

Galaxy cold gas contents in modern cosmological hydrodynamic simulations

Romeel Davé^{1,2,3,4}, Robert A. Crain^{1,4}, Adam R. H. Stevens^{1,5,6}, Desika Narayanan^{1,7,8,9},
Amelie Saintonge^{1,10}, Barbara Catinella^{1,5,6} and Luca Cortese^{1,5,6}

¹*Institute for Astronomy, Royal Observatory, University of Edinburgh, Edinburgh EH9 3HJ, UK*

²*Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town 7535, South Africa*

³*South African Astronomical Observatories, Observatory, Cape Town 7925, South Africa*

⁴*Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK*

⁵*International Centre for Radio Astronomy Research, The University of Western Australia, Crawley, WA 6009, Australia*

⁶*ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Canberra, 2601 Australia*

⁷*Department of Astronomy, University of Florida, 211 Bryant Space Sciences Center, Gainesville, FL 32611, USA*

⁸*University of Florida Informatics Institute, 432 Newell Drive, CISE Bldg E251, Gainesville, FL 32603, USA*

⁹*Cosmic Dawn Center at the Niels Bohr Institute, University of Copenhagen and DTU-Space, Technical University of Denmark, Vibenshuset, Lyngbyvej 2, DK-2100 Copenhagen, Denmark*

¹⁰*Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK*

Accepted 2020 June 26. Received 2020 June 26; in original form 2020 February 17

ABSTRACT

We present a comparison of galaxy atomic and molecular gas properties in three recent cosmological hydrodynamic simulations, namely SIMBA, EAGLE, and IllustrisTNG, versus observations from $z \sim 0$ to 2. These simulations all rely on similar subresolution prescriptions to model cold interstellar gas that they cannot represent directly, and qualitatively reproduce the observed $z \approx 0$ H I and H₂ mass functions (HIMFs and H2MFs, respectively), CO(1–0) luminosity functions (COLFs), and gas scaling relations versus stellar mass, specific star formation rate, and stellar surface density μ_* , with some quantitative differences. To compare to the COLF, we apply an H₂-to-CO conversion factor to the simulated galaxies based on their average molecular surface density and metallicity, yielding substantial variations in α_{CO} and significant differences between models. Using this, predicted $z = 0$ COLFs agree better with data than predicted H2MFs. Out to $z \sim 2$, EAGLE’s and SIMBA’s HIMFs and COLFs strongly increase, while IllustrisTNG’s HIMF declines and COLF evolves slowly. EAGLE and SIMBA reproduce high- $L_{\text{CO}(1-0)}$ galaxies at $z \sim 1$ –2 as observed, owing partly to a median $\alpha_{\text{CO}}(z = 2) \sim 1$ versus $\alpha_{\text{CO}}(z = 0) \sim 3$. Examining H I, H₂, and CO scaling relations, their trends with M_* are broadly reproduced in all models, but EAGLE yields too little H I in green valley galaxies, IllustrisTNG and SIMBA overproduce cold gas in massive galaxies, and SIMBA overproduces molecular gas in small systems. Using SIMBA variants that exclude individual active galactic nucleus (AGN) feedback modules, we find that SIMBA’s AGN jet feedback is primarily responsible by lowering cold gas contents from $z \sim 1 \rightarrow 0$ by suppressing cold gas in $M_* \gtrsim 10^{10} M_\odot$ galaxies, while X-ray feedback suppresses the formation of high- μ_* systems.

Key words: ISM: evolution – galaxies: disc – galaxies: evolution – galaxies: formation – methods: numerical.

1 INTRODUCTION

Galaxies are made up of stars, gas, dust, black holes, and dark matter. Of those, stars represent the most straightforwardly visible component, and so have received the most observational and theoretical attention. However, in the modern baryon cycling paradigm of galaxy evolution, it is the exchange of gas between the interstellar medium (ISM) of galaxies and their surrounding circumgalactic medium (CGM) via inflows, outflows, and recycling that primarily governs how galaxies form and evolve (Tumlinson, Peebles & Werk 2017). As such, it is becoming clear that understanding the gas within and around galaxies is crucial for a full picture of galaxy formation and evolution.

Comprehensive observations of gas in galaxies are challenging because gas is typically diffuse and multiphase, and its emission is often best traced in less accessible portions of the electromagnetic spectrum. For instance, molecular gas (primarily H₂) is canonically traced via heavy element molecular emission lines in the millimetre regime; atomic gas is most straightforwardly traced as radio 21-cm emission; ionized gas appears as ultraviolet and optical emission lines; and gas at high temperatures is mostly evident via X-ray emission. Assembling observations of these various gas phases into a coherent scenario for the role of gas in galaxy evolution is an important goal for current models of galaxy evolution.

Within the ISM of star-forming galaxies, the dominant gas phases are cold ($T \lesssim 100$ K) and warm ($T \sim 10^3$ – 10^4 K), best traced by molecular and atomic hydrogen, respectively. Recent advances in observations of molecular and atomic gases have opened up new windows on understanding the role of ISM gas in galaxy evolution.

* E-mail: rad@roe.ac.uk

This gas not only provides the reservoir for new star formation (SF), but also contains strong signatures of feedback processes from both SF and active galactic nuclei (AGNs). Hence, it is important to situate ISM gas in and around galaxies within the context of hierarchical galaxy formation.

Cosmological gas dynamical simulations provide a comprehensive approach towards elucidating the connection between ISM gas, SF, and feedback. Modern simulations now include sophisticated models for SF and feedback processes, and generally do a good job of reproducing the observed evolution of the stellar component, including suppressing low-mass galaxy growth via stellar feedback and producing massive quenched galaxies via AGN feedback (Somerville & Davé 2015). Simulations such as EAGLE (Schaye et al. 2015), IllustrisTNG (Pillepich et al. 2018a), and SIMBA (Davé et al. 2019) all produce stellar mass functions in reasonable agreement with observations over cosmic time, showing a shallow faint-end slope and a truncation at high masses coincident with the onset of a quiescent galaxy population. Despite this concordance, the detailed physical models for subgrid processes such as SF and feedback implemented in each model are markedly different. Hence, discrimination between such simulations requires comparing to data beyond stellar masses and stellar growth rates.

Emerging observations of cold ISM gas provide a new regime for testing galaxy formation models. Early simulations demonstrated a clear connection between SF-driven feedback and the cold gas contents of galaxies. For instance, Davé, Finlator & Oppenheimer (2011) showed, perhaps counter-intuitively, that increasing the strength of galactic outflows results in an increased gas fraction at a fixed stellar mass in galaxies; while the cold gas mass at a given halo mass is lower, the stellar mass is reduced even further (Crain et al. 2017). The MUFASA simulation used an improved model for SF feedback and obtained good agreement with observations available at the time (Davé et al. 2017). EAGLE (Crain et al. 2015; Schaye et al. 2015) has been successful in reproducing a wide range of galaxy properties, and has been used to investigate the origin of H I in galaxies (Crain et al. 2017). IllustrisTNG has been shown to broadly capture many observed cold gas statistics of galaxies, including trends with galaxy environment, but some curious features in the gas content of massive galaxies that are likely related to AGN feedback remain in tension with observations (Diemer et al. 2018; Stevens et al. 2019a). Despite the increased uncertainty with modelling the observational characteristics of gas in simulated galaxies, and the limited resolution that precludes direct modelling of many detailed ISM processes in cosmological volumes, these results highlight how cold gas observations could potentially provide a valuable test bed for modern galaxy formation models.

Observations of cold gas components within galaxies have also matured in recent years. For atomic hydrogen, large blind H I surveys such as the H I Parkes All-Sky Survey (Barnes et al. 2001; Meyer et al. 2004; Wong et al. 2006) and the Arecibo Legacy Fast ALFA survey (ALFALFA; Giovanelli et al. 2005; Haynes et al. 2018) characterized the properties of H I-selected galaxies over a wide area, but since these galaxies were selected by their H I mass, this meant that these surveys tended to preferentially detect H I-rich systems (Catinella et al. 2012). In order to connect to models, it is more optimal to have a survey that selects on a quantity that is more robustly predicted in models. Ideally, this is stellar mass, since it is the observation that models are most commonly tuned to reproduce, enabling an equal-footing comparison between simulation predictions.

The GALEX Arecibo SDSS Survey (GASS; Catinella et al. 2010) was pioneering in that it measured H I contents for a stellar

mass-selected sample from the Sloan Digital Sky Survey (SDSS) with $M_* \gtrsim 10^{10} M_\odot$, using Arecibo. GASS was able to statistically quantify or place limits on H I-poor galaxies, which become increasingly commonplace towards higher masses. To expand the dynamic range, the GASS-Low survey was done to extend the completeness down to $M_* \gtrsim 10^9 M_\odot$. The aggregate survey, known as extended GASS (xGASS), thus provides H I contents and upper limits for a representative sample of nearly 1200 galaxies with $10^9 < M_* < 10^{11.5} M_\odot$ (Catinella et al. 2018).

Molecular gas measurements have likewise made significant progress in recent years, typically via CO surveys. As a complement to xGASS, the xCOLD GASS survey (Saintonge et al. 2011, 2017) provided CO(1–0) and CO(2–1) measurements for over 500 of the xGASS-observed galaxies. As with xGASS, the careful selection from SDSS enabled reconstruction of a volume-limited sample. xGASS and xCOLD GASS thus provide a benchmark constraint for modern cosmological galaxy formation models, with a well-specified selection function that allows cleaner model–data comparisons, and comprehensive ancillary data that enable tests of the relationship between the cold gas and other components of galaxies.

Moving to higher redshifts, H I surveys are currently quite limited (Catinella & Cortese 2015; Fernández et al. 2016), though, for example, the recently begun LADUMA survey on MeerKAT aims to measure H I directly out to $z \gtrsim 1$ (Blyth et al. 2016). CO surveys at higher redshifts meanwhile have progressed substantially. The pioneering Plateau de Bure High- z Blue Sequence Survey was able to study the molecular content of a well-studied sample including internal kinematics (Tacconi et al. 2013) to $z \sim 2$, but this was not designed to sample a representative volume. The recent CO Luminosity Density at High- z (COLDz) survey using the Jansky Very Large Array (Pavesi et al. 2018) and the Atacama Large Millimetre Array (ALMA) Spectroscopic Survey (ASPECS) in the Hubble Ultra Deep Field (Aravena et al. 2019) have provided a more statistical characterization of the molecular gas contents of galaxies out to $z \sim 2$, albeit with limited samples. Given that such high-redshift cold gas observations are set to improve dramatically in the next few years, the time is ripe to provide a snapshot view of how modern galaxy formation simulations that are successful in reproducing stellar properties fare against available cold gas observations.

This paper compares the predictions from three recent cosmological hydrodynamic simulations, namely SIMBA, EAGLE, and IllustrisTNG, to ALFALFA, xGASS, and xCOLD GASS observations at low redshifts, and the ASPECS and COLDz data at high redshifts. The primary purpose is to assess the range of predictions among state-of-the-art hydrodynamic models in galaxy cold gas contents, and to provide preliminary comparisons to observations. A proper comparison would involve mimicking details of observational selection effects for each survey, which we leave for future work; here we take extant simulation predictions for EAGLE and IllustrisTNG, add SIMBA predictions, and compare to data at face value. We examine H I and H₂ contents that are more directly predicted in these simulations, along with the CO(1–0) luminosity determined via a conversion factor following the recipe based on merger simulations and CO radiative transfer by Narayanan et al. (2012). We find substantial variations among current models in their cold gas predictions, with all models qualitatively reproducing the broad trends but no model quantitatively reproducing all the observations.

This paper is organized as follows. In Section 2, we briefly recap the simulations and key observations used in this work. In Section 3, we first compare stellar properties to show that they are similar among our simulations, and then present our comparisons from $z = 0$ to 2 for the H I and H₂ mass functions (HIMFs and H2MFs), the

H₂-to-CO conversion factor, the CO(1–0) luminosity functions, and the gas content scaling relations at $z = 0$, all compared to a range of relevant observations focused on xGASS and xCOLD GASS at $z \sim 0$ along with other recent surveys from $z \sim 0$ to 2. We further examine different variants of AGN feedback models within SIMBA to better understand the physics driving the evolution of cold gas mass functions. Finally, in Section 4 we summarize our results.

2 SIMULATIONS AND OBSERVATIONS

In this work, we employ the SIMBA, EAGLE, and IllustrisTNG simulations, and will compare these to the xGASS and xCOLD GASS data sets. Here we briefly review these simulations and observations.

Beyond briefly describing each simulation’s input physics, we will focus on each one’s procedure for partitioning gas into ionized, atomic, and molecular phases. As cosmological simulations, each of these models has a spatial resolution of ~ 1 kpc, which is insufficient to directly model physical processes giving rise to the cold ISM phase. Even the onset to self-shielding is not done self-consistently, since this would require radiative transfer, which has a prohibitive computational cost. Instead, each model employs a set of established but approximate subgrid prescriptions in order to determine the atomic and molecular fractions of dense gas. The prescription for self-shielding to form neutral gas is essentially the same among these models, while that for forming molecular gas is not identical but still broadly similar.

2.1 SIMBA

SIMBA (Davé et al. 2019) is the successor to the MUFASA simulation (Davé, Thompson & Hopkins 2016), which was run using a modified version of the gravity plus hydrodynamics solver GIZMO (Hopkins 2015) in its meshless finite mass mode. The simulation evolves a representative $100 h^{-1}$ Mpc comoving volume from $z = 249 \rightarrow 0$ with 1024^3 gas elements and 1024^3 dark matter particles. The mass resolution is $9.6 \times 10^7 M_\odot$ for dark matter particles and $1.82 \times 10^7 M_\odot$ for gas elements, and the minimum adaptive gravitational softening length is $\epsilon_{\min} = 0.5 h^{-1}$ ckpc. Initial conditions are generated using MUSIC (Hahn & Abel 2011) assuming the following cosmology (Planck Collaboration XIII 2016): $\Omega_M = 0.3$, $\Omega_\Lambda = 0.7$, $\Omega_b = 0.048$, $H_0 = 68 \text{ km s}^{-1} \text{ Mpc h}^{-1}$, $\sigma_8 = 0.82$, and $n_s = 0.97$.

SF is modelled using an H₂-based Schmidt (1959) relation, where the H₂ fraction is computed using the subgrid prescription of Krumholz & Gnedin (2011) based on the local metallicity and gas column density, modified to account for variations in resolution (Davé et al. 2016). This H₂ fraction will be directly used in the H₂ results for this paper, so bears further description. For each gas element, the H₂ fraction is computed as

$$f_{\text{H}_2} = 1 - 0.75 \frac{s}{1 + 0.25s}, \quad (1)$$

where

$$s = \frac{\ln(1 + 0.6\chi + 0.01\chi^2)}{0.0396Z(\Sigma/M_\odot \text{ pc}^{-2})}, \quad (2)$$

where Z is the metallicity in solar units, χ is a function of metallicity (see Krumholz & Gnedin 2011), and $\Sigma = \rho^2/|\nabla\rho|$ is the column density calculated using the Sobolev approximation, increased by $\times 30^{2/3}$ to account for subresolution clumping (see Davé et al. 2016, for full discussion). We impose a minimum metallicity of $10^{-3} Z_\odot$ solely for the purposes of this subgrid model.




Figure 1. Phase diagram of hydrogen number density n_{H} versus temperature T for dense, cool gas in SIMBA at $z = 0$. For clarity, a random sample of 0.05 per cent of gas elements is plotted. Points are colour coded by H I fraction in blue/green, and H₂ fraction in red/yellow for the gas elements where $f_{\text{H}_2} > 0$. Abrupt transitions are seen from nearly fully ionized to nearly fully neutral at $n_{\text{H}} \sim 10^{-3} \text{ cm}^{-3}$, while the transition from atomic to molecular is similarly abrupt but depends on metallicity.

The star formation rate (SFR) is calculated from the density ρ and the dynamical time t_{dyn} via $\text{SFR} = \epsilon_* f_{\text{H}_2} \rho / t_{\text{dyn}}$, where $\epsilon_* = 0.02$ (Kennicutt 1998). SF is only allowed to occur in gas with $n_{\text{H}} > 0.13 \text{ H atoms cm}^{-3}$, although the limiting factor is $f_{\text{H}_2} > 0$ for all but supersolar metallicities. SIMBA artificially pressurizes gas above this density by imposing $T = 10^4 (n_{\text{H}}/0.13)^{4/3} \text{ K}$ (Schaye & Dalla Vecchia 2008), in order to prevent numerical fragmentation owing to the Jeans mass being unresolved.

The H I fraction of gas elements is computed self-consistently within the code, accounting for self-shielding on the fly based on the prescription in Rahmati et al. (2013), where the metagalactic ionizing flux strength is attenuated depending on the gas density assuming a spatially uniform ionizing background from Haardt & Madau (2012). This gives the total neutral gas, and subtracting off the H₂ yields the H I. Hence in SIMBA, the H I and H₂ fractions for gas are computed self-consistently, on the fly during the simulation run.

Fig. 1 illustrates the resulting cold gas fractions in phase space in SIMBA at $z = 0$. A random subsample of 0.05 per cent of all particles is plotted in (n_{H}, T) space, with green→blue colours showing the H I fractions, and yellow→red colours indicating H₂ fractions in gas where $f_{\text{H}_2} > 0$. This plot focuses on the cool, dense phase of cosmic gas, since other cosmic phases have tiny neutral fractions; see Christiansen et al. (2019) for a complete phase diagram from SIMBA. For EAGLE, the analogous diagram is shown in fig. 1 of Crain et al. (2017).

There is an abrupt transition at $n_{\text{H}} \approx 10^{-3} \text{ cm}^{-3}$ above which self-shielding kicks in, and the gas goes from highly ionized to mostly neutral. Metal cooling can cool highly enriched gas to fairly low temperatures, but above $n_{\text{H}} > 0.13 \text{ cm}^{-3}$, the ISM pressurization becomes evident as the gas is forced to have a minimum Jeans mass-resolving temperature of $T \propto \rho^{1/3}$ (Davé et al. 2016). Ultimately, it is this ISM pressurization forced by the poor resolution (Schaye & Dalla Vecchia 2008) that prevents direct modelling of the cold molecular ISM phase in SIMBA, as well as in the other cosmological simulations. In this pressurized region, the gas fairly abruptly transitions to molecular, but the density at which this occurs is metal dependent; hence, there is still substantial low-metallicity atomic gas above the threshold (blue points, many hidden underneath the

red points). This illustrates how self-shielding and the Krumholz & Gnedin (2011) prescription work together to transform ionized IGM gas into atomic and molecular phases.

Radiative cooling and photoionization heating are implemented using the GRACKLE-3.1 library (Smith et al. 2017). The chemical enrichment model tracks nine metals during the simulation, tracking enrichment from type II supernovae (SNe), type Ia SNe, and asymptotic giant branch (AGB) stars, including locking some of the metals into dust. A (Chabrier 2003) stellar initial mass function (IMF) is assumed in order to compute stellar evolution. SIMBA includes SF-driven galactic winds as kinetic two-phase, metal-enriched winds with 30 per cent of the wind particles ejected hot and with a mass loading factor that scales with stellar mass, based on the FIRE (Hopkins et al. 2014) zoom simulation scalings from Anglés-Alcázar et al. (2017b). Importantly, hydrodynamics are turned off in the winds (‘decoupled’) until they are well outside the ISM (Springel & Hernquist 2003); hence, they explicitly avoid depositing energy in the ISM on their way out. This is done because in any current hydrodynamic solver, a single fluid element moving at high velocities through an ambient medium is more accurately represented by turning off hydrodynamics rather than allowing the solver to calculate the interactions.

SIMBA’s main improvement on MUFASA is the addition of black hole growth via torque-limited accretion, and AGN feedback via stable bipolar kinetic outflows. For ($T < 10^5$ K) gas, black hole accretion follows the torque-limited accretion model of Anglés-Alcázar et al. (2017a) that is based on the analytic model of Hopkins & Quataert (2011), while for hot gas accretion (Bondi 1952) accretion is employed. AGN feedback in SIMBA is designed to mimic the observed dichotomy in black hole growth modes seen in real AGNs (e.g. Heckman & Best 2014): a ‘radiative’ mode at high Eddington ratios (f_{Edd}) characterized by mass-loaded radiatively driven winds ejected at $\sim 10^3$ km s $^{-1}$, and a ‘jet’ mode at low $f_{\text{Edd}} \lesssim$ few per cent at $\sim 10^4$ km s $^{-1}$. The mass loading is set such that the outflow momentum is $20L/c$, where $L = 0.1\dot{M}c^2$ is the radiative luminosity for a black hole accretion rate of $\dot{M}$. Additionally, we include X-ray heating by black holes based on the model of Choi et al. (2012). This yields a quenched galaxy population (Rodríguez Montero et al. 2019) and galaxy–black hole co-evolution (Thomas et al. 2019) in good agreement with observations.

In addition to the full SIMBA feedback model, we will consider $50 h^{-1}$ Mpc, 2×512^3 variants where we turn off the X-ray feedback (‘No-X’), turn off both X-ray and jet feedback (‘No-jet’), and turn off all AGN feedback (‘No-AGN’). These runs have the same resolution as the $100 h^{-1}$ Mpc run but with 1/8 the volume. All other input physics, as well as the initial conditions, are identical. These models will be described further in Section 3.7. Finally, to assess resolution convergence, we will employ a $25 h^{-1}$ Mpc, 2×512^3 SIMBA run, with identical input physics to the full SIMBA run. This has $8\times$ better mass resolution than the other SIMBA runs. Feedback parameters have not been re-tuned at this higher resolution.

Galaxies are identified using a 6D friends-of-friends (FOF) galaxy finder, using a spatial linking length of 0.0056 times the mean inter-particle spacing (equivalent to twice the minimum softening kernel), and a velocity linking length set to the local velocity dispersion. This is applied to all stars and ISM gas ($n_{\text{H}} > 0.13$ cm $^{-2}$). Haloes are identified using a 3D FOF with a linking length parameter of 0.2. The H I and H $_2$ fractions for individual gas elements are taken directly from the simulation, without any post-processing. To compute galaxies’ H I and H $_2$ contents, we assign each gas particle in a halo to the galaxy that has the highest value of M_{baryon}/R^2 , where M_{baryon} is the total baryonic mass of the galaxy and R is the distance

from the particle to the galaxy’s centre of mass. This enables cold gas, particularly H I, to be assigned to a galaxy even if it is not identified as within the galaxy’s ISM, since the H I can be significant even for gas with $n_{\text{H}} < 0.13$ cm $^{-2}$. For H $_2$, the results are insensitive to whether we consider this low-density gas, since all the molecular gas is located within the ISM (Fig. 1).

For consistency among the models, we will restrict our simulated galaxy samples to $M_* > 10^9 M_{\odot}$, which represents the approximate mass limit for the primary observational sample to which we will compare, even though all our simulations are able to resolve to lower stellar masses.

2.2 EAGLE

The EAGLE simulations (Schaye et al. 2015; Crain et al. 2015, with public data release described by McAlpine et al. 2016) were evolved with a substantially modified version of the N -body Tree-Particle-Mesh (TreePM) smoothed particle hydrodynamics (SPH) solver GADGET3 (last described by Springel 2005). The modifications include significant updates to the hydrodynamics solver, the time-stepping criteria, and the implemented subgrid physics modules. The largest volume EAGLE simulation (Ref-L100N1504 in the nomenclature of Schaye et al. 2015) evolves a region $67.77 h^{-1}$ Mpc (100 Mpc) to the present day, realized with 1504^3 dark matter particles and an (initially) equal number of baryonic particles, yielding particle masses of 1.81×10^6 and $9.70 \times 10^6 M_{\odot}$ for baryons and dark matter, respectively. The Plummer-equivalent gravitational softening length is $\epsilon_{\text{com}} = 2.66$ ckpc, limited to a maximum proper length of $\epsilon_{\text{prop}} = 0.7$ pkpc. We will also show results from the EAGLE-Recal simulation (Recal-L25N752 in Schaye et al. 2015), which evolved a $16.94 h^{-1}$ Mpc volume with $8\times$ higher mass resolution, and $2\times$ higher spatial resolution. Initial conditions were generated with the software described by Jenkins (2010), as detailed in appendix B of Schaye et al. (2015), assuming the Planck Collaboration XVI (2014) cosmogony: $\Omega_{\text{m}} = 0.307$, $\Omega_{\Lambda} = 0.693$, $\Omega_{\text{b}} = 0.04825$, $H_0 = 67.77$ km s $^{-1}$ Mpc $^{-1}$, $\sigma_8 = 0.8288$, and $n_s = 0.9611$.

Interstellar gas is treated as a single-phase star-forming fluid with a polytropic pressure floor (Schaye & Dalla Vecchia 2008), subject to a metallicity-dependent density threshold for SF (Schaye 2004), which reproduces (by construction) the observed Kennicutt–Schmidt relation (Kennicutt 1998) in gas that satisfies local vertical hydrostatic equilibrium. Radiative heating and cooling are implemented element-by-element for 11 species (Wiersma, Schaye & Smith 2009a) in the presence of a time-varying UV/X-ray background radiation field (Haardt & Madau 2001) and the cosmic microwave background. The evolution of the same species due to stellar evolution and mass-loss are tracked during the simulation according to the implementation of Wiersma et al. (2009b). The seeding of BHs and their growth via gas accretion and BH–BH mergers are treated with an updated version of the method introduced by Springel, Di Matteo & Hernquist (2005), accounting for the dynamics of gas close to the BH (Rosas-Guevara et al. 2015). Feedback associated with the formation of stars (Dalla Vecchia & Schaye 2012) and the growth of BHs (Booth & Schaye 2009) are both implemented via stochastic, isotropic heating of gas particles ($\Delta T_{\text{SF}} = 10^{7.5}$ K, $\Delta T_{\text{AGN}} = 10^{8.5}$ K), designed to prevent immediate, numerical radiative losses. Heated particles are not decoupled from the hydrodynamics scheme. The simulations assume the Chabrier (2003) IMF.

Haloes are identified by applying the FoF algorithm to the dark matter particle distribution, with a linking length of 0.2 times the mean inter-particle separation. Gas, stars, and BHs are associated with the FoF group, if any, of their nearest dark matter particle.

Galaxies are equated to bound substructures within haloes, identified by the application of the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009) to the particles (of all types) comprising FoF haloes. Following Schaye et al. (2015), we compute the properties of galaxies by aggregating the properties of the relevant particles residing within 30 kpc of their most-bound particle.

The EAGLE model does not partition hydrogen into its ionized (H II), atomic (H I), and molecular (H₂) forms ‘on-the-fly’, so we estimate the fraction of each on a particle-by-particle basis in post-processing, with the two-step approximation used by Crain et al. (2017), elements of which were also used by Lagos et al. (2015, 2016; Bahé et al. 2016; Marasco et al. 2016). Hydrogen is first partitioned into atomic (H I + H₂) and ionized (H II) components, using the fitting function of Rahmati et al. (2013), which considers both collisional ionization (using temperature-dependent rates collated by Cen 1992) and photoionization by metagalactic UV radiation, calibrated using TRAPHIC radiative transfer simulations (Pawlik & Schaye 2008). Radiation due to sources within or local to galaxies, although likely significant (see e.g. Miralda-Escudé 2005) is not considered explicitly, but is accounted for implicitly via the use of an empirical or theoretical scheme to partition the atomic hydrogen of star-forming gas particles into its atomic and molecular components. Crain et al. (2017) present results using two such schemes, to illustrate the associated systematic uncertainty. The first is the theoretically motivated prescription of Gnedin & Kravtsov (2011), while the second is motivated by the observed scaling of the molecular-to-atomic hydrogen surface density ratio ($R_{\text{mol}} \equiv \Sigma_{\text{H}_2}/\Sigma_{\text{H I}}$) with the mid-plane pressure of galaxy discs (Blitz & Rosolowsky 2006). In the latter case, gas particles are assigned a molecular hydrogen fraction that scales as a function of their pressure (this approach has been widely used elsewhere; see e.g. Popping et al. 2009; Duffy et al. 2012; Davé et al. 2013; Rahmati et al. 2013).

For the EAGLE results in this paper, we will employ the results using the Gnedin & Kravtsov (2011) approach. This is primarily because it is significantly closer to the approaches used in SIMBA and IllustrisTNG. Crain et al. (2017) show that this scheme generally yields higher molecular fractions than the pressure-based approach, and that the partitioning of neutral hydrogen into its atomic and molecular components is a more severe uncertainty on the H I masses of galaxies than, for example, the freedom afforded by present constraints on the amplitude of the present-day UV background. However, the primary systematic influence on atomic hydrogen masses was found to be the resolution of the simulation, with galaxies of fixed mass exhibiting significantly greater ISM masses when simulated at $8\times$ (2 $\times$) greater mass (spatial) resolution, as will be reiterated in the results of this paper.

2.3 IllustrisTNG

IllustrisTNG comprises a suite of cosmological magnetohydrodynamic simulations at various resolutions and volumes, all run with the moving-mesh AREPO code (Springel 2010), assuming a Λ CDM cosmology with parameters based on Planck Collaboration XIII (2016), using a common galaxy formation model (Weinberger et al. 2017; Pillepich et al. 2018a), which was developed from the original Illustris model/simulation (Genel et al. 2014; Vogelsberger et al. 2014). The model includes gas cooling, both primordial and from metal lines, where nine chemical elements are tracked: SF and stellar evolution, which includes kinetic-wind feedback from type-II SNe, and metal enrichment from SNe (types Ia and II) and AGB stars (Pillepich et al. 2018a); massive black hole growth, carrying two modes of feedback: thermal injection for high accretion rates, and

kinetic winds at low accretion rates (Weinberger et al. 2017); and an idealized consideration of magnetic fields (Pakmor & Springel 2013). The main observational constraints used to calibrate the model included the $z = 0$ galaxy stellar mass function (GSMF), the $z = 0$ stellar-to-halo mass relation, and cosmic SFR density history (for details, see Pillepich et al. 2018a).

In this paper, we use the TNG100 simulation, first presented in a series of five papers (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018b; Springel et al. 2018), which has been made publicly available (Nelson et al. 2019). The periodic box length of TNG100 is $75 h^{-1}$ Mpc ≈ 110 comoving Mpc, within which 1820^3 dark matter particles and 1820^3 initial gas cells are evolved. This gives a mass resolution of $7.5 \times 10^6 M_{\odot}$ for dark matter and $\sim 1.4 \times 10^6 M_{\odot}$ for baryons. The smallest resolvable length for gas in galaxy centres is 190 pc. The simulation has been post-processed to decompose atomic gas into its atomic and molecular phases; we follow the methodology of Stevens et al. (2019a), which is based largely on Diemer et al. (2018).

TNG100 galaxy properties in this paper follow the ‘inherent’ definition of Stevens et al. (2019a). This means isolating particles/cells bound to (sub)structures according to SUBFIND (Springel et al. 2001; Dolag et al. 2009) and further making an aperture cut to remove the near-isothermal ‘intrahalo’ component (Stevens et al. 2014).¹ All H I/H₂ properties assume the Gnedin & Draine (2014) prescription, as described in Stevens et al. (2019a). The primary factors in this prescription are the gas density, local UV flux, and local dust content. Critically, it is assumed that UV is generated from sites of SF, with 90 per cent of that UV absorbed in the star-forming cell; the other 10 per cent is propagated through the halo, treating it as a transparent medium (Diemer et al. 2018). A lower limit on the UV flux in any cell is set to the assumed cosmic UV background (Faucher-Giguère et al. 2009). The dust fraction of each cell is assumed directly proportional to its metallicity, in line with Lagos et al. (2015). A comparison of the performance of this prescription with several others has been presented alongside many H I- and H₂-related properties of TNG100 galaxies in a series of recent papers (Diemer et al. 2018, 2019; Stevens et al. 2019a, b).

2.4 Comparison of models

For clarity, we briefly summarize the key differences between our cosmological simulations, in terms of measuring the cold gas properties. All employ the Rahmati et al. (2013) prescription to determine the self-shielded gas, and all employ some variant of the Krumholz & Gnedin (2011) model to compute the H₂ fraction. SIMBA does these on the fly, while the others apply these in post-processing, but since their impact on the dynamics should be relatively minimal, there should not be large variations because of this.

Assigning gas into galaxies has more variations between models. In SIMBA, all H I and H₂ in haloes (which is essentially all cold gas overall) are assigned to the most dynamically important galaxy (i.e. maximum in M_{gal}/r^2 , where r is the distance to the galaxy’s centre of mass). While we identify the galaxies using a 6DFOF and use this to compute properties such as the SFR and M_* , only the locations are used for the cold gas contents, with all cold gas being assigned via proximity. In the case of H₂, this is essentially

¹We note a previously unreported error in the application of this aperture in the results of Stevens et al. (2019a), which meant that the radius always hits its upper limit of R_{200c} . We have corrected for this in this work. The effect is negligible for our results and that of Stevens et al. (2019a).

equivalent to just assigning the H_2 to its own galaxy. The H I, however, can extend well beyond the star-forming region of a galaxy, so this procedure can yield large differences in H I content. The choice for SIMBA is motivated by our comparisons to Arecibo data (see the next section) that have arcminute-scale resolution, and thus is unlikely to be resolving individual galaxies' ISM.

EAGLE, also motivated by matching Arecibo's beam (Bahé et al. 2016), uses a fixed 70 kpc aperture to compute both H I and H_2 contents, within which only gas that is gravitationally bound to a given galaxy is attached to it. IllustrisTNG, meanwhile, uses the variable 'BaryMP' aperture of Stevens et al. (2014). For a fixed stellar mass of $10^9 M_\odot$ at $z = 0$, the aperture radius comes out as 20 ± 10 kpc (rounded 1σ range). For $M_* = 10^{11} M_\odot$, this rises to 70 ± 30 kpc. For the same respective stellar masses at $z = 2$, aperture radii are 15 ± 5 and 40 ± 10 physical kpc. A similar exercise at a fixed $M_{H I}$ of 10^8 and $10^{10} M_\odot$ yields aperture radii of 16 ± 4 and 50 ± 9 kpc, respectively, at $z = 0$. To first order, the aperture scales with the virial radius of each system, where typically $R_{\text{BaryMP}}/R_{\text{vir}} \simeq 0.25 \pm 0.10$. This means that in EAGLE, the largest galaxies may be missing gas that is at large radii compared to what is in TNG. This is not expected to be significant for H_2 where the gas is mostly confined to the dense ISM, but may be important in the case of H I that can extend quite far out. On the flip side, in small galaxies EAGLE may include H I that TNG would not include, but in such systems the H I is not expected to extend very far out. SIMBA does not use aperture masses, but in principle can assign H I from anywhere in the halo, so it is not straightforward to directly compare to aperture masses, but in general SIMBA includes H I from at least as far out as TNG. We emphasize, however, that in all simulations most of the H I is located relatively close to galaxies.

We can compare these radii to the observed H I size–mass relation, $R_1 = \sqrt{M_{H I}/12.88}$, where R_1 in pc is the radius at which the surface density drops to $1 M_\odot \text{ pc}^{-2}$ and $M_{H I}$ is in solar masses (Stevens et al. 2019b). This gives, for $\log M_{H I} = [8, 9, 10, 11]$, $R_1 = [2.8, 8.8, 28, 88]$ kpc. Thus, generally the TNG aperture extends out to at least this radius, while EAGLE's aperture is not expected to miss significant H I except for $M_{H I} \gtrsim 10^{10} M_\odot$. Bahé et al. (2016) showed that EAGLE's choice of 30 kpc can bias the resulting $M_{H I}$ by up to ~ 20 per cent, which is not negligible but does not impact the conclusions of this paper.

Given that the H I distribution is fairly extended and thus more subject to choices regarding aperture or assignment, we thus emphasize that H I comparisons between models and to ALFALFA data should be regarded as preliminary; this could either overestimate or underestimate the H I content depending on the assumed aperture relative to the observed beam. None the less, the differences between models are significantly larger than the expected systematics associated with apertures. Upcoming observations with sensitive radio interferometers such as MeerKAT and ASKAP will avoid blending issues (Elson, Baker & Blyth 2019), but may require even more care when considering H I apertures in simulation comparisons; we leave a fuller investigation of this for future work.

For H_2 , we have a somewhat different issue, because the beam size for our main comparison sample xCOLD GASS, based on the IRAM-30m data, is a factor of ~ 5 smaller than Arecibo, with typical apertures of ~ 10 – 20 kpc. Aperture corrections are applied to these observations, as detailed in section 2.4 of Saintonge et al. (2017) as well as in Saintonge et al. (2012); these indicate that the observations appear to capture the majority of H_2 in these systems. However, if one applies similarly small apertures to the simulations, one can get non-trivially different results, particularly for massive gas-rich galaxies that are quite large in extent. The underlying problem is

probably that all the simulations assume that H_2 is allowed to form in gas with $n_H \gtrsim 0.1$ – 0.2 cm^{-3} , whereas in reality H_2 actually forms only at much higher densities ($n_H \gtrsim 100 \text{ cm}^{-3}$). The choice in these simulations is driven by the numerical resolution, which does not enable the internal structure of the ISM to be fully represented. The result is that H_2 can form in more extended gas in simulations than in real data. Unfortunately, this is an intrinsic limitation even in these state-of-the-art models, and suggests that better subresolution ISM models are required to properly represent the distribution of molecular gas in galaxies.

Ultimately, the models and observations are both including the vast majority of H_2 in the galaxies. However, the spatial distribution in the case of simulations is more extended owing to resolution limitations. Diemer et al. (2019) demonstrated this explicitly, showing that the molecular gas distribution in IllustrisTNG is more extended than observed. Hence, if we were to create mock observations with the same aperture restrictions, we may obtain substantially smaller H_2 contents for the most massive gas-rich galaxies. However, given the limitations in way that H_2 is modelled in the ISM, this would not be a meaningful way to conduct such a comparison. We thus will conduct our comparisons with the apertures as described above, with the caveat that we are comparing only the *global* molecular gas contents, without creating mock CO images tailored to individual surveys.

In truth, none of the decisions surrounding apertures in this paper are necessarily faithful representations of the scale on which gas (or stellar) properties of galaxies are measured observationally. For H I in IllustrisTNG galaxies, this has been explored by Stevens et al. (2019a) and Diemer et al. (2019). In the former, alongside the inherent IllustrisTNG properties, mock-observed properties that were catered to specific H I surveys were presented, namely for ALFALFA and xGASS. Mock and inherent H I measurements differed most for satellite galaxies, primarily due to the much higher probability of 'confusion' (where H I in the mock beam comes from sources other than the galaxy of interest). A similar mocking procedure for H_2 will be presented for TNG100 in Stevens et al. (in preparation). Finally, we note that aperture choices can also be important for stellar mass measures, particularly for massive galaxies (e.g. Bernardi et al. 2017), so these issues are not limited to gas measures.

2.5 xGASS and xCOLD GASS

The twin surveys xGASS and xCOLD GASS were specifically designed to provide a robust and complete census of H I and CO across the $z \sim 0$ Universe for purposes such as comparing and constraining large-scale simulations. Unlike most previous gas surveys, they are neither flux limited nor the result of complex selection functions; the galaxies are selected from SDSS based on stellar mass and redshift alone, and are therefore representative of the local galaxy population.

The H I component of this programme, xGASS, provides 21-cm measurements for 1179 galaxies with $10^9 < M_*/M_\odot < 10^{11.5}$ and $0.01 < z < 0.05$ (Catinella et al. 2010, 2013, 2018). The sample is extracted from the overlap of the SDSS spectroscopic and GALEX imaging surveys, and builds on the ALFALFA H I-blind survey (Giovanelli et al. 2005), which provides the measurements for the most gas-rich systems in the sample.

In addition to the sample selection, the observing strategy employed by xGASS is key to making it fit for purpose here. The observations were designed to provide uniform sensitivity in terms of atomic gas mass fraction, $f_{H I} = M_{H I}/M_*$, which allows us to compare gas contents with global galaxy properties over a wide range of stellar masses. Apart from the galaxies already detected in H I by ALFALFA, each xGASS galaxy was observed with Arecibo until

either a detection of the H I line was obtained or until sensitivity to a gas mass fraction of 2–10 per cent (depending on stellar mass) was reached. For H I non-detections, 5σ upper limits have been computed.

The xCOLD GASS survey targeted 532 galaxies from xGASS with the IRAM-30m telescope to measure their total molecular gas masses via CO(1–0) line emission (Saintonge et al. 2011, 2017). As for the H I, the strategy was to observe each galaxy until either a detection of the CO(1–0) line was obtained or until sensitivity to $f_{\text{H}_2} = M_{\text{H}_2}/M_* = 2$ per cent was reached, allowing stringent upper limits to be placed on the most gas-poor galaxies. The CO line luminosities are corrected for small aperture effects (Saintonge et al. 2012), and then converted into molecular gas masses using the CO-to-H₂ conversion factor α_{CO} of Accurso et al. (2017).

Accurso et al. (2017) characterized α_{CO} using [C II] data from *Herschel* and CO(1–0) data from xCOLD GASS, guided by photodissociation region modelling. They found that α_{CO} depends quite strongly on the galaxy’s metallicity, and weakly on its deviation from the star-forming main sequence:

$$\log \alpha_{\text{CO}}(\pm 0.165) = 14.752 - 1.623(12 + \log(\text{O}/\text{H})) + 0.062 \log \Delta_{\text{MS}}, \quad (3)$$

where $12 + \log(\text{O}/\text{H})$ is the oxygen abundance in conventional notation, and Δ_{MS} is the ratio of the SFR in the galaxy to that in a galaxy on the main sequence at the same M_* . Using $12 + \log(\text{O}/\text{H})_{\odot} = 8.69$ (Asplund et al. 2009), we can rewrite this as

$$\alpha_{\text{CO}} = \frac{4.47^{+2.06}_{-1.41}}{Z_{\text{H}_2}^{1.623} \Delta_{\text{MS}}^{0.062}} M_{\odot} \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}. \quad (4)$$

We will show a comparison of this method for determining α_{CO} versus the simulation-based method we employ from Narayanan et al. (2012) in Section 3.4.

The H I and CO observations are complemented by additional data products. Stellar masses are retrieved from the MPA-JHU catalogue, and stellar mass surface densities were calculated as $\mu_* = M_*/(2\pi r_{50,z}^2)$, where $r_{50,z}$ is the radius encompassing 50 per cent of the z -band flux, in kpc. SFRs were calculated using a combination of UV and IR photometry and an optical SED fitting method (Janowiecki et al. 2017).

3 RESULTS

3.1 Stellar and SFR properties

To set the stage for our gas comparisons, we first present and compare stellar mass and SFR properties in our three simulations. In particular, we look at the GSMF and the SFR–stellar mass relation (‘main sequence’). These have been presented in other papers (Furlong et al. 2015; Pillepich et al. 2018a; Davé et al. 2019), but here we provide a direct comparison to each other as well as to observations. For the main sequence, we will compare to the galaxies in the xGASS sample in order to ensure there are no clear systematic differences in the sample properties.

Fig. 2, top panel, shows the $z = 0$ GSMF in SIMBA (blue), IllustrisTNG (red), and EAGLE (green). The blue shaded region shows the standard deviation computed over 8 subobjects of the SIMBA volume, as an estimate of cosmic variance; since the other simulations have similar volumes, they likely have comparable variance. The observed $z \approx 0$ GSMF is shown from Wright et al. (2017) as the black data points, from the Galaxy And Mass Assembly survey. We do not show EAGLE-Recal here to avoid clutter, but it




Figure 2. The GSMF (top panel) and specific star formation rate versus stellar mass (bottom panel) at $z = 0$ for SIMBA (blue), IllustrisTNG (red), and EAGLE (green) compared with observations. The blue shading for SIMBA shows the variance over 8 simulation subobjects. For the GSMF, representative observations are shown from Wright et al. (2017). For the bottom panel, the hexbin shows the xGASS sample. Running medians are shown for galaxies with $\text{sSFR} > 10^{-1.8} \text{ Gyr}^{-1}$; the dashed black line shows xGASS, while the red, green, and blue lines show IllustrisTNG, EAGLE, and SIMBA, respectively. Error bars show the 1σ spread around the median. All models show good agreement with these stellar-based measures, with some minor discrepancies.

was specifically recalibrated to match the GSMF, so agrees as well as the main EAGLE simulation.

All simulations provide a good match to the observed $z = 0$ GSMF. In part, this is by construction, as the SF and feedback recipes in each have been tuned at some level to reproduce this key demographic. They agree quite well below M_* , but there are some minor differences at the massive end. EAGLE and SIMBA slightly undercut the knee of the GSMF, while IllustrisTNG matches the knee very well but may overproduce the massive end. This highlights the continued difficulty that simulations have in reproducing the sharpness of the exponential cut-off in the $z = 0$ GSMF (e.g. Davé et al. 2016). Note that the massive end is relatively uncertain owing to aperture effects and potentially IMF variations (e.g. Bernardi et al. 2018). Despite these small variations, in general all these simulations reproduce the observed GSMF quite well, within plausible systematic uncertainties.

Fig. 2, bottom panel, shows specific star formation rate ($\text{sSFR} = \text{SFR}/M_*$) versus stellar mass for SIMBA (blue), IllustrisTNG (red), and EAGLE (green). Here, we show a running median for each simulation for star-forming galaxies defined as having $\text{sSFR} > 10^{-1.8} \text{ Gyr}^{-1}$, following Davé et al. (2019). The grey hexbins show the mass-selected xGASS sample, and the black points and line show a similarly selected running median of the xGASS star-forming galaxies. All have error bars representing the 1σ scatter around the median. We focus on a comparison of the star-forming sample since we are primarily interested in gas-rich galaxies in this work; as an aside, there are larger differences in the quenched galaxy fractions among these various models.

All models produce a mildly declining relation of sSFR versus M_* in non-quenched galaxies that is in good agreement with the xGASS

sample, as well as with each other. There are mild differences in the detailed shape of the curves, such as SIMBA producing high sSFR values at low masses, IllustrisTNG having slightly higher sSFR values around M^* , and EAGLE potentially slightly low at low masses. The hexbinned xGASS sample shows a turn-down in sSFR at high masses, which would also be evident in the simulations' running medians if we were to include $\text{sSFR} < 10^{-1.8} \text{ Gyr}^{-1}$ galaxies. The 1σ scatter is typically in the range of ~ 0.3 dex in all models, which is comparable to that seen in the observations (e.g. Kurczynski et al. 2016; Catinella et al. 2018).

Overall, we confirm that SIMBA, IllustrisTNG, and EAGLE all produce stellar mass functions and star-forming main sequences that are in good agreement with observations, and with each other. This is an important check, which sets the baseline for comparisons among their cold gas properties in relation to their M_* and SFR. It is also a non-trivial success for these models, which has only been achieved in the last few years among hydrodynamic galaxy formation simulations. None the less, we will see that the differences among gas properties in these simulations are significantly larger than those seen in stellar properties.

3.2 HIMF

We now examine properties of the neutral gas, starting with the HIMF and its redshift evolution in our three simulations. The HIMF is straightforwardly measurable from 21-cm emission down to quite low masses, albeit currently only at low redshifts owing to the sensitivity limits of radio telescopes. Hence, the HIMF has relatively few concerns regarding selection effects or other observational systematics, modulo confusion issues given the large beam sizes of single-dish surveys (e.g. Elson et al. 2019). That said, there are some non-trivial modelling systematics regarding assumptions about self-shielding and the separation of atomic and molecular hydrogen, although since the neutral gas content tends to be dominated by atomic gas particularly in small systems, this is less of an uncertainty for H I as for H₂. For these reasons, the HIMF is among the more robust constraints on the cold gas content in simulated galaxies.

Fig. 3 shows the HIMF at $z = 0, 1$, and 2 (top to bottom) for SIMBA (blue), IllustrisTNG (red), EAGLE (green solid), and EAGLE-Recal (green dotted). The blue shaded region shows the estimated cosmic variance computed over 8 subobjects of the SIMBA volume. At $z = 0$, we show the observed HIMF from the ALFALFA survey (Jones et al. 2018) as the dark grey band; we repeat this in the $z = 1$ and 2 panels with lighter shading as a reference point to gauge the amount of evolution in models, but current observations of the HIMF are limited to low redshifts, so comparisons to data should only be done at $z = 0$. Finally, for SIMBA, we will show in Section 3.7 that, owing to its relatively low resolution compared to IllustrisTNG and EAGLE, it suffers from incompleteness at the low- $M_{\text{H I}}$ end. To denote this, we have shown the portion of the HIMF that is potentially compromised by resolution effects using a dashed blue line.

At $z = 0$, all simulations show a fairly flat low-mass slope, and a turnover at high masses; broadly, this is similar to that seen in ALFALFA. To quantify this, we fit a Schechter function to each simulated HIMF for galaxies with $M_{\text{H I}} > 10^9 M_\odot$. Given the limited dynamic range, we fix the low-mass slope to the ALFALFA value of -1.25 ; leaving the slope free gives values consistent with this, but with significantly larger uncertainties on all the parameters. We find that the best-fitting characteristic H I mass $M_{\text{H I}}^*$ varies significantly between models: For IllustrisTNG it is $M_{\text{H I}}^* = 10^{10.26} M_\odot$, for SIMBA it is $M_{\text{H I}}^* = 10^{10.07} M_\odot$, and for EAGLE it is $M_{\text{H I}}^* = 10^{9.67} M_\odot$. For comparison, ALFALFA finds $M_{\text{H I}}^* = 10^{9.94} M_\odot$ (Jones et al.




Figure 3. The HIMF in SIMBA (blue), IllustrisTNG (red), and EAGLE (green), at $z = 0, 1$, and 2 (top to bottom). The HIMF from the ALFALFA survey (Jones et al. 2018) at $z \approx 0$ is shown as the grey shading, and these are reproduced at higher z in lighter shading for reference. IllustrisTNG and SIMBA show much more H I than EAGLE, in better agreement with observations, while EAGLE and SIMBA show an increasing HIMF at earlier epochs while IllustrisTNG shows a dropping HIMF.

2018). This confirms the visual impression that IllustrisTNG while generally matching the HIMF as seen in Diemer et al. (2019) mildly overpredicts the HIMF at the high-mass end; we note that Diemer et al. (2019) supplemented TNG-100 with TNG300 and showed that it provided a better match to the high-mass end. Meanwhile, the main EAGLE volume strongly underpredicts the HIMF, while the higher resolution $25 h^{-1} \text{ Mpc}$ EAGLE-Recal simulation produces a significantly higher HIMF. SIMBA provides a very good match to ALFALFA observations in both shape and amplitude. Note that none of these models have been tuned to reproduce the HIMF.

As noted by Crain et al. (2017) and seen in Fig. 3, EAGLE-Recal produces significantly more H I in galaxies, although its volume is too small to probe the high-mass turnover discrepancy. Hence, in EAGLE the H I content appears to be fairly resolution dependent, which we speculate is likely a consequence of EAGLE's subgrid implementation of feedback (intentionally) not incorporating mechanisms to mitigate against resolution sensitivity (as is the case for SIMBA and IllustrisTNG). As noted by Bahé et al. (2016), the thermal energy injected into the ISM by feedback events in EAGLE scales linearly with the baryon particle mass, and at the standard resolution of EAGLE individual heating events can temporarily create $\sim \text{kpc}$ -scale 'holes' in the cold gas distribution. Assuming that higher resolution simulations produce more robust results, the EAGLE-Recal results suggest that the EAGLE feedback model is capable of well reproducing the HIMF.

The HIMF resolution convergence was generally good in the case of SIMBA's predecessor, MUFASA (Davé et al. 2017). Because the SF feedback is quite similar in SIMBA, we expect that this would

also be the case in SIMBA either, and we will demonstrate this in Section 3.7. Thus, to avoid clutter, we choose not to show different resolution versions of these simulations, though we show EAGLE versus EAGLE-Recal. For IllustrisTNG, Diemer et al. (2019) showed the level of convergence in the cold gas mass functions between TNG100 and TNG300,² and found that they were generally in agreement with each other in their overlapping resolved mass ranges, with TNG300 being slightly lower. This suggests that models that use kinetic decoupled winds may fare somewhat better in resolution convergence than thermal feedback models. Note, however, that all these comparisons are subject to aperture effects that could cause more significant changes than resolution (Section 2.4).

Even stronger differences between the simulations are seen when examining redshift evolution. SIMBA and EAGLE both have strong redshift evolution, in the sense that the HIMF shifts to higher M_{HI} at higher redshifts. This occurs owing to the higher inflow rates at higher redshifts, which in the quasi-self-regulated scenario for galaxy growth results in higher gas contents (see e.g. fig. 5 of Crain et al. 2017). Quantified purely in terms of characteristic mass evolution (fixing the low-mass slope) from $z = 0 \rightarrow 2$, M_{HI}^* increases by $\times 2.5$ and $\times 2$ in SIMBA and EAGLE, respectively. Interestingly, EAGLE-Recal does not show as much evolution as EAGLE, and is consistent with the larger volume at $z = 1$ and 2. Meanwhile, IllustrisTNG shows a *reduction* in M_{HI}^* by $\sim \times 2$ between $z = 0 \rightarrow 2$; it is not immediately evident why IllustrisTNG shows this behaviour. In IllustrisTNG and EAGLE, the low-mass slope becomes steeper at high- z . SIMBA shows a turnover at low masses ($M_{\text{HI}} \lesssim 10^9 M_\odot$), but this owes largely to numerical resolution, as we show in Section 3.7. At higher redshifts, the higher M_{HI}/M_* ratios together with the fixed galaxy mass threshold of $M_* = 5.8 \times 10^8 M_\odot$ combine to result in significant incompleteness at $M_{\text{HI}} \lesssim 10^{9.5} M_\odot$ by $z \sim 2$. Modulo these caveats that mostly impact the low- M_{HI} end, it is clear that measurements of the HIMF even out to $z \sim 1$, as is planned with MeerKAT’s LADUMA survey (Blyth et al. 2016), could provide qualitative discrimination between current galaxy formation models.

3.3 H2MF

Stars form from molecular gas, so the molecular gas content provides a connection to the growth rate of galaxies, particularly in star-forming systems where the merger growth rate is subdominant (e.g. Hirschmann et al. 2014). This has been explored in previous works including Lagos et al. (2015) for EAGLE and Davé et al. (2019) for SIMBA. However, as emphasized in analytic ‘equilibrium’ or ‘bathtub’ models of galaxy evolution (e.g. Finlator & Davé 2008; Bouché et al. 2010; Davé, Finlator & Oppenheimer 2012; Lilly et al. 2013), the molecular gas content does not govern the global stellar mass assembly history, but rather represents an evolving balance between gas supply and gas consumption. For a given gas supply into the ISM, if the star formation efficiency ($\text{SFE} = \text{SFR}/M_{\text{H}_2}$) is high, then the gas reservoir will be low, and vice versa, though the time-averaged number of stars formed will not be altered. Meanwhile, the SF history of a galaxy over cosmological time-scales is set primarily by the net gas supply rate (inflows minus outflows), and is globally independent of SFE for reasonable choices (Katz, Weinberg & Hernquist 1996; Schaye et al. 2015). The molecular gas reservoir thus represents a way to characterize this SFE. In observational




Figure 4. The H2MF in SIMBA (blue), IllustrisTNG (red), and EAGLE (green), at $z = 0, 1$, and 2 (top to bottom). Observations from xCOLD GASS (Fletcher et al. 2020) at $z \approx 0$ are shown as the black points, and reproduced in grey at higher z for reference. IllustrisTNG and SIMBA yield significantly larger molecular masses than EAGLE. SIMBA and EAGLE show an increased H2MF at the high-mass end at $z = 2$, while IllustrisTNG’s is lower.

work, this is often presented as measures of its inverse quantity, the molecular gas depletion time.

In cosmological simulations, the SFE is an input parameter. Typically, it is tuned to approximately reproduce the Kennicutt (1998) relation in star-forming galaxies (e.g. Springel & Hernquist 2003), which for instance has been checked in SIMBA (Appleby et al. 2020). In practice, however, the SFE parameter is applied in SIMBA and IllustrisTNG via a volumetric Schmidt (1959) law, which is connected to the Kennicutt (1998) surface density relation via galactic structure; EAGLE uses a scheme based on the local pressure that results in the Kennicutt (1998) relation by construction in the case of vertical hydrostatic equilibrium. In all cases, the molecular gas content in simulations is thus also sensitive to the distribution of molecular gas within galaxies. Furthermore, as mentioned in Section 2.1, SIMBA’s SF prescription uses a molecular gas-based Schmidt (1959) law, while that in IllustrisTNG and EAGLE uses the total gas, which could result in further differences in the internal structure of star-forming gas. For these reasons, the H2MF provides insights into the differences between current SF prescriptions, particularly among simulations that well reproduce the observed growth histories of the galaxy population.

Fig. 4 shows the H2MF at $z = 0, 1$, and 2 (top to bottom) for SIMBA (blue), IllustrisTNG (red), EAGLE (green solid), and EAGLE-Recal (green dotted), with a blue shading on SIMBA for the estimated cosmic variance as before. Observations at $z \approx 0$ are shown as the black points from the xCOLD GASS survey (Fletcher et al. 2020), and reproduced at other redshifts in grey for reference. At higher

²For a full exploration of resolution convergence in IllustrisTNG, see <http://www.benediktdiemer.com/data/hi-h2-in-illustris/>.

redshifts, we will compare more directly to CO luminosity functions (COLFs) in Section 3.5.

As with the HIMF, the general Schechter function shape is found in all simulations, but there are substantive differences between model predictions. At $z = 0$, SIMBA and IllustrisTNG produce nearly identical H2MFs, while EAGLE's is much lower. The differences are particularly dramatic at the high-mass end. A Schechter fit fixing the observed faint-end slope at -1.33 (Saintonge et al. 2017) gives $M_{\text{H}_2}^* = 10^{9.98}$ for SIMBA, $M_{\text{H}_2}^* = 10^{10.07}$ for IllustrisTNG, and $M_{\text{H}_2}^* = 10^{9.21}$ for EAGLE. Comparing to the observed $M_{\text{H}_2}^* = 10^{9.68}$ from xCOLD GASS highlights the discrepancies between all these models and the observations, with SIMBA and IllustrisTNG overproducing the H2MF at the massive end, while EAGLE underpredicts it. We reiterate here that the comparisons at the massive end are potentially subject to uncertainties regarding aperture effects, given that the simulations' apertures are generally significantly larger than those in the observations, as discussed in Section 2.4.

At the low-mass end, SIMBA shows a dip that owes to limited numerical resolution; we denote the low-mass portion of the H2MF that is subject to resolution effects via the dashed blue line (see Section 3.7). However, it agrees around the knee of the H2MF, while IllustrisTNG overpredicts the H2MF somewhat at all masses. These differences may be subject to significant systematics, not the least of which is the assumed CO-to-H₂ conversion factor used to determine M_{H_2} from the observations, as we will explore in Section 3.4.

The shift to higher redshifts again shows a similar pattern as the HIMF: The H2MF is clearly increasing to higher redshifts in SIMBA and EAGLE, but mostly unevolving in IllustrisTNG. For EAGLE, $M_{\text{H}_2}^*$ increases by $\times 4$ between $z = 0$ and 2, for SIMBA it is $\times 2.5$, and there is no clear increase for IllustrisTNG. Comparing to $M_{\text{H}_1}^*$ evolution, we see then that EAGLE and IllustrisTNG both yield a greater increase (or less decrease) in $M_{\text{H}_2}^*$ from $z = 0$ to 2, while for SIMBA the increase is similar in both H₁ and H₂; for SIMBA, this is reflected in the similarity of the evolution of Ω_{H_1} versus Ω_{H_2} as noted in Davé et al. (2019).

Overall, the H2MF shows strong differences between models both at $z = 0$ and in terms of evolution to higher redshifts. This highlights the potential for molecular gas mass measurements to be a key discriminator between models. We will discuss the differences in input physics that may be causing these variations in Section 3.8, but it is interesting that none of the models reproduce the $z \approx 0$ H2MF 'out of the box'. We note that Lagos et al. (2015) found better agreement between EAGLE's H2MF and observations from Keres, Yun & Young (2003) when converting their data to H₂ masses assuming a constant $\alpha_{\text{CO}} = 2 \text{ M}_\odot \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$, but this value is low compared to the canonical value for Milky Way-like galaxies of $\alpha_{\text{CO}} \approx 4.5$; Diemer et al. (2019) found better agreement between IllustrisTNG and the H2MF inferred by Obreschkow et al. (2009), but the recent xCOLD GASS determination is somewhat lower at low masses, resulting in more of an apparent disagreement. At the massive end, we are using a significantly larger aperture than Diemer et al. (2019) in order to capture all the molecular gas and compare the global H₂ content, but this increases the H₂ content of these massive systems into poorer agreement. Popping et al. (2019) likewise found that apertures can have a significant impact on this comparison, and only including mass within a 3.5 arcsec aperture resulted in significantly better agreement. As discussed in Section 2.4, it is not entirely clear which way of doing the comparison is more correct. However, what this does indicate is that systematic uncertainties in both observing and modelling the H2MF may be substantial. Among the most crucial of these is the conversion factor between observed CO luminosity and the H₂ mass. We examine this issue next, and

use this to bring our comparisons into the observational plane of CO(1–0) luminosities.

3.4 The H₂-to-CO conversion factor

Observationally, the molecular hydrogen content is most commonly traced via the CO luminosity. Since simulations most directly model molecular gas, we need to convert the molecular gas mass into a CO luminosity. This conversion factor, known as $\alpha_{\text{CO}} \equiv M_{\text{H}_2}/L_{\text{CO}}$, is the subject of much debate (see the review by Bolatto, Wolfire & Leroy 2013). It is clear that α_{CO} depends on metallicity as well as the local strength of H₂ dissociating radiation, which in turn depends on quantities such as the local SFR and shielding column density. For fairly massive galaxies with close to solar metallicity, it is observed that Milky Way-like galaxies have $\alpha_{\text{CO}} \approx 4.5 \text{ M}_\odot \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$, while starburst-like galaxies have a much lower $\alpha_{\text{CO}} \approx 0.8 \text{ M}_\odot \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$, where the CO measurement here corresponds to the lowest $J = 1-0$ rotational transition. One traditional approach to generating an H2MF is to measure CO(1–0), classify the galaxy into one of these categories, and use an appropriate factor typically assumed to be a constant among all galaxies in a class. Clearly, this is quite simplistic, and it is more likely that there is a continuum of α_{CO} values.

Narayanan et al. (2012) presented a theoretical investigation into how α_{CO} varies with galaxy properties. They used a suite of very high resolution isolated disc galaxy and merger simulations, together with CO line radiative transfer modelling, to directly connect the H₂ measured in their galaxies with the emergent CO(1–0) luminosity. They find that the relation between α_{CO} and galaxy properties is on average reasonably well described by

$$\alpha_{\text{CO}} = \frac{20.6}{Z_{\text{H}_2} \Sigma_{\text{H}_2}^{0.5}} \text{ M}_\odot \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}, \quad (5)$$

where Z_{H_2} is the mass-weighted metallicity of the molecular gas in solar units, and Σ_{H_2} is the molecular mass surface density in $\text{M}_\odot \text{ pc}^{-2}$.

The quantities Z_{H_2} and Σ_{H_2} are calculable for each galaxy in our various simulations. To compute Z_{H_2} , we determine the H₂ fraction-weighted metallicity of all gas particles with a radius containing half the molecular gas (R_{H_2}). For Σ_{H_2} , we compute the projected H₂ surface densities from gas particles within R_{H_2} in the z direction, although we checked in SIMBA that using the average of the x , y , and z projections gives similar results. Given these quantities, we use equation (5) to compute α_{CO} for each galaxy. This enables us to predict CO(1–0) luminosity function and scaling relations for direct comparison to the CO(1–0) luminosities in xCOLDGASS and other CO surveys. We note that Narayanan et al. (2012) recommends applying their formula locally within the ISM, but given the $\sim \text{kpc}$ -scale spatial resolution of our simulations this is often impractical. Hence, we compute these quantities within R_{H_2} to give a single α_{CO} value for each galaxy.

Fig. 5 shows a scatter plot of α_{CO} versus stellar mass (left-hand panels) and metallicity (right-hand panels), at $z = 0$ (top) and $z = 2$ (bottom), in SIMBA, computed using the Narayanan et al. (2012) prescription (equation 5). Individual galaxies are colour coded by their location relative to the star-forming main sequence, computed as a running median of SFR versus M_* , ranging from the reddest points having $\Delta_{\text{MS}} \leq -0.7$ to the bluest points with $\Delta_{\text{MS}} \geq 0.7$. The black hexbin shading in the background shows the values computed instead using the Accurso et al. (2017) method (equation 4). The horizontal dashed lines show reference values typically assumed for starbursts ($\alpha_{\text{CO}} = 0.8$) and Milky Way-like galaxies ($\alpha_{\text{CO}} = 4.5$).




Figure 5. α_{CO} versus stellar mass (left-hand panels) and metallicity (right), at $z = 0$ (top panels) and $z = 2$ (bottom), in SIMBA, using the Narayanan et al. (2012) prescription. Individual galaxies are colour coded by their location relative to the main sequence $\Delta_{\text{MS}} = -0.7 \rightarrow +0.7$ (red to blue). The black hexbins show the values using instead the Accurso et al. (2017) prescription. The horizontal dotted lines demarcate values usually assumed for starbursts ($\alpha_{\text{CO}} = 0.8$) and Milky Way-like galaxies ($\alpha_{\text{CO}} = 4.5$), for reference.

Only galaxies with $M_{\text{H}_2} > 10^8 M_{\odot}$ (as well as our adopted stellar mass resolution limit of $M_* > 10^9 M_{\odot}$) are shown, since otherwise they have too little molecular gas to reliably determine the quantities required to estimate α_{CO} .

Generally, α_{CO} is anticorrelated with both stellar mass and metallicity. The trend with metallicity is stronger, reflecting the inverse metallicity dependence in equation (5). At a given metallicity, low-sSFR galaxies have higher values of α_{CO} , since they are gas poor with lower H_2 surface densities. The trend with mass is shallower than that with metallicity at $z = 0$, because lower mass galaxies have lower metallicity, but this is partly counteracted by their higher gas surface densities. It is notable that SIMBA seems to underpredict the value of α_{CO} in a Milky Way-like star-forming galaxy relative to the nominal value of $\alpha_{\text{CO}} = 4.5 M_{\odot} \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$; this is also true for EAGLE. This would result in an overprediction of CO luminosities for a given M_{H_2} ; we quantify the implications of this for the COLF below.

At $z = 2$, the overall values using the Narayanan et al. (2012) prescription are lower than those at $z = 0$ for a given mass or metallicity. This is because galaxies are more compact and gas rich at high redshifts (e.g. Appleby et al. 2020), leading to higher Σ_{H_2} ; while the metallicities are also lower (Davé et al. 2019), the relatively weak dependence on Z_{H_2} is more than compensated by the increased surface density.

The black hexbins in the background show α_{CO} values computed using equation (4) from Accurso et al. (2017), for comparison. At $z = 0$ for massive ($M_* \gtrsim 10^{10} M_{\odot}$) galaxies, the values from the two methods are similar, though with a slight trend towards higher α_{CO} using this method. At lower masses, the Accurso et al. (2017) prescription yields more significantly higher values, since it is more strongly dependent on metallicity, which is also reflected in its tighter relation with metallicity in the right-hand panel. At high redshifts, equation (4) gives implausibly high values, owing to the significantly lower metallicities that is not mitigated by the higher gas surface densities as in the Narayanan et al. (2012) prescription. This is perhaps not surprising, given that the Accurso et al. (2017) prescription is based on $z \approx 0$ observations. For the remainder of




Figure 6. Histograms of α_{CO} for $M_* > 10^{10} M_{\odot}$ galaxies in SIMBA (blue), IllustrisTNG (red), and EAGLE (green), at $z = 0$ (solid curves) and $z = 2$ (dashed). The vertical lines show values for starbursts ($\alpha_{\text{CO}} = 0.8$) and Milky Way-like galaxies ($\alpha_{\text{CO}} = 4.5$), for reference.

this work, we will use the Narayanan et al. (2012) prescription, as it appears to yield a more plausible redshift evolution owing to accounting for both structural and metallicity changes.

To compare the α_{CO} values among the different simulations, we show in Fig. 6 histograms of α_{CO} at $z = 0$ and 2 (solid and dashed lines, respectively) for galaxies with $M_* > 10^{10} M_{\odot}$ from SIMBA (blue), IllustrisTNG (red), and EAGLE (green). Reference lines for typical starburst and MW values are indicated by the vertical dotted lines.

SIMBA and EAGLE both predict a median value of $\alpha_{\text{CO}} = 2.9 M_{\odot} \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$ at $z = 0$, dropping to ≈ 1 at $z = 2$. A typical dispersion of α_{CO} is ≈ 0.5 dex, which is consistent with the spread seen in Fig. 5 for SIMBA. This shows that the assumption of a constant α_{CO} value even among relatively massive star-forming galaxies may be a poor one. Meanwhile, IllustrisTNG shows generally higher values of α_{CO} , with a median value of $\alpha_{\text{CO}} \approx 6$ at $z = 0$ with rapid evolution to $\alpha_{\text{CO}} \approx 1.5$ at $z = 2$, and an even larger dispersion.

It is worth pointing out that the α_{CO} prescription developed by Narayanan et al. (2012) used simulations with SF and feedback prescriptions that are different to any of the simulations considered here, and also were not cosmologically situated. Since α_{CO} likely depends on the structure and distribution of molecular clouds within the ISM, it could be sensitive to such choices. For instance, the same procedure of running very high resolution zoom versions using our three simulations' own SF and feedback prescriptions and applying a CO radiative transfer code could yield substantially different fitting formulae for α_{CO} in each case. While it is beyond the scope to investigate this here, the variations in α_{CO} among the different simulations even when using the same underlying fitting formula highlight the importance of being able to predict this quantity more accurately in $\sim \text{kpc}$ -scale cosmological simulations if one wants to more robustly compare such simulations to CO observations.

Overall, our computed values of α_{CO} broadly follow expected trends of being around the Milky Way value in massive star-forming galaxies today, shifting towards more starburst-like values at high redshifts. There is, however, no bimodality in the α_{CO} distribution, indicating that using a bimodal α_{CO} value based on galaxy classification may be too simplistic. Moreover, the large spread in α_{CO} at a given M_* or metallicity suggests that using a

single value, regardless of what it is, may be a dangerous assumption. This is particularly true when examining counting statistics such as a mass function, where the scatter in α_{CO} could scatter more numerous low- M_{H_2} galaxies up to high $L_{\text{CO}(1-0)}$ values, thereby increasing $M_{\text{H}_2}^*$ over what one would infer from assuming a constant α_{CO} . In order to examine such effects more quantitatively, we next compare the resulting CO(1–0) luminosity function with α_{CO} computed as above among our various simulations, and compare these to observations from $z \sim 0$ to 2.

3.5 The COLF

With a prescription for computing α_{CO} in hand, albeit with its substantial attendant uncertainties, we can now move the comparison of molecular gas into the observational plane. It is particularly interesting to relate the comparative trends seen for the COLF to the analogous trends seen for the H2MF from the previous section – if α_{CO} was a robust and well-determined quantity, then the general trends between these should mirror each other, but we will see that there are significant differences. Moreover, we can also engage in more direct comparisons to observations out to higher redshifts. Thanks to recent surveys such as ASPECS and COLDz, the CO(1–0) luminosity function (COLF) has now been measured out to $z \gtrsim 2$, to go along with the improved recent low-redshift determination from xCOLD GASS (Saintonge et al. 2017). In this section, we compare our simulated COLFs to each other and to these observations, to assess how well current models do at reproducing data and understand how robust these comparisons are.

Fig. 7 shows the CO(1–0) luminosity function at $z = 0, 1$, and 2 (top to bottom) for SIMBA (blue), IllustrisTNG (red), and EAGLE (green). Also shown are various observational determinations: The grey points at $z = 0$ are the observations from xCOLD GASS (Saintonge et al. 2017), while at $z \approx 1$ and 2 we show observations from ASPECS (Aravena et al. 2019) and COLDz (Pavesi et al. 2018; Riechers et al. 2019).

The $z = 0$ COLF, with α_{CO} computed individually for each simulated galaxy, gives a qualitatively different picture in comparison to observations. First, all the models are now significantly closer to the observations. For instance, EAGLE has gone from being extremely deficient at high M_{H_2} to agreeing very well for $L_{\text{CO}(1-0)}$. IllustrisTNG showed a milk overproduction in the H2MF at all masses, but now agrees quite well with the COLF, thanks to its typically higher values of α_{CO} . This illustrates that the assumptions about α_{CO} qualitatively impact simulation constraints based on the H2MF.

At higher redshifts, the qualitative evolution among the models mimics that seen for the gas mass functions: EAGLE and SIMBA have strong positive luminosity evolution out to high redshifts, while IllustrisTNG’s evolution is also positive (owing to its lower α_{CO} values at high- z) but much weaker than that in the other two simulations. The net result is that IllustrisTNG has a difficult time reproducing the very high $L_{\text{CO}(1-0)}$ values seen in galaxies at $z \sim 1-2$, and tends to overproduce low- $L_{\text{CO}(1-0)}$ systems. A similar failing of galaxy formation models was noticed from a comparison to semi-analytic models done in Riechers et al. (2019), and similarly Popping et al. (2019) found that IllustrisTNG underpredicted the high-luminosity end at $z > 1$. As is often the case, the claimed discrepancy is quite model dependent. SIMBA and EAGLE are well able to produce $L_{\text{CO}(1-0)} \sim 10^{11} \text{ K km s}^{-1} \text{ pc}^2$ systems at $z \sim 2$, as observed. The key is their low α_{CO} values, typically $\alpha_{\text{CO}} \sim 1 \text{ M}_{\odot} \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$. We note that SIMBA reproduces the mass–




Figure 7. CO(1–0) luminosity functions for SIMBA (blue), IllustrisTNG (red), and EAGLE (green) at $z = 0, 1$, and 2 (top to bottom). These are computed using the Narayanan et al. (2012) prescription for α_{CO} . Observations from xCOLD GASS at $z = 0$ (Saintonge et al. 2017), ASPECS at $z \approx 1$ and 2 (Aravena et al. 2019), and COLDz at $z \approx 2$ (Riechers et al. 2019) are shown in black and grey. IllustrisTNG and EAGLE well reproduce the COLF despite widely different H2MFs, while SIMBA overproduces the COLF like the H2MF. SIMBA and EAGLE are able to reproduce observations of high- $L_{\text{CO}(1-0)}$ galaxies at $z \gtrsim 1$.

metallicity relation at $z \sim 2$ (as well as $z \sim 0$; Davé et al. 2019), so the low α_{CO} values are not due to overenriched galaxies.

To illustrate the sensitivity of these predictions to assumptions about α_{CO} , Fig. 8 shows a comparison of these models with different assumptions about α_{CO} , at $z = 0$ (top) and $z = 2$ (bottom). For SIMBA, we show with blue dotted and dashed curves the results of using a constant $\alpha_{\text{CO}} = 2.9 \text{ M}_{\odot} \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$ and the Accurso et al. (2017) α_{CO} prescription, respectively. For IllustrisTNG and EAGLE shown in red and green, the dotted lines show a constant $\alpha_{\text{CO}} = 2.9 \text{ M}_{\odot} \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$. For comparison, the semitransparent solid lines reproduce the results from Fig. 7, and the observations shown there are also reproduced.

For SIMBA at $z = 0$, using the median α_{CO} (dotted line) rather than the full spread in values does not yield a much different COLF. At $z = 2$, however, there is a large difference, as using a constant α_{CO} strongly underpredicts the bright end. This illustrates the importance of including a distribution of α_{CO} values for comparing to observations. Meanwhile, the Accurso et al. (2017) prescription yields a somewhat lower COLF at $z = 0$, owing to its generally higher α_{CO} values. At $z = 2$, the likely unphysically high α_{CO} values predicted in this prescription result in a much poorer agreement with data.

For EAGLE and IllustrisTNG, the story is similar: At $z = 0$, using a constant α_{CO} results in mildly lower COLFs, but at $z = 2$, the difference is very pronounced, and as with SIMBA tends to strongly truncate the bright end of the COLF. The assumption of a constant




Figure 8. CO(1–0) luminosity functions for SIMBA (blue), IllustrisTNG (red), and EAGLE (green) at $z = 0$ and 2 (top to bottom). The different line types show the results when computed using various α_{CO} prescriptions. The faint solid lines show the results using the Narayanan et al. (2012) prescription, reproduced from Fig. 7. The dotted lines for each simulation show the results assuming a constant $\alpha_{\text{CO}} = 2.9 \text{ M}_{\odot} \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$, which is a typical average value at $z = 0$. Finally, the dashed blue line shows the COLF for SIMBA assuming the Accurso et al. (2017) α_{CO} prescription, which produces unphysically high values of α_{CO} at $z = 2$. For reference, observations are shown as in Fig. 7. The assumption of a constant α_{CO} makes an especially large difference for the bright end at $z = 2$.

α_{CO} may thus be a major reason why Riechers et al. (2019) and Popping et al. (2019) found that models could not reproduce the bright end of the high-redshift COLF.

Clearly, independent constraints on α_{CO} at both low and high redshifts would be highly valuable in order to conduct a robust comparison between the observed and simulated COLFs. This could come from direct observations (e.g. Accurso et al. 2017), or else from sophisticated higher resolution simulations including CO line radiative transfer, such as with SÍGAME (Olsen et al. 2017). It is possible that the α_{CO} values coming from the Narayanan et al. (2012) prescription are systematically discrepant in one or more of our simulations, which could then either indicate a failing of that model or the inapplicability of the Narayanan et al. (2012) prescription for that model. There is thus substantial effort still needed in order to be able to utilize the COLF as a robust constraint on galaxy formation models.

Overall, it is encouraging that all our models better reproduce the $z \approx 0$ COLF than the H2MF, since the former is the more direct observable. At higher redshifts, at least some current galaxy formation models have no difficulty forming galaxies with high $L_{\text{CO}(1-0)}$ values at $z \sim 1-2$ – the evolution predicted in EAGLE is in very good agreement with observations, SIMBA’s evolution is still quite reasonable compared to data, and IllustrisTNG has significant difficulties generating high- $L_{\text{CO}(1-0)}$ systems at high redshifts despite its rapid downwards evolution of α_{CO} . EAGLE’s good agreement with the COLF was also noted in Lagos et al. (2015), despite using a different metallicity-dependent α_{CO} prescription. None the less, all these conclusions are highly sensitive to assumptions regarding α_{CO} . To make progress, this must be independently constrained

either observationally and/or theoretically in order to properly assess whether galaxy formation models match observations of molecular gas in galaxies across cosmic time.

3.6 Gas fraction scaling relations

We have seen that our three galaxy formation simulations qualitatively reproduce the distribution functions of cold gas and their measures in galaxies, but there are also significant discrepancies in each case. To investigate the successes and failures in more detail and isolate the galaxy population(s) responsible, it is instructive to examine scaling relations of gas content versus global galaxy properties, which is what we do here.

Fig. 9 shows a montage of scaling relations for our simulations, compared to observations. The y-axis panels show the quantities $f_{\text{H I}} = M_{\text{H I}}/M_*$, $f_{\text{H}_2} = M_{\text{H}_2}/M_*$, and $L_{\text{CO}(1-0)}$ as computed assuming the Narayanan et al. (2012) prescription for α_{CO} . The x-axis quantities are the stellar mass M_* , the sSFR, and stellar mass surface density (computed within the half stellar mass radius) μ_* . In each panel, we show a running median for SIMBA (blue), IllustrisTNG (red), EAGLE (green solid), and EAGLE-Recal (green dashed), with the spread from the 16–84th percentile shown as the shaded blue region for SIMBA. The underlying grey points show the observations from xGASS (for H I) and xCOLD GASS (for H₂ and CO) with downwards arrows indicating upper limits. A running median is shown as the black points with the error bars indicating the 1σ spread around the median. The medians are taken over all data points including non-detections or gas-free galaxies; using the median rather than the mean avoids any ambiguity regarding the values for the upper limits in the observations. The results are not significantly different if we compute the running mean instead. The bins are chosen to have roughly equal numbers of galaxies in each.

The leftmost column shows the relations versus M_* . In the simulations, stellar mass is generally the most accurately predicted quantity; hence, M_* scaling relations are likely the most robust trends predicted by the models. All models predict falling $f_{\text{H I}}$ and f_{H_2} with increasing M_* , and a slow rise in $L_{\text{CO}(1-0)}$ with M_* in the star-forming regime and a quick drop in the most massive (generally quenched) galaxies. The variance around the median in SIMBA is about 0.4 dex in H I, and increases towards high M_* for the molecular and CO trends. These broadly mimics the corresponding observational trends, but there are notable discrepancies.

For $f_{\text{H I}}$, SIMBA and IllustrisTNG agree reasonably well at most masses, reflecting their good agreement with the HIMF as seen in Fig. 3, but somewhat overpredict the atomic fractions at the highest masses. EAGLE, meanwhile, is an interesting case – it strongly underpredicts the HIMF, yet the $f_{\text{H I}}$ is only mildly low, except for a stronger drop at the lowest M_* . This occurs because EAGLE has a significantly larger fraction of galaxies across all masses with little or no H I, which impacts the counts more than the median values. EAGLE-Recal, meanwhile, has higher $f_{\text{H I}}$ at all masses, but in particular follows the observations more closely at $M_* \lesssim 10^{10} \text{ M}_{\odot}$, the combination of which produces an HIMF for EAGLE-Recal that is in good agreement with observations. As discussed in Section 3.2, we speculate that the significant difference between EAGLE and EAGLE-Recal is a consequence of the stronger resolution dependence of EAGLE’s subgrid feedback model relative to those used by SIMBA and IllustrisTNG.

f_{H_2} similarly shows an overall falling trend in all models. IllustrisTNG overproduces the molecular fractions particularly in low- and high- M_* galaxies, and likewise shows an upwards deviation in f_{H_2} towards high M_* like that seen for $f_{\text{H I}}$, indicating that this




Figure 9. Gas fraction scaling relations at $z = 0$ in SIMBA (blue), IllustrisTNG (red), EAGLE (green solid), and EAGLE-Recal (green dashed). The top row shows atomic gas fraction $M_{\text{H I}}/M_*$, the middle shows molecular gas fraction M_{H_2}/M_* , and the bottom row shows $L_{\text{CO}(1-0)}$ with α_{CO} computed using the Narayanan et al. (2012) prescription. These are shown as a function of stellar mass (left-hand panels), specific SFR (middle), and stellar surface density (right). Lines show the running median for each scaling relation; in each case, the bins have been chosen to hold approximately an equal number of galaxies from within each data set. The blue shaded region shows the 16–84 per cent spread around the median in SIMBA. The grey points show results from xGASS for H I data and xCOLD GASS for H₂ and CO data (non-detections shown as downward arrows at their upper limits), with the running median shown as the black points with error bars indicating the 1σ spread about the median including non-detection. All simulations broadly reproduce the trends in gas fractions, but significant discrepancies are seen for SIMBA in molecular gas at low M_* and all sSFRs, for IllustrisTNG in molecular gas at high M_* , and for EAGLE in H I at low masses, which is much improved in EAGLE-Recal.

represents a true bump in the overall cold gas in massive systems as opposed to some artefact of H I–H₂ separation. The overprediction, particularly at the massive end, is subject to uncertainties regarding apertures; Diemer et al. (2019) and Popping et al. (2019) obtained significantly better agreement owing to their use of a smaller aperture.

SIMBA, meanwhile, overproduces f_{H_2} particularly at low masses, suggesting that the excess seen in the H2MF comes from dwarfs

that are overly molecular gas rich. Given SIMBA’s large aperture, it is subject to similar aperture caveats as IllustrisTNG. Curiously, despite matching the H2MF and the GSMF fairly well, SIMBA systematically overproduces the H2 fractions.

EAGLE shows the best agreement in the slope of $f_{\text{H}_2}(M_*)$, although it is slightly low. Hence, for the galaxies that have molecular gas, EAGLE does a good job of reproducing their gas

fractions. EAGLE-Recal is very similar to EAGLE, showing that molecular fractions are less resolution sensitive than atomic fractions in EAGLE.

Both IllustrisTNG and SIMBA produce at least some quite massive galaxies with significant H_2 even though those galaxies are generally quenched, which is also seen in some observed systems (e.g. Davis et al. 2019). It remains to be seen if there is statistical agreement with observations since xCOLD GASS is not a sufficiently large sample to include such rare objects, and current observations of molecular gas in massive galaxies are limited to heterogeneously selected samples. In SIMBA, such gas typically has very low SF efficiency and lies below the Kennicutt–Schmidt relation (Appleby et al. 2020), likely owing to its diffuse distribution.

We can further compare these to the MUFASA gas scaling predictions shown in fig. 5 of Davé et al. (2017). The f_{H_2} relation is fairly similar, since the H_2 formation model has remained the same, but SIMBA produces more molecular gas in low-mass galaxies. The f_{H1} predictions are also fairly similar, but SIMBA produces somewhat more H I at the highest masses. Likely this occurs because MUFASA’s quenching feedback mechanism explicitly heated all the ambient gas in high-mass haloes. MUFASA thus oversuppressed H I in massive galaxies, while SIMBA slightly undersuppresses it.

Looking at $L_{CO(1-0)}$, we see that the impact of the variations in α_{CO} results in somewhat different trends relative to that for f_{H_2} . As in the COLF, SIMBA clearly overproduces $L_{CO(1-0)}$, while EAGLE and IllustrisTNG do generally better, the latter owing to its significantly higher values of α_{CO} . SIMBA and EAGLE show a sharp decline in CO luminosity at $M_* \gtrsim 10^{11} M_\odot$, similar to that seen in xCOLD GASS, while IllustrisTNG continues to show typically high $L_{CO(1-0)}$ out to large stellar masses. In SIMBA, despite there being some molecular gas in these galaxies, the H_2 surface densities are generally low, which increases α_{CO} and thus decreases the CO luminosity. None the less, the very most massive galaxy in SIMBA is relatively CO bright.

The second column depicts the trends versus sSFR. Again, all simulations qualitatively produce the observed trends of increasing cold gas fractions with sSFR, but only a weak trend with $L_{CO(1-0)}$. The trends in individual simulations mirror that seen for M_* : IllustrisTNG and SIMBA match f_{H1} and are somewhat high in f_{H_2} , reflecting the trends seen in the mass functions. For EAGLE, the origins of the discrepancies in the mass functions become clearer in this plot: EAGLE produces too little neutral gas in low-sSFR systems, and generally shows a steeper slope of either H I or H_2 gas fraction versus sSFR versus observations, indicating that EAGLE’s model oversuppresses cold gas in green valley galaxies. EAGLE-Recal simply lacks many galaxies in this green valley regime, but for the few that are there, it seems to produce significantly higher f_{H1} . Meanwhile, the $L_{CO(1-0)}$ trends represent the competing effects of lower sSFR objects generally having lower molecular fractions, but also being larger systems overall. All models achieve this balance naturally, although as before the amplitudes vary somewhat. All models also produce a sharp drop in $L_{CO(1-0)}$ for the quenched systems as observed, more dramatic than that for f_{H_2} . In SIMBA this owes to a higher α_{CO} in such systems from the lower molecular gas surface densities, and this appears to occur naturally in the other simulations as well.

The rightmost column shows trends versus the stellar surface density μ_* . Both data and models show dropping cold fractions with increasing μ_* , as higher μ_* galaxies tend to have lower sSFR and therefore less cold gas relative to stars. The slope of the trend in all models generally reproduces observations. However, the amplitudes vary substantially. IllustrisTNG produces good agreement with all

properties as a function of μ_* , while EAGLE tends to underproduce the gas contents, and SIMBA predicts reasonable gas fractions but generally has a flatter trend with μ_* than observations. The maximum μ_* reached in each model is substantially different, which is directly tied to numerical resolution: IllustrisTNG has the highest resolution, then EAGLE, and then SIMBA, which has $\sim 15\times$ lower mass resolution than IllustrisTNG owing to its larger volume and lower number of gas elements. This is also seen by comparing EAGLE to EAGLE-Recal; the latter produces higher f_{H1} values and also higher f_{H_2} , even though the molecular fractions are a given M_* or sSFRs are not resolution dependent. One interpretation is thus that IllustrisTNG’s resolution is necessary in order to properly resolve the inner structure of the stellar component. If so, this means that predicting inner structural quantities such as stellar mass surface density, bulge fractions, Sersic indices, etc. is quite computationally demanding and subject to a careful assessment of resolution convergence effects. However, we will show in Section 3.7 that SIMBA with AGN feedback off produces a plethora of very high μ_* galaxies, even at the same resolution. Hence, it is certainly possible to produce galaxies with high μ_* at SIMBA’s resolution, which suggests that instead it may be SIMBA’s AGN feedback scheme that prevents galaxies from having sufficiently high stellar densities.

Overall, all our simulations broadly reproduce trends versus stellar mass, surface density, and specific SFR, but there are notable variations and discrepancies that highlight specific differences reflected in the earlier mass and luminosity functions. IllustrisTNG and SIMBA tend to do slightly better on H I fractions, while EAGLE does somewhat better on molecular gas fractions. Stellar surface densities provide a novel constraint on feedback regarding the spatial distribution of stellar growth in simulated galaxies, albeit such measures may be more sensitive to numerical effects. As observations improve both at $z = 0$ and in the distant universe, it is clear that gas scaling relations will provide complementary constraints and insights into the robustness and validity of the input physics in cosmological galaxy formation simulations.

3.7 SIMBA AGN feedback variants

AGN feedback is a key ingredient in modern galaxy formation models in order to quench massive galaxies (Somerville & Davé 2015). In terms of the cold gas content, the predominant effect is to eject and/or prevent gas accretion, owing to energy injection from the AGN. Thus, it is expected that models predict low cold gas contents in massive galaxies, as we have seen above. Yet AGN feedback can come in many different forms (Sturm et al. 2011; Maiolino et al. 2012; Heckman & Best 2014), such as radiatively driven winds off the accretion disc, relativistic jets, and photoheating of surrounding gas. As such, it is interesting to know which type of AGN feedback is responsible for lowering the cold gas content. In SIMBA, we have a particularly interesting model that employs three different types of AGN feedback, broadly following the three feedback modes described in Heckman & Best (2014). Understanding how these different types of AGN feedback impact galaxy cold gas content, particularly in massive galaxies, may provide useful insights that would help guide and constrain quenching models.

In this section, we examine the impact of AGN feedback on the cold gas content of galaxies in SIMBA. As outlined in Section 2.1, the three different forms of AGN feedback in SIMBA are jet feedback at low $f_{Edd} \lesssim 0.02$, radiative feedback at high f_{Edd} , and X-ray feedback at low f_{Edd} and gas fractions (see Davé et al. 2019, for full details). The first two feedback modes are bipolar and purely kinetic, while




Figure 10. HIMF (left) and H2MF (right) at $z = 0$ and 1 (top, bottom) for the four AGN feedback variants: full SIMBA (blue), no-Xray (red), no-jet (green), and no-AGN (cyan). These are run in $50 h^{-1}$ Mpc volumes; the corresponding SIMBA $100 h^{-1}$ Mpc volume results are shown as the blue dashed line. Observations from ALFALFA (Jones et al. 2018) and xCOLD GASS (Fletcher et al. 2020) are shown for reference. The dominant difference comes from the inclusion of AGN jet feedback, which strongly suppresses cold gas in massive galaxies at $z \lesssim 1$.

the X-ray feedback is spherical and mostly kinetic, though typically of much lower strength. The radiative feedback mode corresponds to what is often referred to as ‘quasar mode’ feedback, and is designed to qualitatively model outflows of warm ionized (Perna et al. 2017) and/or cold molecular (Maiolino et al. 2012) gas at speeds of many hundreds of km s^{-1} . The jet mode aims to represent the impact of collimated relativistic radio jets, often called ‘radio mode’ feedback, that eject hot plasma out to tens to hundreds of kpc before imparting their energy into the surrounding medium. X-ray feedback is motivated more fully in Choi et al. (2012) as high-energy radiation pressure that imparts momentum on to the gas surrounding the black hole; unlike the other modes, this mode is spherical, and typically generates a relatively modest outward push of a few hundred km s^{-1} in the gas closest to the black hole, with a strength that scales as the inverse square of the distance.

The SIMBA suite contains $50 h^{-1}$ Mpc box size, 2×512^3 element variants where each of these AGN feedback mechanisms is turned off in turn, as follows:

- (i) ‘Simba’ – Full AGN feedback physics.
- (ii) ‘No-Xray’ – Only X-ray feedback turned off.
- (iii) ‘No-jet’ – X-ray and jet feedback turned off.
- (iv) ‘No-AGN’ – All AGN feedback turned off.

All simulations are run from the same initial conditions, and all other input physics is unchanged. These should thus be regarded as numerical experiments to isolate the impact of each feedback mechanism, rather than realistic variations of galaxy formation models. For instance, only the full SIMBA model accurately reproduces the observed stellar mass function (Davé et al. 2019).

Fig. 10 shows the H I (left-hand panels) and H₂ (right) mass functions, at $z = 0$ (top panels) and $z = 1$ (bottom) for the $50 h^{-1}$ Mpc SIMBA run (blue), No-Xray (green), No-jet (red), and No-AGN (cyan). The blue shading shows the cosmic variance estimated over 8 simulation subobjects in the full physics run; the other models show similar variance but are not shown for clarity. The dotted blue line shows the results from a small-volume ($25 h^{-1}$ Mpc) version of SIMBA with $8 \times$ better mass resolution. The dashed blue line shows the $100 h^{-1}$ Mpc SIMBA results reproduced from Figs 3 and 4. Observations from ALFALFA (left) and xCOLD GASS (right) are also reproduced from those figures, although we will not focus on comparing to data here.

We first compare the 50 and $100 h^{-1}$ Mpc SIMBA runs, which differ only in that the latter has $8 \times$ the volume. In general, they show similar results within uncertainties, but the variance is larger in the $50 h^{-1}$ Mpc case owing to the fewer numbers of galaxies. One notable difference is that the smaller volume produces a higher HIMF, which is within the variance at $z = 0$ but not at $z = 1$. With a larger volume, there are more quenched galaxies with less H I in their outskirts, resulting in a higher mass function for the smaller volume. The difference is less noticeable for H₂, and in this case the large volume has a slightly higher mass function. None the less, these are all variations that are within 1σ expectations of cosmic variance, and since these two volumes are started from different initial conditions, such variations are not unexpected.

We can examine resolution convergence by comparing the $50 h^{-1}$ Mpc (solid) and $25 h^{-1}$ Mpc (dotted) results, which have the same dynamic range but the latter has $8 \times$ better resolution. As such, it resolves galaxies to much lower masses. For the H2MF, there is quite




Figure 11. Gas scaling relations at $z = 0$ as in Fig. 9, for $50 h^{-1}$ Mpc, 2×512^3 variants of SIMBA with different AGN feedback modules included. The blue line shows the running median of these scaling relations for the full SIMBA model, red (‘No-Xray’) shows turning off the X-ray feedback only, green (‘No-jet’) shows turning off the jets and X-ray feedback, and cyan (‘No-AGN’) showing further turning off the radiative feedback. The xGASS and xCOLD GASS data are also shown as in Fig. 9. AGN feedback has no impact on the low- M_* or high-sSFR galaxies, but in particular, jet AGN feedback strongly lowers the gas content in massive galaxies to bring them into better agreement with observations.

good resolution convergence, but it can be seen that the turn-down at $M_{H_2} \lesssim 10^9 M_\odot$ owes primarily to incompleteness, as the $25 h^{-1}$ Mpc results do not show a turnover there. The HIMF, in contrast, does tend to be slightly higher in the higher resolution volume. This general behaviour is also seen between the main EAGLE run and the high-resolution EAGLE-Recal run (Crain et al. 2017), though with a more dramatic difference than that in SIMBA. However, in SIMBA, this might also be a volume effect as described in the previous paragraph. More crucially, it is now evident that the turn-down in the HIMF in SIMBA clearly owes to numerical resolution – with a high-resolution run, the HIMF continues to broadly follow the observed HIMF trend down to $M_{H_1} \lesssim 10^9 M_\odot$. The impact of resolution is particularly evident at $z = 1$, and suggests that the main SIMBA volume can only resolve H I in galaxies down to $M_{H_1} \sim 10^{10} M_\odot$ at these redshifts. Upcoming surveys, however,

are unlikely to probe to lower H I masses, at least for individual detections.

Turning to the AGN feedback variants, at $z = 0$ we see by far the strongest impact owes to the jet feedback mode. This can be seen by noting that turning off X-rays (blue $\rightarrow$ red) only modestly changes the cold gas mass functions. Thus, X-ray feedback does have some effect, but it is barely at the 1σ level relative to cosmic variance. In contrast, further turning off jet feedback (red $\rightarrow$ green) makes a much larger difference, which is true for both H I and H₂. Davé et al. (2019) noted (via analogous tests) that it is jet feedback that is primarily responsible for quenching SF in SIMBA, so it is unsurprising to see this reflected in H₂, but it is perhaps not immediately evident that it would also impact the H I so dramatically. Finally, we note that radiative AGN feedback has very little impact on cold gas content (green $\rightarrow$ cyan). If anything, including this feedback mode tends to

increase the H I, which may be because modest-velocity AGN winds are able to throw more material (albeit a small amount overall) into the CGM that can then remain sufficiently dense to be self-shielded.

Looking at $z = 1$ (bottom panels), it is clear that the differences are much less pronounced. We do not show the results at $z = 2$, but here all models are essentially identical to within cosmic variance. This indicates that the impact of jet feedback is mostly seen at $z \lesssim 1$, which is also reflected in larger scale motions of baryons in the IGM in SIMBA (Borrow, Angles-Alcazar & Dave 2020; Christiansen et al. 2019). Without jets (green and cyan), the cold gas mass functions actually increase at the massive end from $z = 1 \rightarrow 0$, whereas with jets (blue and red) they decrease. Hence, in SIMBA, jet feedback that is responsible for quenching galaxies also dramatically changes late-time cosmic cold gas evolution.

Fig. 11 gives a complementary view of the variations in cold gas properties induced by AGN feedback. This shows cold gas scaling relation comparisons at $z = 0$ as in Fig. 9, but here we compare the suite of $50 h^{-1}$ Mpc AGN feedback variant runs versus xGASS and xCOLD GASS scaling relations. The blue, red, green, and cyan lines show running median scaling relations for the full SIMBA, No-Xray, No-jet, and No-AGN variants as described above. Note that the full SIMBA run here is in a $50 h^{-1}$ Mpc box, rather than the fiducial $100 h^{-1}$ Mpc box shown in Fig. 9, in order to homogenize the comparison to the other runs. The galaxies in the $50 h^{-1}$ Mpc volume have somewhat higher H I contents and slightly lower H₂ contents than the $100 h^{-1}$ Mpc box, as seen in Fig. 10. The observations from xGASS and xCOLD GASS are reproduced from Fig. 9 for reference, though the focus in this plot is a comparison among the AGN feedback variants.

AGN feedback primarily impacts high-mass galaxies in SIMBA. This is seen by the fact that at $M_* \lesssim 10^{9.5} M_\odot$, all models give similar results, even without any AGN feedback at all. There is immediately an increase in gas content by turning on the radiative mode feedback (i.e. cyan No-AGN versus green No-Jet lines), which as we discussed in the mass functions section seems counter-intuitive: Radiative AGN feedback increases the amount of both H I and H₂ (as well as L_{CO}). This is because there is additional cold material being driven out of the galaxy that both mildly lowers the amount of stars formed while also providing more gas to accrete from the CGM.

The largest difference comes from further turning on jet feedback (green versus red lines). This causes a major drop in the gas fractions at $M_* \gtrsim 10^{10} M_\odot$, particularly in molecular gas. It is this feedback, which is also responsible for quenching galaxies (Davé et al. 2019), that is thus responsible for reproducing the strongly negative slope of gas fractions as observed; such a slope is not a trivial outcome of galaxy formation, but rather driven specifically by AGN feedback. Note that the H I fractions still have a negative slope even without jet feedback, since the growth of gravitational hot gaseous haloes suppresses cold gas in the CGM relative to the stars; in contrast, the molecular fractions are nearly mass independent without jet feedback. Adding X-ray feedback has a small but noticeable impact on the molecular content of the most massive systems ($M_* \gtrsim 10^{11} M_\odot$), which comes from the removal of molecular gas from the central regions owing to this feedback mode in SIMBA (Appleyby et al. 2020).

The sSFR scaling relations (middle column) show much less dependence on AGN feedback. This is primarily because the most significant impact of AGN feedback is to quench galaxies and build a population with very low sSFRs; the ones that remain star forming and dominate this plot are not strongly affected. None the less, one can see that turning on radiative mode AGN feedback, for instance, mildly increases the amount of H I and H₂ in galaxies at modest

sSFRs. The slopes of the sSFR- f_{gas} relations are thus in somewhat better agreement with observations when AGN feedback is included versus having all AGN feedback turned off.

The stellar surface density plot provides some insight into the nature of the discrepancies seen in Fig. 9 for the full SIMBA run. Previously, we argued that resolution may be impacting the results to prevent very high surface densities. However, when jets are off (green/cyan lines), it is clear that plenty of galaxies at high stellar surface densities as observed can indeed be achieved at this numerical resolution. Instead, it is primarily the jet feedback that causes the lack of high- μ_* galaxies, with X-ray feedback providing a minor addition. This is seen particularly dramatically in the f_{H_2} plot, where turning on X-ray feedback starkly reduces the number of high- μ_* galaxies. Note that Appleyby et al. (2020) found that this same X-ray feedback is necessary to obtain central depressions in the sSFR profiles of green valley galaxies in accord with observations, yet the impact on the stellar surface densities may be overly strong. It could still be that the results are impacted by numerical resolution; simply producing high- μ_* galaxies does not guarantee that resolution is not a concern. None the less, this comparison illustrates that the stellar surface densities could potentially provide a complementary constraint on X-ray feedback, constrained by the internal build-up of stars within galaxies.

Overall, the comparison between AGN feedback variants in SIMBA demonstrates that the cold gas mass function and its evolution since $z \sim 1$ provide interesting constraints on AGN feedback models. SIMBA's AGN feedback was designed to quench massive galaxies, while keeping the kinetic energy from the jets to be relatively modest, in accord with observations (Whitnam et al. 2018; Davé et al. 2019). Given the large impact of this feedback mode in impacting the HIMF and H2MF, it may be likely that the remaining discrepancies versus e.g. the H2MF could be mitigated with modest tweaking of the jet energy input, while not substantively impacting the quenched population. SIMBA's jet AGN feedback has the largest impact on cold gas properties, and it may actually oversuppress high-stellar surface density systems. None the less, it is a non-trivial success that SIMBA, as well as other simulations, is able to come close to reproducing cold gas mass functions particularly at the massive end via AGN feedback, and highlights such data as a way to independently constrain a poorly understood aspect of modern galaxy formation models.

3.8 Discussion: the interplay of feedback and cold gas

We have seen that despite very similar stellar and SFR properties, EAGLE, IllustrisTNG, and SIMBA have substantially different predictions for cold gas properties. In this section, we briefly speculate on the physical origin of these model differences. A proper study of this would require a systematic parameter space exploration of a range of feedback prescriptions, which is beyond the scope of this work, but there are some qualitative differences in feedback models that may broadly explain some of the variations.

At low masses, the dominant feedback mechanism in these simulations is SF-driven outflows. In both SIMBA and IllustrisTNG, SF feedback is done via decoupled kinetic outflows, meaning that they explicitly do not impact or remove ambient ISM gas. In EAGLE, the feedback is done by raising the temperature of gas particles neighbouring newly formed star particles, by a temperature increment of $10^{7.5}$ K. This drives outflows from the resulting pressure gradient, and also adds heat to the ISM, likely lowering the cold (both molecular and atomic) gas content. Also, it may heat surrounding CGM gas more, which would particularly lower the H I content

associated with the galaxy. This may explain why, at low masses, EAGLE predicts lower H I and particularly H₂ contents, indicating that the low-mass gas content of galaxies is sensitive to the amount of thermal energy deposition from SF feedback into the ISM and CGM. This explanation seems plausible when one considers that the effect is strongly resolution dependent, and that EAGLE’s thermal feedback scheme injects less energy per feedback event at higher resolution.

Larger differences are seen in the HIMF and H2MF at high masses above $\gtrsim M^*$. In this regime, AGN feedback is an important contributor. Again, SIMBA and IllustrisTNG have qualitatively similar AGN feedback approaches, via kinetic jets in massive galaxies at late times. EAGLE continues to use a thermal-based feedback approach, with a higher temperature increment of $10^{8.5}$ K for heating gas around the black hole. Judging from the lack of massive cold gas reservoirs in EAGLE, it may be that this feedback mode is overly aggressive in heating or expelling cold gas. On the other hand, the overproduction of H₂ in SIMBA and IllustrisTNG (modulo uncertainties in α_{CO}) may indicate that a purely kinetic feedback form is insufficient to remove molecular gas from the ISM in a manner in accord with observations. Note that SIMBA generally has weaker jet energy input than IllustrisTNG and does not vary the direction of the jet feedback on short time-scales as IllustrisTNG does, but it does include X-ray feedback that is important for reducing the central molecular gas in quenching high-mass galaxies as observed (Appleby et al. 2020). Thus, the molecular gas content in massive galaxies appears to be impacted by essentially all the different feedback mechanisms, making detailed tests of the impact of each one within each simulation model (as in Section 3.7) important for understanding their impact. We leave such a study for future work.

Among the more striking results is that IllustrisTNG predicts comparable or lower H I and H₂ mass functions at higher redshifts versus $z = 0$, while SIMBA and EAGLE both have increasing mass function to high- z . This is particularly curious since both IllustrisTNG and SIMBA used decoupled kinetic winds and two-mode AGN feedback, yet they yield qualitatively different evolution. One potential difference is that IllustrisTNG uses spherical thermal AGN feedback at high black hole Eddington fractions that are more common at earlier epochs, only going to kinetic at lower Eddington fractions, while SIMBA uses bipolar kinetic AGN feedback in all cases. It could be that such thermal feedback during the peak of black hole growth activity is significantly impacting the gas in the immediate vicinity of star-forming galaxies, thereby modulating the cold gas content. This does not happen in SIMBA owing to the explicit bipolar and decoupled nature of the AGN feedback, even at high Eddington fractions. Since the largest impact is on the H I, it is not immediately clear which approach is in better agreement with data since HIMFs are not yet available at higher redshifts, but this is an interesting prediction that may already be testable using, e.g. the upcoming LADUMA survey on MeerKAT.

4 SUMMARY

We have examined the atomic and molecular hydrogen properties of galaxies in three state-of-the-art cosmological hydrodynamic simulations, namely SIMBA, IllustrisTNG, and EAGLE, and compared them to available observations focusing on the $z \approx 0$ xGASS and xCOLD GASS stellar mass-limited surveys. These simulations employ subgrid models for self-shielding and H₂ formation. In all models, the self-shielding is done via the prescription in Rahmati

et al. (2013). The H₂ fraction in SIMBA is computed on the fly using the Krumholz & Gnedin (2011) prescription, while EAGLE uses a similar method in post-processing, and TNG using a related prescription from Gnedin & Draine (2014). We have studied gas mass functions, CO(1–0) luminosity functions with α_{CO} computed in the simulations using the prescription from Narayanan et al. (2012), and gas scaling relations versus stellar mass, sSFR, and stellar mass surface density. We also looked in SIMBA to examine how the cold gas properties vary when excluding individual AGN feedback processes. We summarize our main results as follows:

(i) SIMBA, EAGLE, and IllustrisTNG all have $z = 0$ stellar mass functions and star-forming main sequences that are in good agreement with each other and with observations, indicating that stellar properties can now be well reproduced in the current generation of hydrodynamic models.

(ii) The simulations’ HIMFs all show a Schechter shape with a flat faint-end slope as observed, but the characteristic mass $M_{\text{H I}}^*$ varies substantially, from $2 \times 10^9 M_\odot$ for EAGLE, to $10^{10} M_\odot$ for SIMBA (in good agreement with the ALFALFA data), to $2 \times 10^{10} M_\odot$ for IllustrisTNG. EAGLE-Recal produces significantly more H I than EAGLE, thus showing some resolution sensitivity for this model.

(iii) SIMBA and EAGLE show a shift towards higher $M_{\text{H I}}^*$ at higher redshifts, increasing by $\sim \times 2$ – 2.5 out to $z = 0$, while IllustrisTNG shows a dropping $M_{\text{H I}}^*$. This qualitative difference will hopefully be discriminated in the upcoming generation of H I surveys.

(iv) The H2MFs for these simulations likewise all show a flat faint-end slope as observed, but with dramatic variations in $M_{\text{H}_2}^*$, from $10^{9.2} M_\odot$ for EAGLE to $10^{10} M_\odot$ for SIMBA and IllustrisTNG; xCOLD GASS finds $M_{\text{H}_2}^* \approx 10^{9.7} M_\odot$, intermediate between these predictions.

(v) As with H I, there is a strong increase in $M_{\text{H}_2}^*$ out to $z = 2$ for EAGLE ($\sim \times 4$) and SIMBA ($\sim \times 2.5$), but no change for IllustrisTNG. Thus, the evolutionary differences are mainly in the total neutral gas content, not in the relative fractions of H I and H₂.

(vi) Examining the more directly observable $L_{\text{CO}(1-0)}$ rather than M_{H_2} gives a substantively different picture, because α_{CO} can vary significantly for galaxies both within a model and between models. EAGLE and SIMBA produce on average $\alpha_{\text{CO}}(z = 0) \approx 3 M_\odot \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$ for $M_* > 10^{10} M_\odot$ star-forming galaxies, while IllustrisTNG has $\alpha_{\text{CO}}(z = 0) \approx 6 M_\odot \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$. There is strong evolution in α_{CO} to higher redshifts, with a median $\alpha_{\text{CO}}(z = 2) \approx 1$ – $1.5 M_\odot \text{ pc}^{-2} \text{ K}^{-1} \text{ km}^{-1} \text{ s}$ in all models.

(vii) The resulting simulated $z = 0$ CO(1–0) luminosity functions are generally in closer agreement with xCOLD GASS observations versus the H2MF, with EAGLE and IllustrisTNG agreeing very well and SIMBA still somewhat too high.

(viii) The evolution of the COLF broadly mimics that of the H2MF. EAGLE and SIMBA show much higher COLFs at $z = 2$ than $z = 0$. These models well reproduce observations of high- $L_{\text{CO}(1-0)}$ galaxies at these epochs in the COLDz and ASPECS surveys, showing that at least some modern galaxy formation models can accommodate these observations.

(ix) Assuming a constant α_{CO} poorly approximates the $z = 0$ COLF for EAGLE and IllustrisTNG, while for SIMBA it is a decent approximation at $z = 0$ but not at $z = 2$. This shows that assuming a constant α_{CO} can significantly bias H₂ comparisons, and that the exact way in which it will be biased depends in detail on the distribution of α_{CO} within the simulated galaxy population. Comparing model predictions using the Accurso et al. (2017) α_{CO} COLF versus the Narayanan et al. (2012) one, we find that the former produces a

somewhat lower COLF at $z = 0$, and a dramatically lower one at $z = 2$ where its predicted α_{CO} values are dubiously high.

(x) Gas fractions drop with M_* and rise with sSFR in all models, with slopes generally as seen in the data. All models broadly reproduce the H I and H₂ fractions but with some discrepancies such as SIMBA being too high at low masses, IllustrisTNG too high at high masses, and EAGLE slightly low overall, though EAGLE-Recal shows very good agreement.

(xi) Comparing to $L_{\text{CO}(1-0)}$ scalings, EAGLE produces reasonable agreement, SIMBA is too high particularly at low M_* and sSFR, and IllustrisTNG may be too high at high M_* .

(xii) Scaling relations versus stellar surface density μ_* generally show good agreement for IllustrisTNG and poorer agreement for EAGLE and SIMBA, which may in part reflect IllustrisTNG's superior numerical resolution for modelling the stellar surface density, but may also reflect failings of the feedback models. EAGLE-Recal shows higher μ_* values than EAGLE, which corroborates at least some of the discrepancy owing to numerical resolution.

(xiii) Comparing AGN feedback variants in SIMBA shows that it is SIMBA's jet mode feedback that is responsible for suppressing cold gas mass functions at the massive end from $z \sim 1 \rightarrow 0$; without jet feedback, the mass functions grow with time.

(xiv) Jet AGN feedback in SIMBA is further responsible for creating a strong anticorrelation between molecular gas fraction and stellar mass; without this, SIMBA predicts roughly constant molecular gas fractions. Concurrently, jet feedback seems to over-suppress the formation of the highest stellar surface density objects, which could be responsible for SIMBA's discrepancies versus these observations.

(xv) Besides α_{CO} and numerical resolution, the dominant systematic for comparisons to observations is the choice of aperture, particularly for H₂ in massive galaxies. While this might be partly mitigated by carefully mocking observations, the H₂ formation threshold and ISM pressurization used in all these simulations intrinsically limit how well the ISM structure can be modelled at cosmological resolutions. Aperture effects are unlikely to qualitatively change our results, but for robust quantitative predictions it will be necessary to model such systematics more carefully.

These results illustrate how cold gas content and its evolution provide strong constraints on key galaxy growth and feedback processes in current models. Even just the mass functions of atomic and molecular gas and their evolution are already qualitatively different in our state-of-the-art simulations that have very similar stellar properties. While H I provides a more straightforward comparison, current radio telescopes provide no constraints at $z \gg 0$, although this will hopefully improve soon with now-online Square Kilometre Array precursors MeerKAT and ASKAP. The molecular gas content is potentially even more constraining, but current comparisons are strongly systematics limited in terms of assumptions regarding α_{CO} that connects the simulated H₂ mass to CO luminosity. More attention should be given to modelling this quantity, which may depend sensitively on many subgrid aspects of ISM modelling. None the less, the large differences between the evolutionary trends in models suggest that even broad constraints on α_{CO} could provide substantial discriminatory power. This would unlock the full potential of the ALMA data to employ the evolution of molecular gas as a constraint on galaxy formation models. The next generation of far-infrared and radio facilities promises to provide novel and complementary constraints on galaxy formation models that will be crucial for building a more comprehensive scenario of galaxy evolution within a cosmological context.

ACKNOWLEDGEMENTS

The authors acknowledge helpful discussions with Daniel Anglés-Alcázar, Sarah Appleby, Katarina Kraljic, and Daniele Sorini, as well as helpful comments from the referee. We thank Robert Thompson for developing CAESAR, and the YT team for development and support of YT. RD acknowledges support from the Wolfson Research Merit Award program of the U.K. Royal Society. RAC is a Royal Society University Research Fellow. ARHS is the Jim Buckee Fellow at UWA. DN is supported in part by NSF Award AST-1715206 and *HST* Theory Award 15043.0001. LC is the recipient of an Australian Research Council Future Fellowship (FT180100066) funded by the Australian Government. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility. The equipment was funded by BEIS capital funding via STFC capital grants ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. Parts of this research were conducted by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013. This research was supported by the Munich Institute for Astro- and Particle Physics (MIAPP) that is funded by the Deutsche Forschungsgemeinschaft (DFG; German Research Foundation) under Germany's Excellence Strategy – EXC-2094-390783311.

DATA AVAILABILITY STATEMENT

The data underlying this work, including simulation snapshots and galaxy catalogues, will be shared on reasonable request to the corresponding author R. Davé.

REFERENCES

- Accurso G. et al., 2017, *MNRAS*, 470, 4750
- Anglés-Alcázar D., Davé R., Faucher-Giguère C.-A., Özel F., Hopkins P. F., 2017a, *MNRAS*, 464, 2840
- Anglés-Alcázar D., Faucher-Giguère C.-A., Kereš D., Hopkins P. F., Quataert E., Murray N., 2017b, *MNRAS*, 470, 4698
- Appleby S., Davé R., Kraljic K., Anglés-Alcázar D., Narayanan D., 2020, *MNRAS*, 494, 6053
- Aravena M. et al., 2019, *ApJ*, 882, 136
- Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, 47, 481
- Bahé Y. M. et al., 2016, *MNRAS*, 456, 1115
- Barnes D. G. et al., 2001, *MNRAS*, 322, 486
- Bernardi M., Meert A., Sheth R. K., Fischer J.-L., Huertas-Company M., Maraston C., Shankar F., Vikram V., 2017, *MNRAS*, 467, 2217
- Bernardi M. et al., 2018, *MNRAS*, 475, 757
- Blitz L., Rosolowsky E., 2006, *ApJ*, 650, 933
- Blyth S. et al., 2016, in *Proc. Sci. Vol. 277*, MeerKAT Science: On the Pathway to the SKA. p. 4
- Bolatto A. D., Wolfire M., Leroy A. K., 2013, *ARA&A*, 51, 207
- Bondi H., 1952, *MNRAS*, 112, 195
- Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
- Borrow J., Anglés-Alcázar D., Davé R., 2020, *MNRAS*, 491, 6102
- Bouché N. et al., 2010, *ApJ*, 718, 1001
- Catinella B., Cortese L., 2015, *MNRAS*, 446, 3526
- Catinella B. et al., 2010, *MNRAS*, 403, 683
- Catinella B. et al., 2012, *A&A*, 544, A65
- Catinella B. et al., 2013, *MNRAS*, 436, 34
- Catinella B. et al., 2018, *MNRAS*, 476, 875
- Cen R., 1992, *ApJS*, 78, 341
- Chabrier G., 2003, *PASP*, 115, 763
- Choi E., Ostriker J. P., Naab T., Johansson P. H., 2012, *ApJ*, 754, 125

- Christiansen J. F., Davé R., Sorini D., Anglés-Alcázar D., 2019, preprint ([arXiv:1911.01343](https://arxiv.org/abs/1911.01343))
- Crain R. A. et al., 2015, *MNRAS*, 450, 1937
- Crain R. A. et al., 2017, *MNRAS*, 464, 4204
- Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
- Davé R., Finlator K., Oppenheimer B. D., 2011, *MNRAS*, 416, 1354
- Davé R., Finlator K., Oppenheimer B. D., 2012, *MNRAS*, 421, 98
- Davé R., Katz N., Oppenheimer B. D., Kollmeier J. A., Weinberg D. H., 2013, *MNRAS*, 434, 2645
- Davé R., Thompson R., Hopkins P. F., 2016, *MNRAS*, 462, 3265
- Davé R., Rafieferantsoa M. H., Thompson R. J., Hopkins P. F., 2017, *MNRAS*, 467, 115
- Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
- Davis T. A., Greene J. E., Ma C.-P., Blakeslee J. P., Dawson J. M., Pandya V., Veale M., Zabel N., 2019, *MNRAS*, 486, 1404
- Diemer B. et al., 2018, *ApJS*, 238, 33
- Diemer B. et al., 2019, *MNRAS*, 487, 1529
- 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
- Elson E. C., Baker A. J., Blyth S. L., 2019, *MNRAS*, 486, 4894
- Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, *ApJ*, 703, 1416
- Fernández X. et al., 2016, *ApJ*, 824, L1
- Finlator K., Davé R., 2008, *MNRAS*, 385, 2181
- Fletcher T. J., Saintonge A., Soares P. S., Pontzen A., 2020, preprint ([arXiv:2002.04959](https://arxiv.org/abs/2002.04959))
- Furlong M. et al., 2015, *MNRAS*, 450, 4486
- Genel S. et al., 2014, *MNRAS*, 445, 175
- Giovanelli R. et al., 2005, *AJ*, 130, 2598
- Gnedin N. Y., Draine B. T., 2014, *ApJ*, 795, 37
- Gnedin N. Y., Kravtsov A. V., 2011, *ApJ*, 728, 88
- Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*. p. 64
- Haardt F., Madau P., 2012, *ApJ*, 746, 125
- Hahn O., Abel T., 2011, *MNRAS*, 415, 2101
- Haynes M. P. et al., 2018, *ApJ*, 861, 49
- Heckman T. M., Best P. N., 2014, *ARA&A*, 52, 589
- Hirschmann M., Dolag K., Saro A., Bachmann L., Borgani S., Burkert A., 2014, *MNRAS*, 442, 2304
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hopkins P. F., Quataert E., 2011, *MNRAS*, 415, 1027
- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Janowiecki S., Catinella B., Cortese L., Saintonge A., Brown T., Wang J., 2017, *MNRAS*, 466, 4795
- Jenkins A., 2010, *MNRAS*, 403, 1859
- Jones M. G., Haynes M. P., Giovanelli R., Moorman C., 2018, *MNRAS*, 477, 2
- Katz N., Weinberg D. H., Hernquist L., 1996, *ApJS*, 105, 19
- Kennicutt R. C., Jr, 1998, *ApJ*, 498, 541
- Keres D., Yun M. S., Young J. S., 2003, *ApJ*, 582, 659
- Krumholz M. R., Gnedin N. Y., 2011, *ApJ*, 729, 36
- Kurczynski P. et al., 2016, *ApJ*, 820, L1
- Lagos C. d. P. et al., 2015, *MNRAS*, 452, 3815
- Lagos C. d. P. et al., 2016, *MNRAS*, 459, 2632
- Lilly S. J., Carollo C. M., Pipino A., Renzini A., Peng Y., 2013, *ApJ*, 772, 119
- McAlpine S. et al., 2016, *Astron. Comput.*, 15, 72
- Maiolino R. et al., 2012, *MNRAS*, 425, L66
- Marasco A., Crain R. A., Schaye J., Bahé Y. M., van der Hulst T., Theuns T., Bower R. G., 2016, *MNRAS*, 461, 2630
- Marinacci F. et al., 2018, *MNRAS*, 480, 5113
- Meyer M. J. et al., 2004, *MNRAS*, 350, 1195
- Miralda-Escudé J., 2005, *ApJ*, 620, L91
- Naiman J. P. et al., 2018, *MNRAS*, 477, 1206
- Narayanan D., Krumholz M. R., Ostriker E. C., Hernquist L., 2012, *MNRAS*, 421, 3127
- Nelson D. et al., 2018, *MNRAS*, 475, 624
- Nelson D. et al., 2019, *Comput. Astrophys. Cosmol.*, 6, 2
- Obreschkow D., Croton D., De Lucia G., Khochfar S., Rawlings S., 2009, *ApJ*, 698, 1467
- Olsen K., Greve T. R., Narayanan D., Thompson R., Davé R., Niebla Rios L., Stawinski S., 2017, *ApJ*, 846, 105
- Pakmor R., Springel V., 2013, *MNRAS*, 432, 176
- Pavesi R. et al., 2018, *ApJ*, 864, 49
- Pawlik A. H., Schaye J., 2008, *MNRAS*, 389, 651
- Perna M., Lanzuisi G., Brusa M., Cresci G., Mignoli M., 2017, *A&A*, 606, A96
- Pillepich A. et al., 2018a, *MNRAS*, 473, 4077
- Pillepich A. et al., 2018b, *MNRAS*, 475, 648
- Planck Collaboration XVI, 2014, *A&A*, 571, A16
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Popping A., Davé R., Braun R., Oppenheimer B. D., 2009, *A&A*, 504, 15
- Popping G. et al., 2019, *ApJ*, 882, 137
- Rahmati A., Pawlik A. H., Raicevic M., Schaye J., 2013, *MNRAS*, 430, 2427
- Riechers D. A. et al., 2019, *ApJ*, 872, 7
- Rodríguez Montero F., Davé R., Wild V., Anglés-Alcázar D., Narayanan D., 2019, *MNRAS*, 490, 2139
- Rosas-Guevara Y. M. et al., 2015, *MNRAS*, 454, 1038
- Saintonge A. et al., 2011, *MNRAS*, 415, 32
- Saintonge A. et al., 2012, *ApJ*, 758, 73
- Saintonge A. et al., 2017, *ApJS*, 233, 22
- Schaye J., 2004, *ApJ*, 609, 667
- Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schmidt M., 1959, *ApJ*, 129, 243
- Smith B. D. et al., 2017, *MNRAS*, 466, 2217
- Somerville R. S., Davé R., 2015, *ARA&A*, 53, 51
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., 2010, *MNRAS*, 401, 791
- 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., 2005, *MNRAS*, 361, 776
- Springel V. et al., 2018, *MNRAS*, 475, 676
- Stevens A. R. H., Martig M., Croton D. J., Feng Y., 2014, *MNRAS*, 445, 239
- Stevens A. R. H. et al., 2019a, *MNRAS*, 483, 5334
- Stevens A. R. H., Diemer B., Lagos C. d. P., Nelson D., Obreschkow D., Wang J., Marinacci F., 2019b, *MNRAS*, 490, 96
- Sturm E. et al., 2011, *ApJ*, 733, L16
- Tacconi L. J. et al., 2013, *ApJ*, 768, 74
- Thomas N., Davé R., Anglés-Alcázar D., Jarvis M., 2019, *MNRAS*, 487, 5764
- Tumlinson J., Peebles M. S., Werk J. K., 2017, *ARA&A*, 55, 389
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
- Weinberger R. et al., 2017, *MNRAS*, 465, 3291
- Whitton I. H., Prescott M., McAlpine K., Jarvis M. J., Heywood I., 2018, *MNRAS*, 480, 358
- 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 O. I. et al., 2006, *MNRAS*, 371, 1855
- Wright A. H. et al., 2017, *MNRAS*, 470, 283

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