulation, which is contrary to the idea that stars only form from molecular gas. However, the amount of SF actually occurring in this neutral gas is very small, only about 3 per cent. This is because the SFR scales as $\rho^{1.5}$, which means that virtually all of the SF occurs in the denser gas that is almost fully molecular.

With each gas particle's H I content determined, we must now associate the gas particles with galaxies. The information from SKID is not sufficient, because SKID only includes star-forming gas and stars in a galaxy, while a significant amount of the H I resides in an extended region around the galaxy, beyond the actively star-forming region. Thus, to account for extended H I, we add up all the H I mass in a sphere around each galaxy that extends to the outermost radius as defined by SKID, i.e. the radius of the farthest SKID particle associated with that galaxy. The outermost radius is typically many times the half-mass radius. While it may seem that this still may not fully account for an extended H I disc, in practice the low threshold density for SF in our simulations means that this choice still encompasses the vast majority of the H I. We tried extending this to 1.5 times the outermost radius, and the total H I mass in the volume increased by only 4 per cent; furthermore, one will start to increasingly 'double count' gas that may be between nearby galaxies as being part of both galaxies. Modestly reducing this radius also has a minimal effect, so our H I masses are not very sensitive to this choice.

2.3 The stellar mass function

As an initial baseline statistic to compare our four wind models, we show in Fig. 2 their galaxy stellar mass functions (GSMFs) at $z = 0$. We compare these to observations from Baldry et al. (2008) using Sloan Digital Sky Survey (SDSS).

The ezw model provides a strikingly good fit to the observed GSMF. At the massive end, this is a direct consequence of tuning the quenching prescription as described previously. In contrast, the low-mass end is unaffected by our quenching prescription, and instead reflects the effect of the hybrid energy–momentum-driven winds. There is a clear upturn in the GSMF at $M_* \lesssim 10^{9.5}\,M_\odot$, which is also seen in the data. As explained in Oppenheimer et al. (2010), in our simulations this arises because above this mass, wind recycling becomes increasingly important and provides extra fuel to higher mass galaxies. Below this mass, the typical recycling time becomes longer than a Hubble time, and the slope begins to steepen towards the dark matter halo mass function's slope.

The vzw model provides not quite as good a fit to the low-mass end, similar to what was seen in a lower resolution version of the same model in Davé et al. (2011a). The differences at the massive end effectively show the impact of the quenching model, since our ezw simulation uses it whereas our vzw simulation does not, while the two outflow prescriptions themselves are identical in this mass regime. Quenching has a substantial effect on the GSMF, but we will show that it has a minimal effect on the HIMF. Meanwhile, the constant-wind model produces a very steep low-mass end slope that looks nothing like the data, while the no-wind model overproduces stars at virtually all masses because it strongly overcools baryons. These results again follow those obtained using lower resolution simulations of the same wind models presented in Oppenheimer et al. (2010) and Davé et al. (2011a), and more discussions of the GSMF can be found there.

This ezw simulation is the first hydrodynamic simulation that we are aware of to yield agreement with the observed GSMF to within statistical uncertainties over the entire mass range probed.
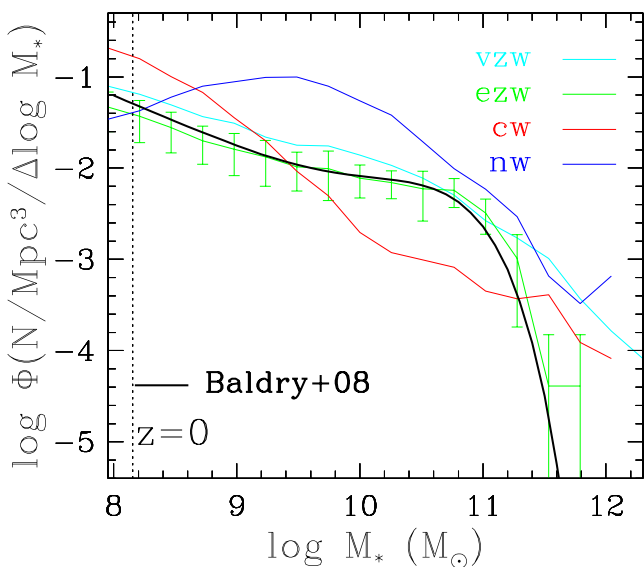

Figure 2. Galaxy stellar mass functions (GSMFs) at $z = 0$ in our four wind models: ezw (green), vzw (cyan), cw (red) and nw (blue). The vertical dotted line shows the stellar mass resolution limit. The thick solid line shows a fit to the observed GSMF from Baldry, Glazebrook & Driver (2008). The ezw outflow model, which includes a heuristic prescription for quenching SF at high masses, matches the observed GSMF to within uncertainties.

This does not imply that this model is fully correct, as it could still be that the SF histories in this model are incorrect, and other constraints may not be as well matched. As a case in point, we will show in Section 6 that the mass–metallicity relation (MZR) in this model looks too steep compared to observations, and that the vzw model provides a better match. Furthermore, the success of this model at high masses owes to including an ad hoc and highly tuned prescription for quenching SF that is not a direct implementation of a physical quenching mechanism. We leave a full comparison of the ezw model to a wide range of observables for future work, and focus here on the H I properties of galaxies in our four wind simulations.

3 THE H I MASS FUNCTION

The most basic counting statistic for characterizing the H I content of galaxies is the HIMF. Improving 21 cm observations have enabled the HIMF to be probed down to H I masses approaching $10^7\,M_\odot$ (e.g. Zwaan et al. 2005; Haynes et al. 2011). However, the redshift evolution of the HIMF remains poorly characterized, awaiting the next generation of facilities. Here we compare our simulated HIMFs with observations, to understand what constraints can be placed on wind models and how the HIMF is expected to evolve.

Fig. 3, top panel, shows the HIMF for our four wind models. The vertical dotted line shows the *stellar* mass resolution limit of our fiducial r32n512ezw simulation. We have no formal H I mass resolution limit, but we will show in Section 8 that this is a reasonable choice for ensuring resolution convergence of the HIMF as well. Error bars (shown for ezw) depict cosmic variance as computed from the error of the mean HIMF in each of eight octants of the simulation volume.

For comparison, we show the observational determination of the HIMF from Haynes et al. (2011) (the first 40 per cent of the ALFALFA survey; dashed line). We note that the ALFALFA observations select by H I mass, while we are effectively selecting by stellar mass; however, the sensitivity of ALFALFA is well below what we can resolve, and hence unless there is a dominant population of small H I-free objects (which we will show later does not occur in our models), this comparison should be robust.

Concentrating on the low-mass end of the HIMF, we see that both the ezw and vzw models do a good job of matching the observed low-mass end slope of $\alpha \approx -1.34 \pm 0.02$ (for the 'whole $\alpha.40$' sample, from table 6 of Haynes et al. 2011). In detail, the predicted best-fitting Schechter function slope for the ezw model is $\alpha = -1.31$, while for the vzw model it is $\alpha = -1.45$ and, therefore, the ezw model provides a slightly better fit. The massive end is relatively unaffected by quenching, as seen from the minor differences between the ezw model that includes quenching and the vzw model that does not include it.

Particularly remarkable is that the amplitudes of the HIMF in these two models are reasonably close to that observed. They are slightly above the observations at all masses, particularly at higher masses. We could in principle tune our values of $N_{\rm H\,I,lim}$ and/or shrink the radial extent to which we associate H I with a galaxy to match the data better, but we prefer to use well-motivated values for these quantities. We have tried several other reasonable values for these quantities, and the net effect is generally to scale the HIMF in amplitude without changing the shape significantly.

The agreement with the low-mass end slope of the HIMF, while simultaneously matching the low-mass end of the galactic stellar mass function, is a stringent constraint that almost all galaxy formation models have difficulty matching (Mo et al.
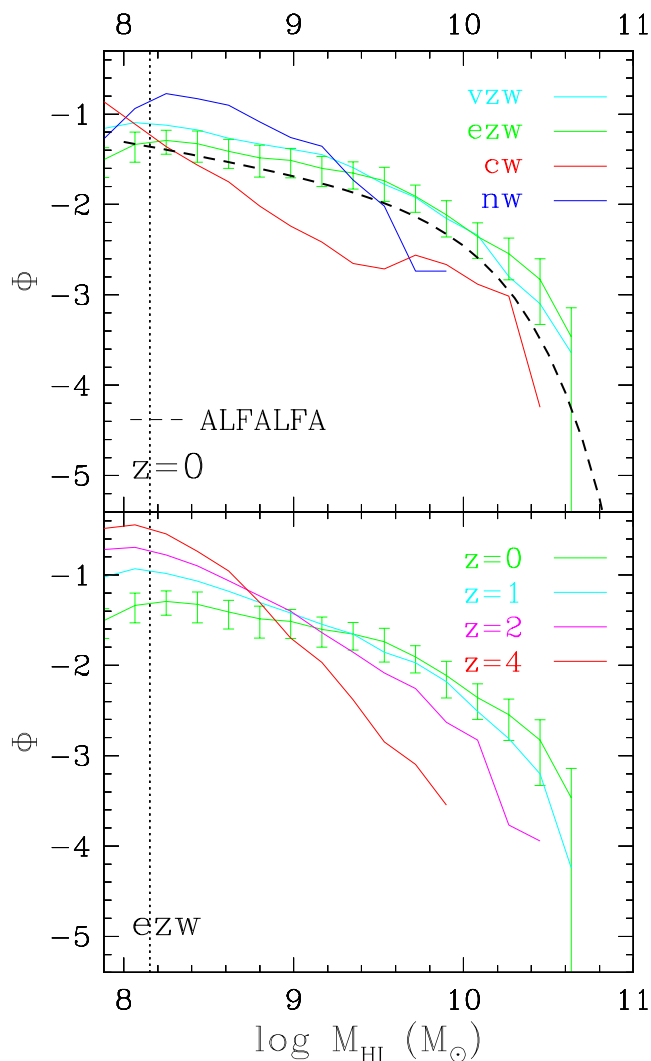
Figure 3. Top: a comparison of the HIMF at $z = 0$ in our four wind simulations – ezw (green), vzw (cyan), cw (red) and nw (blue). The vertical dotted line shows the stellar mass resolution limit of $1.4 \times 10^8\,M_\odot$; this is roughly the H I mass resolution limit as well, since M_* and M_{HI} are comparable at the smallest masses (see Fig. 4). The dashed line shows the HIMF from the $\alpha.40$ sample of the ALFALFA survey (Haynes et al. 2011) for comparison. The ezw and vzw models provide reasonable fits to the data. Bottom: evolution of the HIMF from $z = 4 \rightarrow 0$ in the ezw model. The low-mass end slope increases substantially at high redshifts.

2005; Lu et al. 2012; Lu et al., in preparation). Fundamentally, in most models this arises because the low-mass end slope of the dark matter halo mass function is quite steep, which is exacerbated by low-mass galaxies being more H I rich. In our simulations, the stronger outflows from low-mass galaxies strongly suppress the overall baryon content of galaxies. Even more critical is the impact of wind recycling. As described in Oppenheimer et al. (2010), wind recycling is preferentially stronger in high-mass galaxies because the denser surrounding gas slows outflows more effectively via ram pressure. This results in a higher fraction of ejected gas being recycled back into higher mass galaxies, thereby yielding more SF in more massive systems. Oppenheimer et al. (2010) showed that this flattens the low-mass end slope of the stellar mass function (see Fig. 2), and here we see that this effect is also important for the low-mass end of the HIMF. We will dis-

cuss further the differences between previous semi-analytic galaxy formation model results and our results in Section 9.

Both the constant-wind (cw) and no-wind (nw) models provide significantly poorer fits to the observations than ezw or vzw. Without winds, there is a surplus of low-M_{HI} galaxies, and a deficit of high-M_{HI} systems. This is characteristic of the HIMF in many semi-analytic galaxy formation models (Lu et al. 2012). The former discrepancy arises because there are no outflows to eject material from low-mass galaxies, and hence the baryon fraction in these systems is very large (Davé 2009), in disagreement with observations that show a small baryon content in dwarfs (e.g. McGaugh et al. 2010). The latter arises because, without winds, gas becomes very dense in high-mass galaxies, which means that most of the gas is in molecular form and much of it has been converted into stars. As discussed in Davé et al. (2011a), the stellar content of galaxies in the no-wind model grossly exceeds that observed, and Fig. 2 showed that it overproduces the number density of galaxies at all the stellar masses probed.

For the constant-wind case, the low-mass end slope is quite steep ($\alpha = -1.64$). Moreover, there is a characteristic 'bump' in the mass function at $M_{HI} \sim 10^{10}\,M_\odot$. This bump is reminiscent of a similar bump in the stellar mass function in this model (Davé et al. 2011a), which arises because wind recycling rapidly becomes important around this mass scale as the (constant-velocity) winds are no longer able to escape the galaxy halo's potential well. Such a feature, which is generic to wind models that assume a constant outflow speed, is not observed in the HIMF (nor in the stellar mass function).

Duffy et al. (2012) employed an outflow model that is similar to this constant-wind case. They find that the low-mass end slope is fairly flat down to $M_{HI} \approx 10^9\,M_\odot$, and was too shallow compared to data. In a comparable range, our HIMF also appears fairly flat – formally, the best-fitting low-mass end slope ignoring points below that mass is $\alpha = -1.12$, although a Schechter function is a poor descriptor. Our higher resolution simulation probes further down the mass function, which enables us to see the steep low-mass end. Overall, our results for this wind model agree well with Duffy et al. (2012) in the overlapping mass range, and both show that the constant-wind model fails to match the low-mass end of the HIMF.

The bottom panel shows the redshift evolution of the HIMF out to $z = 4$, focusing on our ezw simulation. The low-mass end slope becomes progressively steeper with redshift, mimicking the behaviour seen in the stellar mass function (Davé et al. 2011a). The low-mass end slopes at $z = 1, 2, 3, 4, 5$ are $-1.54, -1.79, -1.82, -1.99, -2.11$, respectively ($z = 3, 5$ are not shown). This arises because wind recycling becomes increasingly effective at lower redshifts (Oppenheimer et al. 2010), since the recycling time is roughly constant at $\sim 1-3$ Gyr (Oppenheimer & Davé 2008) for $M_* \sim 10^{10}\,M_\odot$ galaxies, which is a small fraction of the Hubble time today but comparable to the Hubble time in the early Universe. Hence, the HIMF steepens rapidly out to $z \sim 2$, and then the steepening becomes more gradual, since at early times the effect of wind recycling, which is responsible for flattening the HIMF, is reduced. Meanwhile, at the massive end, there are fewer galaxies at high redshifts simply because of the hierarchical nature of galaxy assembly.

In summary, the HIMF provides a strong constraint on outflow models. The agreement of the ezw model predictions with the latest observed HIMF from ALFALFA represents a non-trivial success that has not previously been attained in hierarchical models of galaxy formation. The vzw model fares slightly worse but may still be within the overall uncertainties, while the constant-wind and

no-wind cases fare poorly against the observed data. More broadly, this indicates that including a well-motivated model for galactic outflows enables hierarchical structure formation models to produce an HIMF that is in very good agreement with observations down to fairly low H I masses.

4 H I RICHNESS

The H I richness, i.e. the H I mass relative to the stellar mass, provides a complementary characterization of the H I content of galaxies. As we argued earlier, both the H I and stellar masses are reasonably robust predictions of our models, and hence the H I richness should provide a meaningful discriminant between models.

Fig. 4 shows the H I richness in our simulated galaxies, as a function of the stellar mass. The top four panels show our four wind models: ezw (upper left), vzw (middle left), constant (upper right) and no winds (middle right). The points in each panel show
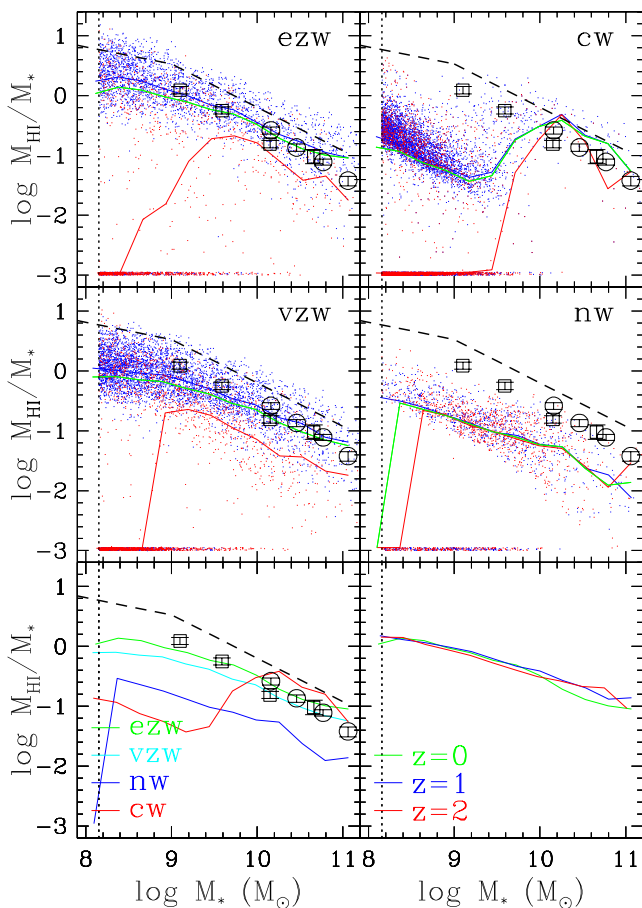

Figure 4. H I richness ($\equiv M_{HI}/M_*$) in galaxies from our four wind models: ezw (upper left), vzw (middle left), cw (upper right) and nw (middle right). The points show individual galaxies at $z = 0$, with central galaxies shown in blue and satellites shown in red. A running median for all galaxies is shown in green, while medians for the centrals and satellites are shown in blue and red, respectively. The circles show H I fractions observed in the GASS of nearby galaxies (Catinella et al. 2010), the squares show results from the Herschel Redshift Survey (Cortese et al. 2011) and the dashed line shows an approximate fit to the locus traced by the H I-selected ALFALFA survey (Huang et al. 2012). The bottom-left panel shows a comparison of all four wind models with observations at $z = 0$, while the bottom-right panel shows the evolution of the median H I richness with redshift out to $z = 2$ in the ezw model.

individual galaxies at $z = 0$, with blue depicting central galaxies and red depicting satellites. Galaxies with zero H I content are plotted along the bottom of each panel; they are almost exclusively satellite galaxies, which we will examine further in Section 5. A binned median of $\log M_{HI}/M_*$ for all galaxies is shown as the green line. We show error bars on the median corresponding to the 1σ spread for the galaxies within each mass bin. We also separately show binned medians for the central and satellite populations. The vertical dotted line shows our galaxy stellar mass resolution limit.

For comparison, mean M_{HI}/M_* observations from the GASS are shown as the large circles (Catinella et al. 2012). GASS has measured H I in a stellar mass-limited sample down to $M_* \sim 10^{10}\,M_\odot$. At lower masses, we show the results from the Herschel Redshift Survey by Cortese et al. (2011, open squares), which uses literature H I data but is approximately stellar mass complete. We also show the fit to results from the ALFALFA survey by Huang et al. (2012, dashed line); since this is H I mass selected, it is biased towards more H I-rich galaxies, as is evident when compared to the M_*-selected data. Our simulation results are most straightforwardly comparable to M_*-limited samples.

All wind models broadly predict that H I richness is anticorrelated with M_*, as observed. However, in detail, the models show distinct differences; for clarity, just the overall medians are plotted in the lower-left panel. Our ezw model produces good agreement with the stellar mass-selected observations down to the lowest probed masses. The trend for the vzw model follows ezw, but shows slightly lowered H I richness particularly in smaller galaxies, indicating that more mass-loaded galactic outflows in low-mass galaxies are favoured. The no-wind case follows the trend of the data, but is too low by a factor of a few in H I richness, showing that there is overly efficient conversion of gas into stars in this model. Finally, the constant-wind model produces a distinct feature in H I richness, mimicking the feature seen in the HIMF, owing to wind recycling. This agrees poorly with the observations, which have no comparable feature. Note that more highly mass-loaded outflows yield *higher* H I richness, not lower as one might naively expect, in part because such outflows suppress M_*.

The redshift evolution of H I richness is shown in the lower-right panel of Fig. 4, for our ezw simulation. Despite the rapid evolution in the HIMF, there is remarkably little evolution in the H I fraction at a given stellar mass from $z = 0 \to 2$. The slow evolution in H I content contrasts the much more rapid evolution observed in the molecular gas fraction (e.g. Tacconi et al. 2010; Geach et al. 2011), although some of that evolution may reflect uncertainties in assessing the molecular gas content from CO emission (e.g. Narayanan, Bothwell & Davé 2012; Bolatto, Wolfire & Leroy 2013). Also, the lack of evolution in H I richness indicates that the evolution in the HIMF discussed in Section 3 mostly reflects the evolution in the stellar and/or halo mass functions. We will see in Section 7 that the lack of evolution in global H I content is also a generic prediction of our models.

5 H I IN SATELLITE GALAXIES

The H I content of galaxies is seen to vary substantially with environment. In the densest environs such as clusters, ram pressure can remove H I from infalling galaxies, and it is long known that cluster galaxies are deficient in H I (e.g. Haynes, Giovanelli & Chincarini 1984). Even in less extreme environments where ram pressure stripping is expected to play less of a role, the H I content appears to be anticorrelated with local density (Robertson, Shields & Blanc 2012); other processes such as strangulation or harassment

Figure 5. HIMF from our fiducial ezw simulation (green line, reproduced from Fig. 3), separated into centrals (blue) and satellites (red). Centrals dominate the HIMF at all H I masses probed.

may also be playing a significant role. In general, such environmental processes are expected to preferentially lower the H I content of satellite galaxies, particularly those within larger haloes that are expected to have hotter ambient gas. In this section, we examine how the H I content of satellites varies with halo mass, which can be considered as a rough proxy for environment.

To begin, in Fig. 5 we separate the total $z = 0$ HIMF in our fiducial ezw simulation (green line) into centrals (blue) and satellites (red). Centrals clearly provide the dominant contribution to the HIMF at all stellar masses probed, being at least three times more abundant than satellites at any given mass. Though we do not show it, this trend holds at all redshifts. It is also consistent with the findings from recent semi-analytic models (SAMs), though at even lower masses ($M_{H I} \lesssim 10^7\,M_\odot$) these models predict that satellites begin to dominate (e.g. Lagos et al. 2011). The basic trend arises because satellites are less abundant than centrals at all masses down to $\sim 10^8\,M_\odot$ (Davé et al. 2011a). As we discuss next, this is further exacerbated by the trends of H I in centrals versus satellites.

As is evident from Fig. 4, satellites (red points and curve) have a lower H I richness than central galaxies (blue) of the same stellar mass. At the high-mass end, the typical difference in median H I richness between centrals and satellites is a factor of ~ 2–3, in all the wind models. Below $M_* \lesssim 10^{9.5}\,M_\odot$, the satellite H I fraction drops very quickly, as there are numerous low-mass H I-free satellites in our simulations (plotted along the bottoms of the panels). While the majority of galaxies of low mass are still centrals (Davé et al. 2011a), and hence the overall median tracks that of central galaxies, it is clear that low-mass satellites in particular are highly deficient in H I for their mass. Hence, not only are satellites less abundant at a given M_*, they also have less H I, and hence contribute very little to the HIMF.

Fig. 6 examines these trends more closely. The top panel shows the H I richness of satellite galaxies in our ezw simulation at $z = 0$, divided into three stellar mass bins of 1.4×10^8–$10^9\,M_\odot$, 10^9–$10^{10}\,M_\odot$ and $>10^{10}\,M_\odot$. Galaxies with H I richness less than 10^{-3} (virtually all of which are H I free) are plotted along the bottom at -3; lines show a running median *not* including these H I-free galaxies.


Figure 6. The top panel shows H I richness ($\equiv M_{H I}/M_*$) for satellite galaxies versus halo mass at $z = 0$, for our ezw model. The satellites are split into low (blue), intermediate (green) and high (red) stellar mass bins, as indicated. Satellites with no H I are plotted at -3, with an artificial spread to aid visualization. The lines show running medians of only the galaxies that have H I (i.e. ignoring the points plotted along the bottom). The bottom panel plots the fraction of H I-free satellites within the same mass bins, colour-coded as above. The solid lines show results from the ezw model that includes an ad hoc quenching model at high halo masses, while the dotted lines show the equivalent plots for the vzw model that does not include such quenching.

This plot shows several key trends. First, at all masses, the satellites that have H I show a mild trend of being more H I rich in lower mass haloes. Since the halo mass traces stellar mass in our models (Davé et al. 2011a), this basically reflects the fact that lower stellar mass galaxies have higher H I richness (as seen in Fig. 4). Similarly, lower mass satellites are more H I rich, again following the trend for centrals discussed in the previous section. Overall these trends appear to reflect the H I content of the satellite galaxies when they were still centrals.

However, a clear difference between central and satellite galaxies is that there are many more satellites that are devoid of H I, particularly at higher halo masses, i.e. the points along the bottom of the plot appear much more frequently in the more massive haloes. In fact, the distribution of H I richness in satellites appears

to be bimodal. Even in large haloes, there are still some satellites that have substantial H I, and their H I content is not grossly different from that of satellites in smaller haloes. However, the *fraction* of satellites that are H I poor increases sharply in higher mass haloes.

To quantify this, we plot in the bottom panel of Fig. 6 the fraction of galaxies with H I richness $<10^{-3}$, as a function of halo mass, divided again into the same three satellite galaxy stellar mass bins as in the top panel. Here, we clearly see that at any halo mass, lower mass satellites are more likely to have had their H I content strongly reduced. Moreover, this trend is a very strong function of halo mass, with satellites of all masses being much more likely to be H I poor if they lie within a more massive halo.

The strong bimodality suggests that the process that renders satellites H I poor occurs on a relatively short time-scale compared to the infall time-scale into the halo (which is roughly comparable to the halo's dynamical time of several Gyr). In future work, we plan to investigate the detailed dynamical processes that remove the H I from satellites in our simulations, though this will likely require some modification to the hydrodynamics algorithm (Agertz et al. 2007; Read & Hayfield 2012; Saitoh & Makino 2013).

The dotted lines in the bottom panel of Fig. 6 show the analogous results for our vzw simulation. In massive haloes where the fraction of H I-poor satellites is substantial, the outflow model is identical between vzw and ezw simulations, but the ezw simulation includes our quenching prescription. The fact that the dotted and solid lines are very similar indicates that the trends seen in the H I content of satellites are not being set by our ad hoc quenching prescription, but more likely by the fact that halo masses above $10^{12}\,M_\odot$ tend to contain much more hot halo gas that can more strongly impact satellites moving through it (Kereš et al. 2005; Gabor et al. 2011). Recall that our quenching prescription applies only to galaxies with a high velocity dispersion, which is typically only the central galaxy in the halo, and hence the only direct impact on smaller satellites would be from the extra heat being added to the halo gas; evidently this has a minimal effect on the satellites in our simulations.

What about the central galaxies? In general, very few of the central galaxies are devoid of H I, as seen in Fig. 4. However, although we do not highlight it, there exists a small population of lower mass ($M_* \lesssim 10^{10}\,M_\odot$) centrals that are H I poor. This is related to what was seen by Gabor & Davé (2012), who found numerous low-mass central galaxies on the red sequence. These turned out to be galaxies that reside just outside, i.e. within several virial radii, of more massive haloes. In the spherical overdensity algorithm we use to identify haloes, such galaxies are identified as centrals, although the influence of the larger galaxy's halo can extend to well beyond its virial radius (e.g. Moldar et al. 2009). Hence, these H I-poor centrals could be impacted by the extended environment of a nearby larger halo, or else they could be former satellites whose orbit has taken them outside the nominal virial radius.

In summary, halo mass plays an increasingly important role in setting the H I content of satellite galaxies. This is particularly seen by the strongly increasing fraction of H I-poor satellites as a function of halo mass. At a given halo mass, low-mass satellites have a greater chance of having their H I removed. Satellites that have not had most of their H I removed lie along similar relations to satellites in lower mass haloes. These results suggest that the process by which H I is removed from satellites in our simulations acts fairly quickly, and preferentially on smaller galaxies. Comparing these predictions to observations can help constrain such H I removal mechanisms.

6 H I DEFICIENCY

The H I richness of galaxies is observed to correlate with a variety of galaxy properties besides stellar mass. In the previous section, we showed that satellite galaxies in higher mass haloes were increasingly stripped of their H I. But even galaxies that still have substantial H I show correlations of their H I content with properties, such as environment (e.g. Cortese et al. 2011) and metallicity (Robertson et al. 2012). Here we examine the second-parameter trends of H I with environment, SFR and metallicity in our simulations to provide insights into the physical drivers that establish the H I content of galaxies.

To set a theoretical context for this discussion, we recall the equilibrium model for galaxy growth as described in Davé et al. (2012), which presents a simple physical scenario that can account for the relationships between key galaxy properties, as well as the scatter around those relationships. In the equilibrium model, accretion on to a galaxy from the cosmic web is fairly quickly processed into either stars or an outflow, resulting in a slowly evolving gas reservoir. This results in fairly tight relations between the stellar mass and SFR (the so-called main sequence; e.g. Davé 2008), gas-phase metallicity (Finlator & Davé 2008) and star-forming gas content (Davé et al. 2011b). The analytically predicted relations are in good agreement with full hydrodynamic simulation results.

Stochasticity in the inflow rate causes scatter about these equilibrium relations. As described in Davé et al. (2011b), an accretion event such as a minor merger will cause an increase in gas content that raises the SFR, while simultaneously lowering the metallicity. Conversely, a lull in accretion, or diminished accretion owing to a galaxy becoming a satellite in a larger halo, will cause the existing gas to be consumed, resulting in a lower gas content, a lower SFR and a higher metallicity. Hence, as an example, departures from the MZR are expected to be inversely correlated with the SFR. Such a trend has been observed (Lara-López et al. 2010; Mannucci et al. 2010) and is known as the fundamental metallicity relation; our simulations naturally predict this (Davé et al. 2011b). Our models analogously predict that deviations in the star-forming (i.e. molecular) gas content will inversely correlate with deviations in metallicity (Davé et al. 2011b).

We now extend this to consider the impact of inflow stochasticity on the H I reservoir of galaxies. The left-hand panels of Fig. 7 show the specific SFR (sSFR), the metallicity and the local galaxy density averaged over 1 Mpc spheres as a function of the stellar mass in our ezw simulation at $z = 0$. Metallicity here is computed as the SFR-weighted oxygen metallicity (Finlator & Davé 2008; Davé et al. 2011b). The observed MZR from SDSS (Tremonti et al. 2004) is shown as the solid black line with contours enclosing 95 per cent of the simulated galaxies indicated by the dashed lines. We have converted these metallicities to solar units assuming a solar oxygen abundance from Asplund et al. (2009), namely $12 + \log[O/H]_\odot = 8.70$. The green line shows a running median in each panel. We also show as the magenta line the median for the vzw simulation (without showing the individual points), for comparison.

The median sSFR drops slowly with M_* at both small and large stellar masses, with a peak at $M_* \sim 5 \times 10^9\,M_\odot$; there are minimal differences between the ezw and vzw models. The vzw line follows that in Davé et al. (2011a), who used lower resolution simulations that only probed down to $M_* \sim 10^9\,M_\odot$, while here the turnover is much more apparent. We note that SDSS observations do not indicate such a turnover (Salim et al. 2007), which reflects a generic problem in hierarchical models that SF in dwarfs peaks too early and is too low today (Weinmann et al. 2012).
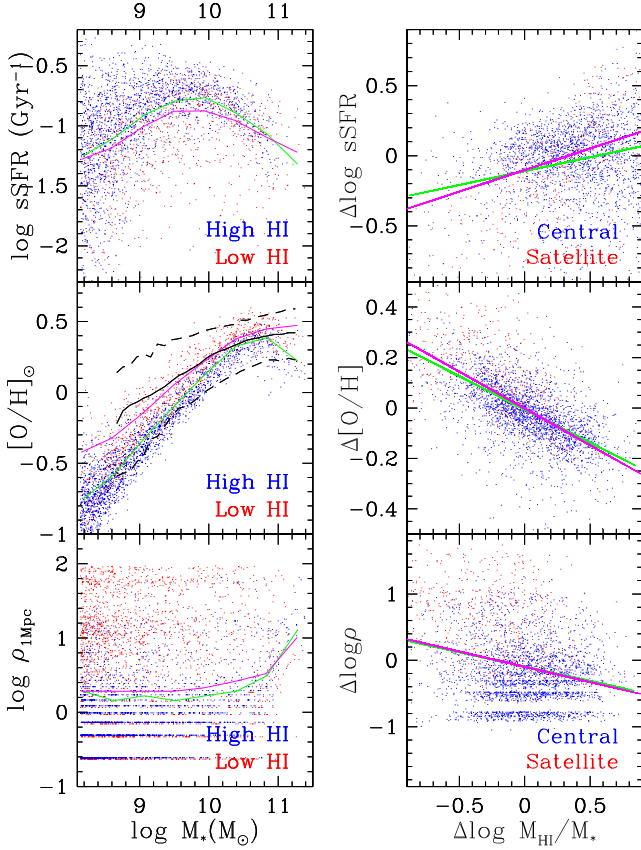
Figure 7. Left-hand panels: specific SFR, metallicity and environment versus stellar mass (top to bottom) from our ezw simulation at $z = 0$, with the blue and red points indicating galaxies that have above and below, respectively, the median H I richness at that given M_*. The green line shows the running median. The cyan line shows the median for the vzw model, although the points are not shown. The black solid line in the middle-left panel plots the median observed relation from Tremonti et al. (2004) with the dashed lines indicating the 95 per cent contours. H I-rich galaxies at a given M_* tend to have higher sSFR, lower metallicity and live in less dense regions. Right-hand panels: the deviation from the median sSFR, metallicity and environment (top to bottom) at a given M_* versus the deviation from the median H I richness. The magenta lines show power-law fits to the deviations, having slopes of 0.31, -0.26 and -0.56, respectively.

The metallicity is tightly correlated and increases with M_* until $M_* \sim 5 \times 10^{10}\,\mathrm{M}_\odot$. The ezw model produces a steeper MZR than the vzw model, for reasons discussed in Finlator & Davé (2008) and Davé et al. (2012). The metallicity $Z \propto 1/(1 + \eta)$ (in the absence of wind recycling), with $\eta \propto 1/\sigma \propto M_*^{-1/3}$ for all galaxies in the vzw model and $\eta \propto 1/\sigma^2 \propto M_*^{-2/3}$ for the ezw model at small masses. For $\eta \gg 1$, as at small masses, this toy model yields $Z \propto M_*^{1/3}$ for the vzw model and $Z \propto M_*^{2/3}$ for the ezw model. This represents an asymptotic slope, and since the ezw model follows the vzw model scalings at high masses, the actual MZR in the ezw model is not as steep as this, but it is clearly steeper than that in the vzw model. The vzw model appears to provide a better fit to the SDSS data, although metallicity measures are sufficiently uncertain that the ezw prediction cannot be ruled out (Ellison et al. 2008). Also, a recent determination of the MZR from direct metallicity measures by Andrews & Martini (2013) suggests a steeper MZR slope at small masses, closer to $Z \propto M_*^{1/2}$.

The simulated galaxies in Fig. 7 are colour-coded by H I richness: the blue points show galaxies that have a higher-than-median H I

richness at that stellar mass, and the red points show the converse. It is clear that H I richness has a strong second-parameter correlation with all these properties: H I-rich galaxies at a given mass also tend to have a higher sSFR, a lower metallicity and a lower local galaxy density.

To reiterate the physical scenario: when accretion occurs, it results in galaxies with higher H I richness than typical. This high H I richness also corresponds to a lower-than-normal gas-phase metallicity. This suggests that the H I content of galaxies can be an indicator of recent accretion, as argued from observations by Moran et al. (2012), who show that the low metallicities in H I-rich systems most strongly appear at the outskirts of discs.

We can quantify these trends using deviation plots, which we show in the corresponding right-hand panels of Fig. 7. This shows the difference between a galaxy's sSFR, metallicity and density, and the median value of these quantities at a given M_*, i.e. approximately the green line in the left-hand panels (although in detail we employ a spline fit to produce more smoothly varying deviations), versus the deviation in H I richness with its median value at that M_*. Points in these panels are colour-coded by centrals (blue) and satellites (red). The green lines show power-law fits to these deviations for the ezw model and the magenta lines show the same for the vzw model (though the individual points are not shown).

These right-hand panels quantify the second-parameter trends from the left-hand panels: galaxies with high H I content also have a high sSFR, a low metallicity and are in low-density environments. A power law provides a reasonable fit to the sSFR and metallicity data, with slopes of 0.31 and -0.26, respectively. The vzw model yields virtually identical slopes, even though the trends (such as the MZR) are noticeably different. This occurs because the trend in the scatter arises from inflow stochasticity, which has little to do with the outflows. The central and satellite galaxies do not show significantly different trends in sSFR or metallicity.

The metallicity dependence can be compared to recent work by Robertson et al. (2012), who looked at the H I deficiency parameter, DEF, relative to the deviation from the expected oxygen abundance for cluster and field galaxies, finding slopes of -0.25 ± 0.12 and -0.41 ± 0.14, respectively. Our simulated galaxies are more comparable to field galaxies, which show a higher slope but are still within their uncertainties. Robertson et al. (2012) compared to results from the simulations in Davé et al. (2011b), which yielded a similar slope. Although the definition of DEF differs from the deviation we plot here, the qualitative agreement in the trend indicates that the basic physical model of a slowly evolving equilibrium in gas content and metallicity appears to be broadly consistent with the observations.

Finally, we consider the environment in the bottom two panels of Fig. 7. The median local galaxy density is independent of mass until the very highest mass systems in our volume at $M_* \gtrsim 5 \times 10^{10}\,\mathrm{M}_\odot$. Unlike in the sSFR and metallicity deviation plots, where the centrals and satellites do not significantly deviate from one other, the environment deviation plot (bottom right) markedly separates the central and satellite galaxies. This environmental dependence arises because the environment can impact the inflow rate into satellites. High-density regions associated with haloes with masses $\gtrsim 10^{12}\,\mathrm{M}_\odot$ will contain substantial amounts of hot gas (Kereš et al. 2005, 2009; Gabor & Davé 2012) that can retard accretion (Dekel & Birnboim 2006; Davé et al. 2012). Hence, galaxies in such regions, particularly satellites, will have lower accretion rates compared to field galaxies at the same mass, and hence less H I. Overall, the full galaxy sample tends towards having less gas-rich galaxies in denser environments, but this is driven almost

entirely by the satellites. Fitting a power law formally yields a slope of -0.56 (shown as the magenta line), though a power law does not appear to be a particularly good fit.

In summary, the correlated deviations between the sSFR, metallicity and H I content provide quantitative constraints in the way in which galaxies oscillate about their equilibrium relations owing to stochastic accretion. The resulting deviation slope reflects the correlation between the instantaneous gas inflow rate, its conversion to neutral hydrogen and the infalling metallicity, as mediated by the local environment. Such constraints provide valuable discriminants between different physical scenarios for galaxy growth. The consistency with available observations, albeit preliminary and qualitative, indicates that the stochasticity in the inflow as expected from cosmological accretion is able to broadly explain the second-parameter trends observed in the relationship between the H I content and other physical galaxy properties.

7 $\Omega_{\mathrm{H\,I}}$ EVOLUTION

Most neutral hydrogen in the cosmos exists in and around galaxies, between the regime where self-shielding happens in the outskirts of discs and where gas becomes molecular dominated within star-forming regions. While it is currently infeasible to probe the evolution of H I in galaxies via 21 cm emission to high redshifts,[3] other avenues have been employed to measure cosmic H I evolution, such as the abundance of damped Lyα systems (DLAs). In this section, we examine the evolution of the cosmic H I density in our simulations directly from the galaxy population, and compare it with various observational determinations.

Fig. 8 shows the evolution of $\Omega_{\mathrm{H\,I}}$ from $z = 5 \rightarrow 0$ in our ezw model, for all galaxies above our approximate H I mass resolution limit of $M_{\mathrm{H\,I}} > 1.4 \times 10^8\,\mathrm{M_\odot}$ (solid blue line), as well as for galaxies with $M_{\mathrm{H\,I}} > 10^9\,\mathrm{M_\odot}$ (dashed) and $M_{\mathrm{H\,I}} > 10^{10}\,\mathrm{M_\odot}$ (dotted). For comparison, we also show the evolution of the stellar mass density Ω_* as the red line. A sample of observations are also shown, as detailed in the caption. Observations of $\Omega_{\mathrm{H\,I}}$ at $z > 2$ from Prochaska & Wolfe (2009) and Noterdaeme et al. (2012) are computed from DLA abundances, while at $0.5 < z < 2$ the data are from Rao et al. (2006), who used strong Mg II systems as a proxy for DLAs. Low-redshift ($z \lesssim 0.1$) data come from 21 cm emission surveys, including Haynes et al. (2011, ALFALFA), Braun (2012) and Delhaize et al. (2013). Note that our prescription for determining self-shielding calibrated to radiative line transfer simulations as described in Section 2 does reasonably well predicting $\Omega_{\mathrm{H\,I}}$ at $z = 0$; we do not adjust any parameters to match these data as was done in Popping et al. (2009).

Our simulations predict that $\Omega_{\mathrm{H\,I}}$ is remarkably constant from $z = 5 \rightarrow 0$, essentially unchanging to within 50 per cent over these 12 billion years. This is in contrast to the dramatic increase of Ω_* over the same time interval; over 98 per cent of all stars in this model form since $z = 4$, with more than 80 per cent since $z = 2$ (in broad agreement with recent observational estimates; e.g. Leitner 2012). It is also quite different from the rapid evolution in the cosmic SFR density (e.g. Hopkins & Beacom 2006; Fardal et al. 2007) and molecular gas fraction (Tacconi et al. 2010, 2012; Geach et al. 2011), which both increase by an order of magnitude or so out to $z = 2$ (though see Narayanan et al. 2012, who argue for less growth in the observations owing to variations in the CO-to-H$_2$

---

[3] This is a key goal of the upcoming LADUMA survey using the MeerKAT array (Holwerda et al. 2011).

Figure 8. $\Omega_{\mathrm{H\,I}}$, the fraction of cosmic mass density in H I, from $z = 5 \rightarrow 0$ from our ezw simulation. The blue solid line shows $\Omega_{\mathrm{H\,I}}$ from all resolved galaxies ($M_* > 1.4 \times 10^8\,\mathrm{M_\odot}$), the blue dashed and dotted lines show $\Omega_{\mathrm{H\,I}}$ in galaxies with $M_* > 10^9\,\mathrm{M_\odot}$ and $M_* > 10^{10}\,\mathrm{M_\odot}$, respectively. The cyan line shows $\Omega_{\mathrm{H\,I}}$ from extrapolating a Schechter function fit to the HIMF at each redshift down to $M_{\mathrm{H\,I}} = 10^7\,\mathrm{M_\odot}$. The magenta line shows the total H I mass density in the entire volume. For comparison, the red line shows the mass density in stars in this model. Observations are shown at $2 \lesssim z \lesssim 4.5$ from Prochaska & Wolfe (2009) (open circles), at $2 \lesssim z \lesssim 3.5$ from Noterdaeme et al. (2012) (open squares), at $0.5 \lesssim z \lesssim 2$ from Rao, Turnshek & Nestor (2006) (filled squares), at $0.02 \lesssim z \lesssim 0.1$ from Delhaize et al. (2013) (filled circles) and at $z = 0$ from Haynes et al. (2011) (open triangle) and Braun (2012) (filled triangle).

---

conversion factor). This emphasizes that H I represents a transient reservoir in the journey of gas from the ionized IGM to stars forming deep within galaxies, and is not directly proportional to the amount of stars being formed or to the amount of molecular gas present. This cautions against overinterpreting quantities such as the 'H I SF efficiency', i.e. the SFR divided by the H I mass, since the two quantities are not directly related.

Subdividing $\Omega_{\mathrm{H\,I}}$ into different H I mass bins, we see that at $z = 0$, most of the cosmic H I is in galaxies with $10^9 < M_{\mathrm{H\,I}} < 10^{10}\,\mathrm{M_\odot}$. Going back in time, H I shifts towards lower mass galaxies; by $z = 2$, only half of the H I is in $M_{\mathrm{H\,I}} > 10^9\,\mathrm{M_\odot}$ systems. In contrast, only a very small portion of the cosmic H I ever resides in galaxies with $M_{\mathrm{H\,I}} > 10^{10}\,\mathrm{M_\odot}$, since higher mass galaxies quickly drop off in their H I richness (Fig. 4). Thus, quenching mechanisms, which primarily affect high-mass galaxies (e.g. Gabor & Davé 2012), are expected to have almost no effect on $\Omega_{\mathrm{H\,I}}$. We also show, as the magenta line, the total H I mass density in the entire simulation volume. The difference between the solid magenta and blue lines reflects H I that is not within resolved galaxies (including the diffuse IGM). This extra contribution is a couple of a per cent at low redshift but rises to ∼50 per cent at $z = 5$, which again reflects the increasing contribution of small (unresolved) galaxies to the global H I budget at high redshifts.

Observations of $\Omega_{\mathrm{H\,I}}$ likewise indicate very little evolution from $z \sim 4 \rightarrow 0$. Different observational tracers generally agree to within a factor of 2, and together indicate essentially no change in $\Omega_{\mathrm{H\,I}}$ for the past 12 billion years, in broad agreement with our predictions. Our simulation begins to underpredict $\Omega_{\mathrm{H\,I}}$ at $z \gtrsim 3$ when compared with the DLA data. The discrepancy could be physical or

numerical. A physical explanation could be that lower metallicity galaxies at early epochs will tend to have less efficient conversion of their atomic gas to molecular gas owing to lower cooling rates and harder interstellar radiation, and hence our locally calibrated prescription for R_{mol} from Leroy et al. (2008) may not be appropriate (see Section 8).

A numerical explanation for the discrepancy may be that we do not fully resolve galaxies with $M_* < 1.4 \times 10^8 \, M_\odot$, and there may be substantial contributions to Ω_{HI} from the very smallest galaxies. At low z, the low-mass end of the HIMF is fairly shallow, so the expected contribution from lower mass galaxies is small. But at higher redshifts, Fig. 3 shows that the slope becomes substantially steeper, meaning that the additional contribution from unresolved systems could be large.

We can crudely correct for this by fitting a Schechter function to the HIMF at each redshift as we did in Section 3, and then integrating down to some chosen lower mass limit. As an illustration, we integrate down to $M_{HI} = 10^7 \, M_\odot$, which is around the lowest mass observable at low z in large surveys such as ALFALFA. The result of this Schechter fit extrapolation to $M_{HI} = 10^7 \, M_\odot$ is shown as the cyan line in Fig. 8. This results in a negligible correction at low z, but at higher redshifts the correction can be up to a factor of 3, which agrees better with the high-z DLA results. We caution that this exercise is intended to be illustrative, since our M_{HI} limit was chosen rather arbitrarily, and it may not correspond to the effective H i masses probed by DLA systems. Furthermore, our Schechter function fits become increasingly uncertain at higher redshifts, owing to the lack of dynamic range in our simulations. Nonetheless, this illustrates that plausible corrections down to lower H i masses can bring our predicted Ω_{HI} into better agreement with the observations by preferentially increasing the high-redshift H i mass density. Given the crudeness of our prescriptions for determining H i content, agreement at this level is encouraging.

In summary, our simulations generally predict a very slowly evolving cosmic H i mass density, in broad agreement with the observations. However, we note that the HIMF actually evolves rather considerably (Fig. 3), with many more low-M_{HI} galaxies and fewer high-M_{HI} ones at high z. It is something of a coincidence that these two variations roughly cancel when summed to give the global H i mass density. Nonetheless, this emphasizes the transient nature of the H i reservoir around galaxies, which does not build up hierarchically in the way that the stellar mass does, but instead responds to evolution in the ionizing background and the ISM physical conditions.

8 NUMERICAL ROBUSTNESS

8.1 Resolution convergence

In this study we have employed some of the highest resolution hydrodynamic simulations of random cosmological volumes ever evolved down to $z = 0$. Nonetheless, it is important to assess whether our resolution is sufficient to robustly predict the H i properties of galaxies. Here we examine our two most basic H i statistics, namely the HIMF and the H i richness versus stellar mass, in a simulation that has 2×384^3 particles in a $48 \, h^{-1}$Mpc volume (r48n384), described in our previous work (e.g. Oppenheimer et al. 2010; Davé et al. 2011a,b). This simulation has eight times poorer mass resolution and two times poorer spatial resolution, albeit in a volume that is 3.4 times larger. We use the identical momentum-


Figure 9. Comparison of the HIMF (top panel) and H i richness (bottom panel) in our fiducial $32 \, h^{-1}$Mpc, 2×512^3 particle simulation (green) compared to a $48 \, h^{-1}$Mpc, 2×384^3 particle simulation (red) with eight times poorer mass resolution. Resolution convergence is good for the HIMF, while the H i richness is decreased in the higher resolution simulation, but is still within the formal uncertainties.

driven wind (vzw) and cooling model in both simulations.[4] We also renormalize the metagalactic background by a flux factor of 1.5, as computed in Davé et al. (2010).

Fig. 9 shows the HIMF (top) and the M_{HI}/M_* ratio versus M_* (bottom) in these two simulations, r32n512vzw in green and r48n384vzw in red. These are analogous to Figs 3 and 4, respectively, and the observations are plotted as in those figures.

The HIMF (top panel) shows very good resolution convergence, at least over this somewhat modest range in mass resolution probed here. The two mass functions are essentially identical until below the resolution limit (indicated by the colour-coded dotted line), where the low-resolution simulation starts to become significantly

---

[4] We use the vzw wind model because we already have the r48n384vzw simulation in hand, and there are usually only minor differences between the vzw and our now-favoured ezw models.
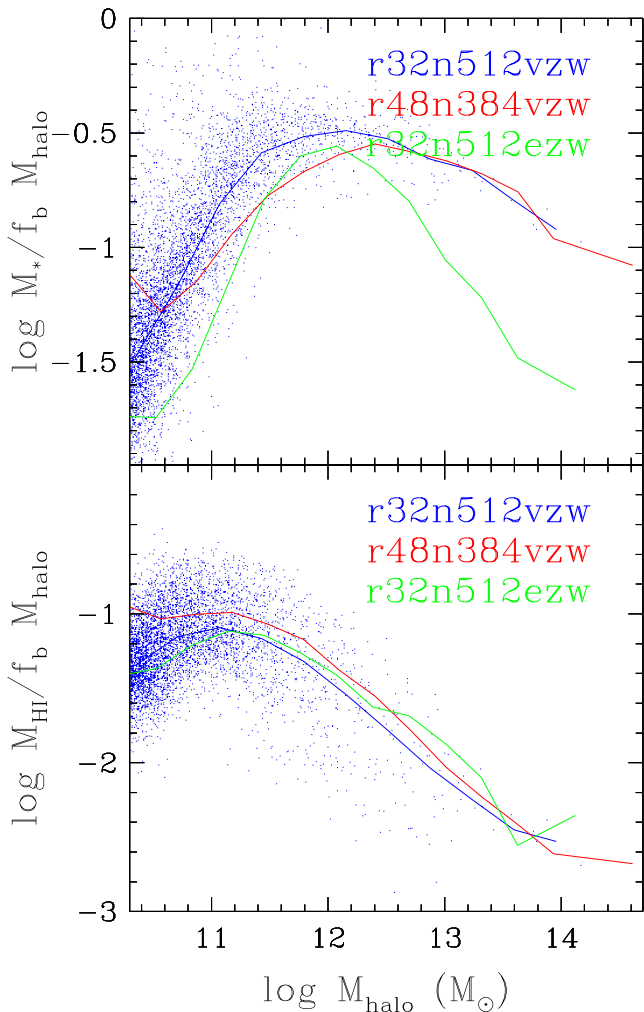
Figure 10. The mass fraction of stars (top panel) and H I (bottom) relative to the expected baryonic mass in the halo ($= f_b M_{halo}$, where $f_b = 0.164$ is the baryon fraction), versus M_{halo}, for central galaxies in our high-resolution vzw simulation (r32n512vzw, blue), low-resolution vzw simulation (r48n384vzw, red) and our fiducial ezw simulation (r32n512ezw, green). Increasing the resolution yields higher stellar masses and lower H I content.

incomplete. This also explicitly demonstrates that our chosen H I mass resolution limit is appropriate when considering the HIMF.

For the H I fraction, the resolution convergence is not quite as good. The running median line for the high-resolution simulation lies slightly below the lower resolution one, indicating that the H I richness decreases with increasing resolution. The effect is roughly ~50 per cent over the factor of 2 spatial (or factor of 8 mass) resolution, and suggests that even higher resolution simulations are required to fully assess the resolution convergence.

To examine the origin of the lack of convergence, we show in Fig. 10 the stellar (top) and H I (bottom) mass relative to the halo baryonic mass (i.e. the halo mass multiplied by the baryon fraction f_b), as a function of the halo mass. We consider only central galaxies since satellite masses are less correlated with their host halo masses. We plot the higher resolution vzw simulation in blue, with points, while for the lower resolution simulation we plot only the running median in red. For reference, we also plot our fiducial ezw simulation in green, showing that it lowers the stellar baryon fraction at

both high and low halo masses but that it has a modest impact on the H I halo fraction at a given halo mass, relative to vzw.

At a given halo mass, the stellar mass is increased and the H I mass is decreased at higher resolution. Both effects contribute roughly equally to lowering the H I richness. It appears that higher resolution simulations are more effective at converting their H I into stars, possibly because the increased resolution enables faster cooling through the H I regime. This suggests that even higher resolution simulations are needed to properly converge both these quantities, though the qualitative trends and interpretations are not impacted by resolution.

Overall, the resolution convergence for H I properties is good but not ideal. For a given M_* the H I content is very robust, but the number density of galaxies at a given H I mass is less well converged. For that reason, Ω_{HI} is also slightly uncertain, as there is an ~30 per cent difference between the $z = 0$ value predicted by the high- and low-resolution simulations. We do not yet know whether these differences reflect a resolution effect or a volume effect; we would need a wider suite of simulations to assess these issues. While these discrepancies are not trivial, they do not significantly affect the main conclusions of this paper.

8.2 Molecular gas prescription

Our prescription for computing the molecular gas fraction relies on an empirical calibration from THINGS (Leroy et al. 2008) fit to their combined spiral sample. It does not explicitly include any dependence on metallicity or redshift. However, there are good theoretical reasons to believe that the conversion of atomic to molecular gas should depend on metallicity, since dust is an important catalyst for H_2 formation, and possibly on redshift since the typical ISM properties in high-redshift galaxies are different from those today. Indeed, several groups have developed theoretical models to investigate this transition, and particularly its dependence on metallicity (Krumholz, McKee & Tumlinson 2008; Gnedin & Kravtsov 2011).

In this section, we employ the analytic prescription developed by Krumholz et al. (2008) to separate atomic and molecular hydrogen, and assess its impact on the HIMF. The prescription follows that detailed in Krumholz & Gnedin (2011):

$$f_{H_2} = 1 - \frac{0.75\,s}{1 + 0.25\,s}, \tag{12}$$

where

$$s = \frac{\ln(1 + 0.6\chi + 0.01\chi^2)}{0.6\tau_c}, \tag{13}$$

$$\chi = 0.756(1 + 3.1Z^{0.365}) \tag{14}$$

and

$$\tau_c = \Sigma\sigma_d/\mu_H. \tag{15}$$

Here, Z is the metallicity in solar units, σ_d is the cross-section for dust, which we assume to be equal to $Z \times 10^{-21}$ cm^2, $\mu_H = 2.3 \times 10^{-24}$ g is the mean mass per H nucleus and Σ is the gas surface density that we take to be the gas density times the SPH smoothing kernel size. This prescription is applied to each star-forming gas particle since molecular gas is only expected to form substantially in star-forming regions; the neutral fraction in non-star-forming gas is not altered and is computed as described earlier (Section 2).

Figure 11. Comparison of star-forming gas particle neutral fractions computed using the observationally derived empirical pressure relation for the H$_2$ content from Leroy et al. (2008) versus the theoretically derived relation from Krumholz & Gnedin (2011), at $z = 0$ (top panel) and $z = 3$ (bottom). Particles are binned and colour-coded by the mean metallicity within each bin. The theoretically determined neutral fraction is higher for lower metallicity gas, and there is more such gas at higher redshifts.

Fig. 11 shows a scatter plot, colour-coded by metallicity, of the neutral fraction computed using our empirical method from Leroy et al. (2008) versus the theoretical method from Krumholz & Gnedin (2011), at $z = 0$ (top) and $z = 3$ (bottom). As expected, lower metallicity gas has a higher neutral fraction in the theoretical prescription, showing an order-of-magnitude change in the neutral fraction at a given ISM pressure over the 2 dex span in metallicity. At $z = 0$, there is little overall systematic shift in the neutral fraction, indicating that the theoretical method is well calibrated to the empirical one for local solar-metallicity spirals from Leroy et al. (2008). However, at $z = 3$ where there is more low-metallicity star-forming gas, the theoretical method yields systematically higher neutral fractions.

Fig. 12 shows the HIMF using the empirical (blue) and theoretical (magenta) methods to compute molecular fractions. The solid lines show the results at $z = 0$. Although there is a systematic metallicity dependence as shown in the previous plot, the overall resulting HIMF at $z = 0$ shows essentially no difference between the two methods, even as a function of mass. Note that the range in metallicities from our most massive to least massive galaxies spans over a dex (see Fig. 7). Nonetheless, the region from where most


Figure 12. Comparison of the HIMF using the observationally derived empirical pressure relation for H$_2$ fraction from Leroy et al. (2008) (blue lines) versus the theoretically derived relation from Krumholz & Gnedin (2011) (red lines). The solid lines show results at $z = 0$, where there is essentially no difference between these prescriptions. The dashed lines show $z = 3$ results, where the metallicity dependence in the theoretical relation generally yields a higher neutral fraction.

of the neutral gas originates is typically well outside the region of vigorous SF that determines a galaxy's metallicity. In these regions, typically at densities of $\lesssim 1$ cm^{-3}, the impact of metallicity on the molecular gas content is evidently not significant.

In contrast, at $z = 3$ the HIMF shows a systematic offset towards higher H I content using the theoretical method. This arises because our simulations show a mild evolution in galaxy metallicity of roughly a factor of 2 at a given stellar mass out to these redshifts (Davé et al. 2011b), which is consistent with observations out to $z \sim 2.5$ (e.g. Erb et al. 2006), although higher redshift data may suggest a more rapid dropoff (Mannucci et al. 2010, though see also Richard et al. 2011). Still, the HIMF is only mildly higher using the metallicity-dependent theoretical prescription, with a maximum increase in H I mass of ~ 1.5 times. We conclude that there is a modest uncertainty in our predictions owing to which the H$_2$ prescription is employed, and it primarily affects higher-z predictions. For instance, it could help to reconcile the predicted $\Omega_{\rm H I}(z)$ with the DLA-derived values at $z \sim 3$.

8.3 Hydrodynamics methodology

Our GADGET-2 simulations employ the EC formulation of SPH (Springel 2005). While this formulation has the advantage of simultaneously conserving entropy and energy, it has been shown to have significant deficiencies particularly when surface instabilities are present (e.g. Agertz et al. 2007), in that it appears to add an artificial surface tension that prevents the disruption of cold clumps moving through a dense medium. This could influence H I content by, for instance, suppressing gas stripping in infalling satellites. Various newer formulations of SPH have been developed that improve this aspect (e.g. Read & Hayfield 2012; Hopkins 2013; Saitoh & Makino 2013), but these are just now being incorporated into full cosmological codes. Furthermore, all forms of SPH tend to be overly dissipative in shear flows, despite explicit attempts to suppress shear viscosity (Balsara 1989). Hence, galaxy discs in

SPH tend to be overly condensed (Torrey et al. 2012) barring the inclusion of strong feedback (Brook et al. 2012), which could affect the conversion of H I to H$_2$ within the disc. These issues beg the question whether H I results using EC SPH might be severely compromised.

Regarding disc sizes, it is worth noting that our simulations do possess the sort of strong feedback that results in larger discs that are more consistent with observations, as we showed using high-resolution cosmological zoom simulations at $z \sim 2$ (Angles-Alcazar et al. 2013). Nonetheless, disc radii in the cosmological simulations we use here are still somewhat smaller than observed at $z = 0$, though not nearly as far off as in simulations without outflows. It is unclear whether this has a substantial effect on the atomic-to-molecular transition, because our prescription relies on the ISM pressure that is determined locally from the subgrid two-phase ISM model (Springel & Hernquist 2003b) rather than from the global properties of the disc. Even when we use the theoretical prescription described in the previous section, the surface density is calculated from individual particles' gas densities and smoothing lengths, rather than global disc properties. Nonetheless, there may be some cumulative effect from the improper density profile of the disc. It is difficult to assess this directly without either doing much higher resolution zoom simulations, in which case it is more difficult to compile statistics, or using a different code such as AREPO (Torrey et al. 2012), in which case one gets discs with larger radii but in other ways still quite unlike observed discs.

Regarding the issue of surface instabilities, the inaccuracies here are expected to manifest primarily in satellite galaxies moving through the intrahalo medium. Such satellites should likely have more gas stripped than our EC SPH algorithm predicts. However, as we showed in Fig. 5, satellites are a strongly subdominant contribution to the HIMF at all masses probed here. Hence, even if two-phase instabilities should be causing all satellites to lose their H I envelope, this will affect the HIMF by 25 per cent at most. A secondary effect is that centrals partly build up their gas mass by accreting gas-rich satellites. However, Kereš et al. (2005) and others have showed that for star-forming galaxies, the contribution to the gas content from infalling satellites (i.e. merging) is small compared to that accreted directly from the IGM or in very small lumps. This has been shown for total gas, not necessarily H I, and in any case assessing the origin of H I is tricky since it is a transient reservoir rather than a steadily accumulating component such as stars. Hence, we leave a detailed examination of the origin of H I for the future, and preliminarily conclude that it seems unlikely that the inaccuracies in EC SPH play a major role in establishing the H I properties of galaxies, though this will require more detailed testing with different codes to establish conclusively.

9 SUMMARY AND DISCUSSION

The H I content of galaxies offers a unique glimpse into baryon cycling processes of inflow and outflow that are believed to govern galaxy evolution. Hydrodynamic simulations can provide a way to connect H I to such baryon cycling processes, as well as to help interpret observations within a hierarchical galaxy formation context. In this work, we have examined the H I properties of galaxies in cosmological hydrodynamic simulations using GADGET-2 with EC SPH employing four different outflow prescriptions. Our key conclusions can be summarized as follows.

(i) The HIMF is strongly affected by the inclusion of galactic outflows. Without outflows, our simulations yield too many galax-ies at small H I masses and too few at large H I masses. Our model with a constant wind speed (cw) introduces a feature in the HIMF around $M_{\rm H\,I} \sim 10^{10}$ M$_\odot$ that is not seen in the observations. Our models with momentum-driven wind scalings produce an HIMF that broadly agrees with the observations, particularly at the low-mass end down to our mass resolution limit of $\sim 10^8$ M$_\odot$. Switching to energy-driven scalings at low masses ($\sigma_{\rm gal} < 75$ km s^{-1}) further improves the agreement and generates a low-mass end HIMF slope of -1.3. The HIMF is slightly above the data at all masses, but likely within the systematic uncertainties regarding how we compute H I masses for our galaxies. The low-mass end slope becomes progressively steeper with redshift out to $z \sim 3$, and then remains constant at around -2 at higher redshifts. As an aside, we showed that the ezw model also provides an excellent match, within the quoted uncertainties, to the GSMF.

(ii) All models show that low-mass galaxies are more H I rich, relative to their stellar content, than high-mass galaxies. The exact shape and amplitude of the relation depend on the wind model. All the models broadly agree with the GASS data that probe down to $M_* \sim 10^{10}$ M$_\odot$, but the different wind models predict different trends to lower masses. The constant-wind model predicts significantly lower gas richness than our vzw (momentum-driven wind scalings) or ezw models at $M_* \lesssim 10^{10}$ M$_\odot$, reflecting a strong suppression of inflow owing to energetic winds (van de Voort et al. 2011) that likewise suppresses the sSFRs (Davé et al. 2011a). Current observations favour a continuing increase of gas richness to lower masses, and our ezw model again fares slightly better than vzw at low masses, and both match much better than the cw model. There is essentially no redshift evolution in the H I richness as a function of the stellar mass, showing that the evolution of the HIMF is interconnected with the evolution of the GSMF.

(iii) Galaxies with a high H I content tend to have lower gas-phase metallicities and higher SFRs at a given M_*. The deviation in sSFR and metallicity versus the deviation in H I richness can be fitted by power laws with slopes of 0.31 and -0.26, respectively. The latter slope compares favourably with existing observations of the H I deficiency parameter versus metallicity deviation. These deviation trends can be understood within the context of a model in which galaxies tend to have equilibrium relations in SFR, metallicity and gas content, perturbed by stochastic accretion events that result in correlated deviations in SFR, metallicity, star-forming gas mass and H I mass. This suggests that higher H I richness at a given M_* is an indicator of recent accretion events in the outskirts of galaxies that stimulates SF.

(iv) Environment plays an important role in governing H I content, causing an increasing suppression of the H I content of satellites in more massive haloes. The suppression is rapid, as indicated by the bimodal distribution of the H I fraction at a given halo mass. Lower mass satellites are more likely to have their H I removed. The deviation of H I with environment is anticorrelated with the deviation in H I richness, driven by H I-poor satellite galaxies in denser regions.

(v) The global H I mass density evolves slowly from $z \sim 5 \rightarrow 0$, in broad agreement with the observations of DLAs and other measures. This is in contrast to the stellar mass density, which grows substantially over that interval, highlighting the transient nature of the H I reservoir around galaxies. Today, the majority of the cosmic H I mass is in galaxies with $10^9 \lesssim M_{\rm H\,I} \lesssim 10^{10}$ M$_\odot$. Our fiducial simulations predict $\Omega_{\rm H\,I}$ that is lower than observed at $z \gtrsim 2$, which could reflect either the inability of these models to resolve the low-mass galaxies that are a significant contribution to global H I at high redshifts or an evolving efficiency

of converting H I to H_2 owing to generally lower metallicities at high z.

(vi) We briefly examine numerical issues, including resolution convergence over a modest span of a factor of 8 in mass resolution between two simulations with momentum-driven winds. The HIMF is well converged, but the H I richness is not converged at the ~50 per cent level between the two simulations. Using a theoretically based molecular gas criterion also results in up to ~50 per cent difference in the HIMF at high redshifts. There may be further uncertainties owing to our use of EC SPH, though it is expected to be small if the differences are limited to satellite galaxies. In general, it appears that our predictions are not free of numerical uncertainties, but seem to be robust at the factor of 2 level, and hence the qualitative conclusions presented above are generally robust.

Simultaneously matching the observed low-mass end of the galactic stellar mass function and the low-mass end slope of the HIMF has been challenging in simulations as well as in SAMs of galaxy formation, where one has much more freedom to adjust the physical model (Mo et al. 2005; Lagos et al. 2011; Lu et al. 2012; Lu et al., in preparation). Fundamentally, in most models this arises because the low-mass end slope of the dark matter halo mass function is quite steep, which is exacerbated by the fact that low-mass galaxies are more H I rich. Mo et al. (2005) argue that it is not possible to match the low-mass end of the HIMF in a very broad range of CDM-based galaxy formation models if one only makes two basic assumptions: first, that the original gas distribution has the same specific angular momentum as the dark matter halo, which it conserves as it forms an exponential disc (as in Mo, Mao & White 1998), and that the angular momentum distribution of the gas and stars within the disc does not change subsequent to their accretion on to the disc; secondly, that stars only form in gas above a critical surface density Σ_{crit}, as observed (Kennicutt 1989). Many SAMs make these same assumptions. Mo et al. (2005) populate some fraction of the CDM haloes expected in a ΛCDM model with exponential gas discs and assume that any region originally above Σ_{crit} has its surface density lowered to Σ_{crit}. Making the (conservative) assumption that half of the gas is neutral, they overproduce the observed HIMF at the low-mass end by a factor of >5. Lu et al. (2012) confirm this result by adjusting all their SAM parameters using a Bayesian inference technique to match the observed low-redshift K-band galaxy luminosity function, and find that they overpredict the observed HIMF by a large factor.

Lu et al. (in preparation) expand on this work, by adjusting their parameters to simultaneously match both the observed low-z K-band galaxy luminosity function and the HIMF. They find that this is not possible using standard models that include the above two assumptions, but that it is possible if one does any one of the following: (1) reduce Σ_{crit} by about an order of magnitude, (2) allow the gas disc to maintain an exponential profile while stars are forming, i.e. allow gas from large radii beyond the star-forming radius to lose angular momentum by some process and move inwards, (3) include some process that preheats the gas before it can be accreted into a dark matter halo or have the formula by which one determines that the expected gas accretion rate be different for small mass haloes. The first option violates observations by a substantial margin, and furthermore for smaller, lower metallicity galaxies one observes Σ_{crit} to increase, and not to decrease (Bolatto et al. 2013). Our simulations assume a three-dimensional critical SF density of $n_H = 0.13 \, cm^{-3}$, which approximately matches the observed Σ_{crit} (Springel & Hernquist 2003a), and hence this is not likely to be the difference between our models and the SAMs. The third possibility

is also unlikely in our simulations, since we do not explicitly add any sort of preheating, nor do we explicitly vary the dark matter accretion rates from the usual ΛCDM expectations (e.g. Neistein & Dekel 2008; Faucher-Giguere, Kereš & Ma 2011). However, it is possible that large-scale structure formation itself creates its own form of preheating (Mo et al. 2005). Hence, if the arguments presented in Lu et al. (in preparation) apply to our simulations, then since we simultaneously match the HIMF and the stellar mass function, we suspect that some angular momentum transport processes, either physical or numerical, causes the gas in our simulated discs to move inwards.

SAMs with different recipes from those in Lu et al. (in preparation) have had some success matching the HIMF. Obreschkow et al. (2009), Power, Baugh & Lacey (2010), Cook et al. (2010) and Lagos et al. (2011) all present various independent SAMs that are broadly successful, though curiously they all tend to overpredict the observed HIMF by ~3 times in the vicinity of $M_{H I} \sim 10^8 - 10^9 \, M_\odot$. It is beyond the scope of this paper to determine which ingredients in these SAMs differ from the simple assumptions in Lu et al. (2012), or whether the persistent discrepancy in the dwarf mass range is a fundamental issue or simply a matter of more parameter tuning. Unfortunately, this illustrates a difficulty with SAMs that, owing to the many plausible choices to achieve similar successes at matching data, it can be difficult to extract robust interpretations for the key physical ingredients required. This is particularly true since the dynamics of the gas is crucial for modelling a transient reservoir such as H I, but SAMs do not model the gas dynamics directly, and must rely on assumptions to track this. Nonetheless, SAMs offer the advantage of exploring a wide range of parameter space to better understand all the various avenues by which models can reproduce observations. Hence, SAMs and hydrodynamical simulations provide complementary information about how to build a successful model to match H I and other properties of galaxies and their associated gas.

With the emergence of numerous major new radio facilities such as the JVLA, MeerKAT (Karoo Array Telescope) and ASKAP, and eventually the SKA itself in the next decade, the future of H I studies looks promising. Simulations such as the ones presented here, hopefully improved substantially in the coming years, will be crucial for providing a context to interpret these observations within the broader multiwavelength landscape of hierarchical galaxy evolution. At this point, it seems equally important both to probe the evolution of the H I content of galaxies and to go deeper around nearby galaxies to better understand how H I connects to the cosmic web. Such data promise to provide interesting new constraints on the key processes of galaxy evolution such as inflows, outflows and wind recycling. This paper represents a first step towards providing a comprehensive framework for interpreting the wealth of forthcoming H I data.

REFERENCES

Agertz O. et al., 2007, MNRAS, 380, 963

Andrews B., Martini P., 2013, ApJ, 765, 140

Anglés-Alcázar D., Davé R., Özel F., Oppenheimer B. D., 2013, ApJ, preprint (arXiv:1303.6959)

Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, ARA&A, 47, 481

Baldry I. K., Glazebrook K., Driver S. P., 2008, MNRAS, 388, 945

Balsara D. S., 1989, PhD thesis, Univ. of Illinois

Battisti A. J. et al., 2012, ApJ, 744, 93

Becker G. D., Hewett P. C., Worseck G., Prochaska J. X., 2012, MNRAS, MNRAS, 430, 2067

Bolatto A. D., Wolfire M., Leroy A. K., 2013, ARA&A, preprint (arXiv:1301.3498)

Booth C. M., Schaye J., 2009, MNRAS, 398, 53

Bouche N. et al., 2010, ApJ, 718, 1001

Braun R., 2012, ApJ, 749, 87

Brook C. B., Stinson G., Gibson B. K., Roskar R., Wadsley J., Quinn T., 2012, MNRAS, 419, 771

Catinella B. et al., 2010, MNRAS, 403, 683

Catinella B. et al., 2012, A&A, 544, 65

Chabrier G., 2003, PASP, 115, 763

Cook M., Evoli C., Barausse E., Granato G. L., Lapi A., 2010, MNRAS, 402, 941

Cortese L., Catinella B., Boissier S., Boselli A., Heinis S., 2011, MNRAS, 415, 1797

Dalla Vecchia C., Schaye J., 2008, MNRAS, 387, 1431

Davé R., 2008, MNRAS, 385, 147

Davé R., 2009, in Jogee S., Marinova I., Hao L., Blanc G. A., eds, ASP Conf. Ser. Vol. 419, Galaxy Evolution: Emerging Insights and Future Challenges. Astron. Soc. Pac., San Francisco, p. 347

Davé R., Oppenheimer B. D., Katz N., Kollmeier J. A., Weinberg D. H., 2010, MNRAS, 408, 2051

Davé R., Oppenheimer B. D., Finlator K. M., 2011a, MNRAS, 415, 11

Davé R., Finlator K. M., Oppenheimer B. D., 2011b, MNRAS, 416, 1354

Davé R., Finlator K. M., Oppenheimer B. D., 2012, MNRAS, 421, 98

Dekel A., Birnboim Y., 2006, MNRAS, 368, 2

Dekel A. et al., 2009, Nat, 457, 451

Delhaize J., Meyer M., Stavely-Smith L., Boyle B., 2013, MNRAS, 433, 1398

Di Matteo T., Springel V., Hernquist L., 2005, Nat, 433, 604

Duffy A. R., Kay S. T., Battye R. A., Booth C. M., Dalla Vecchia C., Schaye J., 2012, MNRAS, 420, 2799

Ellison S. L., Patton D. R., Simard L., McConnachie A. W., 2008, ApJ, 672, L107

Erb D. K., Shapley A. E., Pettini M., Steidel C. C., Reddy N. A., Adelberger K. L., 2006, ApJ, 644, 813

Fardal M. A., Katz N., Weinberg D. H., Davé R., Katz N., 2007, MNRAS, 379, 985

Faucher-Giguere C. A., Kereš D., Dijkstra M., Hernquist L., Zaldarriaga M., 2010, ApJ, 725, 633

Faucher-Giguere C. A., Kereš D., Ma C.-P., 2011, MNRAS, 417, 2982

Finlator K., Davé R., 2008, MNRAS, 385, 2181

Finlator K., Davé R., Papovich C., Hernquist L., 2006, ApJ, 639, 672

Gabor J. M., Davé R., 2012, MNRAS, 427, 1816

Gabor J. M., Davé R., Oppenheimer B. D., Finlator K. M., 2011, MNRAS, 417, 2676

Geach J. E., Smail I., Moran S. M., MacArthur L. A., Lagos C. d. P., Edge A. C., 2011, ApJ, 730, 19

Giovanelli R. et al., 2005, AJ, 130, 2598

Gnedin N. Y., Kravtsov A., 2011, ApJ, 728, 88

Haardt F., Madau P., 2001, in Neumann D. M., Van J. T. T., eds, Proc. XXXVIth Rencontres de Moriond, Recent Results of XMM–Newton and Chandra, available at: http://wwwdapnia.cea.fr/Conferences/Morion astro 2001/index.html

Haynes M. P., Giovanelli R., Chincarini G. L., 1984, ARA&A, 22, 445

Haynes M. P. et al., 2011, AJ, 142, 170

Hinshaw G. et al., 2009, ApJS, 180, 225

Holwerda B. W. et al., 2012, in Tuffs R. J., Popescu C. C., eds, Proc. IAU Symp. 284, The Spectral Energy Distribution of Galaxies. Cambridge Univ. Press, Cambridge, p. 496

Hopkins P. F., 2013, MNRAS, 428, 2840

Hopkins A. M., Beacom J. F., 2006, ApJ, 651, 142

Hopkins P. F., Quataert E., Murray N., 2012, MNRAS, 412, 3522

Huang S., Haynes M. P., Giovanelli R., Brinchmann J., Stierwalt S., Neff S. G., 2012, AJ, 143, 133

Katz N., Weinberg D. H., Hernquist L., 1996, ApJS, 105, 19

Kennicutt R. C., 1989, ApJ, 344, 685

Kennicutt R. C., 1998, ApJ, 498, 541

Kereš D., Katz N., Weinberg D. H., Davé R., 2005, MNRAS, 363, 2

Kereš D., Katz N., Fardal M., Davé R., Weinberg D. H., 2009, MNRAS, 395, 160

Kirkman D., Tytler D., Lubin D., Charlton J., 2007, MNRAS, 376, 1227

Krumholz M. R., Gnedin N. Y., 2011, ApJ, 729, 36

Krumholz M. R., McKee C. F., Tumlinson J. T., 2008, ApJ, 689, 865

Lagos C. D. P., Baugh C. M., Lacey C. G., Benson A. J., Kim H.-S., Power C., 2011, MNRAS, 418, 1649

Lara-López M. A. et al., 2010, A&A, 521, L53

Leitner S. N., 2012, ApJ, 745, 149

Leroy A. K., Walter F., Brinks E., Bigiel F., de Blok W. J. G., Madore B., Thornley M. D., 2008, AJ, 136, 2782

Lu Y., Mo H. J., Katz N., Weinberg M. D., 2012, MNRAS, 421, 1779

Mannucci F., Cresci G., Maiolino R., Marconi A., Gnerucci A., 2010, MNRAS, 408, 2115

McGaugh S. S., Schombert J. M., de Blok W. J. G., Zagursky M. J., 2010, ApJ, 708, L14

McKee C. F., Ostriker J. P., 1977, ApJ, 218, 148

Meyer M. J. et al., 2004, MNRAS, 350, 1195

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

Mo H. J., Yang X., van den Bosch F. C., Katz N., 2005, MNRAS, 363, 1155

Moldar S. M., Hearn N., Haiman Z., Bryan G., Evrad A. E., Lake G., 2009, ApJ, 696, 1640

Moran S. M. et al., 2012, ApJ, 745, 66

Murray N., Quataert E., Thompson T. A., 2005, ApJ, 618, 569

Murray N., Quataert E., Thompson T. A., 2010, ApJ, 709, 191

Narayanan D., Bothwell M., Davé R., 2012, MNRAS, 426, 1178

Neistein E., Dekel A., 2008, MNRAS, 388, 1792

Noterdaeme P. et al., 2012, A&A, 547, L1

Obreschkow D., Croton D., De Lucia G., Khochfar S., Rawlings S., 2009, ApJ, 698, 1467

Oppenheimer B. D., Davé R., 2008, MNRAS, 387, 577

Oppenheimer B. D., Davé R., Kereš D., Katz N., Kollmeier J. A., Weinberg D. H., 2010, MNRAS, 406, 2325

Popping A., Davé R., Braun R., Oppenheimer B. D., 2009, A&A, 504, 15

Power C., Baugh C. M., Lacey C. G., 2010, MNRAS, 406, 43

Prochaska J. X., Wolfe A. M., 2009, ApJ, 696, 1543

Rao S. M., Turnshek D. A., Nestor D. B., 2006, ApJ, 636, 610

Read J. I., Hayfield T., 2012, MNRAS, 422, 3037

Richard J., Jones T., Ellis R., Stark D. P., Livermore R., Swinbank M., 2011, MNRAS, 413, 643

Robertson P., Shields G. A., Blanc G. A., 2012, ApJ, 748, 48

Rupke D. S., Veilleux S., Sanders D. B., 2005, ApJS, 160, 115

Saitoh T. R., Makino J., 2013, ApJ, 768, 44

Salim S. et al., 2007, ApJS, 173, 267

Schmidt M., 1959, ApJ, 129, 243

Springel V., 2005, MNRAS, 364, 1105

Springel V., Hernquist L., 2003a, MNRAS, 339, 289

Springel V., Hernquist L., 2003b, MNRAS, 339, 312

Tacconi L. J. et al., 2010, Nat, 463, 781

Tacconi L. J. et al., 2012, ApJ, 768, 74

Torrey P., Vogelsberger M., Sijacki D., Springel V., Hernquist L., 2012, MNRAS, 427, 2224

Tremonti C. A. et al., 2004, ApJ, 613, 898

van de Voort F., Schaye J., Booth C. M., Haas M. R., Dalla Vecchia C., 2011, MNRAS, 414, 2458

Verner D. A., Ferland G. J., 1996, ApJS, 103, 467

Weinmann S. M., Pasquali A., Oppenheimer B. D., Finlator K., Mendel J. T., Crain R. A., Maccio A. V., 2012, MNRAS, 426, 2797

Wiersma R. P., Schaye K., Theuns T., Dalla Vecchia C., Tornatore L., 2009, MNRAS, 399, 574

Zwaan M. A., Meyer M. J., Staveley-Smith L., Webster R. L., 2005, MNRAS, 359, L30

This paper has been typeset from a T_EX/L^AT_EX file prepared by the author.