

The distribution of atomic hydrogen in EAGLE galaxies: morphologies, profiles, and H I holes

Yannick M. Bahé,^{1*} Robert A. Crain,^{2,3} Guinevere Kauffmann,¹ Richard G. Bower,⁴ Joop Schaye,³ Michelle Furlong,⁴ Claudia Lagos,^{5,6} Matthieu Schaller,⁴ James W. Trayford,⁴ Claudio Dalla Vecchia^{7,8} and Tom Theuns⁴

¹Max-Planck-Institut für Astrophysik, Karl-Schwarzschild Str. 1, D-85748 Garching, Germany

²Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK

³Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands

⁴Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK

⁵European Southern Observatory, Karl-Schwarzschild-Str. 2, D-85748 Garching, Germany

⁶International Centre for Radio Astronomy Research (ICRAR), M468, University of Western Australia, 35 Stirling Hwy, Crawley, WA 6009, Australia

⁷Instituto de Astrofísica de Canarias, C/Vía Láctea s/n, E-38205 La Laguna, Tenerife, Spain

⁸Departamento de Astrofísica, Universidad de La Laguna, Av. del Astrofísico Francisco Sánchez s/n, E-38206 La Laguna, Tenerife, Spain

Accepted 2015 November 12. Received 2015 September 29; in original form 2015 May 11

ABSTRACT

We compare the mass and internal distribution of atomic hydrogen (H I) in 2200 present-day central galaxies with $M_{\text{star}} > 10^{10} M_{\odot}$ from the 100 Mpc EAGLE ‘Reference’ simulation to observational data. Atomic hydrogen fractions are corrected for self-shielding using a fitting formula from radiative transfer simulations and for the presence of molecular hydrogen using an empirical or a theoretical prescription from the literature. The resulting neutral hydrogen fractions, $M_{\text{H I}+\text{H}_2}/M_{\text{star}}$, agree with observations to better than 0.1 dex for galaxies with M_{star} between 10^{10} and $10^{11} M_{\odot}$. Our fiducial, empirical H₂ model based on gas pressure results in galactic H I mass fractions, $M_{\text{H I}}/M_{\text{star}}$, that agree with observations from the GASS survey to better than 0.3 dex, but the alternative theoretical H₂ formula from high-resolution simulations leads to a negative offset in $M_{\text{H I}}/M_{\text{star}}$ of up to 0.5 dex. Visual inspection of mock H I images reveals that most H I discs in simulated H I-rich galaxies are vertically disturbed, plausibly due to recent accretion events. Many galaxies (up to 80 per cent) contain spuriously large H I holes, which are likely formed as a consequence of the feedback implementation in EAGLE. The H I mass–size relation of all simulated galaxies is close to (but 16 per cent steeper than) observed, and when only galaxies without large holes in the H I disc are considered, the agreement becomes excellent (better than 0.1 dex). The presence of large H I holes also makes the radial H I surface density profiles somewhat too low in the centre, at $\Sigma_{\text{H I}} > 1 M_{\odot} \text{pc}^{-2}$ (by a factor of $\lesssim 2$ compared to data from the Bluedisk survey). In the outer region ($\Sigma_{\text{H I}} < 1 M_{\odot} \text{pc}^{-2}$), the simulated profiles agree quantitatively with observations. Scaled by H I size, the simulated profiles of H I-rich ($M_{\text{H I}} > 10^{9.8} M_{\odot}$) and control galaxies ($10^{9.1} M_{\odot} > M_{\text{H I}} > 10^{9.8} M_{\odot}$) follow each other closely, as observed.

Key words: methods: numerical – galaxies: formation – galaxies: ISM – galaxies: structure.

1 INTRODUCTION

Radio observations have revealed the presence of atomic hydrogen (H I) in the Milky Way (see Dickey & Lockman 1990), as well as in many other galaxies (see Walter et al. 2008 and references therein). Although the median H I mass fraction $M_{\text{H I}}/M_{\text{star}}$ is only ~ 0.1

for Milky Way mass galaxies ($M_{\text{star}} \approx 10^{10.5} M_{\odot}$), the presence of substantial scatter means that the ratio can exceed unity in individual cases (Catinella et al. 2010). This H I reservoir is believed to be fuel for future star formation (e.g. Prochaska & Wolfe 2009; Davé et al. 2010; van de Voort et al. 2012), which makes the ability to correctly model its structure and evolution an integral part of the wider quest to better understand galaxy formation.

The long time-scales of galaxy formation ($\gtrsim 1$ Gyr) imply that observations are effectively limited to one point in time for any

*E-mail: ybahe@mpa-garching.mpg.de

individual galaxy, so that studying the evolution of galactic gas necessarily involves theoretical modelling. In the ‘Semi-Analytic Modelling’ (SAM) approach (e.g. Kauffmann, White & Guiderdoni 1993; Guo et al. 2011; Fu et al. 2013), the evolution of baryonic galaxy components is described by analytic equations that are combined with an underlying dark matter distribution from N -body simulations (e.g. Springel et al. 2005b; Boylan-Kolchin et al. 2009) or the extended Press–Schechter formalism (Bond et al. 1991; Bower 1991). SAMs have become increasingly refined over time, and a number of authors have used them to study various aspects of $H\text{I}$ in galaxies such as its evolution (Lagos et al. 2011; Popping, Somerville & Trager 2014), radial distribution (Fu et al. 2013; Wang et al. 2014) and origin in early-type galaxies (Lagos et al. 2014).

However, SAMs are not able to predict the detailed structure of gas within and around galaxies: its accretion, for example, is typically modelled in an ad hoc way without fully accounting for the filamentary structure of the intergalactic medium (but see Benson & Bower 2010 for a counter-example). This motivates the use of cosmological hydrodynamical simulations, which model the accretion and outflows of gas from galaxies self-consistently, and at (potentially) high spatial resolution. Further benefits include the ability to trace the thermodynamic history of individual fluid elements (e.g. Kereš et al. 2005; van de Voort et al. 2011; Nelson et al. 2013), and that they permit the study of satellite galaxies without additional assumptions (e.g. Bahé & McCarthy 2015).

A number of authors have studied low-redshift $H\text{I}$ with cosmological hydrodynamical simulations in the past. Popping et al. (2009) successfully reproduced the observed distribution of $H\text{I}$ column densities over seven orders of magnitude, as well as the $H\text{I}$ two-point correlation function (see also Duffy et al. 2012; Rahmati et al. 2013a). The sensitivity of $H\text{I}$ to supernova feedback, and its evolution over cosmic time, was explored by Davé et al. (2013) and Walker et al. (2014), while Cunnamea et al. (2014) and Rafiee-antsoa et al. (2015) investigated the influence of the group/cluster environment on $H\text{I}$ (see also Stinson et al. 2015).

However, a common problem of these simulations has been their inability to produce galaxies whose stellar component agrees with observations. In particular, angular momentum from infalling gas was typically dissipated too quickly and too severely to form realistic discs (Steinmetz & Navarro 1999), and ‘overcooling’ (e.g. Katz, Weinberg & Hernquist 1996) manifested itself in galaxy stellar mass functions that are too high at the massive end (e.g. Crain et al. 2009; Oppenheimer et al. 2010; Lackner et al. 2012).

In the recent past, several groups have developed simulations which are able to avoid these problems. Incorporation of efficient supernova feedback – in a physical and numerical sense – and/or increased resolution has led to the formation of realistic disc galaxies (e.g. Governato et al. 2007, 2010; McCarthy et al. 2012; Aumer et al. 2013; Marinacci, Pakmor & Springel 2014). The inclusion of additional feedback from accreting supermassive black holes (‘AGN feedback’), on the other hand, has reduced the overcooling problem at the high-mass end and led to more accurate stellar masses of simulated galaxies (Rosas-Guevara et al. 2015, see also Springel, Di Matteo & Hernquist 2005a, Sijacki et al. 2007, Booth & Schaye 2009, and Vogelsberger et al. 2013). With these and other improvements, the Evolution and Assembly of GaLaxies and their Environments (EAGLE) project (Schaye et al. 2015, see also Crain et al. 2015) has yielded a cosmologically representative population of galaxies with realistic properties such as stellar masses and sizes (see also Vogelsberger et al. 2014 for the *ILLUSTRIS* simulation). EAGLE has also been shown to broadly reproduce e.g. the observed colour distribution of galaxies (Trayford et al. 2015) at $z \sim 0$, as

well as the redshift evolution of the stellar mass growth and star formation rates (SFRs; Furlong et al. 2015a).

The Lagrangian smoothed particle hydrodynamics (SPH) formalism adopted by EAGLE makes it, in principle, possible to study directly the physics governing the accretion and outflow of atomic hydrogen in simulated galaxies. Unlike the $z \approx 0$ stellar mass function and sizes, $H\text{I}$ properties were not taken into account when calibrating the EAGLE galaxy formation model. It is therefore uncertain whether the distribution of $H\text{I}$ is modelled correctly: as Crain et al. (2015) have shown, even stellar masses and sizes are sufficiently independent of each other that reproducing observations of one does not necessarily imply success with the other. Comparing the $H\text{I}$ properties of simulated galaxies to observations therefore offers an opportunity to directly test the galaxy formation model, as well as being a necessary step to ascertain the extent to which simulation predictions are trustworthy.

At high redshift ($z \geq 1$), Rahmati et al. (2015) have shown that the column density distribution function and covering fractions of $H\text{I}$ absorbers in EAGLE agree with observations, while Lagos et al. (2015) demonstrated that EAGLE galaxies contain realistic amounts of molecular hydrogen (H_2) both at $z = 0$ and across cosmic history. Here, we conduct a series of detailed like-with-like comparisons between EAGLE and recent low-redshift $H\text{I}$ observations including the *Galex* *Arecibo* SDSS Survey (GASS; Catinella et al. 2010, 2013) and the *Bluedisk* project (Wang et al. 2013, 2014). Our aim is to analyse the distribution of $H\text{I}$ within individual $z = 0$ galaxies; the cosmological distribution of $H\text{I}$ in EAGLE will be investigated separately (Crain et al., in preparation).

The remainder of this paper is structured as follows. In Section 2, we review the key characteristics of the EAGLE project, and give an overview of the GASS (Catinella et al. 2010) and CO Legacy Database (COLD GASS; Saintonge et al. 2011) surveys in Section 3. Our $H\text{I}$ modelling scheme is then described in Section 4, followed by a comparison of galaxy-integrated neutral hydrogen and $H\text{I}$ masses to observations in Section 5. Section 6 analyses the internal distribution of $H\text{I}$ in the simulated galaxies, including a comparison of $H\text{I}$ surface density profiles to *Bluedisk* data. Our results are summarized and discussed in Section 7. All masses and distances are given in physical units unless specified otherwise. A flat Λ CDM cosmology with Hubble parameter $h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1}) = 0.6777$, dark energy density parameter $\Omega_\Lambda = 0.693$ (dark energy equation-of-state parameter $w = -1$), and matter density parameter $\Omega_M = 0.307$ as in Planck Collaboration XVI (2014) is used throughout this paper. The EAGLE simulations adopt a universal Chabrier (2003) stellar initial mass function with minimum and maximum stellar masses of 0.1 and $100 M_\odot$, respectively.

2 THE EAGLE SIMULATIONS

2.1 Simulation characteristics

The Evolution and Assembly of GaLaxies and their Environments (EAGLE) project consists of a large suite of many cosmological hydrodynamical simulations of varying size, resolution, and sub-grid physics prescriptions. They are introduced and described in detail by Schaye et al. (2015) and Crain et al. (2015); here we only summarize the main characteristics that are particularly relevant to our study.

The largest simulation (Ref-L100N1504 in the terminology of Schaye et al. 2015), upon which our analysis here is based, fills a cubic box of side length 100 comoving Mpc (‘cMpc’) with $N = 1504^3$

dark matter particles ($m_{\text{DM}} = 9.7 \times 10^6 M_{\odot}$) and an initially equal number of gas particles ($m_{\text{gas}} = 1.81 \times 10^6 M_{\odot}$). The simulation was started at $z = 127$ from cosmological initial conditions (Jenkins 2013), and evolved to $z = 0$ using a modified version of the GADGET-3 code (Springel 2005). These modifications include a number of hydrodynamics updates collectively referred to as ‘Anarchy’ (Dalla Vecchia in preparation, see also Hopkins 2013, appendix A of Schaye et al. 2015, and Schaller et al. 2015) which eliminate most of the problems associated with ‘traditional’ SPH codes related to the treatment of surface discontinuities (e.g. Agertz et al. 2007; Mitchell et al. 2009) and artificial gas clumping (e.g. Nelson et al. 2013).

The gravitational softening length is 0.7 proper kpc (‘pkpc’) at redshifts $z < 2.8$, and 2.66 ckpc at earlier times. In the warm interstellar medium, the Jeans scales are therefore marginally resolved, but the same is not true for the cold molecular phase. For this reason, the simulation imposes a temperature floor $T_{\text{eos}}(\rho)$ on gas with $n_{\text{H}} > 0.1 \text{ cm}^{-3}$, in the form of a polytropic equation of state $P \propto \rho^{\gamma}$ with index $\gamma = 4/3$ and normalized to $T_{\text{eos}} = 8 \times 10^3 \text{ K}$ at $n_{\text{H}} = 10^{-1} \text{ cm}^{-3}$ (see Schaye & Dalla Vecchia 2008 and Dalla Vecchia & Schaye 2012 for further details). In addition, gas at densities $n_{\text{H}} \geq 10^{-5} \text{ cm}^{-3}$ is prevented from cooling below 8000 K.

The EAGLE simulation code includes significantly improved sub-grid physics prescriptions. These include element-by-element radiative gas cooling (Wiersma, Schaye & Smith 2009a) in the presence of the cosmic microwave background and an evolving Haardt & Madau (2001) UV/X-ray background, reionization of hydrogen at $z = 11.5$ and helium at $z \approx 3.5$ (Wiersma et al. 2009b), star formation implemented as a pressure law (Schaye & Dalla Vecchia 2008) with a metallicity-dependent density threshold (Schaye 2004), stellar mass-loss and chemical enrichment on an element-by-element basis (Wiersma et al. 2009b), as well as energy injection from supernovae (Dalla Vecchia & Schaye 2012) and accreting supermassive black holes (AGN feedback; Rosas-Guevara et al. 2015; Schaye et al. 2015) in thermal form.

For a detailed description of how these sub-grid models are implemented in EAGLE, the interested reader is referred to Schaye et al. (2015). However, three aspects in the implementation of energy feedback from star formation merit explicit mention here. First, because the feedback efficiency cannot be predicted from first principles, its strength was calibrated to reproduce the $z \approx 0$ galaxy stellar mass function and sizes. Secondly, the feedback parametrization depends only on local gas quantities, in contrast to e.g. the widely-used practice of scaling the parameters with the (global) velocity dispersion of a galaxy’s dark matter halo (e.g. Okamoto et al. 2005; Oppenheimer & Davé 2006; Davé et al. 2013; Vogelsberger et al. 2013; Puchwein & Springel 2013). Finally, star formation feedback in EAGLE is made efficient not by temporarily disabling hydrodynamic forces or cooling for affected particles (e.g. Springel & Hernquist 2003; Stinson et al. 2006), but instead by stochastically heating a small number of particles by a temperature $\Delta T = 10^{7.5} \text{ K}$ (Dalla Vecchia & Schaye 2012). These details can be expected to influence in non-trivial ways the galactic distribution of H I (see e.g. Davé et al. 2013), so that an examination of this diagnostic also informs our understanding of the impact of this scheme on the structure of the simulated ISM.

2.2 Galaxy selection

From the 100 cMpc EAGLE simulation Ref-L100N1504, we select our $z = 0$ target galaxies as self-bound subhaloes – identified using the SUBFIND algorithm (Dolag et al. 2009, see also Springel

et al. 2001) – with a stellar mass of $M_{\text{star}} \geq 10^{10} M_{\odot}$. This limit ensures that individual galaxies are well resolved ($\gg 1000$ baryon particles) and that our sample is directly comparable to the observational GASS and Bluedisk surveys. Crain et al. (in preparation) will present the full H I mass function in EAGLE extending down to much smaller galaxies. Stellar masses are computed as the total mass of all gravitationally bound star particles within a spherical aperture of 30 kpc, centred on the particle for which the gravitational potential is minimum. Schaye et al. (2015) showed that this definition mimics the Petrosian mass often used by optical surveys.

Note that we only select central galaxies – i.e. the most massive subhalo in a friends-of-friends halo – because satellites are subject to additional complex environmental processes that can impact upon their H I content (e.g. Fabello et al. 2012; Catinella et al. 2013; Zhang et al. 2013, see also Bahé et al. 2013). We focus here on testing the arguably more fundamental accuracy of the simulations for centrals; the H I properties of EAGLE satellites will be discussed elsewhere (Marasco et al., in preparation). In total, we have a sample of 2200 galaxies, the vast majority of which (2039) have stellar masses below $10^{11} M_{\odot}$.

3 THE GASS AND COLD GASS SURVEYS

Before describing our H I analysis as applied to EAGLE, we now give a brief overview of the GASS and COLD GASS surveys, which will be compared to our simulations below. We also describe our approach for comparing EAGLE in a consistent way to these observations. For clarity, we will describe the third main survey used in our work, Bluedisk (Wang et al. 2013), in Section 6.4.2 where its results are compared to predictions from EAGLE.

3.1 The Galax Arcibo SDSS Survey (GASS)

The Galax Arcibo SDSS Survey (GASS)¹ (Catinella et al. 2010, 2013) was designed to provide an unbiased census of the total H I content in galaxies with stellar mass $M_{\text{star}} > 10^{10} M_{\odot}$. Measuring this observationally does not require very high spatial resolution and can therefore be achieved with single-dish observations on e.g. the Arcibo telescope. However, an important issue is that of galaxy selection: blind surveys such as ALFALFA (Giovanelli et al. 2005) are naturally more likely to detect abnormally H I-rich than -poor galaxies because the volume over which the latter can be detected is small (Catinella et al. 2010; Huang et al. 2012). The galaxies in GASS are therefore selected only by stellar mass, and observed until either the 21-cm line from H I is detected or an upper limit of $M_{\text{HI}}/M_{\text{star}} \approx 0.015$ has been reached.² Out of the 760 galaxies in the full GASS sample, centrals are selected by cross-matching to the Yang et al. (2012) SDSS group catalogue (see Catinella et al. 2013 for details), which leaves us with 386 galaxies with $10^{10} M_{\odot} \leq M_{\text{star}} \leq 10^{11} M_{\odot}$ (and 522 at $M_{\text{star}} \geq 10^{10} M_{\odot}$). We note that the equivalent EAGLE sample is almost an order of magnitude larger ($N = 2083$ and 2200, respectively), because of the larger effective volume.

In order to make the comparison between EAGLE and GASS fair, it is important to compute M_{HI} for EAGLE galaxies as done in observations, i.e. by integrating over the same range in projected radius and line-of-sight distance. For the former, we use a fixed

¹ Data available at <http://www.mpa-garching.mpg.de/GASS>.

² This limit is fixed to $M_{\text{HI}} = 10^{8.7} M_{\odot}$ for galaxies with $M_{\text{star}} < 10^{10.5} M_{\odot}$, so that the gas fraction detection threshold increases towards lower stellar masses.

value of 70 kpc, which roughly corresponds to the Arecibo L-Band Feed Array (ALFA) full width at half-maximum (FWHM) beam size of ~ 3.5 arcmin (Giovanelli et al. 2005) at the median redshift of the GASS sample, $\bar{z} = 0.037$ (Catinella et al. 2010). The line of sight is taken as the simulation z -coordinate; we include all particles (including those outside haloes) with peculiar velocity relative to the mass-weighted velocity of the galaxy subhalo in the range $[-400, +400]$ km s $^{-1}$ to approximately match what was done in GASS. A comparison of this integration range with simple spherical shell apertures can be found in Appendix A2, which confirms that masses obtained with this ‘GASS-equivalent’ method agree well with the mass of H I inside a (3D) aperture of 70 kpc, but exceed those measured inside a 30 kpc aperture at both the high- and low- M_{HI} end by typically up to a factor of 2.

3.2 CO Legacy Database for GASS (COLD GASS)

To complement the GASS data base with information on the molecular hydrogen content of galaxies, the COLD GASS³ survey (Saintonge et al. 2011) observed a randomly selected subset of ~ 250 galaxies from the GASS sample in CO with the IRAM 30-m telescope. Similarly to GASS, galaxies were observed until either the CO (1–0) line was detected or an upper limit equivalent to an H₂ mass fraction of ~ 1.5 per cent (for galaxies with $M_{\text{star}} > 10^{10.6} M_{\odot}$) or an absolute H₂ mass of $10^{8.8} M_{\odot}$ (for galaxies with $M_{\text{star}} < 10^{10.6} M_{\odot}$) was achieved.

A detailed comparison of EAGLE to results from COLD GASS is presented by Lagos et al. (2015). Here, we combine the results from GASS and COLD GASS to obtain observational constraints on the *total* neutral hydrogen mass in galaxies, which we compare to predictions from EAGLE in Section 5.1. For simplicity, we adopt the same particle selection as described above for GASS: this is justified because H₂ is concentrated more strongly towards the galaxy centre than H I and the COLD GASS survey is designed to measure the total H₂ masses of its galaxies (see Saintonge et al. 2011 for more details). Neutral hydrogen masses obtained from the simulations with the relatively large aperture matched to GASS can therefore be meaningfully compared to the sum of H I and H₂ masses from GASS and COLD GASS, respectively.

4 H I MODELLING

The EAGLE simulation output itself only contains the mass of hydrogen in each gas particle, but not how much of this is in ionized (H II), atomic (H I) or molecular (H₂) form.⁴ Although it is possible to separate these self-consistently using radiation transport and detailed chemical network modelling (e.g. Pawlik & Schaye 2008; Altay et al. 2011; Christensen et al. 2012; Rahmati et al. 2013a; Richings, Schaye & Oppenheimer 2014a,b; Walch et al. 2015), the computational expense of dynamically coupling these techniques to the simulation and e.g. calculate SFRs directly from the H₂ phase is unfeasibly high for a 100 cMpc simulation like EAGLE. Although we cannot, therefore, make truly self-consistent predictions for the individual hydrogen phases, we can still gain insight by employing

an approximation scheme in post-processing to calculate the H I mass of gas particles, as follows.

4.1 Neutral and ionized hydrogen

First, we compute the fraction of hydrogen in each gas particle that is *neutral* (H I and H₂). For this, we use the ionization fitting formula of Rahmati et al. (2013a), which was calibrated using (smaller) simulations with detailed radiation transport modelling.⁵ Their prescription relates the total ionization rate (photo- plus collisional ionization) to that from the UV background, for which we adopt a value⁶ of $\Gamma_{\text{UVB}} = 8.34 \times 10^{-14}$ s $^{-1}$ (Haardt & Madau 2001), accounting for self-shielding. Not taken into account, however, is the (difficult to constrain) effect of local stellar radiation, which Rahmati et al. (2013b) found to affect dense H I systems even at $z = 0$. In Section 5.1, we show that the resulting neutral gas masses are in good agreement with observational constraints.

4.2 Atomic and molecular hydrogen (H I/H₂)

In a second step, we then model the fractions of neutral hydrogen in molecular (H₂) and atomic form (H I) with two different approaches. Our fiducial method, similar to what was done by Altay et al. (2011), Duffy et al. (2012), and Davé et al. (2013), is to exploit the empirical relation between gas pressure and molecular fraction (Wong & Blitz 2002; Blitz & Rosolowsky 2006) which can be measured observationally on scales comparable to the resolution of EAGLE. This approach is approximately self-consistent, because star formation is also implemented based on pressure (Schaye & Dalla Vecchia 2008) and observational evidence strongly suggests a link between the two (e.g. Leroy et al. 2008; Krumholz, McKee & Tumlinson 2009; Bigiel et al. 2011; Huang & Kauffmann 2014, but see the theoretical work of Glover & Clark 2012).

From observations of 11 nearby, non-interacting galaxies spanning almost one decade in total metallicity and three decades in pressure, Blitz & Rosolowsky (2006, hereafter BR06) derived the H₂/H I fraction in terms of the mid-plane gas pressure P as

$$R_{\text{mol}} \equiv \frac{\Sigma_{\text{H}_2}}{\Sigma_{\text{H I}}} = \left(\frac{P}{P_0} \right)^{\alpha}, \quad (1)$$

with best fit parameters⁷ $P_0/k_B = 4.3 \times 10^4$ cm $^{-3}$ K and $\alpha = 0.92$. Assuming that the molecular and atomic phases have the same scaleheight, R_{mol} is also equal to the ratio between the *volume* densities of H₂ and H I. We furthermore assume that neutral hydrogen only has a contribution from H₂ in particles with a non-zero SFR, the density threshold for which is motivated by whether physical conditions allow the formation of a cold molecular phase (Schaye 2004).

⁵ For particles within 0.5 dex of the imposed equation of state, we assume a fixed temperature of $T = 10^4$ K when calculating collisional ionization and recombination with this prescription.

⁶ This value is larger by a factor of ~ 3 than the more recent determination by Haardt & Madau (2012). We have tested both, and found no significant impact on the H I results presented here.

⁷ Leroy et al. (2008) studied a somewhat larger sample of 23 galaxies, and found a best-fitting normalization $P_0/k_B = 1.7 \times 10^4$ cm $^{-3}$ K and exponent $\alpha = 0.8$; this parametrization was used by Duffy et al. (2012) and Davé et al. (2013). We show in Appendix A1 that the difference between applying these two parameterisations to our simulated galaxies is negligible. The same is true for the effect of including three additional interacting galaxies in the analysis of BR06.

³ Data available at http://www.mpa-garching.mpg.de/COLD_GASS.

⁴ The radiative cooling prescription of EAGLE takes into account that only a fraction of the gas is neutral. However, these ratios are computed without accounting for self-shielding (see Rahmati et al. 2013a and Schaye et al. 2015 for further details) and can therefore not be used directly in our present work, where we study the highly self-shielded regime of galaxy interiors.




Figure 1. Neutral hydrogen mass fractions for our simulated galaxies as predicted by the Rahmati et al. (2013a) fitting formula (blue line, shaded bands show the 1σ uncertainty and 50 per cent scatter, respectively). For comparison, observational data from the combined GASS and COLD GASS surveys are shown as grey symbols, upward (downward) facing triangles differ in that non-detections are set to zero (upper limits). Large triangles show the observed medians, small ones the 75th percentile of the distribution. The light blue triangle shows an additional lower limit from H I masses in the full GASS survey (see text for details). The neutral hydrogen masses in EAGLE agree with observational constraints to within 0.1 dex, although there are large uncertainties on the observational median at $M_{\text{star}} > 10^{11} M_{\odot}$.

As an alternative, we also consider the theoretically motivated H_2 partitioning scheme of Gnedin & Kravtsov (2011, hereafter GK11), which is based on high-resolution simulations with an explicit treatment of the formation and destruction of H_2 . For more details on this scheme and its implementation in EAGLE, we refer the interested reader to Lagos et al. (2015), where this prescription was shown to yield good agreement between the H_2 content of EAGLE galaxies and observations. Further approaches of modelling the H I in EAGLE are explored in Appendix A1; these include simple prescriptions such as ignoring H_2 altogether or assuming a fixed ratio of $m_{H_2}/m_{H I} = 0.3$ for each particle (as in Popping et al. 2009), both of which give similar results as our fiducial empirical BR06 method described above. In the same place, we also test the alternative theoretical prescription by Krumholz (2013) as implemented into EAGLE by Lagos et al. (2015).

5 NEUTRAL AND ATOMIC HYDROGEN FRACTIONS COMPARED TO OBSERVATIONS

5.1 Neutral hydrogen fractions

We begin by showing in Fig. 1 the total neutral hydrogen fraction, i.e. $M_{H I+H_2}/M_{\text{star}}$, of our simulated galaxies as a function of stellar mass M_{star} (blue). The solid dark blue line shows the running median of the distribution. The dark and light shaded bands indicate the statistical 1σ uncertainty on the median and the 50 per cent scatter, respectively, i.e. they extend from f_{low} to f_{high} where $f = M_{H I+H_2}/M_{\text{star}}$ and $f_{\text{low}(\text{high})} = \tilde{f} + (P_{15.9(84.1)} - \tilde{f})/\sqrt{N}$; \tilde{f} here denotes the median and P_n the n th percentile of the distribution in a bin with N galaxies. This prediction is compared to observational constraints from the intersection of the GASS and COLD

GASS surveys shown as grey symbols. Both have a large fraction of non-detections: only 46 per cent of central galaxies targeted in both surveys are detected in H I and CO, although the majority of galaxies (83 per cent) are detected in at least one component. To bracket the resulting uncertainty on the observed median, we have computed it with non-detections set both to zero (giving lower limits, shown by upward facing triangles) and the observational upper limit (downward facing triangles). At $M_{\text{star}} < 10^{11} M_{\odot}$, both approaches differ by less than 0.2 dex. The 75th percentile of the observed distribution is analogously shown by small triangles. The impact of non-detections is much smaller here (< 0.1 dex).

The median neutral hydrogen fraction predicted by EAGLE agrees remarkably well with observational constraints, deviating by $\lesssim 0.1$ dex in the regime $\log_{10}(M_{\text{star}}/M_{\odot}) = [10.0, 11.0]$ in the sense that the simulated galaxies contain, in general, slightly too little neutral gas. The 75th percentiles agree at a similar level, but without a consistent sign of the deviation. For the most massive galaxies ($M_{\text{star}} > 10^{11} M_{\odot}$), the large observational uncertainties induced by frequent non-detections prevent strong statements on the accuracy of the simulation prediction for the median neutral fraction, but the 75th percentiles are well-constrained observationally and show a significant shortfall of the simulation, by ~ 0.3 dex.

We note that in the second most massive M_{star} bin, $\log_{10}(M_{\text{star}}/M_{\odot}) = [11.0, 11.25]$, less than 50 per cent of (central) galaxies in the COLD GASS sample are detected in either H I or CO, so the log-scaling of Fig. 1 prevents us from showing a lower limit on the observed median here. However, in the (larger) GASS sample, the H I detection fraction in the same bin is 52 per cent, so we can place at least a (conservative) lower limit on the neutral gas fraction in this bin from the GASS H I median alone (light blue triangle in Fig. 1). Including this additional constraint, the median EAGLE neutral gas fractions are consistent with observations at the 0.2 dex level over the range $\log_{10}(M_{\text{star}}/M_{\odot}) = [10.0, 11.25]$.

5.2 Atomic hydrogen fractions compared to GASS

Having established that the neutral hydrogen content of EAGLE galaxies agrees with observations, we now turn to analysing the atomic hydrogen subcomponent. Fig. 2 presents a comparison of the atomic hydrogen mass fractions, $M_{H I}/M_{\text{star}}$ in EAGLE with data from the GASS survey (Catinella et al. 2010; see Section 3). We show here the distribution of $M_{H I}/M_{\text{star}}$ for galaxies in four narrow bins of stellar mass (individual panels, mass increases from left to right). Blue/red histograms show the distribution for simulated EAGLE galaxies: in the top row, we adopt the empirical BR06 formula to account for the presence of H_2 (blue histograms), while this is achieved following the theoretical GK11 formula in the bottom row (red). In both cases, GASS data are represented by black lines. The vertical orange dash-dotted line marks the (maximum) GASS detection threshold in each stellar mass bin: for consistency, we combine all galaxies with $M_{H I}/M_{\text{star}}$ lower than this into a single ‘non-detected’ bin (blue/red open square and black open diamond in the shaded region on the left).

Both EAGLE and GASS show a decrease in $M_{H I}/M_{\text{star}}$ with increasing M_{star} (see also Catinella et al. 2010). While both H_2 models (top/bottom row) lead to broad agreement with the observed distribution in shape and normalization, the match is considerably better with the empirical H_2 formula of BR06 (top/blue): the median H I mass fractions (vertical dotted lines) differ by < 0.2 dex in all four bins of stellar mass and show no systematic deviation from the observed median. This level of agreement is considerably better than

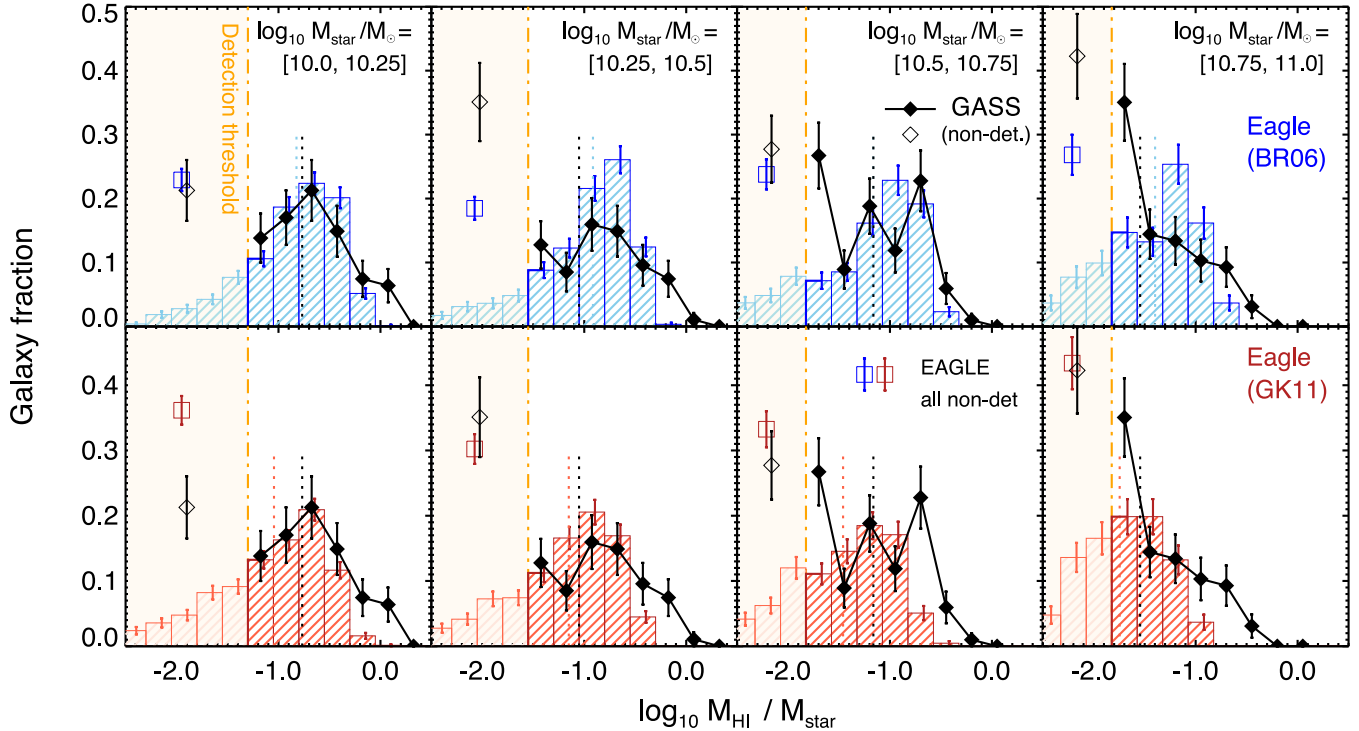


Figure 2. Comparison of the H I mass of EAGLE galaxies (blue/red histograms) with GASS observations (black lines); both samples include only central galaxies. In the top panel, the presence of H₂ in EAGLE is accounted for with the empirical Blitz & Rosolowski (2006) pressure-law prescription, while the bottom panel shows the corresponding results from the theoretical Gnedin & Kravtsov (2011) partition formula. The shaded region on the left is below the (maximum) GASS detection threshold in each panel; all simulated galaxies in this regime (light blue/red) are combined into the blue/red open square for comparison to the observations. Vertical black (blue/red) dotted lines indicate the median H I mass fraction of all GASS (EAGLE) galaxies per stellar mass bin; in the third panel of the top row both lie on top of each other. Error bars show statistical Poisson uncertainties. Both H₂ prescriptions lead to broad agreement of the predicted H I masses with observations, but the detailed match is considerably better for the Blitz & Rosolowski (2006) H₂ formula (top).

obtained by other recent hydrodynamic simulations (e.g. Aumer et al. 2013, whose H I fractions are higher than observed by ~ 0.5 dex; see Wang et al. 2014). In contrast, the median H I fraction obtained with the theoretical GK11 approach is consistently too low by ~ 0.1 – 0.4 dex.

It is possible that the differences between these two H₂ schemes are driven by inaccurate gas-phase metallicities in EAGLE galaxies, to which the BR06 pressure law is by construction insensitive; further work is required to test whether this is indeed the case. It is also important to keep in mind that dense gas is modelled in a highly simplified way in EAGLE, with the primary aim of circumventing numerical problems that would arise if gas were allowed to cool below $\sim 10^4$ K at the resolution of EAGLE (see e.g. Schaye & Dalla Vecchia 2008). It is plausible that the BR06 and GK11 prescriptions simply reflect this imperfect ISM model in different ways, leading to different predictions about the H₂ fractions.

As mentioned above, Lagos et al. (2015) obtained good agreement between the H₂ content of EAGLE galaxies and observations with the GK11 prescription, which we find to yield too low H I fractions. The likely reason for this apparent contradiction is that Lagos et al. (2015) focus their analysis on the sub-sample of galaxies above the COLD GASS detection threshold, and include both centrals and satellites: here on the other hand, we consider centrals only and calculate overall medians (from both detections and non-detections). Our results are therefore not directly comparable.

The match to GASS in Fig. 2 is not quite perfect even with the empirical BR06 model, however: on close inspection, the scatter in M_{HI} at fixed M_{star} is slightly smaller in EAGLE than GASS,

which manifests itself in a relative deficiency of non-detections (26 per cent versus 42 per cent in the highest stellar mass bin) and very H I-rich galaxies (a difference of -0.15 dex in the 90th percentile of $M_{\text{HI}}/M_{\text{star}}$ in the highest stellar mass bin). The latter discrepancy is also seen with the GK11 H₂ model, and we confirm in Appendix A1 that it is still present even when we ignore the presence of H₂ completely and assign all neutral gas as ‘H I’. Although the observational scatter may be overestimated due to uncertainties in the stellar mass measurements, we demonstrate below that a more likely cause is the presence of spuriously large H I holes in the simulated galaxies. Overall, however, we can conclude that (central) EAGLE galaxies acquire approximately realistic amounts of H I by $z = 0$, with only relatively minor uncertainties introduced by the choice of model to account for the presence of H₂.

6 THE INTERNAL STRUCTURE OF H I IN SIMULATED GALAXIES

We now investigate the internal distribution of atomic hydrogen in the simulated galaxies. Even in light of the good match between total H I masses in EAGLE and observations as demonstrated above, there is no guarantee that the former is modelled in an equally realistic way: many previous hydrodynamical simulations have suffered from ‘overcooling’ which leads to an artificially enhanced gas density in the central region, especially in massive galaxies (e.g. McCarthy et al. 2012, see also Crain et al. 2015). In combination with a deficit of gas in the outskirts, the total H I mass in simulated galaxies could agree with observations even in this case.

6.1 Visual inspection of H I morphologies

As a first qualitative step, we have created mock H I images of all simulated galaxies satisfying our selection criteria, by assigning the H I mass of particles within a line-of-sight interval of $[-70, 70]$ kpc relative to the galaxy centre to an $x \times y$ grid with pixels of 0.5 kpc, and smoothing with a Gaussian FWHM of 1 kpc. For simplicity, we only do this with the empirical BR06 H₂ correction. All images were then inspected visually, and the galaxies assigned to one of three broad morphological categories: (a) ‘Irregular’ (no disc-like structure), (b) ‘Disturbed H I discs’ (which are not flat when edge-on, and instead show e.g. prominent warps), and (c) ‘Clean H I discs’. Their relative abundances will be discussed in Section 6.2 below. While this is inevitably a subjective classification, it can still offer valuable insight into the H I structure that may not be apparent from a simple quantitative analysis. A typical example galaxy from each category is shown in the first two rows of Fig. 3; each has similar total H I mass ($\log_{10} M_{\text{HI}}/M_{\odot} = [9.9, 10.1]$) but rather different appearance. Galaxies have been rotated to face-on in the top row, and to edge-on in the middle; the disc plane is defined to be perpendicular to the angular momentum axis of all H I within a spherical 50 kpc aperture which corresponds roughly to the edge of the largest H I discs in our sample (see Fig. 6 below). The scaling is linear from 0 to 10 $M_{\odot} \text{ pc}^{-2}$ with darker shades of blue indicating denser gas; the smoothing scale of 1 kpc is indicated with a purple circle in the middle-right panel.

‘Irregular’ galaxies in particular (left-hand panels) typically contain a large number of H I ‘blobs’ that are typically representing gas in the process of accreting on to the galaxy. Although the ‘disc’ galaxies (middle and right-hand columns) have their H I predominantly in a more or less thin disc, many of these also show pronounced substructure with dense clumps as well as large (~ 10 – 20 kpc) H I holes. To illustrate the varying degree to which the latter are apparent in different galaxies, we show in Fig. 4 mock face-on H I images of three galaxies: one with clearly visible holes (left), one with only tentative hole identifications (middle), and one without visible large holes (right).

We have verified that these holes are also found in the total gas density maps, so they are not an artefact of our H I modelling (i.e. they are not regions where most of the neutral gas is H₂). From inspection of the high time resolution ‘snapshot’ outputs in EAGLE (see Schaye et al. 2015), they form rapidly and can reach sizes of ~ 10 kpc within only 20 Myr. This suggests that they are the result of heating events associated with star formation that are (individually) orders of magnitude more energetic than supernova explosions in the real Universe as a result of the limited resolution of EAGLE. Their detailed formation and survival is almost certainly more complex, however: as we show below, a clear correlation between star formation and occurrence of holes is not observed in the EAGLE galaxies.

For comparison, the bottom row of Fig. 3 displays three observed H I maps of nearby spiral galaxies of comparable M_{HI} – NGC 5457, NGC 6946, and NGC 5055 – from The H I Nearby Galaxies Survey (THINGS; Walter et al. 2008) smoothed to the same spatial resolution as our simulated maps (1 kpc). These clearly look different from even the simulated ‘clean disc’ galaxies and have a much smoother but also more intricate structure (see also Braun et al. 2009) with clear spiral arms. This difference in appearance can be attributed to the imperfect modelling of the ISM in EAGLE, which does not explicitly include a cold phase, and must therefore impose a pressure floor for high-density gas (Schaye & Dalla Vecchia 2008; Schaye et al. 2015). In addition, the observed galaxies do not show

H I holes comparable to those in the simulation, although such features do occur on smaller scales (Boomsma et al. 2008): the small inset in the bottom row shows a zoom of a 20×20 kpc region of the almost face-on galaxy NGC 5457 which clearly shows numerous holes up to scales of ~ 5 kpc (note that the inset uses a square-root scaling for improved clarity, and is therefore shown in a different colour). Rather than the *existence* of H I holes being an artefact of the simulation, it is rather their size and hence covering fraction in the disc that is in tension with observations.

Another apparent disagreement between the simulations and observations is the thickness of the H I discs: as the middle-right panel shows, even a ‘clean disc’ extends several kpc in the vertical direction. The exponential scaleheight of this disc is ~ 1.5 kpc, several times larger than the values indicated by observations of H I in the Milky Way (Dickey & Lockman 1990). We note, however, that this is only a factor of ~ 2 larger than the gravitational softening length of the simulation, which may largely explain this discrepancy. Another plausible contributor is the imposed temperature floor of $\sim 10^4$ K, which corresponds to a Jeans length of ~ 1 kpc. It is conceivable that the artificial thickening of the disc and the unrealistically large size of H I holes noted above are in fact related: a thicker disc implies a lower volume density and hence less mass swept up per unit distance. Despite these differences in detail, we show below that the azimuthally averaged distribution of H I in our simulated galaxies agrees quantitatively with observations, both in terms of the H I sizes and surface density profiles.

6.2 Correlations of H I morphology with other galaxy properties

The relative abundance of each morphological type as a function of, respectively, the total H I mass, stellar mass, H I mass fraction and specific star formation rate (sSFR) is shown in the top panels of Fig. 5. H I morphology correlates strongly with all four of these parameters: irregular distributions are most common at low H I mass, high M_{star} and low sSFR, whereas the fraction of discs is highest in the opposite regimes. Interestingly, the fraction of *clean* discs (blue line in Fig. 5) shows no such simple behaviour: they are most common at intermediate $M_{\text{HI}} \approx 10^{9.4} M_{\odot}$, while most galaxies at the high M_{HI} end show a ‘disturbed disc’ morphology. As H I content and sSFR are correlated, it is not surprising to see similar trends with the latter (rightmost panel): the fraction of clean discs drops sharply for the most actively star-forming galaxies.

Note that the left-hand column in particular (trends with M_{HI}) suffers from incompleteness due to our imposed stellar mass threshold of $M_{\text{star}} \geq 10^{10} M_{\odot}$: there is a large population of less massive galaxies, many of which with M_{HI} high enough to fall within the range plotted here. With a median $M_{\text{HI}}/M_{\text{star}} = 0.1$ at $M_{\text{star}} = 10^{10} M_{\odot}$ (see Fig. 2), this mostly affects the range $M_{\text{HI}} \lesssim 10^9 M_{\odot}$ which is therefore shaded grey in Fig. 5. This should be kept in mind when comparing to H I limited surveys, but insofar as only *massive* galaxies are concerned ($M_{\text{star}} \geq 10^{10} M_{\odot}$), it does not affect our results.

The bottom panels of Fig. 5 show the fraction of galaxies in the two ‘disc’ categories with one or more H I holes. They are generally more common in clean than disturbed discs and at higher M_{HI} (e.g. 55 per cent at $M_{\text{HI}} \approx 10^{10} M_{\odot}$) and lower M_{star} . Including tentative hole identifications (dotted lines), their occurrence increases by a factor of ~ 2 to ~ 80 per cent at both the low- M_{star} and high- M_{HI} ends. Perhaps surprisingly, the hole fraction shows no clear increase with increasing sSFR, and when tentative detections are included, there is a clear increase towards *lower* sSFR, at least in clean discs

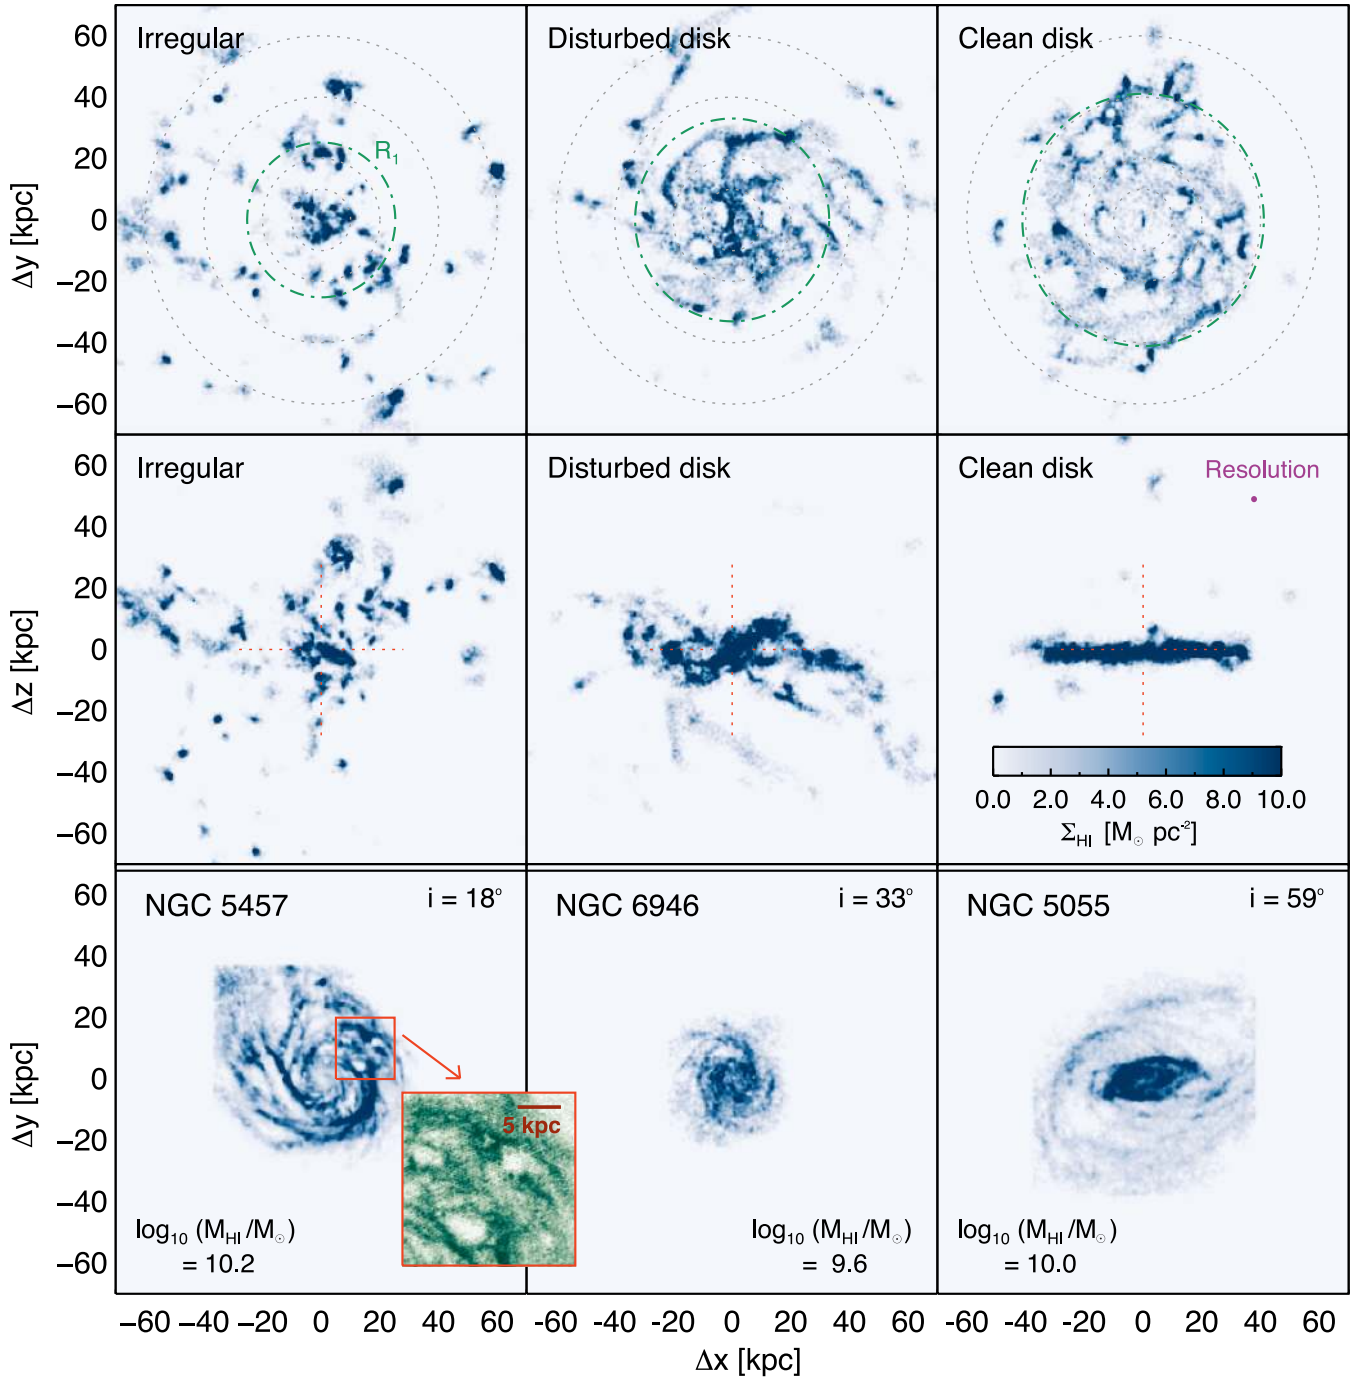


Figure 3. Top and middle row: Three typical examples of EAGLE galaxies with different H I morphologies, but similar H I mass $\log_{10}(M_{\text{HI}}/M_{\odot}) = [9.9, 10.1]$: Irregular (left), Disturbed disk (middle) and Clean disk (right). Face-on images are shown in the top row, edge-on equivalents below. For comparison, the bottom row shows three observed H I images from THINGS (Walter et al. 2008) at the same physical scale; these do not correspond to the same morphology categories as the simulated galaxies above. All images use the same linear scaling, as given in the middle-right panel, and are Gaussian-smoothed to a (FWHM) resolution of 1 kpc (purple circle in the middle-right panel). The green dash-dotted rings in the top row show the characteristic radius R_1 for each galaxy (see Section 6.3); the grey dotted circles indicate radii of 10, 20, 40, and 60 kpc, respectively. The red dotted ‘cross-hairs’ in the middle row indicate the best-fitting H I disc plane and axis. The inset in the bottom row shows a 20×20 kpc zoom-in of NGC 5457, revealing H I holes similar to what is seen in EAGLE but on a smaller scale.

(blue). Even though we here show current sSFR – which may well already have been lowered by the presence of low-density holes – we have tested for correlation between holes and star formation in the recent past (as well as total star formation and SFR density), which yields a similar result. This suggests a complex connection between

star formation (and the associated feedback) and the occurrence of holes. It is possible, for example, that disc instabilities can prolong the lifetime of holes and therefore make them a more prominent feature in H I-rich galaxies (see also Mitchell et al. 2012; Agertz, Romeo & Grisdale 2015).

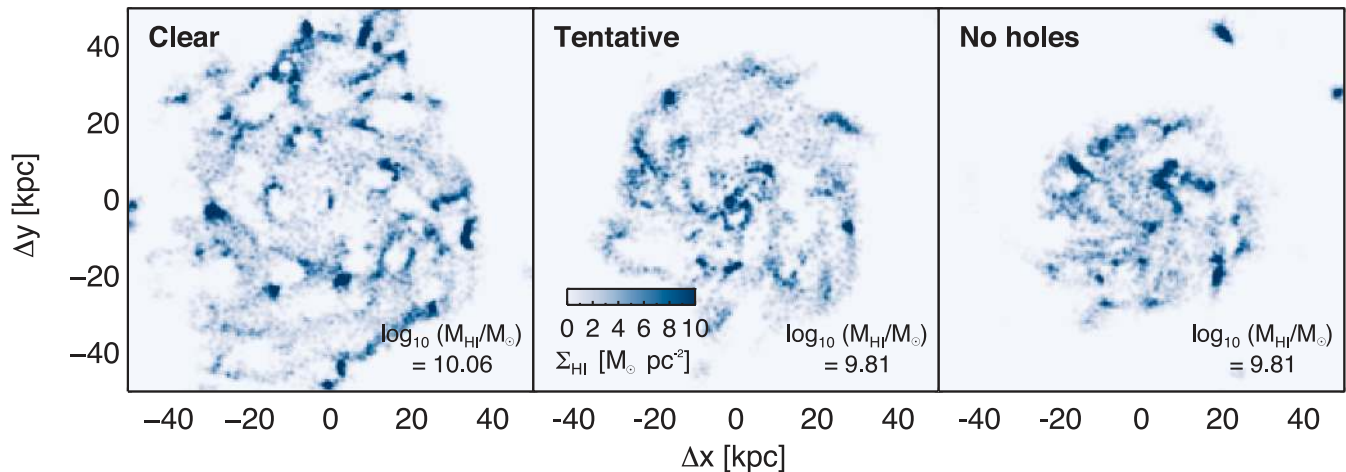


Figure 4. Examples of simulated face-on galaxy H I images with clear (left), tentative (middle) and no (right) H I holes as found by visual inspection; all of them have similar H I mass. The scaling is the same as in Fig. 3. Note that even the ‘no holes’ galaxy in the right-hand panel shows some hole-like structures, but these are much smaller than those seen in the left-hand panel and therefore not classified as ‘holes’ here.

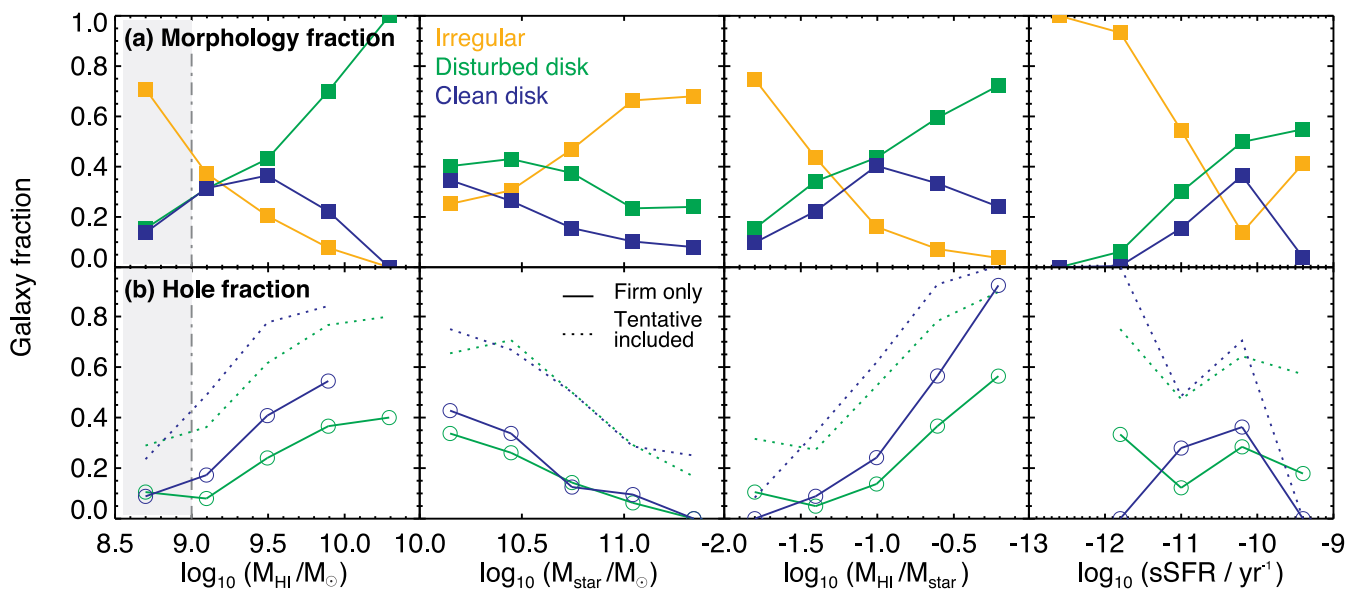


Figure 5. Fraction of galaxies with different H I morphologies (top), and with visible H I holes (bottom, dotted lines include tentative identifications). From left to right, the individual columns show the fractions as a function of H I mass, stellar mass, H I mass fraction and specific star formation rate (sSFR). The fraction of disturbed discs increases strongly with H I mass, in qualitative agreement with observations; the same is true for the fraction of galaxies with visible H I holes. Trends are also seen with the other galaxy parameters, as discussed in the text. The grey region in the left-hand column is affected by our imposed stellar mass limit.

In summary, a large fraction of our simulated galaxies (64 per cent) show a disc morphology of their H I content (particularly at high M_{HI}), but nearly two thirds of these are evidently disturbed. All H I discs are a factor of several too thick and lack the intricate spiral structure seen in observed H I maps, which is a direct consequence of the simplified ISM modelling in EAGLE. Particularly at high M_{HI} and low M_{star} , many galaxies furthermore show H I holes that are larger than what is observed (43 and 19 per cent of all central galaxies, respectively, with and without tentative hole identifications).

6.3 H I size–mass relation

As a simple one-parameter proxy for the internal gas distribution, we next investigate the ‘characteristic’ size of the H I discs, which

we define as the radius R_1 at which the azimuthally averaged surface mass density drops below $1 \text{ M}_{\odot} \text{ pc}^{-2}$; this definition is commonly encountered in the observational literature (e.g. Broeils & van Woerden 1994; Wang et al. 2013, 2014). To find R_1 , we align all galaxies to face-on (see above) and generate two-dimensional H I surface density maps Σ_{HI} with a pixel size of 0.5 kpc, integrating along the line of sight over the range $[-70, +70]$ kpc. We compare our data to the mass–size relation of Broeils & Rhee (1997); their galaxies were observed at the Westerbork Synthesis Radio Telescope (WSRT) with a beam size of ~ 12 arcsec and are at a median distance of ~ 35 Mpc (Broeils & van Woerden 1994) which corresponds to a physical resolution of ~ 2 kpc. Note that this relation has recently been verified and extended to very H I-rich galaxies by Wang et al. (2013). We convolve our H I maps with a Gaussian kernel of $\text{FWHM} = 2$ kpc, and extract radial profiles in a set of

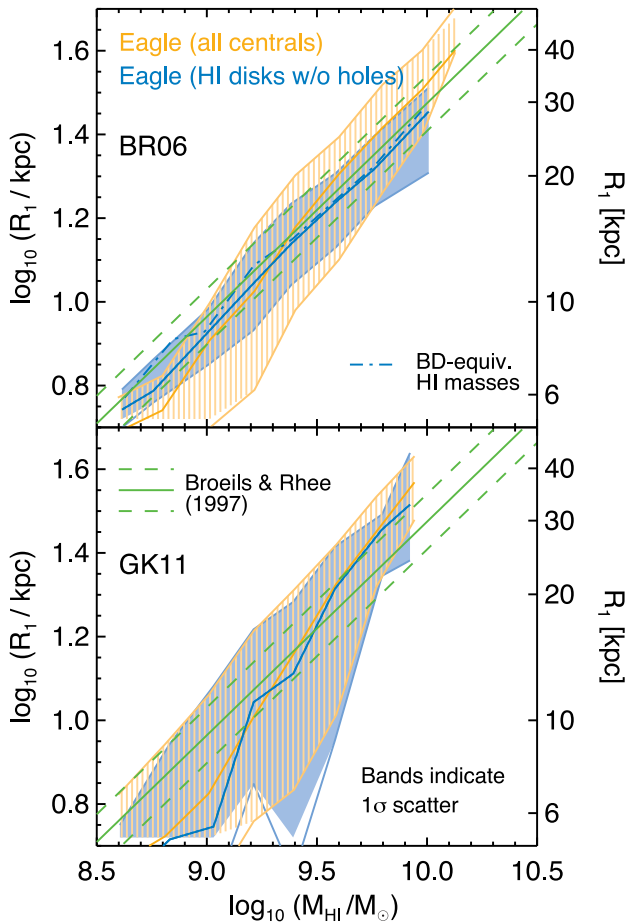


Figure 6. The H I mass–size relation for EAGLE galaxies (blue/orange bands), compared to the observational data from Broeils & Rhee (1997) (green lines). Orange includes all (central) EAGLE galaxies whereas blue only includes those with a visually confirmed H I disc that did not show prominent H I holes. The top panel shows results with the empirical Blitz & Rosolowski (2006) H₂ model, whereas the theoretical Gnedin & Kravtsov (2011) approach is used in the bottom panel. For the former, the agreement of the median trends is already quite good in the full sample, and even better for the ‘clean’ subset (blue) over more than an order of magnitude in H I mass, especially when using ‘Bluedisk’ equivalent H I masses (dash–dotted line; see text). The theoretical model (bottom) predicts an H I mass–size relation that is significantly too steep.

(circular) annuli of width 2.5 kpc. R_1 is then determined by interpolating linearly between the outermost bin with density above the threshold of $1 M_{\odot} \text{pc}^{-2}$, and the one beyond this.⁸ In the top row of Fig. 3, R_1 is shown as a dark green dash–dotted circle, and coincides approximately with what one would visually identify as the ‘outer edge’ of the H I disc.

In Fig. 6, we show the resulting relation between R_1 and M_{HI} for EAGLE galaxies. As before, we select only central galaxies;⁹

⁸ We note that this procedure is not strictly self-consistent, because such an interpolation yields, in general, a cumulative mass profile that differs from the true profile. However, we have experimented with more elaborate methods such as linear or quadratic spline fits, or narrower profile bins. None of these alternatives differ substantially from the simple method adopted here.

⁹ We have verified that the result is virtually unchanged when potentially interacting galaxies are excluded, i.e. those where a neighbouring galaxy within 150 kpc has a stellar mass exceeding one tenth of its own.

we show results for both the empirical BR06 (top) and theoretical GK11 prescription (bottom) to account for H₂. The running median and 1σ scatter (defined as the 15.9/84.1 percentile) of these is shown in orange. To test the influence of H I morphology, we also show (in blue) the size–mass relation for only the subset of simulated galaxies that have been visually classified as containing a well-aligned H I disc without prominent holes.¹⁰ The green solid line represents the best fit of Broeils & Rhee (1997), with 1σ standard deviation indicated by the green dashed lines (note that our plot has the x- and y-axis swapped relative to their fig. 4, and that they show $D_1 \equiv 2R_1$ instead).

In general, the EAGLE galaxies follow the observed relation quite well. The full sample (orange) has a slope that is somewhat steeper than observed with both H₂ recipes (0.59 versus 0.51, i.e. a 16 per cent difference with BR06, and 52 per cent with GK11). With the empirical BR06 method (top), the ‘disc-only, no holes’ distribution (blue) agrees with the observations to better than 0.1 dex over the full range of M_{HI} that we probe here, $\log_{10}(M_{\text{HI}}/M_{\odot}) = [8.5, 10.0]$. The agreement gets better still when we calculate the H I masses in analogy to the Bluedisk survey (see below), the observations for which were also conducted at the WSRT: this reduces M_{HI} slightly at the low- M_{HI} end (see top panel of Fig. A2) and therefore improves the agreement of EAGLE with the Broeils & Rhee (1997) relation to <0.05 dex. We note, however, that Broeils & van Woerden (1994) find close agreement between the interferometry-derived H I masses used by Broeils & Rhee (1997) and single-dish measurements, so it is not clear whether this is indeed a more fair comparison to the data. We also note that the scatter in the overall EAGLE sample is somewhat too large, by a factor of ~ 2 . Again, the agreement here is better in the sample that excludes galaxies with prominent H I holes (blue). With the GK11 H₂ formula (bottom), excluding holes makes no appreciable difference, and the H I size–mass relation remains steeper and broader than observed.

6.4 Density profiles for H I-rich and ‘normal’ galaxies

A more detailed quantitative test of the H I structure is to compare the radial (surface) density profiles to observations. A particularly interesting question in this respect is how atypically H I-rich galaxies – which have likely been particularly efficient at accreting H I recently – compare to those with average H I content. Motivated by this consideration, the Bluedisk survey (Wang et al. 2013) has recently observed a set of 25 galaxies expected to be H I-rich, and a similar number of ‘control’ galaxies, generating resolved H I maps with a resolution of ~ 10 kpc. One key discovery of this study has been that the H I surface density profiles of *all* galaxies, both H I-rich and normal, follow a ‘universal’ shape as long as they are normalized by the characteristic H I disc size R_1 (Wang et al. 2014). We will now compare the EAGLE galaxies to these observations. For simplicity, we focus first on results obtained with the empirical BR06 H₂ model which, as shown above, leads to total H I masses in good agreement with observations. Profiles obtained with the theoretical H₂ formula of GK11 will be presented in Section 6.4.3 below.

6.4.1 Sample definitions

From our 2083 central EAGLE galaxies with $\log_{10}(M_{\text{star}}/M_{\odot}) = [10.0, 11.0]$ – the same range as in Bluedisk – we select those

¹⁰ Dropping the requirement for the galaxies to have an H I disc makes no noticeable difference, but the Broeils & Rhee (1997) sample includes only spiral and irregular galaxies (by optical morphology).

with a Petrosian half-light radius $R_{50,z} \geq 3$ kpc. This radius is defined as enclosing 50 per cent of the Petrosian flux in the SDSS z -band¹¹ obtained from stellar population synthesis (SPS) modelling (Trayford et al. 2015). As we show in Appendix C1, this simple size cut approximately reproduces the more complex original sample selection in the Bluedisk survey (Wang et al. 2013). We have verified that the profiles shown below are insensitive to the exact value of the size cut, and are actually almost unchanged when all galaxies are included, regardless of size.

This simulated sample with 607 members is then divided into ‘H I-rich’ galaxies with $M_{\text{HI}} \geq 10^{9.8} M_{\odot}$, and a ‘control’ sample with $10^{9.1} M_{\odot} \leq M_{\text{HI}} \leq 10^{9.8} M_{\odot}$. As opposed to an – equally plausible – set of cuts in $M_{\text{HI}}/M_{\text{star}}$, these limits approximately correspond to the sample division in Bluedisk (see Wang et al. 2013 and Appendix C2). For consistency with the observations, we calculate H I masses here in a ‘Bluedisk-equivalent’ fashion (Wang et al. 2013): a two-dimensional H I surface density map with pixel size 0.5 kpc was created (with the simulation z -coordinate as the line of sight, i.e. random galaxy orientations) and then smoothed with an elliptical Gaussian of FWHM = 14 (9) kpc major (minor) axes. This corresponds approximately to the WSRT beam size at the median redshift of the Bluedisk galaxies ($\bar{z} \approx 0.027$). From these maps, we then sum over all pixels with Σ_{HI} above the median Bluedisk detection threshold of 4.6×10^{19} atoms cm^{-2} ($=0.37 M_{\odot} \text{pc}^{-2}$ in H I) to obtain the total H I mass of the galaxy. In Appendix A2, we show that the resulting M_{HI} are typically less than 0.1 dex below the ‘GASS-equivalent’ mass for $M_{\text{HI}} \geq 10^{9.1} M_{\odot}$.

Recall from Section 5 that EAGLE has a deficiency of galaxies at the H I-rich end. This is unfortunate for the present purpose, because it means that our simulated H I-rich galaxies are typically not quite as extreme as those in the corresponding Bluedisk sample. However, the two populations are still clearly different and enable us to test the universality of the H I density profile: The 406 simulated ‘control’ galaxies have a median $\log_{10}(M_{\text{HI}}/M_{\odot}) = 9.61$, whereas the 133 ‘H I-rich’ counterparts¹² have a median $\log_{10}(M_{\text{HI}}/M_{\odot}) = 9.93$. For comparison, the Bluedisk H I-rich galaxies have a median $\log_{10}(M_{\text{HI}}/M_{\odot}) = 10.09$ ($N = 23$) and the control sample $\log_{10}(M_{\text{HI}}/M_{\odot}) = 9.59$ ($N = 18$). The difference between the two EAGLE samples (0.32 dex) is therefore approximately two thirds of that in Bluedisk (0.5 dex).

6.4.2 Density profile comparison

In Wang et al. (2014), density profiles were extracted using elliptical annuli with orientation and axis ratio (b/a) taken from the best-fitting ellipse to the stellar r -band light, and then multiplied by a factor of $b/a \approx \cos(\theta)$ to correct the profiles to face-on. For the analysis of our simulated galaxies, we simply rotate them to face-on by aligning the angular momentum axis of the H I in the central 50 kpc with the line of sight (as in Fig. 3). For each galaxy, an H I image was then created as described above, and the surface density

¹¹ As detailed in Appendix C1, the radii $R_{50,z}$ obtained for our simulated galaxies are systematically too large compared to observations from SDSS; this is investigated in more detail by Furlong et al. (2015b). For our purpose, we simply re-scale the distribution of $R_{50,z}$ to enforce a match to the observational data. This ad hoc fix does not invalidate our results below, because we are only concerned with relative size comparisons, and only use the stellar sizes to select the overall sample to compare to Bluedisk.

¹² There are some additional galaxies with $M_{\text{HI}} < 10^{9.1} M_{\odot}$, which are not included in either sample.




Figure 7. The scaled H I surface density profiles for galaxies in the EAGLE simulation (red/green rectangles) and in the Bluedisk survey (yellow/blue shaded bands). Galaxies are split into ‘H I-rich’ and ‘control’ samples based on their H I mass, as explained in the text. The Bluedisk profiles are identical in both panels, but different methods are applied to EAGLE: Top: galaxies are rotated to face-on. Bottom: profiles are extracted in random orientation in elliptical bins with position angle and axis ratio determined from the stellar r -band light (as in the Bluedisk analysis). In agreement with observations, both methods yield similar profiles for simulated H I-rich and control galaxies. However, the EAGLE profiles deviate from the observations in the central region, and with the ‘elliptical bin’ inclination correction (bottom), they are also too shallow in the outskirts.

profile extracted with 20 equally spaced bins from 0 to $2R_1$ (recall that R_1 is defined as the radius at which the H I surface density drops to $1 M_{\odot} \text{pc}^{-2}$). Galaxies were then median-stacked to obtain the average profile for the H I-rich and control galaxies. The same procedure was applied to the Bluedisk profiles from Wang et al. (2014).

The result is shown in the top panel of Fig. 7. Yellow and blue lines trace the median profiles for H I-rich and control galaxies in Bluedisk. As in Fig. 1, shaded bands indicate the statistical 1σ uncertainty on the median, i.e. they extend from σ_{low} to σ_{high} where $\sigma_{\text{low (high)}} = \bar{\Sigma}_{\text{HI}} + (P_{15.9(84.1)} - \bar{\Sigma}_{\text{HI}})/\sqrt{N}$; $\bar{\Sigma}_{\text{HI}}$ here denotes the median and P_n the n th percentile of the Σ_{HI} distribution in a bin with N galaxies. The simulated profiles are shown with green (red) boxes for H I-rich (control) galaxies whose vertical extent gives the 1σ uncertainty on the median calculated in the same way; they are small due to the comparatively large sample size. In addition, the galaxy-to-galaxy scatter is indicated by thin error bars for EAGLE




Figure 8. Cumulative H I mass profiles for EAGLE galaxies (green/red) compared to observations from Bluedisk (blue/yellow); different colours denote control/H I-rich galaxies as in Fig. 7, but here the x -axis is not normalized by R_1 . For Bluedisk, dotted lines show the profiles corrected for the offset in total H I mass in each sample relative to EAGLE (see text). Bands and thick boxes indicate the 1σ uncertainty on the median. There is a clear lack of H I in the central region ($R \lesssim 20$ kpc) in EAGLE, which manifests itself in the lower-than-observed H I masses of H I-rich galaxies (blue/green). For the ‘control’ sample, the H I deficit in the centre is largely compensated by the slightly too shallow decline of the profile in the outer parts (yellow/red).

and by the grey hatched region for Bluedisk, both of which show the interval occupied by 50 per cent of galaxies (i.e. P_{25} to P_{75}).

In the outer region ($R \gtrsim R_1$), simulated and observed profiles generally agree to within the statistical uncertainties. Both simulated control and H I-rich galaxies follow an exponential profile (straight line in the log-linear plot), although the gradient is slightly steeper for the H I-rich galaxies (green, discrepancy at $\sim 2\sigma$ level). The observed galaxy profiles show an approximately equal, but opposite difference (steeper profiles for control galaxies), although the small number of Bluedisk galaxies means that this difference is not statistically significant.

In the inner parts ($R \lesssim R_1$), there is a much more pronounced discrepancy between simulations and observations: the former ‘undercut’ the observed profile. Interestingly, the simulated H I-rich and control profiles are still almost identical, to within <0.1 dex (with a minor excess, significant at $\sim 2\sigma$, in H I-rich galaxies, as in observations). To test whether this is a result of the slightly different analysis for the simulated and observed galaxies, we have reproduced the Bluedisk analysis on our simulated galaxies exactly (i.e. we extracted the profiles in elliptical rings with position angle and axis ratio given by the stellar r band, and multiplied with a correction factor of b/a), the result of which is shown in the bottom panel of Fig. 7. However, this fails to ameliorate the tension in the central region, and adds another disagreement in the outer parts, where the simulated profiles are now far too shallow; we shall return to this shortly.

It is worth keeping in mind that, due to the scaling of the x -axis in Fig. 7 by R_1 , one cannot directly infer the actual mass distribution from the profiles. We therefore explicitly show a comparison between the cumulative mass profiles in EAGLE and Bluedisk in Fig. 8; the symbols have the same meaning as in Fig. 7 with the addition of the grey hatched region denoting the (small) zone influenced

by resolution effects ($R \leq 3\epsilon$, the gravitational softening length). Note that the x -axis here shows the actual galactocentric radius R in kpc, and is not normalized by R_1 . As expected, simulated galaxies show a deficit of mass in the inner region ($R \lesssim 30$ kpc). However, for the control galaxies (red) this is almost completely compensated by the outer parts, where the surface density profile is slightly too shallow. In the H I-rich sample, on the other hand, the central deficit manifests itself in a lower total H I mass in simulated galaxies, as already seen in Fig. 2. To confirm that this interpretation also holds quantitatively, the dashed lines in Fig. 8 show the Bluedisk profiles re-scaled by the ratio of median total $M_{\text{H I}}$ in EAGLE and Bluedisk, 0.95 and 0.69, respectively, for the H I-rich and control samples. As expected, this shows much better agreement with the simulated H I-rich profile at $R \gtrsim R_1$. The lack of extremely H I-massive galaxies in EAGLE is therefore directly connected to the missing H I in the central galaxy regions.

6.4.3 Why are the simulated profiles different from observations?

We are therefore faced with two puzzles: (i) Why does the inclination correction through elliptical bins work so poorly for the outskirts of the simulated galaxies? (ii) Why is the H I density too low in the inner regions of simulated galaxies?

For the first question, one natural explanation might be that the inclination of the H I disc does not correspond exactly to the ellipticity of the r -band light (e.g. Serra et al. 2012; Wang et al. 2013). We test this in Appendix D and there is indeed good evidence that this is the case: observationally, the distribution of b/a in the Bluedisk sample is far from flat, with a deficit at both small ($b/a \lesssim 0.3$) and high ($b/a \gtrsim 0.9$) axis ratios, which is in conflict with the simple assumption $\cos\theta = b/a$ because the distribution of $\cos\theta$ should be uniform. Likewise, a direct comparison between the inclination angle and stellar b/a in EAGLE shows both significant scatter and a systematic offset for galaxies with $b/a \lesssim 0.6$. However, we have repeated the ‘elliptical bin’ analysis on our simulated galaxies, with axis ratio and orientation angle derived directly from the orientation of the H I angular momentum axis instead of fits to the stellar light, and the resulting profiles are almost identical to those shown in the bottom panel of Fig. 7.

There must therefore be a genuine difference between the simulated and observed galaxies, and one obvious candidate for this is the artificially increased thickness of the H I disc in EAGLE (see discussion in Section 6.1), which may lead the elliptical-bin inclination correction, based on an infinitesimally thin H I disc, to fail. Any line of sight will intersect a thick disc not only at one point, but over a finite interval, which effectively smears out the resulting profile: an inclined line of sight that intersects the disc mid-plane in the galaxy outskirts (at a radius R_a) will actually pick up most H I at $R < R_a$ where the H I density is higher, thus leading to a shallower outer profile. By conservation of total mass, the density must appear lower in the inner regions, exactly as seen in Fig. 7. This interpretation would also explain why the (outer) face-on profiles in EAGLE are still a good match to observations, because they are by definition insensitive to the vertical structure of the H I discs.

It is conceivable, of course, that a similar effect is also present in observed galaxies. However, we have tested this by comparing the (inclination-corrected) profiles for Bluedisk galaxies with axis ratio b/a below and above the median of 0.6. If the inclination correction for observed galaxies were to suffer from the same bias as in EAGLE, the more edge-on sample with $b/a < 0.6$ should show a systematically shallower outer profile, but we did not find evidence for such an offset. The profiles presented by Wang et al. (2014) can

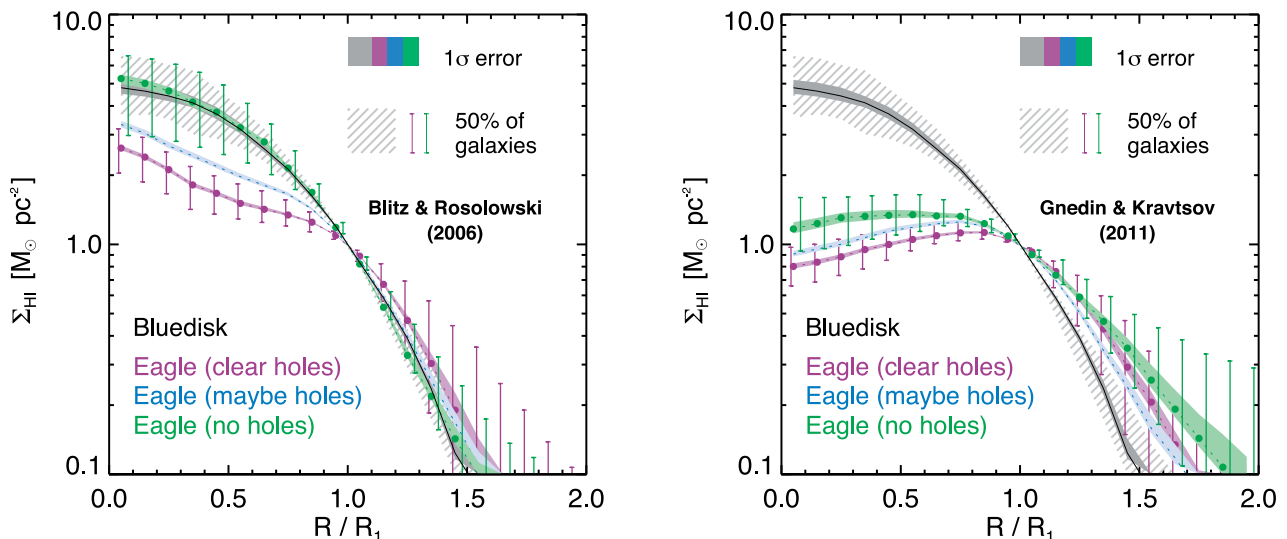


Figure 9. Dependence of H I surface density profiles on the presence of visible H I holes. The left-hand panel shows results obtained with the empirical Blitz & Rosolowski (2006) H₂ model, whereas the theoretical Gnedin & Kravtsov (2011) formula is used in the right-hand panel. Green points show the profile for galaxies with no visual evidence of holes, while purple points show those galaxies where holes are clearly visible. Galaxies with tentative hole detections are shown in blue; for clarity we have omitted the error bars here. With the empirical H₂ formula (left), the hole-free sample (green) is in good agreement with the Bluedisk data over the entire radial range we probe here, while galaxies with hole detections show a deficit in the central H I profile. Using the theoretical formula from Gnedin & Kravtsov (2011), galaxies with and without holes have central $\Sigma_{\text{H I}}$ profiles that are significantly too low compared to observations.

therefore be regarded as truly face-on, and the comparison to the face-on profiles from EAGLE in the top panel of Fig. 7 as the most meaningful test of the simulated H I surface density profiles.

The discrepancy in the *central* regions must have a different origin, as it is present in both the face-on and elliptical-bin-corrected profiles. One possibility here is that this is related to the spurious H I holes that we had already noted in the discussion of Fig. 3. To test this hypothesis, we show in Fig. 9 the median-stacked profiles (now again normalized by R_1) for galaxies with clear visual hole detections (purple), as well as those without (green) and those with tentative-only identifications (blue). For simplicity, we here combine H I-rich and control galaxies. The left-hand panel shows the profiles obtained with the empirical BR06 pressure-law to account for the presence of H₂; for comparison we also show the equivalent profiles obtained when using the theoretical GK11 formula in the right-hand panel. In both cases, the observed profile from Bluedisk (H I-rich and control combined) is shown in black.

It is evident that the discrepancy in the inner region seen with the BR06 model as discussed above is indeed connected to the presence of H I holes. Simulated galaxies without holes (green) follow the observed profile almost exactly over the entire radial range that we consider here (discrepancy < 0.07 dex). The scatter remains slightly larger than in the observations, but only by typically $\lesssim 50$ per cent. In contrast, galaxies with clearly visible holes (purple) have a much shallower central profile (by up to a factor of ~ 2), with tentative identifications (blue) lying in-between. It is worth pointing out that the ‘hole’ and ‘no-hole’ populations have profiles that differ by more than the typical scatter in each: there is a clear difference between *individual* galaxies in the two categories, and not just in a statistical sense.

Although smaller, there is also a slight effect in the outer profiles where the H I surface density is slightly higher in galaxies with holes than without (by $\lesssim 0.1$ dex). This may seem counter-intuitive, but is explained by the fact that the profiles are scaled by R_1 which is affected by the presence of holes as well (see Fig. 6). It is, of course, no surprise that galaxies with holes lack H I and therefore

show shallower inner profiles than those without visible evidence for them. What is far less trivial, however, is that the ‘hole-free’ sample agrees so well with the observations: *the discrepancy in the surface density profiles between EAGLE and Bluedisk can thus be fully attributed to the existence of these holes.*

The GK11 profiles (right-hand panel of Fig. 9) show a similar general trend – a higher central H I surface density for hole-free galaxies – but the difference is much smaller here than with the BR06 formula. Moreover, all three profiles are significantly too low within R_1 , with a discrepancy by a factor of $\gtrsim 5$ in the very centre. Evidently, the theoretically based GK11 formula assigns an unphysically large fraction of the neutral gas in the central EAGLE galaxy regions as H₂, which would explain why the GK11 total H I masses as shown in Fig. 2 are biased low. In addition, the *outer* $\Sigma_{\text{H I}}$ profiles are actually slightly more discrepant (i.e. shallower) for hole-free galaxies than those with holes, indicating that the normalization radius R_1 is somewhat too small at fixed M_{star} (see also Fig. 6). The lack of H I in the central region also explains the smaller effect of holes, compared to BR06: there is simply not enough H I even in hole-free galaxies for these features to have a significant impact.

We note that it is, in principle, conceivable that the profile agreement between H I-rich and control galaxies is not actually a physical feature, but rather a result of the comparatively large beam size: recall that the H I maps from EAGLE had been artificially reduced in resolution to the same level as in Bluedisk. However, we explicitly check for this in Appendix E, where we calculate the EAGLE H I surface density profiles from higher-resolution maps. Although the detailed shape of the profiles does indeed vary with the beam size, the close agreement between H I-rich and control galaxies remains. This strongly suggests that it is a genuine physical feature of the simulated – and observed – galaxies, rather than a smoothing artefact.

The conclusion that the discrepancy between EAGLE and the Bluedisk observations can be attributed to feedback-related H I holes, which are generated *after* accretion of the gas, combined

with the good match to the observed outer H I profile, supports the hypothesis that the accretion of H I on to galaxies is overall modelled well in EAGLE. At the same time, it highlights the fact that gas properties of galaxies are particularly sensitive to the adopted feedback implementation in hydrodynamical simulations (see also Walker et al. 2014; Crain et al. 2015).

7 SUMMARY AND DISCUSSION

We have conducted a series of detailed like-with-like comparisons between the atomic hydrogen (H I) content of present-day galaxies with $M_{\text{star}} \geq 10^{10} M_{\odot}$ from the 100 cMpc EAGLE simulation and various observational data sets. Our main aim has been to test the EAGLE galaxy formation model in a regime that was not considered during calibration, and to assess the usefulness of these simulations to better understand galactic H I evolution and the origin of observed scaling relations. Our main findings may be summarized as follows.

(i) The EAGLE simulations, combined with the Rahmati et al. (2013a) fitting formula for self-shielding and collisional ionization, predict median neutral hydrogen fractions for central galaxies that agree with observations to better than 0.1 dex at galaxy stellar masses in the range 10^{10} to $10^{11} M_{\odot}$. For the most massive galaxies ($M_{\text{star}} > 10^{11} M_{\odot}$) large observational uncertainties due to frequent non-detections prevent strong statements on the accuracy of EAGLE.

(ii) Accounting for molecular hydrogen (H₂) in EAGLE with the empirical Blitz & Rosolowski (2006) pressure law (BR06) pressure law leads to atomic hydrogen masses that are generally in good agreement with observations from GASS; the medians differ by <0.2 dex over a decade in stellar mass from 10^{10} to $10^{11} M_{\odot}$. A minor deficiency is the failure of the simulation to produce galaxies as H I-rich as the richest observed galaxies, as well as a moderate shortfall of very H I-poor galaxies.

(iii) An alternative, theoretically based H₂ model of Gnedin & Kravtsov (2011, GK11) applied to EAGLE results in galaxies having overall too little H I, particularly in their centres.

(iv) Using visual classification of H I morphologies, we have shown that H I discs are increasingly common in simulated galaxies of increasing M_{HI} and decreasing M_{star} , but at $M_{\text{HI}} \gtrsim 10^{9.4} M_{\odot}$ the majority of discs appear vertically disturbed.

(v) Many simulated galaxies contain large (up to ~ 20 kpc) H I holes which are a factor of several larger than seen in observations. They are more common at high M_{HI} and low M_{star} , but show no clear correlation with the specific star formation rate of galaxies.

(vi) Simulated galaxies match the observed H I mass–size relation reasonably well (the slope is too steep by 13 per cent), and the agreement becomes excellent (better than 0.1 dex) when only galaxies with visually confirmed H I discs without holes are included and the empirical BR06 prescription is used to account for the presence of H₂.

(vii) The H I surface density profiles of H I-rich ($M_{\text{HI}} > 10^{9.8} M_{\odot}$) and control galaxies ($10^{9.1} M_{\odot} \leq M_{\text{HI}} \leq 10^{9.8} M_{\odot}$) in EAGLE follow each other closely when scaled by the characteristic H I radius R_1 , as observed. While the outer profiles ($R > R_1$) also agree quantitatively, the surface density around $0.5 R_1$ is too low by a factor of ~ 2 . This tension can be fully attributed to the presence of H I holes. Galaxies without holes follow the observed density profile almost perfectly, to better than 0.07 dex.

To our knowledge, this is the first time that such relatively detailed agreement of H I properties with observations has been demonstrated in self-consistent cosmological hydrodynamical simulations.

Our results bode well for theoretical studies aiming to use EAGLE to obtain a better understanding of a wide variety of physical processes relevant for galaxy evolution, such as H I accretion and recycling, as well as its stripping in the dense environment of groups and clusters, for which realistic initial conditions are a significant advantage. In companion papers we show that this success can be predominantly attributed to the calibrated model for energetic feedback from star formation in EAGLE (Dalla Vecchia & Schaye 2012; Crain et al. 2015; Schaye et al. 2015).

The discrepancy between the empirical and theoretical H I/H₂ partition schemes that we have identified (see also Appendix A1) is the largest systematic uncertainty in our results. Although the pressure law parametrization of BR06 is based directly on observational data, its simplicity may hide detailed problems of the EAGLE simulations that have a strong impact on the other partition models. One possibility here is that the metallicities are somewhat too high in the central galaxy region, which would lead to overestimation of the H₂ content by the GK11 scheme.

An unambiguous shortcoming of the EAGLE galaxies in terms of their gas distribution appears to be the common occurrence of unphysically large H I holes with low surface density. Although they appear to be seeded by heating events which are included in EAGLE to model star formation feedback, the fact that holes are (slightly) more common in galaxies with low star formation activity suggests that their formation and survival depends on other galaxy properties as well. Efforts to improve the simulations in this respect must therefore likely involve improvements to both the star formation feedback scheme and ISM model. Another important area for improvement would be to account for the (as yet uncertain) effect of local stellar radiation on the neutral hydrogen fractions: Rahmati et al. (2013b) showed that this can significantly affect H I column densities even at $z = 0$.

In a follow-up paper, we will study the H I accretion on to H I-rich EAGLE galaxies directly by tracing galaxy progenitors back in time to find out when and how the $z = 0$ H I reservoirs were built up. As well as providing new insight into the details of galaxy formation, this will also lead to new predictions that can be tested against future observations, and contribute to our theoretical understanding of galaxies.

ACKNOWLEDGEMENTS

RAC is a Royal Society University Research Fellow. We thank the THINGS, Bluedisk, GASS and COLD GASS teams for making their observations publicly available, Mei-Ling Huang for technical help with the data, and Barbara Catinella for providing the match of GASS galaxies to the Yang et al. (2012) halo catalogue. We gratefully acknowledge helpful discussions with Ali Rahmati, Diederik Kruijssen, Thorsten Naab, Bernhard Röttgers, Jing Wang, and Simon White. This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grant ST/H008519/1, and STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the National E-Infrastructure. We also gratefully acknowledge PRACE for awarding the EAGLE project with access to the Curie facility based in France at Très Grand Centre de Calcul. Support was also received via the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office ([AP P7/08 CHARM]), the National Science Foundation under Grant No. NSF PHY11-25915,

and the UK Science and Technology Facilities Council (grant numbers ST/F001166/1 and ST/I000976/1) via rolling and consolidated grants awarded to the ICC, and the European Research Council (grant agreement 278594 – GasAroundGalaxies). Several figures in this paper were produced using the `ASTROPY` (Astropy Collaboration et al. 2013) and `MATPLOTLIB` (Hunter 2007) PYTHON packages.

REFERENCES

- Agertz O. et al., 2007, *MNRAS*, 380, 963
 Agertz O., Romeo A. B., Grisdale K., 2015, *MNRAS*, 449, 2156
 Altay G., Theuns T., Schaye J., Crighton N. H. M., Dalla Vecchia C., 2011, *ApJ*, 737, L37
 Astropy Collaboration et al., 2013, *A&A*, 558, A33
 Aumer M., White S. D. M., Naab T., Scannapieco C., 2013, *MNRAS*, 434, 3142
 Bahé Y. M., McCarthy I. G., 2015, *MNRAS*, 447, 969
 Bahé Y. M., McCarthy I. G., Balogh M. L., Font A. S., 2013, *MNRAS*, 430, 3017
 Baldry I. K. et al., 2012, *MNRAS*, 421, 621
 Benson A. J., Bower R., 2010, *MNRAS*, 405, 1573
 Bigiel F. et al., 2011, *ApJ*, 730, L13
 Blitz L., Rosolowsky E., 2006, *ApJ*, 650, 933 (BR06)
 Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, *ApJ*, 379, 440
 Boomsma R., Oosterloo T. A., Fraternali F., van der Hulst J. M., Sancisi R., 2008, *A&A*, 490, 555
 Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
 Bower R. G., 1991, *MNRAS*, 248, 332
 Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, *MNRAS*, 398, 1150
 Braun R., Thilker D. A., Walterbos R. A. M., Corbelli E., 2009, *ApJ*, 695, 937
 Broeils A. H., Rhee M.-H., 1997, *A&A*, 324, 877
 Broeils A. H., van Woerden H., 1994, *AAPS*, 107, 129
 Catinella B. et al., 2010, *MNRAS*, 403, 683
 Catinella B. et al., 2013, *MNRAS*, 436, 34
 Chabrier G., 2003, *PASP*, 115, 763
 Christensen C., Quinn T., Governato F., Stilp A., Shen S., Wadsley J., 2012, *MNRAS*, 425, 3058
 Crain R. A. et al., 2009, *MNRAS*, 399, 1773
 Crain R. A. et al., 2015, *MNRAS*, 450, 1937
 Cunnamea D., Andrianomena S., Cress C. M., Faltenbacher A., Gibson B. K., Theuns T., 2014, *MNRAS*, 438, 2530
 Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
 Davé R., Oppenheimer B. D., Katz N., Kollmeier J. A., Weinberg D. H., 2010, *MNRAS*, 408, 2051
 Davé R., Katz N., Oppenheimer B. D., Kollmeier J. A., Weinberg D. H., 2013, *MNRAS*, 434, 2645
 Dickey J. M., Lockman F. J., 1990, *ARA&A*, 28, 215
 Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
 Duffy A. R., Kay S. T., Battye R. A., Booth C. M., Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 420, 2799
 Dutton A. A. et al., 2011, *MNRAS*, 416, 322
 Fabello S., Kauffmann G., Catinella B., Li C., Giovanelli R., Haynes M. P., 2012, *MNRAS*, 427, 2841
 Fu J. et al., 2013, *MNRAS*, 434, 1531
 Furlong M. et al., 2015a, *MNRAS*, 450, 4486
 Furlong M. et al., 2015b, preprint ([arXiv:1510.05645](https://arxiv.org/abs/1510.05645))
 Giovanelli R. et al., 2005, *AJ*, 130, 2598
 Glover S. C. O., Clark P. C., 2012, *MNRAS*, 421, 9
 Gnedin N. Y., Kravtsov A. V., 2011, *ApJ*, 728, 88 (GK11)
 Governato F., Willman B., Mayer L., Brooks A., Stinson G., Valenzuela O., Wadsley J., Quinn T., 2007, *MNRAS*, 374, 1479
 Governato F. et al., 2010, *Nature*, 463, 203
 Guo Q. et al., 2011, *MNRAS*, 413, 101
 Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*. *Frontieres*, Paris, p. 64
 Haardt F., Madau P., 2012, *ApJ*, 746, 125
 Hopkins P. F., 2013, *MNRAS*, 428, 2840
 Huang M.-L., Kauffmann G., 2014, *MNRAS*, 443, 1329
 Huang S., Haynes M. P., Giovanelli R., Brinchmann J., 2012, *ApJ*, 756, 113
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Jenkins A., 2013, *MNRAS*, 434, 2094
 Katz N., Weinberg D. H., Hernquist L., 1996, *ApJS*, 105, 19
 Kauffmann G., White S. D. M., Guiderdoni B., 1993, *MNRAS*, 264, 201
 Kereš D., Katz N., Weinberg D. H., Davé R., 2005, *MNRAS*, 363, 2
 Krumholz M. R., 2013, *MNRAS*, 436, 2747 (K13)
 Krumholz M. R., McKee C. F., Tumlinson J., 2009, *ApJ*, 699, 850
 Lackner C. N., Cen R., Ostriker J. P., Joung M. R., 2012, *MNRAS*, 425, 641
 Lagos C. D. P., Baugh C. M., Lacey C. G., Benson A. J., Kim H.-S., Power C., 2011, *MNRAS*, 418, 1649
 Lagos C. D. P., Davis T. A., Lacey C. G., Zwaan M. A., Baugh C. M., Gonzalez-Perez V., Padilla N. D., 2014, *MNRAS*, 443, 1002
 Lagos C. D. P. et al., 2015, *MNRAS*, 452, 3815
 Leroy A. K., Walter F., Brinks E., Bigiel F., de Blok W. J. G., Madore B., Thornley M. D., 2008, *AJ*, 136, 2782 (L08)
 McCarthy I. G., Schaye J., Font A. S., Theuns T., Frenk C. S., Crain R. A., Dalla Vecchia C., 2012, *MNRAS*, 427, 379
 Marinacci F., Pakmor R., Springel V., 2014, *MNRAS*, 437, 1750
 Mitchell N. L., McCarthy I. G., Bower R. G., Theuns T., Crain R. A., 2009, *MNRAS*, 395, 180
 Mitchell N. L., Bower R. G., Theuns T., Vorobyov E. I., 2012, in Capuzzo-Dolcetta R., Limongi M., Tornambè A., eds, *ASP Conf. Ser. Vol. 453, Advances in Computational Astrophysics: Methods, Tools, and Outcome*. *Astron. Soc. Pac.*, San Francisco, p. 19
 Nelson D., Vogelsberger M., Genel S., Sijacki D., Kereš D., Springel V., Hernquist L., 2013, *MNRAS*, 429, 3353
 Okamoto T., Eke V. R., Frenk C. S., Jenkins A., 2005, *MNRAS*, 363, 1299
 Oppenheimer B. D., Davé R., 2006, *MNRAS*, 373, 1265
 Oppenheimer B. D., Davé R., Kereš D., Fardal M., Katz N., Kollmeier J. A., Weinberg D. H., 2010, *MNRAS*, 406, 2325
 Pawlik A. H., Schaye J., 2008, *MNRAS*, 389, 651
 Planck Collaboration XVI, 2014, *A&A*, 571, A16
 Popping A., Davé R., Braun R., Oppenheimer B. D., 2009, *A&A*, 504, 15
 Popping G., Somerville R. S., Trager S. C., 2014, *MNRAS*, 442, 2398
 Prochaska J. X., Wolfe A. M., 2009, *ApJ*, 696, 1543
 Puchwein E., Springel V., 2013, *MNRAS*, 428, 2966
 Rafieferantsoa M., Davé R., Anglés-Alcázar D., Katz N., Kollmeier J. A., Oppenheimer B. D., 2015, *MNRAS*, 453, 3980
 Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013a, *MNRAS*, 430, 2427
 Rahmati A., Schaye J., Pawlik A. H., Raičević M., 2013b, *MNRAS*, 431, 2261
 Rahmati A., Schaye J., Bower R. G., Crain R. A., Furlong M., Schaller M., Theuns T., 2015, *MNRAS*, 452, 2034
 Richings A. J., Schaye J., Oppenheimer B. D., 2014a, *MNRAS*, 440, 3349
 Richings A. J., Schaye J., Oppenheimer B. D., 2014b, *MNRAS*, 442, 2780
 Rosas-Guevara Y. M. et al., 2015, *MNRAS*, 454, 1038
 Saintonge A. et al., 2011, *MNRAS*, 415, 32
 Schaller M., Dalla Vecchia C., Schaye J., Bower R. G., Theuns T., Crain R. A., Furlong M., McCarthy I. G., 2015, *MNRAS*, 454, 2277
 Schaye J., 2004, *ApJ*, 609, 667
 Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
 Schaye J. et al., 2015, *MNRAS*, 446, 521
 Serra P. et al., 2012, *MNRAS*, 422, 1835
 Shen S., Mo H. J., White S. D. M., Blanton M. R., Kauffmann G., Voges W., Brinkmann J., Csabai I., 2003, *MNRAS*, 343, 978
 Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877
 Springel V., 2005, *MNRAS*, 364, 1105
 Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
 Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
 Springel V., Di Matteo T., Hernquist L., 2005a, *MNRAS*, 361, 776
 Springel V. et al., 2005b, *Nature*, 435, 629

Steinmetz M., Navarro J. F., 1999, *ApJ*, 513, 555
 Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, *MNRAS*, 373, 1074
 Stinson G. S. et al., 2015, *MNRAS*, 454, 1105
 Trayford J. W. et al., 2015, *MNRAS*, 452, 2879
 van de Voort F., Schaye J., Booth C. M., Haas M. R., Dalla Vecchia C., 2011, *MNRAS*, 414, 2458
 van de Voort F., Schaye J., Altay G., Theuns T., 2012, *MNRAS*, 421, 2809
 Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
 Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
 Walch S. et al., 2015, *MNRAS*, 454, 238
 Walker A. P., Gibson B. K., Pilkington K., Brook C. B., Dutta P., Stanimirovic S., Stinson G. S., Bailin J., 2014, *MNRAS*, 441, 525
 Walter F., Brinks E., de Blok W. J. G., Bigiel F., Kennicutt R. C., Jr, Thornley M. D., Leroy A., 2008, *AJ*, 136, 2563
 Wang J. et al., 2013, *MNRAS*, 433, 270
 Wang J. et al., 2014, *MNRAS*, 441, 2159
 Wiersma R. P. C., Schaye J., Smith B. D., 2009a, *MNRAS*, 393, 99

Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, *MNRAS*, 399, 574
 Wong T., Blitz L., 2002, *ApJ*, 569, 157
 Yang X., Mo H. J., van den Bosch F. C., Zhang Y., Han J., 2012, *ApJ*, 752, 41
 Zhang W., Li C., Kauffmann G., Xiao T., 2013, *MNRAS*, 429, 2191

APPENDIX A: SENSITIVITY TO MODELLING ASSUMPTIONS

A1 H I modelling

We here explore in more detail the impact of adopting alternative prescriptions to model H I in post-processing. Four groups of variations are considered in turn: (i) a different parametrization of the H₂ pressure law; (ii) theoretically motivated H I/H₂ partition schemes; (iii) more simplistic partition schemes; and (iv) a simple temperature–density cut to identify ‘H’ gas particles.

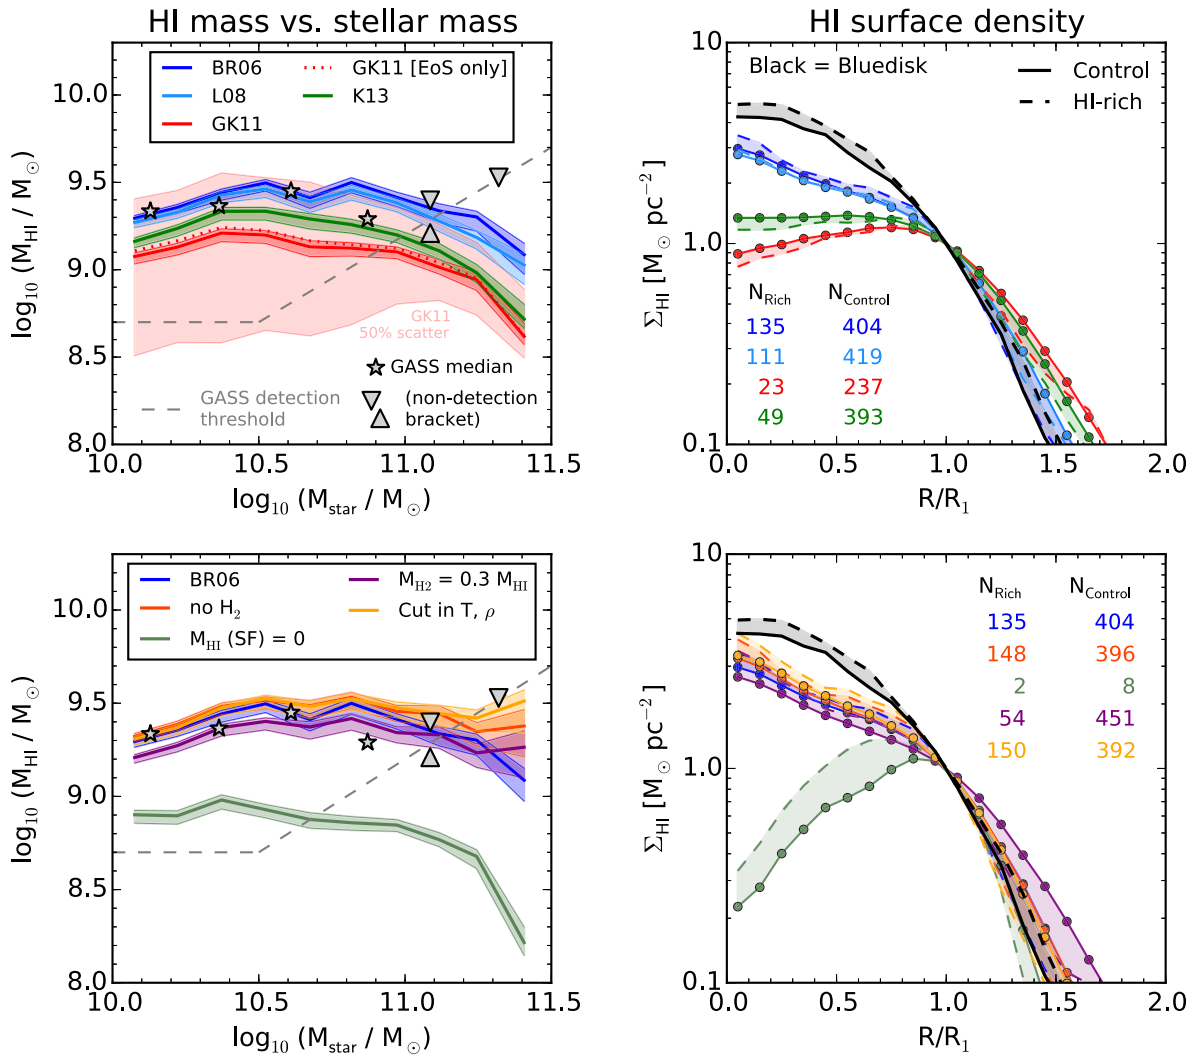


Figure A1. Comparison of the $M_{\text{HI}}-M_{\text{star}}$ relation (left) and H I surface density profiles (right) predicted by different H I models. Top row: empirical pressure law fit of Leroy et al. (2008, light blue) as well as the theoretically motivated formulae of Gnedin & Kravtsov (2011, red) and Krumholz (2013, green). The bottom row shows more simplistic models: ignoring the presence of H₂ (orange), assuming all star-forming gas particles are exclusively H₂ (green), a fixed ratio of $m_{\text{H}_2} = 0.3 m_{\text{HI}}$ as used by Popping et al. (2009, purple), and a temperature–density cut (yellow). For ease of comparison, the empirical Blitz & Rosolowski (2006) pressure law model is shown in both panels (dark blue). Lines show median values, the dark shaded bands in the right-hand panels indicate their 1σ uncertainty. For the Gnedin & Kravtsov (2011) model in the top (red), we additionally show the 50 per cent scatter with a light-shaded band as explained in the text; this is similar for the other models and hence not shown for clarity. Numbers in the right-hand panels show how many H I-rich and control galaxies are predicted by each model.

For each of these, we show in Fig. A1 the resulting $M_{\text{star}}-M_{\text{HI}}$ relation (left-hand panels) and scaled H I surface density profiles in analogy to Fig. 7 (right-hand panels). Solid lines represent running medians; the Blitz & Rosolowski (2006) pressure law (BR06) is shown in dark blue in both rows for ease of comparison.

In the right-hand plots, the dark shaded bands represent the 1σ uncertainty on the median (see Section 6.4.2). The 50 per cent galaxy-to-galaxy scatter is shown for the Gnedin & Kravtsov (2011) model (GK11) in the top row as a light red band; for the other models this scatter is similar and hence not shown for clarity. Observed median H I masses from GASS (Catinella et al. 2013) are overplotted as grey symbols. In analogy to Fig. 1, we consider setting non-detections in GASS both to zero and to the observational upper limit. Where the median is the same in both approaches (i.e. where its value is unaffected by the presence of non-detections), we show it as a star-symbol, otherwise the two differing values which bracket the true median are shown by downward and upward facing triangles.¹³ The grey dashed line shows the nominal GASS detection threshold.

In the right-hand plots, ‘control’ ($\log_{10}(M_{\text{HI}}/M_{\odot}) = [9.1, 9.8]$) and ‘H I-rich’ galaxies ($\log_{10}(M_{\text{HI}}/M_{\odot}) \geq 9.8$) are differentiated by solid and dashed lines, respectively; their numbers are shown in the bottom-left or top-right corners. For comparison, the observational data from Bluedisk (Wang et al. 2014) are shown in black. Shaded bands highlight the extent of discrepancy between the H I-rich and control profiles.

A1.1 Different parametrizations of the H_2 pressure law

The light blue lines in the top panels represent the results obtained with the alternative pressure law parametrization of Leroy et al. (2008, hereafter L08). Their analysis is based on the THINGS survey (Walter et al. 2008) with a larger sample of galaxies (23) than that of BR06 (11 galaxies). The best-fitting parameters determined by L08 are both slightly lower than in BR06: $P_0/k_B = 1.7 \times 10^4 \text{ cm}^{-3} \text{ K}$ and $\alpha = 0.8$ (BR06: $P_0/k_B = 4.3 \times 10^4 \text{ cm}^{-3} \text{ K}$ and $\alpha = 0.92$). However, the impact on our results is negligible, and both methods agree to better than 0.05 dex except at the high- M_{star} end (where the molecular fraction is highest; Saintonge et al. 2011). Likewise, the central Σ_{HI} profiles are lower in L08 (especially for H I-rich galaxies), but only by $\lesssim 10$ per cent. We therefore conclude that our results are not significantly affected by uncertainties in the H_2 pressure law parametrization.

A1.2 Theoretically motivated H I/ H_2 partition schemes

In addition to the GK11 formula, Lagos et al. (2015) also implemented the theoretically motivated H I/ H_2 partition schemes of Krumholz (2013, hereafter K13) into EAGLE, which is derived from analytic modelling; as with GK11, we refer to Lagos et al. (2015) for details of the implementation. In the top row of Fig. A1, the K13 model is shown in dark green, and the corresponding result from GK11 in red.

¹³ The reason for being able to compute a non-zero lower limit in the second-highest stellar mass bin which differs from the upper limit is that 52 per cent of these galaxies are detected in H I by GASS, but some non-detected galaxies have upper limits higher than the lowest detected H I masses; see Catinella et al. (2010) for details. In the highest stellar mass bin, only 36 per cent of galaxies are detected so the lower limit on the median is zero.

In contrast to the L08 parametrization discussed above, the K13 model predicts H I masses that are too low compared to observations (by ~ 0.2 dex at $M_{\text{star}} \approx 10^{10} M_{\odot}$ and ~ 0.5 dex at the high- M_{star} end). They are also lower than the predictions from either of the two pressure law parametrizations (see above), but consistently higher than what is predicted by the GK11 formula (red, a difference of ~ 0.1 – 0.2 dex). As the top-right panel of Fig. A1 shows, this discrepancy is predominantly driven by the inner galaxy regions ($R \lesssim R_1$) where K13 (as well as GK11) combined with EAGLE fails to predict (median) surface densities above $1.5 M_{\odot} \text{ pc}^{-2}$ (lower by a factor of ~ 4 than observed by Bluedisk).

We also note that the Lagos et al. (2015) implementation of GK11 and K13 allows for the existence of a (small) molecular fraction in all gas particles, including those not forming stars. In contrast, the pressure law models of BR06 and L08 are only applied to star-forming particles whose properties admit the formation of a significant amount of cold gas (Schaye 2004; Schaye & Dalla Vecchia 2008; Schaye et al. 2015). To demonstrate that this is of little significance for the resulting H I content (< 0.05 dex), the median M_{HI} obtained with the GK11 prescription but H_2 restricted to star-forming particles is shown in the top-left panel as a red dotted line.

A1.3 Simple H I/ H_2 partition schemes

The H I/ H_2 partition schemes considered so far all have a particle-level physical basis, either on theoretical or observational grounds. Three more simple separation prescriptions are considered in the bottom panel of Fig. A1: ignoring the presence of H_2 altogether (red), assuming a fixed ratio of $m_{H_2}/m_{HI} = 0.3$ (purple), which approximately corresponds to the cosmic average at $z = 0$ for the galaxy masses considered here (Saintonge et al. 2011) and was used by Popping et al. (2009); and an extreme prescription of assuming H I is only found in non-star-forming particles (which, in turn, contain no H_2 ; pale green).

As expected, the latter prescription strongly underpredicts (by ~ 0.5 – 1 dex) the H I content of simulated galaxies, even more than the GK11 or K13 partition schemes, and shows an unphysical, strong central decline in Σ_{HI} (because most gas inside R_1 has a non-zero star formation rate in the simulation). In contrast, the other two agree relatively closely with BR06 (dark blue), deviating by less than 0.1 dex in M_{HI} except for the few galaxies with $M_{\text{star}} > 10^{11.2} M_{\odot}$. By construction, the model ignoring H_2 predicts slightly higher M_{HI} whereas the fixed particle-level m_{H_2}/m_{HI} ratio (purple) generally yields slightly lower integrated M_{HI} except for the comparatively H_2 rich most massive galaxies. Both are compatible with the observational constraints from GASS.

These trends are mirrored by the Σ_{HI} profiles, except that the profiles for ‘control’ galaxies are slightly too shallow with fixed particle-level m_{H_2}/m_{HI} ratio (purple), plausibly because this assumption breaks down in the (molecule-poor) outer parts and therefore underestimates the true R_1 radius.

A1.4 Temperature–density cut for H I

Finally, we test the arguably simplest model of assuming a cut in temperature¹⁴ ($T \leq 10^{4.5} \text{ K}$) and density ($n_{\text{H}} \geq 0.01 \text{ cm}^{-3}$); we then set $m_{\text{HI}} = m_{\text{H}}$ for these particles and $m_{\text{HI}} = 0$ for all others. In contrast to all other models explored here, this prescription does not

¹⁴ We use a fixed value of $T = 10^4 \text{ K}$ for star-forming gas here.




Figure A2. Difference between the ‘GASS-equivalent’ H I measurement (within a 2D aperture of 70 kpc, and $|\Delta v| \leq 400 \text{ km s}^{-1}$), and the mass obtained from H I images at the Bluedisk resolution with a threshold H I surface density of $\Sigma_{\text{thresh}} = 0.37 M_{\odot} \text{ pc}^{-2}$ (top panel). Thin shaded bands enclose 50 per cent of galaxies in random orientation (green) or face-on (blue), while the darker shaded bands indicate the statistical 1σ uncertainty on the median (solid line). The difference is less than 10 per cent for galaxies with H I-masses above the ‘control’ limit ($M_{\text{HI}} \geq 10^{9.1} M_{\odot}$; red line). For reference, the light blue vertical line indicates the division between H I-rich and control galaxies in Section 6. Bottom: Comparison to H I masses in spherical apertures of 30 kpc (red), 70 kpc (orange), and 150 kpc (yellow).

rely on the Rahmati et al. (2013a) fitting formula for calculating the neutral hydrogen fraction, and is shown in yellow in the bottom row of Fig. A1. Although both the integrated M_{HI} and surface density profiles are slightly higher than with the empirical BR06 approach (dark blue), this difference is similarly small as for the ‘ignoring H_2 ’ model (orange) and deviates by less than 0.1 dex from our default prescription except for the most massive galaxies.

A2 Integration aperture

In Fig. A2, we test to what extent the H I masses obtained using our default ‘GASS-equivalent’ mass definition ($R_{2\text{D}} \leq 70 \text{ kpc}$ and $|\Delta v_z| \leq 400 \text{ km s}^{-1}$) agree with other choices. The top panel compares to the ‘Bluedisk-equivalent’ definition (see Section 6.4): we make a $140 \times 140 \text{ kpc}$ mock H I image from all particles with $|z| \leq$

70 kpc, convolve this with a Gaussian $14 \times 9 \text{ kpc}$ FWHM beam, and then integrate over all pixels with $\Sigma_{\text{HI}} \geq 0.37 M_{\odot} \text{ pc}^{-2}$ (as in Wang et al. 2013). Blue and green lines show the median H I mass difference for galaxies in face-on and random orientation, respectively (dark and light shaded bands give the 1σ uncertainty on the median and 50 per cent scatter). Although the ‘Bluedisk’ method gives somewhat lower H I masses, the difference is only ~ 0.1 dex (25 per cent) at the lower limit of our ‘control’ sample ($M_{\text{HI}} = 10^{9.1} M_{\odot}$, vertical red dash-dotted line), and less than 0.05 dex (12 per cent) for ‘H I-rich’ galaxies ($M_{\text{HI}} \geq 10^{9.8} M_{\odot}$). Our results presented above are therefore not significantly affected by the difference between these two integration methods (see also Fig. 6).

In the bottom panel, we compare to simple 3D radial cuts at 30 kpc (red), 70 kpc (orange), and 150 kpc (yellow). The second of these agrees very well with our ‘GASS-equivalent’ aperture down to H I masses as low as $10^{8.5} M_{\odot}$: only 13 per cent of all galaxies deviate by more than 25 per cent. At intermediate H I masses ($M_{\text{HI}} \approx 10^{9.5} M_{\odot}$) the spherical cut leads to a slightly *higher* H I mass: this is because some particles within the sphere have peculiar z -velocities larger than 400 km s^{-1} and are therefore excluded in the ‘GASS-equivalent’ measurement (the excess is only of the order of a few per cent, however). Adopting a smaller aperture (30 kpc, red) leads to much more significant mass deficits of ~ 0.2 dex (60 per cent) at both the high- and low- M_{HI} ends. Perhaps surprisingly, the masses still agree well at intermediate M_{HI} . This behaviour is likely due to an interplay of two effects: at masses of $M_{\text{HI}} \gtrsim 10^{9.5} M_{\odot}$, R_1 is typically 20 kpc or more (Fig. 6), so a 30 kpc aperture misses a non-negligible amount of H I in the galactic outskirts (see Fig. 8). For H I poor galaxies ($M_{\text{HI}} \lesssim 10^{9.0} M_{\odot}$), on the other hand – which are predominantly of ‘irregular’ H I type (Fig. 5) – a significant contribution to the ‘GASS-equivalent’ H I mass comes from H I clumps along the line of sight which are not directly connected to the galaxy. This would also explain why these galaxies show slightly smaller-than-expected R_1 radii in Fig. 6: it is really their H I masses that are slightly too high.

Overall, we conclude that our H I masses are not just compatible with the respective observations, but also physically meaningful at least for galaxies with $M_{\text{HI}} > 10^9 M_{\odot}$ upon which the majority of our analysis here is based.

APPENDIX B: TESTS OF NUMERICAL CONVERGENCE

In this appendix, we test the numerical convergence of our results, by comparing three additional simulations from the EAGLE suite run in a 25 cMpc periodic box. Two of these (Ref-L025N0752 and Recal-L025N0752 in the terminology of Schaye et al. 2015) were run at eight times higher mass resolution (i.e. $m_{\text{gas}} \simeq 2.2 \times 10^5 M_{\odot}$), while the third one (Ref-L025N0376) uses the same resolution as in the large Ref-L100N1504 simulation that we have analysed in the main part of this paper. The difference between the two high-resolution simulations is that Ref-L025N0752 uses the same sub-grid physics parameters as run Ref-L100N1504 (and Ref-L025N0376), whereas in Recal-L025N0752, the parameters were re-calibrated to improve the match to the observed galaxy stellar mass function. For more details, the interested reader is referred to Schaye et al. (2015).

The rationale behind this re-calibration is that, as explained in detail by Schaye et al. (2015), the interpretation of the numerical sub-grid physics parameters in a hydrodynamical simulation is in general resolution dependent. This makes it unlikely that a

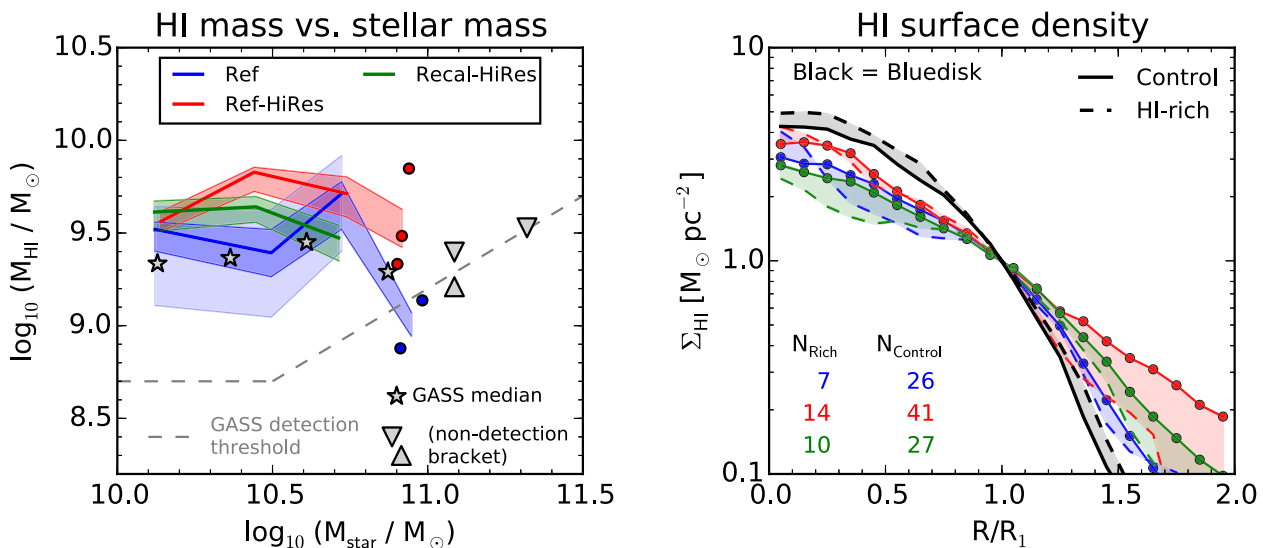


Figure B1. Strong and weak convergence tests for EAGLE H I properties. The layout is analogous to Fig. A1 above. We compare three 25 cMpc simulations: Ref-L025N0376 (blue; same resolution and parameters as the large run analysed in the main part of the paper), Ref-L025N0752 (red, higher resolution but same parameters), and Recal-L025N0752 (green, higher resolution and re-calibrated parameters). See text for details. In both higher resolution simulations, galaxies are more H I rich, particularly in the galaxy outskirts (shallower profiles in right-hand panel). As expected, this difference is less strong in the re-calibrated simulation (green).

simulation such as EAGLE can achieve ‘strong convergence’ – i.e. the calculation result being unaffected by a change in resolution while keeping the sub-grid parameters fixed. However, one may still obtain ‘weak convergence’ by re-calibrating the sub-grid parameters at the new resolution level.

We will test both the strong convergence of our results, by comparing the two Ref-L025 simulations (standard and high resolution), and weak convergence by comparing Ref-L025N0376 and Recal-L025N0752. Both these tests are presented in Fig. B1, the structure of which is identical to Fig. A1 above. Note, though, that due to the much smaller volume of the 25 cMpc boxes compared to the large 100 cMpc run (by a factor of $4^3 = 64$), the number of galaxies is significantly smaller here, resulting in larger statistical uncertainties.

Looking first at the total H I masses (left-hand panel), the general trend is that the high-resolution simulations contain galaxies that have higher M_{HI} at a given M_{star} . This difference is larger for simulation Ref-L025N0752 (red), indicating that the weak convergence (blue/green curves; offset by $\lesssim 0.2$ dex) is better than strong convergence (blue/red; offset by $\lesssim 0.4$ dex), as expected. However, even the recalibrated high-resolution simulation is significantly too H I-rich at least at stellar masses between $10^{10.0}$ and $10^{10.5} M_{\odot}$ compared to observations.

These trends are mirrored by the H I surface density profiles shown in the right-hand panel of Fig. B1. In the central region ($R \lesssim R_1$), agreement between all three simulations is quite good, with the biggest discrepancy again between the two Ref simulations (strong convergence test; high-resolution simulation higher by ~ 0.2 dex). In the outskirts, both high-resolution simulations show profiles that are shallower than the standard resolution run Ref-L025N0376, indicating extra H I at increased resolution. Again, this difference is stronger for the non-recalibrated high-resolution run Ref-L025N0752 (red), in particular for ‘control’ galaxies with $\log_{10}(M_{\text{HI}}/M_{\odot}) = [9.1, 9.8]$ where the offset reaches ~ 0.5 dex. Although the agreement is better for the recalibrated high-resolution run (green), an offset of ~ 0.2 dex remains at $1.5 R_1$. Furthermore, neither higher resolution simulation reproduces the similarity be-

tween the outer H I profiles of H I-rich and control galaxies seen in Bluedisk, although here, as well, the effect is ameliorated by the parameter recalibration: at $1.5 R_1$, the difference between Σ_{HI} in H I-rich and control galaxies is ~ 0.3 dex in Ref-L025N0752, and only ~ 0.1 dex in Recal-L025N0752.

In summary, both H I masses and surface density profiles in EAGLE show reasonably good weak convergence, with differences at a level of ~ 0.1 – 0.2 dex in the sense that higher resolution galaxies contain more H I. Strong convergence is considerably less good, with differences up to ~ 0.5 dex. It is conceivable that further parameter fine-tuning might improve the weak convergence in terms of H I properties, but with typical observational uncertainties at a level of ~ 0.1 – 0.2 dex, it is questionable whether this would actually be justified. A more detailed discussion of convergence of H I masses in EAGLE will be presented by Crain et al. (in preparation).

APPENDIX C: BLUEDISK EQUIVALENT SAMPLE SELECTION

Unlike GASS, the Bluedisk sample is by construction biased. The actual selection procedure is somewhat complex (see Wang et al. 2013 for details) and involves a large range of galaxy parameters including stellar mass, stellar surface mass density, NUV and r -band colours as well as colour gradients. Although it is, in principle, possible to mimic all these for our simulated galaxies, the risk of many small deviations adding up to major inconsistencies is considerable, especially because of the complex influence of dust on the NUV luminosities. However, we demonstrate here that the selection can be reproduced with a simple size cut of $R_{50,z} = 3$ kpc, and a threshold at $\log_{10}(M_{\text{HI}}/M_{\odot}) = 9.8$ to separate ‘H I-rich’ and ‘control’ galaxies.

C1 Overall sample selection by galaxy size

Bluedisk targeted 25 galaxies expected to be H I-rich, selected out of those for which the photometric gas fraction plane of Catinella




Figure C1. Stellar surface density versus stellar mass for Bluedisk galaxies (green) and the GASS sample (purple). Compared to the unbiased GASS sample, Bluedisk galaxies have preferentially lower surface densities, i.e. larger sizes.

et al. (2010) with an additional colour gradient correction (Wang et al. 2013) predicted an H I mass fraction ≥ 0.6 dex above the median at a given stellar mass (Catinella et al. 2010). The H I mass fraction is strongly anticorrelated with stellar surface mass density $\mu_* \equiv M_{\text{star}} / (2 R_{50,z}^2)$, so these target galaxies have preferentially low μ_* .¹⁵ An equal number of control galaxies were selected to match the H I-rich targets in μ_* (amongst other properties), so these are also preferentially of low stellar surface density.

In Fig. C1 we plot μ_* against stellar mass, M_{star} , for all Bluedisk galaxies (large green circles), and also for the (μ_* -unbiased) GASS sample in purple. As expected, the Bluedisk galaxies are congregated at the lower end of the GASS distribution, with an upper boundary corresponding approximately to a fixed size of $R_{50,z} = 3$ kpc (only two out of 50 galaxies are slightly smaller).

We have therefore calculated the half-light radius $R_{50,z}$ for our simulated galaxies, using luminosities calculated for each star particle using stellar population synthesis (Trayford et al. 2015). However, as we show in Fig. C2, the resulting sizes are systematically larger than in SDSS, of which we take the GASS parent sample from Catinella et al. (2010) with almost 12 000 galaxies as a large, unbiased subset. The origin of this discrepancy is not entirely clear, because Schaye et al. (2015) have demonstrated that the $z = 0.1$ stellar sizes in EAGLE, when calculated based on Sérsic profile fits to stellar mass profiles as described by McCarthy et al. (2012), are in good agreement with the SDSS-based results of Shen et al. (2003). We note here that both the observational size–mass relation of Shen et al. (2003) and the EAGLE relation shown in Schaye et al. (2015) are based on galaxies with Sérsic index $n_s < 2.5$. On the other hand, Baldry et al. (2012) find a size–mass relation for blue galaxies that is ~ 0.2 dex higher than the result of Shen et al. (2003),

¹⁵ $R_{50,z}$ is defined as the projected radius enclosing 50 per cent of the Petrosian flux in the z -band. The Petrosian flux, in turn, is the total flux within two Petrosian radii r_p . Note that, in SDSS, Petrosian fluxes in all five bands are based on r_p as measured in the r band, to avoid aperture bias effects.




Figure C2. SDSS z -band Petrosian half-light radii $R_{50,z}$ in EAGLE (blue) and SDSS (GASS parent sample; purple). The sizes of the simulated galaxies appear systematically larger than in SDSS, the difference being comparable to the typical size excess of the Bluedisk galaxies (green). Note that the size parameter $R_{50,z}$ used here is not the same as that shown by Schaye et al. (2015).

and Dutton et al. (2011) have shown that the Shen et al. (2003) sizes are biased low due to their use of circular apertures. Our Fig. C1, on the other hand, does not select galaxies by any other criterion than stellar mass.

It is also possible that the observational analysis in the SDSS pipeline slightly underestimates the true Petrosian radius (e.g. due to limited depth of the images), that the z -band profiles differ from the stellar mass profiles for EAGLE, or otherwise that the (outer) stellar light profiles in the simulations are somewhat too shallow. For the present purpose, we simply re-scale the $R_{50,z}$ radii of the EAGLE galaxies so that its distribution function matches that of the SDSS sample; from Fig. C2 it can be seen that this typically corresponds to a reduction by a factor of ~ 40 per cent. The EAGLE galaxy sizes and their evolution are studied in detail by Furlong et al. (2015b). We emphasize, however, that this discrepancy has virtually no impact on our results here, because we *only* use stellar sizes to select the overall subsample to compare to Bluedisk and have verified that our results are virtually unchanged when we instead select all EAGLE galaxies, regardless of size.

C2 Division into H I-rich and control galaxies by total H I mass

After observation, the Bluedisk galaxies were (re-)classified into H I-rich and control samples based on a comparison between the actual observed H I mass ($M_{\text{H I, observed}}$) and that predicted by the original Catinella et al. (2010) gas fraction plane ($M_{\text{H I, C10-prediction}}$). Galaxies with H I mass larger than predicted are classified as H I-rich, while those whose H I mass is less than predicted were assigned to the control sample (compared to all galaxies at a given stellar mass, these are still slightly H I enhanced because of the aforementioned bias towards low stellar surface densities).¹⁶

¹⁶ As Wang et al. (2013) show, this effect is largely cancelled out by a bias in the Catinella et al. (2010) gas fraction plane, which overpredicts the H I content of most galaxies targeted as the ‘control’ sample. Their (actual) $M_{\text{H I}}$ is therefore not far from the overall median.




Figure C3. Definitions of ‘H I-rich’ and ‘control’ galaxies in the Bluedisk survey (blue and yellow points) and the simplified sample division adopted here (horizontal dotted lines and shaded regions). The observational division is reproduced almost exactly by these simple cuts.

In Fig. C3, we plot the observed and predicted H I masses of Bluedisk galaxies, with H I-rich galaxies ($M_{\text{H I, observed}} > M_{\text{H I, C10-prediction}}$) shown in blue and control galaxies ($M_{\text{H I, observed}} < M_{\text{H I, C10-prediction}}$) in yellow. As can be seen, the two samples can be relatively cleanly separated by a single cut at $M_{\text{H I, observed}} = 10^{9.8} M_{\odot}$: There are only two H I-rich galaxies whose H I mass is below this threshold (one of them only marginally), and conversely only three galaxies in the Bluedisk control sample lie above the threshold. At the lower end, the Bluedisk control sample is well limited by $M_{\text{H I, observed}} \geq 10^{9.1} M_{\odot}$, with again only two galaxies falling outside this range (one of which is excluded from analysis because of a nearby companion; see Wang et al. 2013).

To test our overall sample selection (see above), we also plot the GASS galaxies as purple circles, large ones representing ‘big’ GASS galaxies ($R_{50,z} \geq 3$ kpc, in the same range as Bluedisk), and all others shown as small points. Encouragingly, the first set is also located in the same region as the Bluedisk galaxies in this plot, whereas the second set (galaxies smaller than Bluedisk) form a tail towards the lower left (i.e. towards lower H I masses).

APPENDIX D: INCLINATION OF THE H I DISC AND OPTICAL ELLIPTICITY

We had noted in the discussion of Fig. 7 that differences in the outer H I surface density profiles emerge depending on whether the inclination correction is performed using the angular momentum axis of the H I, or the optical r -band ellipticity. We here test the possibility that this is due to incorrect alignment of the H I discs in the latter approach.

In the top panel of Fig. D1, we show the distribution of axis ratios b/a fit to the 25 mag arcsec $^{-2}$ r -band isophote of our EAGLE galaxies (blue) and those for the Bluedisk sample (magenta), as well as the GASS parent sample (yellow). Because the orientation of the galaxies in the simulation box is random, the distribution of inclination angles $\cos(\theta)$ is flat, and the same should be true for b/a if the two are equal in a statistical sense. However, this is clearly not the case for either simulations or observations: both show a marked deficit at both small ($b/a \lesssim 0.3$, i.e. very elongated isophotes)

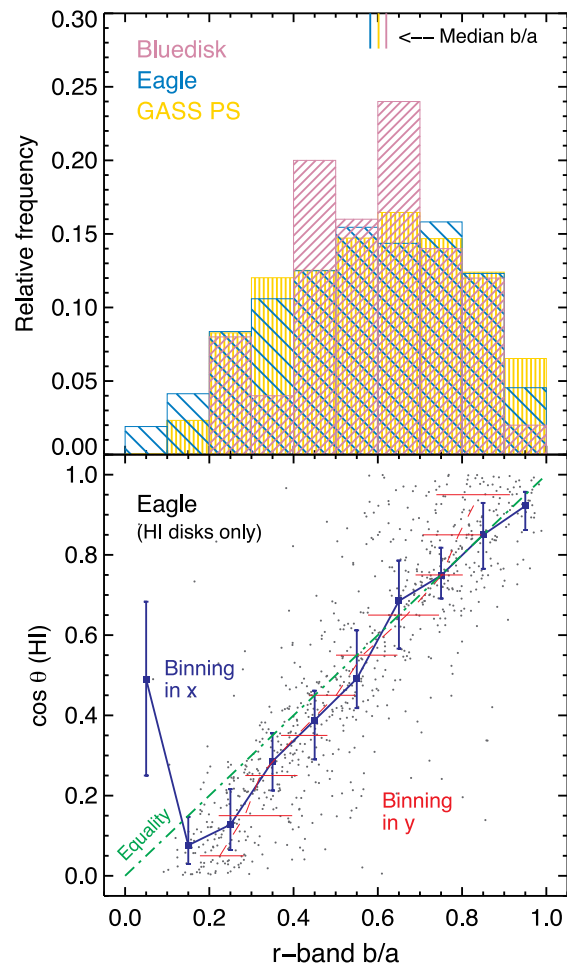


Figure D1. Correlation between the H I disc inclination and stellar r -band ellipticity. The top panel shows the distribution of the latter quantity (b/a) for Bluedisk (magenta), the GASS parent sample (yellow), and EAGLE (blue): This should be uniform if it actually measured the (random) inclination, but is clearly concentrated towards intermediate values of ~ 0.6 in all samples (vertical ticks at the top indicate the medians, which agree well). Bottom: direct comparison between b/a and H I inclination angle $\cos(\theta)$ for EAGLE galaxies with an H I disc. There is a fair amount of scatter for individual galaxies, but for $b/a \gtrsim 0.6$ the approximation that $\cos(\theta) = b/a$ is true on average.

and large ellipticities ($b/a \gtrsim 0.9$, almost circular). In principle, the Bluedisk sample selection could lead to a subtle selection bias towards certain axis ratios, but the fact that the distribution is very similar to the stellar-mass-only selected EAGLE and GASS samples makes this very unlikely. A two-sided KS test shows a likelihood of 0.26 for the EAGLE and Bluedisk distributions being drawn from the same parent population, and the medians – vertical lines at the top of the plot – are also very similar for all three data sets (0.58, 0.60, and 0.62 for EAGLE, GASS, and Bluedisk, respectively). Instead, the uneven distribution simply reflects the fact that extreme axis ratios are rare because even if a (stellar) disc is aligned perfectly edge-on (face-on), any deviation from a circular, infinitely thin disc will increase (decrease) the measured axis ratio and therefore drive the b/a distribution towards intermediate values.

In the bottom panel, we directly compare r -band axis ratio and H I disc inclination angle for EAGLE galaxies with an H I disc (i.e. excluding irregular H I morphologies; grey points), and also show the running median and 25th/75th percentile binned both by b/a

(blue) and $\cos(\theta_{\text{HI}})$ (red). There is clearly substantial scatter in the relation, of typically ~ 0.1 around the median, but in general there is a clear correlation between the two. As a function of $\cos\theta$ (red), the relation is slightly steeper than one-to-one, which explains the uneven shape of the distributions in the top panel. More importantly for the interpretation of observational data, however, is the opposite trend: at $b/a \gtrsim 0.6$, the median $\cos\theta$ (blue) is almost exactly equal to b/a (green line). Although care must evidently be taken when applying this relation to individual galaxies, the estimation of inclination angles from optical ellipticities should at least be reliable in a statistical sense for galaxies with moderately elliptical isophotes.

APPENDIX E: SENSITIVITY OF H I PROFILES TO THE ADOPTED BEAM SIZE

In the profile comparison in Section 6, we had artificially reduced the resolution of the simulated H I maps to match the beam size of the Westerbork Synthesis Radio Telescope (WSRT). We now briefly investigate how this smoothing has affected our results, by analysing the profiles obtained from higher resolution images with a (circular) beam of FWHM = 2 kpc, more similar to what is shown in Fig. 3. Based on our discussion in Section 6, we choose to




Figure E1. Variation of H I surface density profiles with beam size. The default low-resolution profiles with beam size 14×9 kpc FWHM (see Section 6) are shown with dashed lines, and the profiles obtained from higher-resolution maps with beam size 2×2 kpc FWHM as filled circles connected with dotted lines. In both cases, red represents control galaxies and green H I-rich galaxies. For guidance, the combined median Bluedisk profile is shown in black. Error bars indicate 1σ uncertainties on the median. Although there is some change in the detailed profile shape with increased resolution, the close match between H I-rich and control galaxies is robust to this change.

rotate galaxies to face-on as the most physically meaningful way of extracting surface density profiles, and present the result in Fig. E1. Particle H I masses are here calculated with the empirical BR06 H_2 formula.

The higher resolution profiles are shown with filled circles connected by dotted lines, red for control galaxies and green for those which are H I-rich. For comparison, we also re-produce the low-resolution (14×9 kpc) profiles shown in Fig. 7 with dashed lines. Statistical 1σ uncertainties on the medians are indicated with error bars in the former case, and the width of the shaded bands in the latter. For guidance, we also include the overall Bluedisk median profile (black).

Unsurprisingly, the increase in resolution does lead to some change in the detailed shape of the profile, but overall the effect is rather small. In the outer regions ($R \gtrsim 0.9 R_1$), the high-resolution profile is slightly steeper (reaching $0.1 M_\odot \text{pc}^{-2}$ at 1.4 instead of $1.5 R_1$); the same is true for the very centre ($R \lesssim 0.3 R_1$). More interestingly, however, Fig. E1 clearly demonstrates that the close agreement between H I-rich and control galaxies discovered by Wang et al. (2014) is *not* sensitive to the beam size, and is a genuine physical feature of our simulations. This strongly suggests that the same should also be true for the real Universe, a prediction that can be tested directly in future with high-resolution H I surveys such as from the Square Kilometre Array.