



THE UNIVERSITY *of* EDINBURGH

## Edinburgh Research Explorer

### **MUFASA: Galaxy star formation, gas, and metal properties across cosmic time**

**Citation for published version:**

Davé, R, Rafieferantsoa, MH, Thompson, RJ & Hopkins, PF 2017, 'MUFASA: Galaxy star formation, gas, and metal properties across cosmic time', *Monthly Notices of the Royal Astronomical Society*, vol. 467, no. 1, pp. 115-132. <https://doi.org/10.1093/mnras/stx108>

**Digital Object Identifier (DOI):**

[10.1093/mnras/stx108](https://doi.org/10.1093/mnras/stx108)

**Link:**

[Link to publication record in Edinburgh Research Explorer](#)

**Document Version:**

Publisher's PDF, also known as Version of record

**Published In:**

Monthly Notices of the Royal Astronomical Society

**General rights**

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [openaccess@ed.ac.uk](mailto:openaccess@ed.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# MUFASA: galaxy star formation, gas and metal properties across cosmic time

Romeel Davé,<sup>1,2,3★</sup> Mika H. Rafieferantsoa,<sup>1,4</sup> Robert J. Thompson<sup>1,5</sup>  
and Philip F. Hopkins<sup>6</sup>

<sup>1</sup>University of the Western Cape, Bellville, Cape Town 7535, South Africa

<sup>2</sup>South African Astronomical Observatories, Observatory, Cape Town 7925, South Africa

<sup>3</sup>African Institute for Mathematical Sciences, Muizenberg, Cape Town 7945, South Africa

<sup>4</sup>Max-Planck-Institut für Astrophysik D-85741, Garching, Germany

<sup>5</sup>National Center for Supercomputing Applications, Champaign-Urbana, IL 61801, USA

<sup>6</sup>California Institute of Technology, Pasadena, CA 91125, USA

Accepted 2017 January 12. Received 2017 January 5; in original form 2016 October 1

## ABSTRACT

We examine galaxy star formation rates (SFRs), metallicities and gas contents predicted by the MUFASA cosmological hydrodynamic simulations, which employ meshless hydrodynamics and novel feedback prescriptions that yield a good match to observed galaxy stellar mass assembly. We combine 50, 25 and 12.5  $h^{-1}$  Mpc boxes with a quarter billion particles each to show that MUFASA broadly reproduces a wide range of relevant observations, including SFR and specific SFR functions, the mass–metallicity relation, H I and H<sub>2</sub> fractions, H I (21 cm) and CO luminosity functions, and cosmic gas density evolution. There are mild but significant discrepancies, such as perhaps too many high-SFR galaxies, overly metal-rich and H I-poor galaxies at  $M_* \gtrsim 2 \times 10^{10} M_\odot$ , and specific star formation rates that are too low at  $z \sim 1$ –2. The H I mass function increases by  $\times 2$  out to  $z \sim 1$ , then steepens to higher redshifts, while the CO luminosity function computed using the Narayanan et al. conversion factor shows a rapid increase of CO-bright galaxies out to  $z \sim 2$  in accord with data.  $\Omega_{\text{HI}}$  and  $\Omega_{\text{H}_2}$  both scale roughly as  $\propto (1+z)^{0.7}$  out to  $z \sim 3$ , comparable to the rise in H I and H<sub>2</sub> fractions. MUFASA galaxies with high SFR at a given  $M_*$  have lower metallicities and higher H I and H<sub>2</sub> fractions, following observed trends; we make quantitative predictions for how the fluctuations in the baryon cycle drive correlated scatter around galaxy scaling relations. Most of these trends are well converged with numerical resolution. These successes highlight MUFASA as a viable platform to study many facets of cosmological galaxy evolution.

**Key words:** methods: numerical – galaxies: abundances – galaxies: evolution – galaxies: formation – galaxies: ISM – galaxies: star formation.

## 1 INTRODUCTION

Observations of galaxy properties from today back to the early Universe are improving at a remarkable pace, thanks to advancing multiwavelength photometric and spectroscopic galaxy surveys. Progress has been particularly impressive in the near-infrared (IR) and longer wavelengths, which provides more robust constraints on stellar and metal content at high redshifts and gas content across all redshifts. Models for galaxy formation thus find it increasingly challenging to be able to reproduce such observations within a physically motivated concordance cosmology framework.

Recent cosmological hydrodynamic simulations have been impressively successful at broadly reproducing key galaxy demographic observables over cosmic time (see Somerville & Davé 2015, and references therein). A primary benchmark used to test galaxy formation models is the observed galaxy stellar mass function (GSMF). Many modern simulations can now match this to within a factor of several over the majority of cosmic time and mass (Davé et al. 2013; Genel et al. 2014; Crain et al. 2015; Davé, Thompson & Hopkins 2016; Kaviraj et al. 2016), which is typically within the range of current systematic uncertainties in the data. To do so, all cosmological-scale simulations incorporate heuristic models for feedback processes associated with star formation that suppress galaxy formation at the low-mass end, combined with feedback often associated with active galactic nuclei that suppresses massive galaxy growth. However, the precise physical mechanisms invoked

\* E-mail: romeeld@gmail.com

for feedback can vary substantially amongst simulations, despite their predicted GSMFs being similar. To further test and discriminate between models, and thereby constrain the physical mechanisms giving rise to feedback, it is thus important to move beyond the GSMF and consider other aspects of galaxy demographics.

Advancing multiwavelength observations have made impressive progress at characterizing the gas and metal content of galaxies across cosmic time. Metallicity measures at higher redshifts have been aided by new near-IR spectroscopic capabilities that have enabled the same optical emission line measures used at low redshifts to be applied to  $z \sim 2\text{--}3$  galaxies (Steidel et al. 2014; Sanders et al. 2015). Molecular gas contents have now been measured out to similar redshifts thanks to deep millimetre-wave data that can detect redshifted carbon monoxide (CO) emission lines (Geach et al. 2011; Tacconi et al. 2013). Direct measures of atomic gas (H I) remain confined to low redshifts ( $z \lesssim 0.5$ ) as of yet owing to the sensitivity of current instruments (Delhaize 2013; Fernández et al. 2016), but the Square Kilometre Array (SKA) and its precursors such as MeerKAT aim to probe H I out to  $z \sim 1$  and beyond (e.g. Holwerda, Blyth & Baker 2012). These observations provide a direct glimpse into the gaseous fuel for star formation, as well as products of massive star formation as traced by chemical enrichment, hence they can more directly probe the baryon cycle of gaseous inflows and outflows that are viewed as being the central driver of cosmological galaxy evolution.

Cosmological galaxy formation simulations have utilized these observations to provide additional constraints on feedback mechanisms and other physical processes of galaxy formation (e.g. Vogelsberger et al. 2014; Schaye et al. 2015; Davé et al. 2016). For instance, the slope of the mass–metallicity relation strongly suggests that low-mass galaxies preferentially eject more of their gas in outflows versus forming it into stars (e.g. Finlator & Davé 2008). The high gas fraction in low-mass galaxies is likewise a reflection of strong outflows that prevents the gas from forming into stars (e.g. Davé, Finlator & Oppenheimer 2011). This broadly agrees with the notion that low-mass galaxies must have stronger feedback in order to suppress the faint end of the GSMF (e.g. Somerville et al. 2008; Davé et al. 2011). While these trends generally point towards a qualitatively similar picture (Somerville & Davé 2015), it remains highly challenging for a single model to quantitatively reproduce all the relevant observed relations across a wide range of mass scales and cosmic epochs.

In this paper, we present a further analysis on the suite of cosmological hydrodynamic simulations of galaxy formation using GIZMO, called the MUFASA simulations, introduced in Davé et al. (2016, hereafter Paper I). MUFASA uses updated state of the art feedback modules, including two-phase kinetic outflows with scalings taken from the Feedback In Realistic Environments (FIRE) simulations (Muratov et al. 2015), an evolving halo mass-based quenching scheme (Gabor & Davé 2015; Mitra, Davé & Finlator 2015), 11-element chemical evolution and molecular gas-based star formation (Krumholz, McKee & Tumlinson 2011; Thompson et al. 2014). We run three volumes, each with  $512^3$  dark matter particles and  $512^3$  gas elements, having box sizes of 50, 25 and  $12.5 h^{-1}$  Mpc, in order to cover halo masses from  $\sim 10^{10}\text{--}10^{14} M_{\odot}$  and stellar masses from  $\sim 10^7$  to  $10^{12} M_{\odot}$ .

In Paper I, we showed that MUFASA does an excellent job at reproducing the observed evolution of the GSMF over most of cosmic time. Here we compare MUFASA to a wider suite of observations encompassing galaxy star formation rates (SFRs), gas and metal content, in order to quantitatively examine whether a model that accurately reproduces stellar mass growth can also match these independent properties. One significant discrepancy seen in Paper I

was that specific SFRs (sSFRs) at  $z \sim 1\text{--}2$  were well below the observations, even though galaxy growth rates as measured by GSMF evolution seemed to be in accord with the data. Here we further investigate this issue using SFR and sSFR functions over cosmic time. Since MUFASA directly tracks  $H_2$  within galaxies using a sub-grid prescription (Krumholz et al. 2011), we investigate H I and  $H_2$  contents separately, along with their evolution. Galaxy metallicities provide a crucial barometer for feedback, so we compare our predictions to emerging observations out to Cosmic Noon. Simulations naturally predict that deviations from the mean galaxy scaling relations are correlated in that galaxies at a given stellar mass that are high in SFR are also low in metallicity (Davé et al. 2011) and gas content (Rafieferantsoa et al. 2015). Here we generalize this analysis across all quantities considered, showing that the deviations from the mean relations in SFR, metallicity, H I, and  $H_2$  versus  $M_*$  are all correlated, and we quantify these correlations.

Taken together, these results extend the overall success of the MUFASA simulations as a reasonably faithful reproduction of the real universe, thereby highlighting MUFASA’s utility as a platform to study of the physics of galaxy evolution across cosmic time. This paper is outlined as follows: In Section 2, we briefly recap the key ingredients of the MUFASA simulations. Section 3 discusses predicted SFRs and sSFRs, Section 4 presents the mass–metallicity relation, and Section 5 shows gas fractions and gas mass functions. In Section 6, we quantify the second-parameter dependences of the scatter around key scaling relations. We summarize our findings in Section 7.

## 2 SIMULATION DESCRIPTION

We employ a modified version of the gravity plus hydrodynamics solver GIZMO (Hopkins 2015), which uses the GADGET-3 gravity solver (Springel 2005), along with the meshless finite mass (MFM) hydrodynamics solver. We use adaptive gravitational softening throughout for all particles (Hopkins 2015), with a minimum (Plummer-equivalent) softening length set to 0.5 per cent of the mean interparticle spacing. For more details on these aspects as well as the feedback choices summarized below, see Paper I.

We include radiative cooling from primordial (non-equilibrium ionization) and heavy elements (equilibrium ionization) using the GRACKLE-2.1 chemistry and cooling library (Bryan et al. 2014; Kim et al. 2014). A spatially uniform photoionizing background is assumed, namely the 2011 update of the determination in Faucher-Giguere et al. (2009). Gas above a threshold density is assumed to have an equation of state given by  $T \propto \rho^{1/3}$  (Schaye & Dalla Vecchia 2008), and for the primary run employed in this paper, the threshold density is taken to be  $0.13 \text{ cm}^{-3}$ . Stars are formed using a molecular gas-based prescription following Krumholz, McKee & Tumlinson (2009), which approximates the  $H_2$  fraction based on the local density, the Sobolev approximation in which the optical depth is given by  $\rho/|\nabla\rho|$ , where  $\rho$  is the particle’s density, and the particle’s metallicity scaled to solar abundance based on Asplund et al. (2009). We vary the assumed clumping factor with resolution, as described in Paper I.

Young stellar feedback is modelled using decoupled, two-phase winds. Winds are ejected stochastically, with a probability that is  $\eta$  times the SFR probability. The formula for  $\eta$  is taken to be the best-fitting relation from the FIRE suite of zoom simulations Muratov et al. (2015), namely

$$\eta = 3.55 \left( \frac{M_*}{10^{10} M_{\odot}} \right)^{-0.351}, \quad (1)$$

where  $M_*$  is the galaxy stellar mass determined using an on-the-fly friends-of-friends galaxy finder. The ejection velocity  $v_w$  scaling is also taken to follow that predicted by FIRE, but with a somewhat higher amplitude:

$$v_w = 2 \left( \frac{v_c}{200 \text{ km s}^{-1}} \right)^{0.12} v_c + \Delta v_{0.25}, \quad (2)$$

where  $v_c$  is the galaxy circular velocity estimated from the friends-of-friends baryonic mass, and  $\Delta v_{0.25}$  accounts for the potential difference between the launch location and one-quarter of the virial radius where Muratov et al. (2015) measured the scalings from FIRE. Winds are also ejected with a random 30 per cent fraction being ‘hot’, namely at a temperature set by the difference between the supernova energy and the wind launch energy (if this is positive), with the remaining 70 per cent launched at  $\ll 10^4 \text{ K}$ . Wind fluid elements are allowed to travel without hydrodynamic forces or cooling until the time when its relative velocity versus surrounding (non-wind) gas is less than 50 per cent of the local sound speed, or alternatively if it reaches limits in density of 0.01 times the critical density for star formation, or a time given by 2 per cent of the *Hubble* time at launch. We further include energy Type Ia supernovae (SNIa) and asymptotic giant branch (AGB) stars, implemented as a delayed component using stellar evolution as tracked by Bruzual & Charlot (2003) models with a Chabrier (2003) initial mass function (IMF). See Paper I for full details.

Chemistry is tracked for hydrogen, helium and nine metals: C, N, O, Ne, Mg, Si, S, Ca and Fe, comprising over 90 per cent of metal mass in the universe. Type II SN yields are taken from Nomoto et al. (2006), parametrized as a function of metallicity, which we multiplied by 0.5 in order to more closely match observed galaxy metallicities. Type II yields are added instantaneously to every star-forming gas particles at every time-step, based on its current SFR. For SNIa yields, we employ the yields from Iwamoto et al. (1999), assuming that each SNIa yields  $1.4 M_\odot$  of metals. For AGB stars, we employ enrichment as a function of age and metallicity from various sources as described in Oppenheimer & Davé (2008), further assuming a 36 per cent helium fraction and a nitrogen yield of 0.001 18. The enrichment, like the energy, is added from stars to the nearest 16 gas particles, kernel-weighted, following the mass-loss rate as computed assuming a Chabrier (2003) IMF.

We note that ISM gas ejected from our simulated galaxies is done so without any modification to its metallicity. We do not employ a separate ‘metal loading factor’ parameter (i.e. the metallicity of the ejected gas relative to the ISM metallicity) that preferentially ejects enriched (or de-enriched, as in Illustris; Vogelsberger et al. 2014) ISM material; in other words, we assume a metal loading factor of unity. The physical justification for this is that particularly in low-mass galaxies where the mass loading factor  $\eta$  is high, direct supernovae ejectae represent only a very small portion of the total outflowing material; hence, it makes sense that the outflow metallicity is dominated by ambient ISM gas (surrounding the launch site). In higher mass galaxies where  $\eta$  is low, this assumption can break down and it may be more appropriate to include a metal loading factor greater than unity. Without more detailed modelling, it is difficult to determine exactly what the appropriate metal loading factor is, so we eschew this complication for the present. Note that the FIRE simulations with self-consistently generated outflows find metal loading factors around unity for all galaxies (Ma et al. 2016), supporting our assumption.

To quench massive galaxies, we employ an on-the-fly halo mass-based quenching scheme that follows Gabor & Davé (2012, 2015). Above a halo quenching mass  $M_q$ , we maintain all halo gas at a

temperature above the system virial temperature, by continuously adding heat. This is intended to mimic the effects of ‘radio mode’ or ‘jet mode’ quenching (Croton et al. 2006), where jets inflate superbubbles in surrounding hot gas that approximately sphericalizes the jet energy and counteracts gas cooling (McNamara & Nulsen 2007). We only add heat to gas that is not self-shielded, defined as having a neutral (atomic+molecular) fraction above 10 per cent after applying a self-shielding correction following Rahmati et al. (2013). We take  $M_q$  as determined from the analytic ‘equilibrium model’ constraints required to match the observed evolution of the galaxy population from  $z = 0$  to 2 (Mitra et al. 2015), namely

$$M_q = (0.96 + 0.48z) \times 10^{12} M_\odot. \quad (3)$$

As demonstrated in Paper I, this evolving quenching mass is nicely consistent with the observations during early epochs ( $z \sim 2$ ) and today, while providing a sharp turnover in the stellar mass function at late epochs that closely matches the observations.

Paper I focused on the  $50 h^{-1} \text{ Mpc}$  MUFASA simulation using  $512^3$  gas fluid elements (i.e. mass-conserving cells),  $512^3$  dark matter particles and  $0.5 h^{-1} \text{ kpc}$  minimum softening length. Table 1 of Paper I lists the details for two higher resolution runs with the identical input physics and number of particles, having box sizes of 25 and  $12.5 h^{-1} \text{ Mpc}$ , and proportionally smaller softening lengths. At that time, these simulations were only evolved to  $z = 2$ , but since then we have evolved the  $25 h^{-1} \text{ Mpc}$  volume to  $z = 0$  and the  $12.5 h^{-1} \text{ Mpc}$  run to  $z = 1$ . We will use these to extend the dynamic range of our predictions and to test resolution convergence.

We generate initial conditions at  $z = 249$  using MUSIC (Hahn & Abel 2011) assuming a cosmology consistent with Planck Collaboration XIII (2016) ‘full likelihood’ constraints:  $\Omega_m = 0.3$ ,  $\Omega_\Lambda = 0.7$ ,  $\Omega_b = 0.048$ ,  $H_0 = 68 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ,  $\sigma_8 = 0.82$  and  $n_s = 0.97$ . We output 135 snapshots down to  $z = 0$  (105 to  $z = 1$ ). For these parameters, the 32-star particle galaxy mass resolution limit is  $5.8 \times 10^8 M_\odot$  for the fiducial  $50 h^{-1} \text{ Mpc}$  volume (as listed in table 1 of Paper I), and a factor of 8 and 64 smaller for the two smaller volumes.

We analyse the snapshots using SPHGR<sup>1</sup> (Thompson 2015) that identifies galaxies using SKID and haloes using ROCKSTAR (Behroozi, Wechsler & Conroy 2013), links them via their positions and outputs a catalogue of properties required for all the analyses in this paper. SKID identifies bound groups of stars and cold ( $T < 10^{4.5} \text{ K}$ ), dense (i.e. above the threshold density for star formation) gas. The stellar mass is then the total mass of all star particles in each galaxy, while the  $H_2$  mass is the sum over the  $H_2$  fraction times the gas mass of gas particles. We will discuss the computation of  $H_1$  in Section 5.1, as this requires some post-processing for self-shielding. The SFR is computed as the sum of all gas particles’ SFRs in each galaxy, and thereby represents an instantaneous measure. The metallicity is taken to be the SFR-weighted metallicity among all gas particles, which most closely mimics the weighting associated with emission-line measures of gas-phase metallicity to which we will compare. There is no aperture or surface brightness cut applied, so there may be some minor systematics when comparing to the observations, but we do not expect this to affect any of our conclusions. Also, we consider all galaxies except in some cases such as gas-phase metallicity measures where only star-forming galaxies are utilized; we leave a study of the central versus satellite population for future work.

<sup>1</sup> <http://sphgr.readthedocs.org/en/latest/>

### 3 STAR FORMATION RATES


Paper I compared MUFASA to the evolution of the stellar mass function, showing general agreement with the growth of the stellar content of galaxies across much of cosmic time. However, it also reiterated a longstanding discrepancy in predictions of sSFRs at a given  $M_*$ , i.e. the main sequence, during the peak epoch of cosmic star formation, in which simulated galaxies have  $\sim \times$  two to three lower SFRs compared to the observations at  $z \sim 2$ . This is a generic issue in cosmological models, both hydrodynamic and semi-analytic (e.g. Davé 2008; Genel et al. 2014; Sparre et al. 2015; Furlong et al. 2015; Somerville & Davé 2015), which could reflect an issue with the formation history of low-mass galaxies (e.g. Weinmann et al. 2012; White, Somerville & Ferguson 2015), though MUFASA seems to reproduce this at least qualitatively (Paper I). Here we explore the distribution of SFRs in more detail, by comparing MUFASA to two other SFR observables, namely the star formation rate function and the SSFR function.

#### 3.1 Star formation rate function

Fig. 1 shows SFRFs at  $z = 0, 1, 2$  from our suite of MUFASA simulations. The red solid, green dashed and blue dotted curves show the results from our 50, 25 and 12.5  $h^{-1}$  Mpc (at  $z \geq 1$ ) simulations. The hatched region shows cosmic variance as computed over the eight sub-octants within each simulation volume. The vertical dotted line indicates the typical SFR at the stellar mass resolution limit of 32 gas particle masses from a fit to the  $M_*$ -SFR relation; below this, the distribution of SFRs is expected to be significantly compromised by numerical resolution, and even above this SFR, there may be some galaxies that are affected by poor resolution owing to the scatter in the  $M_*$ -SFR relation. Hence, this line should be regarded as an approximate rather than a strict resolution limit. Indeed, one can see from comparing the various simulations' SFR functions at the same SFR that the lack of resolution convergence begins significantly above the dotted line, perhaps around  $0.5 M_\odot \text{yr}^{-1}$  at  $z = 0$  for the  $50 h^{-1}$  Mpc volume where it drops below the cosmic variance of the smaller volume. From this simulation's SFR- $M_*$  relation (fig. 6 of Paper I), this corresponds to a stellar mass of about  $10^{9.5} M_\odot$  for main sequence (i.e. star-forming) galaxies, well above our nominal stellar mass resolution limit.


Observations are shown in the various panels from H $\alpha$  luminosity functions, converted to SFR using the relation taken from Kennicutt (1998), adjusted for a Chabrier IMF. At  $z \sim 0$ , we show data from Bothwell et al. (2011, dotted black) and Gunawardhana et al. (2013, dashed black), at  $z \sim 1$  from Colbert et al. (2013), and at  $z \sim 2$  from Mehta et al. (2016). All these observations account for extinction based on considering H $\beta$  and sometimes more, but there is still uncertainty in such corrections.

At  $z = 0$ , the simulated SFRFs are in good agreement with Bothwell et al. (2011), but overpredict by up to  $\sim \times 3$  the more recent Gunawardhana et al. (2013) data from the Galaxy and Mass Assembly (GAMA) survey. The former uses panchromatic measures of SFR including the far-IR, while the latter is from extinction-corrected H $\alpha$  measures. It is beyond the scope of this paper to determine which observation is more accurate; panchromatic measures may be more robust from the perspective that they more accurately capture the bolometric output of young stars, yet the integrated SFR density from Bothwell et al. (2011) of  $0.025 M_\odot \text{yr}^{-1} \text{Mpc}^{-3}$  is noticeably above the compilation of SFRD observations from Madau & Dickinson (2014) at  $z \approx 0$ . If the GAMA data are more accurate, this would imply that there are several times more SFGs



**Figure 1.** Star formation rate functions at  $z = 0, 1, 2$  (top to bottom), in our suite of MUFASA simulations. Results for the 50, 25 and 12.5  $h^{-1}$  Mpc (at  $z \geq 1$ ) are shown in red solid, green dashed and blue dotted lines, respectively, with the hatched region showing the cosmic variance computed over the eight sub-octants of the simulation volume. The vertical dotted line is an approximate resolution limit, taken as the mean SFR at our stellar mass resolution limit; there is likely some incompleteness even above this value, which can be seen by comparing the different volumes. Observations are shown as dashed black lines in each panel, using H $\alpha$  luminosity functions converted to SFRFs from Gunawardhana et al. (2013,  $z \approx 0$ ), Colbert et al. (2013,  $z \approx 2$ ) and Mehta et al. (2016,  $z \approx 2$ ).

with  $\text{SFR} \sim 1\text{--}10 M_\odot \text{yr}^{-1}$  in MUFASA than in the real Universe. However, it is possible that H $\alpha$  surveys miss the most highly star-forming galaxies since they are typically highly obscured. This discrepancy is consistent with the finding in Paper I (see their fig. 3) that the cosmic SFR density is overpredicted by  $\sim 50$  per cent at  $z = 0$  in MUFASA compared to Madau & Dickinson (2014), although it is in better agreement with the Bothwell et al. (2011) measurement. Hence MUFASA's predictions for the SFRF today are broadly consistent with data, although perhaps overpredicting galaxies with SFRs comparable to or exceeding that of the Milky Way.



**Figure 2.** SSFR functions at  $z = 0.25, 1, 2$  (top to bottom rows) in our  $50$  (black lines) and  $25 h^{-1}$  Mpc (grey lines) MUFASA simulations, in four bins of increasing stellar mass (left to right). Hatched regions show the cosmic variance computed over the eight sub-octants in the volume. Observations from Ilbert et al. (2015) are shown, which include data from COSMOS at  $z = 0.2\text{--}0.4$  (blue in upper panels), and in the middle panels from COSMOS at  $z = 0.8\text{--}1$  (blue), COSMOS at  $z = 1\text{--}1.2$  (green), GOODS from  $z = 0.8$  to  $1$  (red) and  $1$  to  $1.2$  (magenta). The predicted sSFR functions match the observations very well at  $z \sim 0.25$ , showing that MUFASA reproduces the distribution of sSFRs quite well at low- $z$ . At  $z = 1$ , MUFASA matches well the shape of the distribution but is shifted to slightly lower sSFR. This indicates that MUFASA is properly capturing the physical causes of the fluctuations around the main sequence, as well as the number of galaxies transitioning to quiescence.

At  $z = 1$ , the predicted SFRF is similar to that at  $z = 0$  at low SFRs, but shows an excess at high SFRs, such that now we start to see galaxies with  $\text{SFR} \gtrsim 100 M_{\odot} \text{yr}^{-1}$  in our  $50 h^{-1}$  Mpc volume. The SFRF shows less of a strong truncation at high-SFR as it does at low- $z$ . Generally, MUFASA exceeds the observations at high SFRs. This trend continues at  $z = 2$ , with the low-SFR end mostly unevolving but the high-SFR end more heavily populated.

The discrepancy of overpredicting the SFRF is actually quite difficult to reconcile with the results from Paper I (and as generically found in cosmological simulations) that MUFASA underpredicts the SFR- $M_*$  relation at  $z \sim 1\text{--}2$ . The fact that MUFASA also well reproduces the  $z = 2$  GSMF (and if anything slightly underpredicts that as well at high  $M_*$ ) would suggest that the convolution of those two predictions should yield a clear underprediction of the SFRF. Yet the SFRF appears to be overpredicted at these epochs. At the risk of blaming observations for discrepancies with models, it appears that perhaps the  $\text{H}\alpha$  measures are systematically low at high- $z$ , which could occur if such samples are missing highly obscured galaxies that become more prevalent at high- $z$ . It is beyond the scope of this work to fully examine all the relevant systematics, but it highlights that leaving aside the models, there appears to be some consistency issues purely among the observational measures of SFRs during Cosmic Noon.

In summary, the SFRFs predicted by MUFASA generally show the observed shape from  $z = 0\text{--}2$ , though with an amplitude that may be somewhat too high at low redshifts depending on the comparison

data set. The general agreement is encouraging and may be within current systematic uncertainties in measuring a complete sample of star-forming galaxies across all these epochs. There is no obvious discrepancy in the SFRF at  $z = 2$  that would explain the discrepancy in the SFR- $M_*$  relation. Resolution convergence in the SFRFs between the various MUFASA volumes is reasonable, though not ideal, partly because our simulations yield a fixed stellar mass resolution that translates into a soft SFR resolution owing to the scatter in the SFR- $M_*$  relation.

### 3.2 SSFR function

A separate test of SFRs is whether our simulations reproduce the correct distribution of specific SFRs at a given stellar mass. Qualitatively, at high redshifts, the spread in sSFRs measures the fluctuations around the main sequence owing to inflow fluctuations (e.g. Mitra et al. 2017), while at lower redshifts, a substantial low-sSFR population appears corresponding to quenched galaxies. Matching the amplitude and evolution of the distribution of sSFRs in stellar mass bins is thus a stringent test of whether the predicted MUFASA galaxy population is in accord with the rate at which galaxies are fluctuating around the main sequence, and eventually quenched (e.g. Tacchella et al. 2016).

Fig. 2 shows the specific star formation rate function (sSFRF) in four bins of stellar mass in the range  $10^{9.5} < M_* < 10^{11.5} M_{\odot}$  (left to right), at  $z = 0.25, 1, 2$  (top to bottom). We will primarily

consider the  $50 h^{-1}$  Mpc volume here for clarity, particularly since we want to well sample the rate of galaxy quenching for which we prefer our largest volume containing the most massive haloes. However, we also show the  $25 h^{-1}$  Mpc volume to indicate the level of resolution convergence. Lines show the predicted sSFRF (black for  $50 h^{-1}$  Mpc, grey for  $25 h^{-1}$  Mpc), while the hatched region shows the cosmic variance computed among the eight sub-objects only for the  $50 h^{-1}$  Mpc case. Observations are shown from a compilation by Ilbert et al. (2015) at  $z = 0.2$ – $0.4$  and  $0.8$ – $1.2$ , from various sources as described in the caption, generally from extinction-corrected UV measures or spectral energy distribution fitting. Note that the observations only consider galaxies that have a measurable SFR that we mimic in our simulations by excluding galaxies with  $\log \text{sSFR} < -3$  (which would lie off this plot in any case).

At  $z = 0.25$ , the sSFRF shows a peak at the median sSFR within that  $M_*$  bin, a sharp truncation to higher sSFR and a broader extension to low sSFR corresponding to green valley galaxies. MUFASA provides a remarkably good match (i.e. within cosmic variance) to the observed sSFRF in every stellar mass bin. This is true in both the  $50$  and  $25 h^{-1}$  Mpc volumes, though here and at all redshifts, there is a slight systematic tendency for the smaller volume to shift the sSFR distribution to lower values, likely owing to the slightly increased stellar masses produced at higher resolution (Paper I). Statistically, both runs agree about equally well with the observations at this epoch.

This new test demonstrates that the scatter in sSFRs, and hence the fluctuations around the main sequence as well as the rate at which the green valley is being populated, is being well modelled in MUFASA. In particular, the amplitude and shape match in the most massive bin would suggest that MUFASA is not overproducing the number of galaxies with high sSFRs, even if Fig. 1 suggested that it might be doing so. These can be reconciled if MUFASA is producing a few too many massive galaxies, which is indeed a trend noted in the  $z = 0$  GSMF shown in Paper I, albeit with large cosmic variance.

At  $z = 1$ , the shape of the sSFRF is well reproduced, but there is clearly an offset in the distribution such that the predicted values are lower by  $\sim \times 2$ . This is simply reflecting the fact that the median sSFR is underproduced at this epoch, as shown in Paper I, continuing a trend generically seen in cosmological galaxy formation models. It appears that the discrepancy in the median sSFR is not reflective of the emergence of some new population of galaxies in the observations that do not appear in the models, but rather an overall systematic shift in the measured sSFR values at that epoch. We would expect that these trends would continue on to  $z = 2$ , but we do not know of sSFRFs published at this epoch.

Overall, MUFASA does an excellent job of reproducing the low- $z$  distribution of sSFRs, including the peak value, the sharp truncation to high sSFRs that highlights the rarity of starbursts locally and the gradual decline towards low-sSFR that reflects the population of galaxies likely in the process of quenching. There are still a non-trivial number of SFGs even at the highest masses in MUFASA, which is in agreement with the observations. This suggests that MUFASA does a good job reproducing the SFR fluctuations and quenching rate of galaxies, which provides some empirical support for the implemented subgrid models for star formation and quenching.

#### 4 METALLICITY

Chemical enrichment provides a key tracer for star formation and feedback activity in and around galaxies. Within a simple equilibrium or bathtub-type model, the mass–metallicity relation directly

reflects the mass-loss rate in outflows together with the recycling of previously ejected (enriched) material back into the ISM (e.g. Finlator & Davé 2008; Somerville & Davé 2015). Galaxy metallicities are thus a crucial test for how accurately a particular model is representing the baryon cycle.

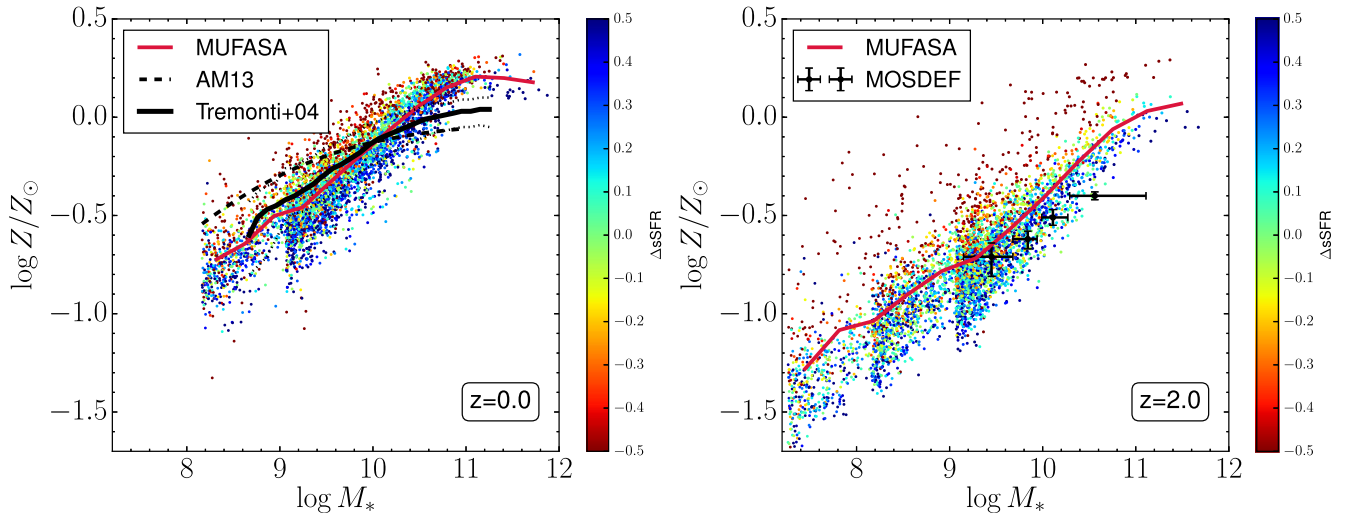
The stellar mass–gas phase metallicity relation (MZR) is one of the tightest observed correlations between any two galaxy properties, with a scatter typically around 0.1 dex (Tremonti et al. 2004). Unfortunately, calibration issues may add significant systematic uncertainties (Kewley & Ellison 2008), but, none the less, the shape of the MZR is likely to be reasonably robust even if the amplitude is less certain. In this section, we present predictions for the MZR from MUFASA, along with the comparisons to key observations at the present epoch and in the early Universe. Note that we only consider star-forming galaxies in this section, since we are measuring the gas-phase metallicities.

Fig. 3 shows the MZR at  $z = 0, 2$  (left-hand, right-hand panels) in our MUFASA simulation suite. At  $z = 2$ , we have overplotted all three volumes down to each of their galaxy stellar mass resolution limit; these are the three ‘groupings’ of points, with the  $12.5 h^{-1}$  Mpc volume extending to the lowest masses, and the  $50 h^{-1}$  Mpc volume dominating at high masses. At  $z = 0$ , we only have the  $50$  and  $25 h^{-1}$  Mpc volumes. The colour-coding shows the deviation in  $\log \text{SFR}$  for each galaxy off of the global  $M_*$ –SFR relation at that redshift (Paper I). The thick red line shows a running median for the combined sample of simulated galaxies; while we do not show the individual volumes’ medians separately, it is evident that the agreement between them is reasonable in the overlapping mass ranges, as there is no significant break in the median fit when crossing over a mass resolution threshold, though higher resolution simulations tend to predict slightly higher metallicities at a given mass. We employ here a 64-particle resolution limit, because we have found that even though the stellar masses are well converged down to 32 particles, the metallicities are only well converged to a higher mass limit, below which the MZR scatter starts to grow substantially in a manner that is inconsistent with higher resolution runs (this is typical in simulations; see e.g. Davé et al. 2011). Observations at  $z = 0$  are shown from the Sloan Digital Sky Survey (SDSS), via nebular line fitting (T04, black solid line) and ‘direct’ abundance measures from stacked spectra (AM13, black dashed line). At  $z = 2$ , we show the observations from the Mosfire Deep Evolutionary Field survey (Kriek et al. 2015) using O3N2 abundances obtained from near-IR Keck spectroscopy (Sanders et al. 2015, points with error bars).

Broadly, the agreement between MUFASA and the observations is fairly good. The faint-end slope is generally consistent with data at both redshifts, and at high- $z$ , it can be seen that the simulated MZR slope extends unabated to much lower masses than can be observed prior to the *James Webb Space Telescope*. At the massive end, there is clearly a turnover at low redshifts above  $M_* \gtrsim 10^{11} M_\odot$ , and even at  $z = 2$ , there is a hint of a similar turnover, though even the  $50 h^{-1}$  Mpc volume does not adequately probe the very high-mass end at that epoch.

At low masses, there is  $\sim 0.2$ – $0.3$  dex increase in the metallicity at a fixed  $M_*$  from  $z = 2 \rightarrow 0$ . The evolution is slightly less at high masses, creating a more prominent flat portion of the MZR. This amount of evolution, and the trend of a more prominent turnover at low masses, is generally consistent with the observations (Steidel et al. 2014; Zahid et al. 2014; Sanders et al. 2015).

A more careful comparison to MZR data reveals some notable discrepancies. Most obviously, there is a clear overprediction of the metallicity at  $M_* \gtrsim 10^{10.3} M_\odot$  at  $z = 0$ . It appears that the high-mass flattening begins at a lower mass scale in the data as compared to in MUFASA, which continues with an unabated power



**Figure 3.** MUFASA galaxy mass–metallicity relations at  $z = 0, 2$ , computed from the predicted SFR-weighted oxygen abundance assuming the solar oxygen abundance from Vagnozzi, Freese & Zurbuchen (2016). Displayed points are combined from the 50, 25 and 12.5  $h^{-1}$  Mpc runs, and for each run, every galaxy with gas is plotted down to the 64-particle stellar mass resolution limit where a break is evident. Points are colour-coded by their distance from the  $M_*$ –SFR relation; bluer points have higher SFR for their  $M_*$ , as indicated by the colour bar. Observations at  $z = 0$  are shown from Tremonti et al. (2004, hereafter T04; solid black) and Andrews & Martini (2013, hereafter AM13; grey dashed), while  $z \approx 2$  data are shown from Sanders et al. (2015).

law up to nearly  $10^{11} M_{\odot}$  before flattening. There is even a hint of such an overproduction at  $z = 2$ ; while the overall amplitude is slightly too large compared to these observations at all masses (e.g. by about 0.1 dex at  $M_* = 10^{10} M_{\odot}$ ), this is particularly exacerbated for the highest mass bin where the discrepancy reaches 0.3 dex. One possibility for reconciling this in the models would be that the metal loading factor at  $M_* \gtrsim 10^{10.3} M_{\odot}$  should be greater than unity, which would preferentially eject a higher fraction of metals out of the high-mass galaxies. Alternatively, it could be that the models have excess wind recycling at high masses; we will examine mass flows and recycling in detail in future work. Finally, this could be an outgrowth of the quenching algorithm that begins to prevent accretion in galaxies around this mass, which results in a closed-box metallicity evolution with gas consumption and commensurate rapid enrichment from stellar mass-loss.

One can also see that the low-mass end of the MZR is in better agreement with the T04 nebular line MZR than the direct abundances measures by AM13. Such discrepancies between observational analyses highlight the difficulty in robustly calibrating metallicity indicators (Kewley & Ellison 2008). Moreover, at high redshifts, it is possible that the typical stellar population in  $z \sim 2$  star-forming galaxies may be substantially different than that at low redshifts (Steidel et al. 2016), which could alter the usual metallicity calibrations applied to nebular emission-line measures. In light of this, the disagreements between MUFASA MZR predictions and observed may be regarded as preliminary.

Finally, the colours of the points show a clear trend that galaxies with low sSFR at a given mass will have high metallicity and vice versa. This has been noted in data (Ellison et al. 2008; Lara-López et al. 2010; Mannucci et al. 2010; Salim et al. 2014; Telford et al. 2016), and Mannucci et al. (2010) dubbed this the fundamental metallicity relation (FMR) because they further argued that the SFR– $M_*$ – $Z$  relation was also redshift-independent. More recent results have called into question whether the FMR is truly redshift-independent (Salim et al. 2015; Brown, Martini & Andrews 2016; Grasshorn Gebhardt et al. 2016), and also whether it is even seen at high redshift (Steidel et al. 2014; Sanders et al. 2015). However,


it appears that the samples at  $z \sim 2$  may not be sufficient for such a trend to have been apparent, and moreover calibration issues can mask such subtle correlations (Salim et al. 2015). It is thus unclear whether the FMR exists at  $z \sim 2$  observationally. We will discuss this second-parameter dependence of the MZR on the sSFR further in Section 6.

In MUFASA, the general trend of the SFR– $M_*$ – $Z$  relation is apparent at both  $z = 0$  and 2. However, the predicted MZR is notably tighter at  $z = 0$  (typical variance of  $\sigma \approx 0.1$  dex around the mean relation) than at  $z = 2$  ( $\sigma \approx 0.2$  dex). By  $z = 2$ , the most metal-rich galaxies already have metallicities comparable to the most metal-rich objects at  $z = 0$ , across all  $M_*$ , while the most metal-poor objects are much less enriched.

The physical explanation for the second-parameter correlations with SFR is that an increase in gas accretion will bring in metal-poor gas while fuelling new star formation, and conversely, a lull in accretion will result in an evolution more similar to a closed box that will raise the metallicity quickly by consuming its gas (e.g. Finlator & Davé 2008). As pointed out in Davé et al. (2011), the lull is permanent for satellite galaxies, causing them to reach a slightly higher metallicity at a given mass before running out of fuel, as observed (Pasquali et al. 2010); though we do not show it here, this is true in MUFASA as well, and similar results are seen in Illustris (Genel 2016) and Evolution and Assembly of GaLaxies and their Environments (EAGLE) (Bahé et al. 2017), with these papers pointing out the additional impact from concentrating the gas owing to harassment and stripping of low-metallicity gas in the outskirts. Hence, in the fluctuating ‘smooth accretion’ scenario for galaxy fuelling (Kereš et al. 2005; Dekel et al. 2009), the FMR is a natural outcome, and the scatter about the relation reflects the frequency and impact of accretion fluctuations such as mergers. Confirming the reality of the FMR at  $z \sim 2$  is thus a crucial test of this scenario. In Section 6, we will quantify predictions for this second-parameter correlation that can be tested against present and future observations.

Fig. 4 shows the evolution of the MZR from  $z = 6 \rightarrow 0$ , computed as a running median from the combined sample of three runs. The





**Figure 4.** MUFASA median galaxy mass–metallicity relations at  $z = 0, 0.5, 1, 2, 3, 4, 6$ . Displayed relations combine galaxies from the 50, 25 and  $12.5 h^{-1}$  Mpc runs down to each of their 64-particle resolution limit. There is steady upwards evolution of the MZR over time since  $z \sim 4$ , with  $\sim \times 2$  increase in metallicity at a given  $M_*$ , since  $z = 2$ .

MZR shows a constant low- $M_*$  power-law slope of  $\approx 0.5$  at all redshifts. At the lowest redshifts, there is the onset of a flattening in the MZR at  $M_* \gtrsim 10^{10.7} M_\odot$ . The MZR rises steadily but slowly with time. At a given mass (below the flattening), the evolution is  $\sim 0.2$  dex out to  $z \sim 1$ , and then 0.1 dex per unit redshift out to  $z \sim 4$ , and no further evolution to  $z = 6$ . In Finlator & Davé (2008), it was argued that barring any evolution in  $\eta$  with  $M_*$  (MUFASA assumes none), then the evolution of the MZR, must reflect the enrichment level of accreted material, i.e. wind recycling. It remains to be seen if such a scenario is consistent with a mass-independent increase in the metallicity down to quite low masses.

Qualitatively, the generally slow evolution and the onset of a high-mass flattened portion at lower redshifts is consistent with observations (e.g. Zahid et al. 2014), as well as the data-constrained analytic models of galaxy evolution (Mitra et al. 2015). However, the mass at which the flattening occurs is generally much higher in MUFASA than in such data, where the onset of flattening is typically below  $M_* \lesssim 10^{10} M_\odot$ . This again reflects the fact that MUFASA appears to produce too steep an MZR at  $10^{10} \lesssim M_* \lesssim 10^{11} M_\odot$ .

Overall, the slope and evolution of the MZR is in broad agreement with the observations, showing mild evolution out to  $z \sim 4$ . However, there is a key discrepancy around  $L^*$  galaxies that bears further investigation. In future work, we will examine the detailed origin for the evolution of the MZR, highlighting contributions from *in situ* enrichment versus pre-enriched accreted gas.

## 5 GAS CONTENT

The gas content of galaxies provides a measure of the fuel available for new star formation. Molecular gas ( $H_2$ ) directly traces material that is forming into stars, while atomic gas ( $H I$ ) typically resides in a more extended reservoir that connects the ionized IGM with the molecular ISM. Hence the gas content of galaxies represents a combination of the effects of how gas is converted into stars within the ISM, as well as the processes that fuel new star formation via gas from the IGM.

Observationally, it is generally believed that the atomic gas in galaxies evolves out slowly to high redshifts, while molecular gas evolves more rapidly upwards. The canonical explanation for this

is that  $H I$  represents a transient reservoir that does not directly trace star formation, while  $H_2$  traces star-forming gas much more closely and hence drops with time in a manner similar to what is seen for the cosmic SFR.

In actuality, the story is more subtle. In simple terms, one can rewrite the ratio of star-forming gas to stars as

$$\frac{M_{\text{gas}}}{M_*} = \frac{M_{\text{gas}}}{SFR} \frac{SFR}{M_*} = t_{\text{dep}} sSFR, \quad (4)$$

where the first term is the depletion time and the second term is the specific SFR (e.g. Davé, Finlator & Oppenheimer 2012). Given a fixed depletion time, one then expects the gas content of high-redshift galaxies to be increased. However, one also expects the depletion time to be reduced to higher redshifts, since galaxies typically form a relatively fixed fraction of their gas into stars per dynamical time (Kennicutt 1998), and disc dynamical times are expected to scale approximately with the *Hubble* time (Mo, Mao & White 1998). If  $sSFR \propto (1+z)^{2.5}$  and  $t_{\text{dyn}} \propto H^{-1}(z)$ , then one gets approximately  $f_{H_2} \propto (1+z)$ . Hence, galaxies are expected to have higher star-forming gas fractions at earlier epochs.

Meanwhile, the evolution of atomic hydrogen is not so straightforward to predict. In the simplest model where the time-scale to pass through the atomic phase also scales with the halo (or, equivalently, disc) dynamical time,  $H I$  should follow  $H_2$ . But physically, atomic gas occurs when gas can self-shield against ionizing radiation, yet is not dense enough to be molecular (i.e. to self-shield against  $H_2$  dissociating radiation). At high redshifts, gas is physically denser and accretion is more filamentary (Dekel et al. 2009), but the ionizing background is stronger. Which effect wins will depend on the detailed interplay of how gas is accreted around galaxies.

In this section, we examine how the atomic, molecular and total neutral (atomic+molecular) gas evolves within galaxies, as a function of stellar mass, in terms of mass functions and globally as a cosmic mass density.

### 5.1 Gas fractions

MUFASA, like many recent simulations of galaxy formation, tracks the amount of molecular gas formed in galaxies. Owing to the limitations of resolution, this is done via a sub-resolution prescription as described in Paper I, broadly following Krumholz et al. (2009) with minor additions.

Meanwhile, the atomic gas fraction is typically significant only in regions that are able to self-shield against the cosmic metagalactic flux (ignoring, as we do here, ionizing radiation emitted locally by the galaxy itself). Hence, we must account for self-shielding in order to separate the neutral gas from the ionized gas.

We follow the prescription in Rahmati et al. (2013) for determining the self-shielded fraction. They provide a fitting formula to the attenuation in the cosmic metagalactic flux as a function of local density, based on full radiative transfer simulations. Given the attenuated ionizing flux impinging on each gas particle, we then compute the rate balance equations to determine the equilibrium atomic fraction following Popping et al. (2009). For particles at low densities ( $n_H \lesssim 10^{-3} \text{cm}^{-3}$ ), the gas is generally optically thin, but above this density, one quickly gets more self-shielded gas, increasing the fraction to unity typically above  $n_H \gtrsim 10^{-2} \text{cm}^{-3}$ . From this self-shielded gas, we then subtract the molecular fraction as tracked directly in the simulation, which yields the atomic fraction. We compute a galaxy's  $H I$  content by summing all atomic gas that is more gravitationally bound to that galaxy relative to any other

galaxy, using the total baryonic mass to compute the gravitational binding. In practice, we do not consider gas with  $n_{\text{H}} < 10^{-4} \text{ cm}^{-3}$  since this is never self-shielded and thus contributes negligibly to the total H I content.

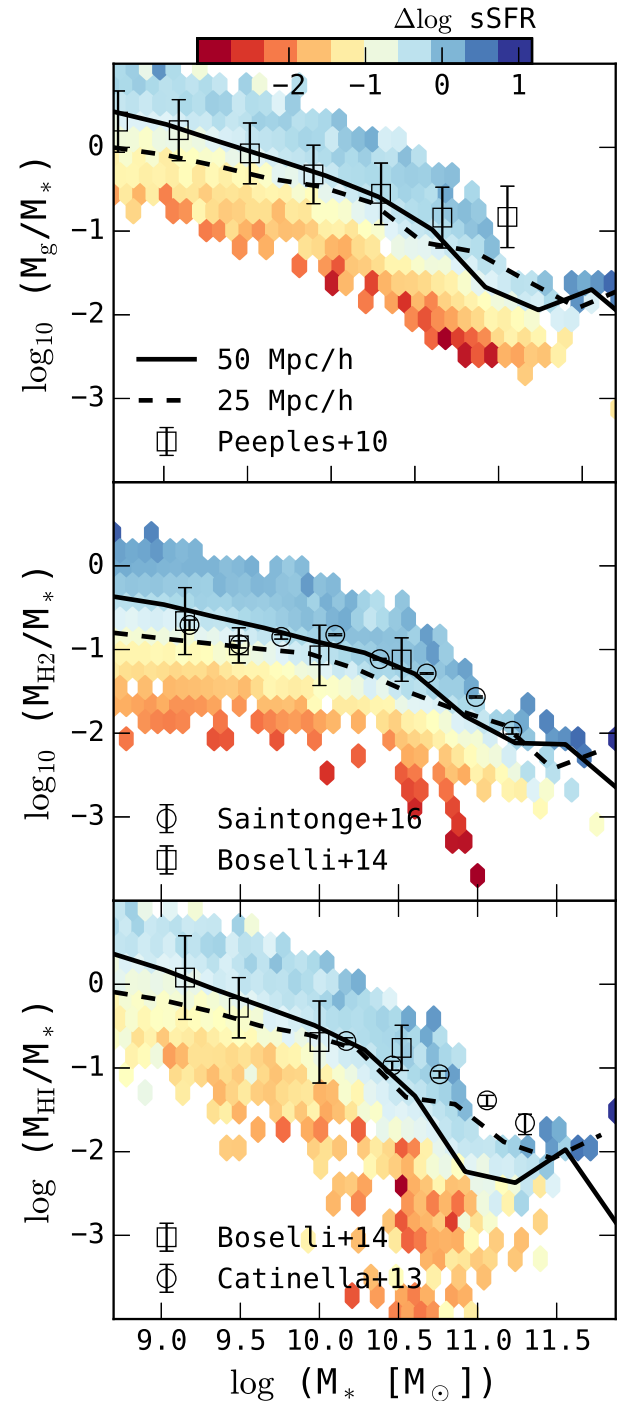
Fig. 5 shows the total (H I+H<sub>2</sub>) (top panel), molecular (middle) and atomic (bottom) gas fractions as a function of stellar mass at  $z = 0$ . The solid black line shows a running median for the fiducial  $50 h^{-1}$  Mpc volume. The overlaid hexbins are colour-coded by the average sSFR at that gas fraction relative to the global average sSFR at the given  $M_*$ . The dashed line shows a similar running median for the  $25 h^{-1}$  Mpc run, to illustrate the level of resolution convergence. This plot, and subsequent ones, includes all (resolved) galaxies.

In the top panel, the total gas fraction as a function of  $M_*$  in the  $50 h^{-1}$  Mpc run is in excellent agreement with a compilation of the observations by Peeples & Shankar (2011) over most of the mass range. At the highest masses, the observations lie above the model predictions. While these data only include galaxies where gas was detected, and many of the simulated galaxies have such low gas fractions that they would likely evade detection, since there are no predicted galaxies at all at the median total gas fraction, it appears at face value that the discrepancy is real. For  $M_* \lesssim 10^{10.5} M_{\odot}$ , however, galaxy samples are quite complete, and hence the agreement is a robust success of the models. To quantify the agreement, we interpolate the simulated gas fractions to the observed mass bins and compute the difference in predicted and observed values in units of the observational uncertainty  $\sigma$ , and combine these in quadrature. For  $M_* < 10^{10.7} M_{\odot}$ , this yields a  $0.4\sigma$  deviation between MUFASA and the Peeples & Shankar (2011) data.

The middle panel shows that the molecular gas fractions are likewise in good agreement with the observations from the COLDGASS survey (Saintonge et al. 2016), as well as the Herschel Reference Survey (HRS; Boselli et al. 2014). COLDGASS (Saintonge et al. 2011) is an  $M_*$ -complete survey and hence is quite directly comparable to our simulated galaxies. MUFASA even traces the slight turn-down in  $f_{\text{H}_2}$  at  $M_* \gtrsim 10^{10.5}$  relative to an extrapolated trend from lower masses, which is indicative of a typical mass scale at which quenching kicks in. The combined deviation versus HRS is  $\lesssim 0.8\sigma$ , and that versus COLDGASS is large, owing to the very small statistical uncertainties, but the true uncertainties are likely dominated by systematics that are difficult to quantify; none the less, the agreement in  $f_{\text{H}_2}$  is within  $\sim 0.2$  dex at all  $M_*$ .

The atomic gas fractions are compared to data from the Galex Arcibo SDSS Survey (GASS; Catinella et al. 2010) that is the parent survey of COLDGASS and hence also a  $M_*$ -selected sample of SDSS galaxies down to very low H I fractions. At low masses ( $M_* \lesssim 10^{10.5} M_{\odot}$ ), there is a quite good agreement with the GASS data, which again is a non-trivial success. However, our  $50 h^{-1}$  Mpc volume predicts a sharp drop in  $f_{\text{H I}}$  above this mass, whereas the data show a more gradual trend. This is likely the origin of the discrepancy in the total gas fraction at these masses, since  $f_{\text{H}_2}$  shows good agreement in this mass range. Statistically, in versus HRS that covers masses below the drop, the combined deviation is just  $\sim 0.8\sigma$ . However, in versus GASS, the statistical agreement is poor, owing to the discrepancy at  $M_* \gtrsim 10^{10.5} M_{\odot}$ .

The  $25 h^{-1}$  Mpc volume (dashed lines) consistently shows lower gas content at  $M_* \lesssim 10^{10} M_{\odot}$ , and thus a shallower trend with  $M_*$  that results in an  $\sim \times 2$  deficit with respect to the  $50 h^{-1}$  Mpc volume at the lowest probed masses. The deficit is essentially identical in both H I and H<sub>2</sub>, which suggests that gas consumption is more rapid in the  $25 h^{-1}$  Mpc volume, likely owing to its higher resolution that achieves higher densities where more rapid star formation can occur.



**Figure 5.** Total neutral (H I+H<sub>2</sub>; top panel), molecular (middle) and atomic (bottom) gas fractions as a function of stellar mass predicted from the MUFASA  $50 h^{-1}$  Mpc simulation at  $z = 0$ . Running medians are shown as the solid black lines. Colour-coding shows the mean sSFR deviation from the main sequence ( $\Delta \log$  sSFR) in each hexbin. Dashed line shows the running median from the  $25 h^{-1}$  Mpc run to assess resolution convergence. Data are shown in the top panel from the compilation by Peeples & Shankar (2011), in the middle panel from COLDGASS (Saintonge et al. 2016) and HRS (Boselli et al. 2014), and in the bottom panel from HRS and GASS (Catinella et al. 2013). The  $50 h^{-1}$  Mpc box shows good agreement with the observations over the mass range probed by the data, but the  $25 h^{-1}$  Mpc run tends to show lower gas fractions at low masses. At a given  $M_*$ , galaxies with higher gas content have higher SFR, and the trend appears tighter for H<sub>2</sub>.

Interestingly, this volume shows no ‘dip’ in the H I and hence total gas content at  $M_* \sim 10^{11} M_\odot$ , indicating that the disagreement in the  $50 h^{-1}$  Mpc may be a peculiarity in that simulation or else some issue with resolution convergence in terms of the way it interacts with the quenching model. One possibility is that the  $25 h^{-1}$  Mpc is able to self-shield gas in massive haloes more effectively owing to its ability to resolve clumpier structures, and thus the quenching model is less impactful here, since, by construction, it only operates on non-self-shielded gas. In any case, at high and low masses, it appears that resolution convergence is not ideal for predicting gas fractions, and the resulting systematic uncertainties are of the order of a factor of 2.

Other models that track molecular and atomic gas generally predict H I and H<sub>2</sub> fractions broadly in accord with the observations, be they semi-analytic (Lagos et al. 2011; Popping, Somerville & Trager 2014) or state-of-the-art hydrodynamic models such as EAGLE (Lagos et al. 2015; Crain et al. 2017). For instance, Lagos et al. (2015) finds good agreement in EAGLE with H I and H<sub>2</sub> gas fractions as well as mass functions, particularly when adopting the Gnedin & Kravtsov (2011) prescription for computing  $f_{\text{H}_2}$ . Alongside MUFASA’s success, these results suggest that the overall gas content is a fairly robustly predictable quantity in models, at least at  $z = 0$ . We note that all these models (including ours) have been tuned at varying levels in order to match the present-day stellar mass function. It may be that predicting the GSMF correctly, plus having a molecular gas-based prescription for converting gas into stars, generically leaves the proper amount of gas in galaxies. If so, this represents a non-trivial success for current models of galaxy formation.

The coloured hexbins show that at a given  $M_*$ , both molecular and atomic gas content are highly correlated with ongoing star formation. In both cases, galaxies with enhanced gas content for their  $M_*$  also have higher sSFR. The trend appears to be qualitatively stronger in the molecular case, which is unsurprising since stars form out of molecular gas in our simulations. None the less, it is also clearly present in the atomic gas, indicating that the H I reservoir plays a role in regulating star formation even if it is not directly forming stars. This is qualitatively consistent with the observations that show more low-metallicity gas in the outskirts of bluer (i.e. higher sSFR) galaxies (Moran et al. 2012). Simulations have also found this trend; Davé et al. (2011) showed that this is generic outcome of models regardless of stellar feedback implementation – it is even present without any feedback – and Lagos et al. (2016) found a similar trend in EAGLE.

As with the metallicity, the qualitative explanation of this is that a temporary enhancement (lull) of accretion results in both increased (decreased) gas content and star formation, along with the aforementioned reduction (increase) in metallicity. Since it takes some time for the inflowing gas to first turn into atomic gas and then molecular and finally stars, the molecular gas content is expected to be more highly correlated with the SFR. Hence, as with the FMR, the second parameter dependence of gas content on SFR most directly reflects the fluctuations in the inflow rate (Mitra et al. 2017). This idea of self-regulation of gas content has been presented previously in various forms (e.g. Davé et al. 2011), but was explored comprehensively in a suite of EAGLE simulations by Lagos et al. (2016), who argued that the  $f_{\text{H I}} - M_* - \text{SFR}$  relation formed a ‘Fundamental Plane of star formation’ whose detailed shape is set by the star formation law. Overall, it appears that the correlated behaviour of gas fraction and star formation with stellar mass is a generic outcome of the galaxy growth process. We will quantify this behaviour in Section 6.

## 5.2 Gas fraction evolution

Galaxies at a given mass are observed to be more molecular gas-rich at earlier epochs (e.g. Geach et al. 2011; Tacconi et al. 2013). The amount of evolution is subject to some uncertainties regarding the conversion between the observed CO intensity and the molecular gas mass (see e.g. Bothwell et al. 2013), but this is unlikely to erase the qualitative trend. Far-IR dust continuum measures can also be used to probe gas content evolution, though likewise subject to some uncertainties regarding the conversion of dust to gas mass (e.g. Scoville et al. 2016); in general, such studies tend to find less strong evolution than CO-based studies.


Neutral gas above  $z \gtrsim 0.3$  is currently only observable in absorption line studies such as with Mg II absorbers (Rao, Turnshek & Nestor 2006) or Damped Lyman Alpha (DLA) systems (Prochaska & Wolfe 2009); it is not obvious how these systems trace galaxies, as it is usually challenging to identify the individual galaxy giving rise to such absorbers, though clustering measures offer some general guide that they typically live in  $10^{11-12} M_\odot$  haloes (e.g. Bouché et al. 2005; Font-Ribera et al. 2012). Measuring H I fractions requires having a measure of the stellar mass from optical or near-IR data for individual galaxies, for which 21 cm emission can be observed. While current H I-21 cm surveys only probe to  $z \sim 0.4$  (Fernández et al. 2016), upcoming radio telescopes promise to push direct H I-21 cm gas content measures in optically selected samples out to  $z \gtrsim 1$ , for example, using the new MeerKAT telescope (Holwerda et al. 2012), and will be further advanced with the SKA. Here we make testable predictions for gas fractions that can guide such efforts.

Fig. 6 shows the evolution at  $z = 0, 1, 2, 4$  of the median total (H I+H<sub>2</sub>) gas fraction versus stellar mass from the  $50 h^{-1}$  Mpc MUFASA run in the top panel, and the next two panels show this subdivided into molecular and atomic gas fractions. The  $z = 0$  curve is identical to that in Fig. 5, but here we also show with shading the cosmic variance estimated via jackknife resampling over eight volume sub-octants. Here we do not show the second-parameter dependence on SFR as we did in Fig. 5, but a similar trend persists at all redshifts. We do not explicitly show any observations on this plot, since molecular gas observations span some range depending on the type of data, while atomic gas fraction measures do not yet exist at  $z \gtrsim 0.4$ .

The total gas content of galaxies at a given  $M_*$  is higher at earlier epochs. There appears to be some mild mass dependence to this statement, as high-mass galaxies lose their gas more quickly than low-mass galaxies, with an overall effect of steepening the  $f_{\text{gas}} - M_*$  relation. Much of the evolution occurs from  $z \sim 1 \rightarrow 0$ , prior to which the evolution was somewhat slower.

Neutral hydrogen (bottom panel) represents the majority of the cold gas content of galaxies at almost all epochs and masses, except at high masses today. Hence, the evolutionary trends in H I fraction tend to drive those of the total gas content. The strong evolution particularly at high masses out to  $z \sim 1$  is good news for the upcoming H I surveys designed to measure 21 cm emission from galaxies out to this epoch such as Looking At the Distant Universe with the MeerKAT Array (LADUMA), and will figure prominently in the evolution of the H I mass function (HIMF) discussed in Section 5.3.

Although we do not show it, the  $25 h^{-1}$  Mpc box actually shows quite good resolution convergence with the  $50 h^{-1}$  Mpc box shown here for all redshifts except  $z = 0$ . At  $z = 0$ , the  $25 h^{-1}$  Mpc volume shows a flatter relation (as seen in Fig. 5), but at higher  $z$ , the relations are similar, which implies less mass dependence to the



**Figure 6.** MUFASA median gas fractions as a function of  $M_*$  at  $z = 0$  (solid black), 1 (dashed blue), 2 (dot-dashed green) and 4 (dotted red) from the  $50 h^{-1}$  Mpc run. Cosmic variance over the eight sub-octants is shown as the hatched region around the  $z = 0$  curve. The H I fraction (bottom panel) decreases with time for massive galaxies, while low-mass galaxies are always H I-dominated. The molecular gas fraction (middle panel) increases steadily across all masses, but is typically sub-dominant to H I. The total molecular plus atomic content (top panel) is thus driven by the H I evolution.

evolution. Hence one should regard the detailed mass dependence of the evolution as a less robust prediction.

The trend to earlier epochs for the molecular gas (middle panel) is broadly similar to that for the atomic gas, in that it is increasing at all masses. There is a steady decrease in  $f_{\text{H}_2}$  with time across all  $M_*$  of about 0.2–0.3 dex between  $z = 2 \rightarrow 0$ , with only a very slight trend for more evolution at the highest masses. Predicted gas fractions continue to increase at a given  $M_*$  out to  $z = 4$ , so we expect even more gas-rich galaxies at high masses, but unfortunately, even Atacama Large Millimeter/submillimeter Array (ALMA) will have difficulty measuring the molecular content at these epochs except in the very largest systems (Decarli et al. 2016).

Comparing to the observations, it appears that MUFASA predicts  $\text{H}_2$  fractions that are too low versus data at  $z \sim 1$ –2. CO-based gas fractions from Tacconi et al. (2013) show a large scatter but generally lie between 20 and 40 per cent for the most massive galaxies, and 50 per cent for moderate-mass galaxies. To lower masses, fractions up to 90 per cent are inferred for the smallest  $z \sim 2$  galaxies by inverting the Kennicutt (1998) relation (e.g. Erb et al. 2006). The dust continuum-based measures from Scoville et al. (2016) also show typically molecular fractions of 20–40 per cent for the main-sequence galaxies at  $z \sim 1$ –2, and even higher for starbursts. In contrast, MUFASA predicts  $z = 2$  gas fractions of  $\sim 10$  per cent for massive ( $M_* \sim 10^{11}$ ) galaxies, and only up to  $\sim 40$  per cent (relative to the molecular+stellar mass) for the smallest galaxies that are well below what can be probed directly with the observations. Hence, in general, it appears that MUFASA high- $z$  gas fractions are too low by  $\sim \times 2$ . Given the uncertainties, this is not a gross failure, but it is notable.

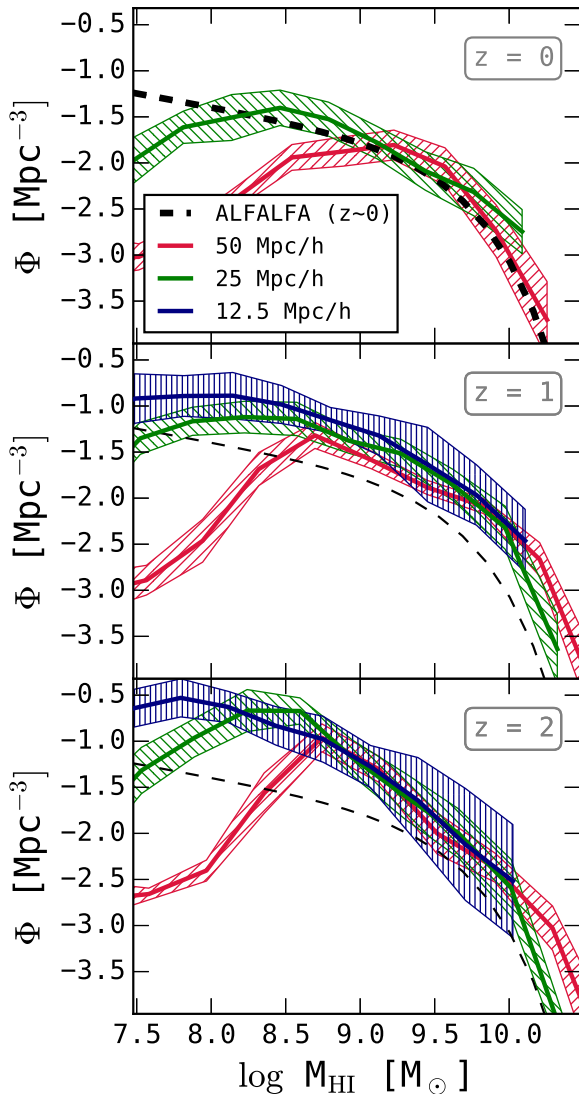
Such low high- $z$  molecular gas fractions are predicted in other simulations and SAMs as well (Popping et al. 2015b; Lagos et al. 2015). For example, using EAGLE, Lagos et al. (2015) found that the typical molecular gas fraction at  $z \sim 2$  for  $M_* > 5 \times 10^9 M_\odot$  galaxies is  $\sim 20$  per cent, and actually even drops slightly to higher redshifts. The GADGET-3 simulations by Davé et al. (2011) also found modest gas fractions of  $\sim 30$  per cent for intermediate-redshift massive galaxies. These discrepancies with the observations may be partially but not completely explainable by selection effects in which targeted CO observations tend to select highly star-forming (and thus gas-rich) galaxies; Tacconi et al. (2013) accounted for this and still found generally higher  $f_{\text{H}_2} \sim 40$  per cent than predicted in simulations. Another possibility is that locally calibrated CO-to- $\text{H}_2$  conversion factors may not be correct at high- $z$ ; we will explore this issue in more depth in Section 5.4. None the less, at face value, it appears that many models including MUFASA struggle to reproduce quite as high gas fractions as inferred for massive high- $z$  galaxies.

### 5.3 H I mass function

The HIMF combines information from the galaxy mass function and H I fractions to provide a complementary constraint on models. Observations of the HIMF extend to quite low masses locally thanks to deep surveys with the Arecibo telescope such as the H I-selected The Arecibo Legacy Fast Alfa survey (ALFALFA; Haynes et al. 2011) and the stellar mass-selected GASS (Catinella et al. 2010). However, the sensitivity of current instrumentation precludes characterization of the HIMF at significantly higher redshifts. The SKA and its precursors aim to improve on this, and hence predictions for the evolution of the HIMF are useful for quantifying expectations for upcoming surveys such as the LADUMA survey (Holwerda et al. 2012).

Fig. 7 shows the predicted HIMF from MUFASA, showing the three volumes in different colours with cosmic variance (shading) estimated as before, from the variance of the HIMF in each of the eight sub-octants within each simulation volume. As discussed in Paper I, this is likely to somewhat underestimate the true cosmic variance. The three panels show the HIMF at  $z = 0, 1, 2$  from top to bottom. The ALFALFA mass function at  $z \approx 0$  is shown as the thick black dashed line, and this is reproduced in the other redshift panels to better visualize the amount of evolution in the models; however, direct comparison to MUFASA should only be made at  $z = 0$ .

The top panel shows that MUFASA provides a reasonable match to the observed HIMF, in the resolved range. At  $M_{\text{HI}} \lesssim 10^9 M_\odot$ , the  $50 h^{-1}$  Mpc volume shows a departure from the data, but the higher



**Figure 7.** HIMFs in the 50 (red solid), 25 (green dashed) and  $12.5 h^{-1}$  Mpc (blue dotted, for  $z \geq 1$ ) MUFASA simulations. Hatched region shows the cosmic variance computed over the eight sub-octants in each volume. Black dashed line shows  $z \approx 0$  observations from the ALFALFA survey, reproduced at  $z = 1, 2$  in order to better depict the evolution. MUFASA does reasonably well reproducing the HIMF over the  $M_{\text{HI}}$  range of convergence; at  $z \sim 0$ , the  $50 h^{-1}$  Mpc box begins to deviate from the  $25 h^{-1}$  Mpc at  $M_{\text{HI}} \lesssim 10^9 M_{\odot}$ , while at higher redshifts, this occurs at  $M_{\text{HI}} \sim 10^{8.7} M_{\odot}$ . The HIMF generally increases in amplitude to higher redshifts, and also steepens noticeably at  $z = 2$ . There is fair resolution convergence down to  $M_{\text{HI}} \sim 10^9 M_{\odot}$  in the  $50 h^{-1}$  Mpc volume, and to  $\approx 8 \times$  lower in the  $25 h^{-1}$  Mpc volume.

resolution run continues unabated, suggesting that the turnover at low masses is an artefact of numerical resolution. Indeed, if one combines the stellar mass resolution limit of  $10^{8.7} M_{\odot}$  with the fact that galaxies at that  $M_*$  have an H I fraction of around two (Fig. 5), this suggests that galaxies with  $M_{\text{HI}} \lesssim 10^9 M_{\odot}$  will suffer from incompleteness in our simulations. The  $25 h^{-1}$  Mpc volume extends another factor of almost 8 lower in mass before turning over, as expected from its  $8 \times$  higher mass resolution. Computing the combined deviation for the  $50 h^{-1}$  Mpc box for  $M_{\text{HI}} > 10^9 M_{\odot}$  gives a deviation of  $0.8\sigma$  between the MUFASA and ALFALFA data, where  $\sigma$  is taken to be our cosmic variance estimate. For the  $25 h^{-1}$  Mpc box, pushing to a factor of 8 lower in  $M_{\text{HI}}$ , we obtain a combined

deviation of  $1.2\sigma$ . While there is a mild trend for the predictions to be systematically higher than the data, the differences are not statistically significant, and hence MUFASA provides a good match to the ALFALFA HIMF.

The agreement of MUFASA with both the stellar mass function (Paper I) and the HIMF is an important success. Previous simulations by Davé et al. (2013) also showed good agreement with data for both, even when sub-divided into stellar mass bins (Rafieeantsoa et al. 2015). The EAGLE simulation likewise shows good agreement with both (Crain et al. 2017), including sub-divided into  $M_*$  bins (Bahé et al. 2016). However, semi-analytic models constrained to match the stellar mass function do not necessarily agree well with the HIMF (e.g. Benson 2014). The SAMs of Popping et al. (2014) do fairly well at  $M_{\text{HI}} \gtrsim 10^9 M_{\odot}$ , but predict a significant upturn to lower masses, which is not observed. Such an upturn is also seen in the older Overwhelmingly Large Simulations HIMF as well (Duffy et al. 2008). The semi-empirical model of Popping, Behroozi & Peeples (2015a) likewise produces a steep faint end of the HIMF, deviating strongly at  $M_{\text{HI}} \lesssim 10^9 M_{\odot}$ . In simulations such as MUFASA, H I represents a transient reservoir of cold gas infalling into a galaxy, as demonstrated in Crain et al. (2017); such a dynamic origin suggests that fully dynamical models are best suited to make predictions for the nature and evolution of H I in galaxies. It appears that the low-mass ( $M_{\text{HI}} \lesssim 10^9 M_{\odot}$ ) HIMF may be a key discriminant for the dynamics of gas infall.

Looking at the evolution to  $z = 1$ , we see that the HIMF is best described by an overall increase in the mass of H I in each galaxy by a factor of  $\sim 2$  to 3, particularly at the massive end. This is consistent with the evolution seen in Fig. 6. This is good news for surveys such as LADUMA that will probe the bright end of the HIMF at these redshifts; in future work, we will make more specific predictions for LADUMA. Interestingly, this is somewhat contrary to the trend predicted by our previous simulations in Davé et al. (2013), which showed a steepening of the HIMF to higher redshifts, but the massive end was generally unchanged or lowered. This is because the H I fraction in the Davé et al. (2013) simulations was invariant with redshift, whereas in MUFASA galaxies, they are substantially more H I-rich, particularly at high masses.

From  $z = 1$  to 2, the main trend is that the HIMF is steeper at low masses, while the massive end does not evolve significantly. This is driven by the steepening of the stellar mass function, since the H I fraction, if anything, has a shallower trend with stellar mass at higher redshifts (Fig. 6). This general trend agrees better with that in Davé et al. (2013).

Overall, the HIMF in MUFASA at  $z = 0$  is a reasonable match to the observations, even though the dynamic range is limited compared to other simulations such as EAGLE. By combining various box sizes, we can span a similar dynamic range, and the HIMF shows good resolution convergence in the overlapping H I mass range. MUFASA predicts a noticeably higher HIMF at  $z \sim 1$ , and then a steepening trend to  $z \sim 2$ , which must await future SKA and pathfinder telescope data for testing.

#### 5.4 CO luminosity function

The mass function of molecular gas is more complicated to determine than that of atomic gas, since the observations typically do not directly trace  $\text{H}_2$  but rather some proxy such as CO. For ordinary (non-starburst) galaxies, canonically, the best proxy for  $\text{H}_2$  is the  $J = 1-0$  rotational transition of CO. None the less, this still requires a conversion factor ( $X_{\text{CO}}$ ) to obtain the  $\text{H}_2$  mass, and the dependence of  $X_{\text{CO}}$  on the intrinsic properties of galaxies, such as SFR and

metallicity, is uncertain. This becomes particularly problematic at high redshifts, where the ISM conditions in typical main-sequence galaxies vary substantially from that today.

A typical assumption is that galaxies that are near the main sequence have ‘Milky Way-like’  $X_{\text{CO}} \approx 4$ , whereas starbursts have  $X_{\text{CO}} \approx 0.8$  (Tacconi et al. 2013). However, substantial work has gone into predicting  $X_{\text{CO}}$  based on galaxy properties from detailed simulations, yielding a continuous rather than bimodal trend. In particular, Narayanan et al. (2012) used zoomed simulations together with a CO line radiative transfer code to develop an approximate fitting function for  $X_{\text{CO}}$  as a function of  $\text{H}_2$  surface density and metallicity:

$$X_{\text{CO}} = \frac{1.3 \times 10^{21}}{Z' \times \Sigma_{\text{H}_2}^{0.5}}, \quad (5)$$


where  $Z'$  is the metallicity in solar units.

Here, we use this formula to compute  $X_{\text{CO}}$  individually for each galaxy, obtaining  $\Sigma_{\text{H}_2}$  by dividing the  $\text{H}_2$  half-mass of each galaxy by the area computed from the  $\text{H}_2$  half-mass radius. Using this  $X_{\text{CO}}$ , we then convert our simulated  $\text{H}_2$  masses into CO luminosities ( $L_{\text{CO}}$ ) that can be compared more directly against the observations. In this way, we specifically account for the metal and gas content evolution in CO-to- $\text{H}_2$  conversions when comparing to the observations. This is analogous to the approach in Narayanan, Bothwell & Davé (2012), except that here we convert simulated galaxies to get  $L_{\text{CO}}$ , while they took the converse approach of converting the observations into  $M_{\text{H}_2}$  to compare with models. However, we will see that our conclusions are similar.

Fig. 8 shows the CO luminosity function (COLF) from our MUFASA simulations, showing once again our available simulation volumes at each redshift  $z = 0, 1, 2$  (top to bottom). At  $z = 0$ , it is possible to directly observe CO 1–0 down to very low  $L_{\text{CO}}$ , and such observations by Kereš, Yun & Young (2003) are shown as the data points. To higher redshifts, blind CO surveys where the survey volume can be robustly estimated are difficult, so one typically uses another proxy for this. The dashed lines show observations from Vallini et al. (2016), which used far-IR luminosity as a proxy for CO luminosity; at  $z = 0$ , they agree with the Kereš et al. (2003) data. At higher redshifts, we plot their observations down to their approximate completeness limit. We note that recent direct CO measures from ALMA by Decarli et al. (2016) indicate a somewhat higher number of high- $L_{\text{CO}}$  objects at  $z \sim 2$  than Vallini et al. (2016), but the statistics are small and the cosmic variance is large, so the discrepancy is only marginally significant.


At  $z = 0$ , MUFASA generally predicts a reasonable COLF, with a hint of an excess at high  $L_{\text{CO}}$ . The  $50 h^{-1}$  Mpc volume shows a turnover at low- $L_{\text{CO}}$  owing to numerical resolution, while the  $25 h^{-1}$  Mpc continues to agree well with the observations down to the lowest probed  $L_{\text{CO}}$ . The observations of Vallini et al. (2016) generally find an increase in the number of high- $L_{\text{CO}}$  galaxies with redshift, and the simulations follow this trend, generally agreeing with data with still a hint of a high- $L_{\text{CO}}$  excess. By  $z = 2$ , the observations only probe the brightest CO galaxies, where only the  $50 h^{-1}$  Mpc volume has comparably bright systems, but these are in very good agreement with the data. Statistically, the agreement with the Vallini et al. (2016) data is marginal; the  $50 h^{-1}$  Mpc box typically agrees to within  $\sim 2\sigma$  in cosmic variance at all redshifts in the overlapping  $L_{\text{CO}}$  range, while the  $25 h^{-1}$  Mpc agrees to  $\lesssim 1.5\sigma$ . Given the systematic uncertainties involved in converting between CO luminosity and  $\text{H}_2$  mass, this level of agreement is encouraging.

It is interesting that despite the relatively mild evolution of  $\text{H}_2$  fractions in Fig. 6 and a putative underprediction of  $f_{\text{H}_2}$  at  $z \sim 2$ ,



**Figure 8.** CO luminosity functions in the  $50$  (red solid),  $25$  (green dashed) and  $12.5 h^{-1}$  Mpc (blue dotted, for  $z \geq 1$ ) MUFASA simulations. Hatched region shows the cosmic variance computed over the eight sub-octants in each volume. We compute the CO 1–0 luminosity from our predicted  $\text{H}_2$  mass based on the prescription in Narayanan et al. (2012) derived from zoom simulations coupled with CO radiative transfer. We show the observations from Vallini et al. (2016) at  $z = 0, 1, 2$  as the black dashed lines, down to their approximate completeness limit at each redshift; note that these are based on an  $L_{\text{IR}}$ -to- $L_{\text{CO}}$  conversion. Resolution convergence amongst the volumes is generally quite good.

MUFASA reproduces well the evolution of the COLF out to  $z = 2$ , and shows significantly more high- $L_{\text{CO}}$  objects at high redshifts. This suggests that using a physically motivated prescription for converting CO into  $\text{H}_2$  (or vice versa) can lead to inferring a different amount of evolution in the gas fractions, and, in general, could potentially reconcile the relatively low amount of evolution in simulations versus the stronger evolution inferred using standard assumptions regarding  $X_{\text{CO}}$ ; this broadly echoes the conclusions of Narayanan et al. (2012). Generically, metallicity-dependent  $X_{\text{CO}}$  prescriptions such as Narayanan et al. (2012) and Feldmann, Gnedin & Kravtsov (2012) tend to predict more  $\text{H}_2$  at high masses and less at low masses owing to enhanced  $\text{H}_2$  production at high



**Figure 9.** Cosmic mass density in H I (blue) and H<sub>2</sub> (red) as a function of redshift, in our 50 (solid) and 25  $h^{-1}$  Mpc (dashed) MUFASA simulations.  $\Omega_{\text{H I}}$  observations are shown from Delhaize (2013,  $z < 0.1$ ), Rao et al. (2006,  $0.5 < z < 1.3$ ) and Noterdaeme et al. (2012,  $2.1 < z < 3.35$ ). The predicted trend of approximately  $\Omega_{\text{H I}} \propto (1+z)^{0.74}$  is a good match to the compiled observations, as is the normalization, although it is somewhat sensitive to numerical resolution.  $\Omega_{\text{H 2}}$  shows a similar redshift trend as  $\Omega_{\text{H I}}$ , which is substantially slower than the evolution of the cosmic SFR density as shown in fig. 3 of Paper I.

metallicities, which serves to increase the bright end of the COLF and flatten the faint end (Popping et al. 2014), thus yielding better agreement with the COLF. Empirical luminosity-dependent  $X_{\text{CO}}$  calibrations have a qualitatively similar effect (Boselli et al. 2014). Hence MUFASA may be plausibly reproducing the evolution of the molecular gas content in galaxies in spite of its modest evolution of  $\sim \times 2$  in the H<sub>2</sub> content at a fixed  $M_*$  out to  $z \sim 2$ .

### 5.5 Cosmic gas mass evolution

A synthesis of all the above evolutionary measures is provided in the evolution of the global cosmic gas density, typically parametrized in units of the critical density (i.e. as  $\Omega_{\text{gas}}$ ). The slow evolution of  $\Omega_{\text{H I}}$  relative to the overall cosmic star formation rate density has been noted as evidence that H I is not directly physically associated with star formation, while the more rapid evolution of H<sub>2</sub> fractions can explain at least part of the rapid evolution in the cosmic SFRD. However, such interpretations are complicated by detailed assumptions regarding  $X_{\text{CO}}$  as discussed in the previous section, and how H I gas traces galaxies. Here we examine predictions for the evolution of the cosmic H I and H<sub>2</sub> mass densities, in the context of the evolutionary trends we have discussed above.

Fig. 9 shows the evolution of the cosmic density in atomic gas (blue) and molecular gas (red) as a function of  $\log(1+z)$ . Solid and dashed lines show the results from our 50 and 25  $h^{-1}$  Mpc, respectively. This is obtained by summing over all SKID-identified galaxies; using instead the sum over all H I or H<sub>2</sub> in the volume (which includes the IGM) makes a negligible difference.

Data points with the blue error bars correspond to various observational measures of  $\Omega_{\text{H I}}$ : from 21cm emission (Delhaize 2013,  $z < 0.1$ ), using Mg II absorbers as a proxy for DLAs (Rao et al. 2006,  $0.5 < z < 1.3$ ), and DLA absorbers selected from the SDSS (Noterdaeme et al. 2012,  $2.1 < z < 3.35$ ). For H<sub>2</sub>, no data are shown; at  $z = 0$ , Kereš et al. (2003) inferred  $\Omega_{\text{H 2}} \approx 2 \times 10^{-4}$ , which is above the predictions, but given the good agreement shown versus

the  $z \approx 0$  COLF from the previous section, this could be subject to uncertainties regarding  $X_{\text{CO}}$ .

$\Omega_{\text{H I}}$  roughly follows a power law in  $(1+z)$ ; a best-fitting relation to the 50  $h^{-1}$  Mpc run is given by  $\Omega_{\text{H I}} = 10^{-3.53} (1+z)^{0.74}$ , and is higher in amplitude by 20 per cent for the 25  $h^{-1}$  Mpc volume. Generally, this provides a good fit to the trend seen in the compilation of the observations from various sources and techniques, particularly for the higher resolution volume. The difference between the volumes, while only about 0.1 dex, none the less suggests that there is a sub-optimal-resolution convergence in this quantity, likely driven by the fact that the 50  $h^{-1}$  Mpc volume does not resolve many low- $M_{\text{H I}}$  galaxies as seen in Fig. 7. Semi-analytic models tend to predict that  $\Omega_{\text{H I}}$  rises somewhat out to intermediate redshifts, but then falls at  $z \gtrsim 1-2$  (Obreschkow et al. 2009; Popping et al. 2015a), in clear disagreement with a continued rise in  $\Omega_{\text{H I}}$  out to  $z \sim 3.5$ . Hence, the broad agreement in the redshift evolution of  $\Omega_{\text{H I}}$  is highly encouraging, and suggests that H I in and around galaxies is being viably modelled by MUFASA across a range of epochs.

In contrast, the evolution of  $\Omega_{\text{H 2}}$  is predicted to be substantially slower than often believed. MUFASA predicts essentially the same redshift evolution for  $\Omega_{\text{H 2}}$  as for  $\Omega_{\text{H I}}$ , with an increase of a factor of  $\times 3$  from  $z = 0 \rightarrow 3$ . The SAMs of Lagos et al. (2011) predict almost no evolution for  $\Omega_{\text{H I}}$ , but a  $\times 7$  increase from  $z = 0 \rightarrow 3$  for  $\Omega_{\text{H 2}}$ . More recently, Lagos et al. (2015) found slower evolution of  $\Omega_{\text{H 2}}$  in the EAGLE simulation, more similar to MUFASA. Observations cannot yet clearly distinguish between these predictions.


In summary, MUFASA predicts mild evolution in both the total H I and H<sub>2</sub> cosmic mass densities, scaling approximately as  $(1+z)^{0.8}$ . Such a scaling roughly follows from the simple equilibrium model arguments outlined at the start of this section. The evolution of  $\Omega_{\text{H I}}$  is in good agreement with the observations, but the predictions for  $\Omega_{\text{H 2}}$  are not currently robustly testable. As CO and far-IR surveys improve with ALMA and other facilities, such constraints will provide important tests of these and other models.

## 6 FLUCTUATIONS AROUND SCALING RELATIONS

In the prevalent baryon cycling paradigm, quasi-continuous gas inflows drive galaxy growth, modulated by feedback (Somerville & Davé 2015). The net result is that galaxies live on fairly tight scaling relations between stellar mass, SFR, metallicity and gas content (e.g. Finlator & Davé 2008; Davé et al. 2012; Lilly et al. 2013; Lagos et al. 2015). Fluctuations in the inflow rate owing to e.g. mergers can cause the fluctuations around these scaling relations. Indeed, inflow fluctuations owing to stochastic dark matter infall alone yield a scatter that is in good agreement with the observed scatter in the SFR- $M_*$  relation (Forbes et al. 2014; Mitra et al. 2017).

In addition to scatter in SFR, such fluctuations also give rise to correlated scatter in the metallicity and gas content. For instance, a boost in inflow will enhance the gas content and lower the metallicity, while boosting the SFR owing to the abundance of fresh fuel. Hence one expects that at a given  $M_*$ , high gas content should correlate with low metallicity and high SFR. This results in the second-parameter dependences with SFR in the scatter around these scaling relations. In this section we quantify these second-parameter dependences in MUFASA, which provides predictions for baryon cycling that are testable with current and future observations.

Figs 3 and 5 already showed clear second-parameter dependences on the SFR in MUFASA galaxies: Galaxies that have higher SFR for their  $M_*$  also have lower metallicities and higher gas fractions in



**Figure 10.** Plots showing the deviation from mean scaling relations versus  $M_*$  in our  $50 h^{-1}$  Mpc volume at  $z = 0$  for four quantities (in logarithm): H I fraction, H<sub>2</sub> fraction, specific SFR and metallicity. Scatter plots show these deviations plotted against each other, depicting how the fluctuations in these quantities are correlated. The trends can be reasonably represented by the best-fitting power laws shown as the solid lines, with the slope indicated in the upper right of each panel. Overall, galaxies at a given  $M_*$  with higher sSFR have higher H I and H<sub>2</sub> fractions and lower metallicity. The panels along the diagonal show histograms of the scatter around each scaling relation. While generally Gaussian, there is a tail to low sSFR and gas content arising from green valley galaxies.

both H I and H<sub>2</sub>. To quantify this, we use ‘deviation plots,’ i.e. we plot the deviation away from the mean scaling relation in two quantities versus  $M_*$  against each other. This isolates the second-order aspects of baryon cycling-driven galaxy evolution by directly quantifying how fluctuations drive correlated scatter, while removing the dependence on the overall inflow rate that sets the first-order (mean) scaling relationship between quantities.

As an example, in order to make a deviation plot for sSFR versus H I, we begin with the sSFR– $M_*$  and  $f_{\text{HI}}-M_*$  relations. For each galaxy, we then compute the difference between  $\log$  sSFR of that galaxy and the median of all galaxies’  $\log$  sSFR at that  $M_*$ ; we call this  $\Delta \log$  sSFR. Similarly, we compute the difference between  $\log f_{\text{HI}}$  for that galaxy and the median  $\log f_{\text{HI}}$  at that galaxy’s  $M_*$ ; this is  $\Delta \log f_{\text{HI}}$ . We can analogously compute  $\Delta \log f_{\text{H}_2}$  and  $\Delta \log Z$  for the molecular gas and metallicity, respectively. Note that here we are always using  $M_*$  as our independent variable, because this quantity is stable on the (relatively) short time-scales over which deviations are occurring; in principle, it is possible to use any property as the independent variable, but we leave such explorations for future work.

Fig. 10 shows 2D histograms of the deviations  $\Delta \log f_{\text{HI}}$ ,  $\Delta \log f_{\text{H}_2}$ ,  $\Delta \log$  sSFR and  $\Delta \log Z$  plotted against each other. Only star-forming galaxies are included, and for simplicity, we only show the  $50 h^{-1}$  Mpc volume at  $z = 0$ , but the trends are similar in the other volumes. The panels along the diagonal show the histograms of deviation values for each quantity, which illustrate the shape of the scatter around the median scaling relation versus  $M_*$ . The solid line in each panel shows the best-fitting power law to the deviations shown, and the number in the upper right corner is the best-fitting slope.

Fig. 10, at its most basic level, shows that deviations in the SFR, H I and H<sub>2</sub> all correlate positively with each other, while metallicity deviations (bottom row) anticorrelate with all of the others. This quantifies the amount by which galaxies that lie above the mean MZR also tend to lie below the mean relations in the sSFR,  $f_{\text{HI}}$  and  $f_{\text{H}_2}$  versus  $M_*$ . Such trends arise naturally in a ‘gas regulator’-type model, which is an ISM mass-balance formalism in which the gas content is allowed to vary (Lilly et al. 2013).

The slope of the best-fitting line contains information about how well quantities track each other. For instance, consider  $\Delta \log f_{\text{H}_2}$



versus  $\Delta \log$  sSFR: the slope is close to linear, which means that the fluctuations in  $H_2$  are directly tracking the fluctuations in SFR. This is unsurprising, since in our simulations, it is assumed that the SFR of any given gas element is proportional to its  $f_{H_2}$ ; none the less, it is not trivial that this translates into a similar trend in galaxy-integrated quantities. The slope versus  $\Delta \log f_{H_1}$ , in contrast, is somewhat sub-linear for  $\Delta \log f_{H_2}$  and  $\Delta \log$  sSFR, indicating that the fluctuations in  $H_2$  and sSFR do not perfectly reflect the fluctuations in  $H_1$ .

The deviations in metallicity versus the MZR, i.e. the plots along the lowermost row, have garnered much attention in the literature. For instance, the panel showing  $\Delta \log$  sSFR versus  $\Delta \log Z$  (lower rightmost) corresponds to the FMR, showing that galaxies with higher SFR at a given  $M_*$  have lower  $Z$ . The best-fitting line has a slope of  $-0.16$ , which represents a higher order testable prediction of MUFASA’s ability to depict the fluctuations in baryon cycle that give rise to scatter around the scaling relations.

The bottom leftmost panel corresponds to the observational trend noted by Hughes et al. (2013), Lara-López et al. (2013) and Bothwell et al. (2013), which the latter dubbed the  $H_1$ -FMR: galaxies with higher  $H_1$  content at a given  $M_*$  are seen to have lower metallicities. Resolved spectroscopy by Moran et al. (2012) indicated that the excess in  $H_1$  tends to be accompanied by a drop in the outer metallicity, strongly suggesting that this trend is driven by accretion in the outskirts of galaxies. MUFASA predicts a slope for  $\Delta \log f_{H_1}$  versus  $\Delta \log Z$  of  $-0.18$ , similar to but slightly stronger than that versus  $\Delta \log$  sSFR. Robertson et al. (2013) measured this deviation slope to be  $-0.41 \pm 0.14$  for field galaxies ( $-0.31$  for cluster galaxies). This is steeper than our current predictions, but this was done at a fixed sSFR rather than  $M_*$ , which likely accounts for some of the difference. Metallicity is formally most strongly tied to  $H_2$ ; the slope of this deviation relation is  $-0.22$ .

One can also examine the spread of points around the best-fitting linear relation within the deviation plots. This is another measure of how tightly any given two quantities fluctuate. One can quantify this by measuring the mean deviation in, say, metallicity, from the best-fitting relations involving  $\Delta \log Z$ . For the metallicity relations, the mean departure in  $\Delta \log Z$  is 0.062, 0.061 and 0.064 with respect to  $\Delta \log f_{H_1}$ ,  $\Delta \log f_{H_2}$  and  $\Delta \log$  sSFR, respectively. This again suggests, at a very marginal level, that metallicity more strongly follows  $H_1$  than sSFR, which is a conclusion also reached in the observational analysis by Bothwell et al. (2013). Still, metallicity tracks  $H_2$  slightly better than either of these quantities (as also found by Lagos et al. 2015), which is perhaps not surprising since  $Z$  here is the SFR-weighted gas-phase metallicity, and thus it is effectively  $f_{H_2}$ -weighted.

Finally, the diagonal panels show the scatter of each quantity around the mean scaling relation versus  $M_*$ . The shape is generally Gaussian, with a spread that is slightly smaller in  $H_2$  relative to sSFR and  $H_1$ . Metallicity has quite small scatter, consistent with  $\sim 0.1$  dex as observed T04. In detail, there is a longer tail to low- $\Delta$ sSFR and correspondingly low gas fraction deviations, which arises from galaxies on their way to quenching.

This deviation plot represents the global view over all galaxies down to the resolution limit of our  $50 h^{-1}$  Mpc volume at  $z = 0$ . Clearly, it is instructive to examine this plot using galaxies binned by mass, or colour or at different redshifts. We do not show this here, but we have checked that the trends depicted in Fig. 10 are generally well converged with resolution in the overlapping mass range, and they are qualitatively similar at higher redshifts. One can also examine trends by fixing other quantities besides  $M_*$ , such as SFR. In future work, we will explore the implications of these

deviation plots in terms of baryon cycling, and present more detailed comparisons to the relevant observations.

In summary, deviations plots quantify how galaxies respond to the fluctuations in the baryon cycle. By examining only the departures around the mean relations (with respect to  $M_*$ ), we remove the ‘first-order’ component of galaxy growth (along with many associated systematics) and isolate the impact of the ‘second-order’ fluctuations on observable quantities. We thus quantify the correlations in scatter among these various quantities, thereby presenting a new and higher order test of galaxy formation models. While there currently exist various forms of the observational characterizations for these trends, we plan to conduct a more thorough and direct comparison to data regarding the second-parameter trends for both gas fractions and metallicities in future work.

## 7 SUMMARY

We have presented predictions of the MUFASA simulations and compared to observations of the SFR, metal and gas content of galaxies. MUFASA uses state-of-the-art feedback modules and hydrodynamics methodology taken from the high-resolution zoom simulations and analytic models. To further extend our dynamic range, we employ several simulations using identical input physics but varying in volume (box sizes of 50, 25,  $12.5 h^{-1}$  Mpc), and check that generally simulations at different numerical resolution make similar predictions in their overlapping mass ranges.

Following Paper I where we showed that MUFASA performed creditably at reproducing the observed stellar mass function over a range of cosmic epochs, here we further show that it also fares well against a number of other key barometers, including many that have not been examined extensively versus previous models such as the sSFRF. We also make novel and testable predictions for the correlations in the fluctuations around mean scaling relations in SFR, metallicity and gas content, as a direct means to quantify how galaxies respond to the fluctuations in the baryon cycle.

Our main results are summarized as follows.

(i) The sSFRF in MUFASA shows a Schechter shape with a relatively shallow faint end, in broad agreement with the observations out to  $z \sim 2$ , albeit with a hint that MUFASA overpredicts high-SFR galaxies. This is curious, given that in Paper I, we demonstrated that MUFASA matches the stellar mass function well but strongly underpredicts the  $z \sim 2$  sSFR- $M_*$  relation, which implies that, if anything, MUFASA should underpredict the SFR function. This highlights the continued difficulty in reconciling current SFR measures during Cosmic Noon with models and, in some cases, among the various data sets themselves.

(ii) The sSFRF provides a more detailed test of how well models reproduce the scatter around the main sequence. MUFASA reproduces the observed sSFR function at low- $z$  quite well, indicating that this simulation is nicely reproducing the number of galaxies in the green valley, is correctly capturing the spread around the main sequence and is not missing a large population of starbursts. At  $z \sim 1$ , the entire sSFR function is shifted by  $\sim \times 2$  with respect to the observations although the shape matches well, reiterating the result from Paper I, showing that the mean sSFR at that epoch is underpredicted by a similar factor.

(iii) The mass-metallicity relation shows a reasonable low-mass slope and amplitude versus the observations at both  $z \approx 0$  and 2. In contrast,  $M_* \gtrsim 2 \times 10^{10} M_\odot$  star-forming galaxies at  $z = 0$  continue to show a strong rise in the MZR that does not agree well with the observations, and then abruptly flattens at roughly the

appropriate metallicity. We conjecture that wind recycling, which plays a key role in setting the MZR at these masses, may be too vigorous in our simulations at these masses, or else these galaxies should have a metal loading factor above our assumed value of unity. Finally, the MZR clearly shows a second-parameter trend that galaxies with high SFR at a given  $M_*$  have lower metallicity.

(iv) MUFASA directly tracks molecular gas; hence, we can separate the gas content into atomic, molecular and ionized. MUFASA well reproduces observations of the total cold gas fraction ( $H\text{I}+H_2$ ) as a function of  $M_*$ , and provides a fair match to the  $H\text{I}$  fraction individually, with a notable deficit at high masses in only our lowest resolution run. Like with the metallicity,  $H\text{I}$  and  $H_2$  content also show a second-parameter trend that galaxies with high SFR tend to have higher gas fractions.

(v) Gas fractions are broadly predicted to increase with redshift, which, at least qualitatively, agrees with the observations. However, the predicted rate of evolution for  $H_2$  ( $\sim \times 2-3$  out to  $z \sim 2$ ) is slower than canonically observed for molecular gas.  $H_2$  evolves similarly across all masses, while  $H\text{I}$  evolves slightly faster at higher masses; there is an order of magnitude more  $H\text{I}$  in an  $M_* = 10^{11} M_\odot$  galaxy at  $z \sim 1-2$  versus today.

(vi) As a result of the rapid  $H\text{I}$  evolution at high masses, the bright end of the HIMF evolves fairly rapidly as well. The predicted HIMF agrees well with the observations at  $z \sim 0$ , and evolves upwards at all masses by  $\sim \times 2-3$  by  $z = 1$ . At  $z = 2$ , we predict a steeper faint end, although this may not be accessibly observationally in 21 cm prior to the full SKA.

(vii) In order to explore the potential discrepancy in molecular gas evolution further, we use the simulation-based prescription from Narayanan et al. (2012) to convert MUFASA molecular gas masses to a CO luminosity based on the metallicity and molecular gas content, and compare to inferred COLFs observed out to  $z \sim 2$ . We find surprisingly good agreement at all masses for the COLF, despite the mild evolution in  $f_{H_2}$ . This highlights that systematic uncertainties in  $X_{CO}$  can be an overriding factor in making robust comparisons to molecular gas content data at intermediate redshifts.

(viii) The cosmic mass density in  $H\text{I}$  is predicted to evolve mildly upwards out to high- $z$ , such that  $\Omega_{H\text{I}} \propto (1+z)^{0.7-0.8}$ . This evolution is in good agreement with the observations from a variety of techniques. The amplitude is somewhat sensitive to resolution, and our  $25 h^{-1}$  Mpc volume has 10–20 per cent higher  $\Omega_{H\text{I}}$  than our  $50 h^{-1}$  Mpc cube, which agrees slightly better with data.

(ix) In accord with our predicted mild evolution for  $f_{H_2}$ , we also predict mild evolution for  $\Omega_{H_2}$ , with a similar redshift scaling as  $\Omega_{H\text{I}}$  but lower by  $\sim \times 3$ . We note that this is much less steep than the evolution of the cosmic SFR density; it has been suggested that the drop in cosmic SFR owes directly to the drop in molecular gas mass, but MUFASA does not support this interpretation, as the cosmic SFRD (see Paper I) evolves significantly more rapidly than  $\Omega_{H_2}$ .

(x) An independent and higher order test of models is whether they reproduce the observed scatter around the mean scaling relations. We make predictions for this using deviation plots, where we correlate the deviations for each galaxy relative to scaling relations in the sSFR, metallicity,  $f_{H_2}$  and  $f_{H\text{I}}$  versus  $M_*$ . We show that MUFASA qualitatively reproduces observed trends that indicate that at a given  $M_*$ , galaxies with high SFR have low metallicity and high gas content. We make predictions for the power-law slopes between deviations in the sSFR, metallicity,  $f_{H_2}$  and  $f_{H\text{I}}$  that can be tested against the observations.

As in Paper I, MUFASA continues to demonstrate good agreement with now a wider range of galaxy observables across cosmic time,

indicating that it provides a viable platform to study the physics of galaxy evolution in a cosmological context on  $\gtrsim$  kiloparsec scales. This implies, among other things, that employing scalings taken from the FIRE simulations into cosmological-scale runs satisfyingly reproduces some of the same data-concordant trends as individual FIRE zoom runs. Our current heuristic quenching model seems to populate the green valley and lower their gas contents relative to blue cloud galaxies approximately as observed, though we are working towards a more self-consistent black hole growth and feedback model that may substantially affect these predictions particularly at the massive end. Together with recent results from other cosmological simulations such as EAGLE, these successes showcase the emerging promise of cosmologically situated galaxy formation simulations in helping to understand the Universe as mapped through large-scale multiwavelength galaxy surveys probing the various constituents of galaxies back to early epochs.

## ACKNOWLEDGEMENTS

The authors thank D. Anglés-Alcázar, F. Durier, K. Finlator, S. Huang, N. Katz, D. Narayanan and R. Somerville for helpful conversations and comments. RD, MR and RJT acknowledge support from the South African Research Chairs Initiative and the South African National Research Foundation. Support for MR was also provided by the Square Kilometre Array post-graduate bursary program. Support for RD was provided by National Aeronautics and Space Administration (NASA) ATP grant NNX12AH86G to the University of Arizona. Support for PFH was provided by an Alfred P. Sloan Research Fellowship, NASA ATP Grant NNX14AH35G and National Science Foundation Collaborative Research Grant #1411920 and CAREER grant #1455342. The MUFASA simulations were run on the Pumbaa astrophysics computing cluster hosted at the University of the Western Cape, which was generously funded by UWC's Office of the Deputy Vice Chancellor. These simulations were run with revision e77f814 of GIZMO hosted at <https://bitbucket.org/rthompson/gizmo>.

## REFERENCES

- Andrews B., Martini P., 2013, ApJ, 765, 140 (AM13)  
 Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, ARA&A, 47, 481  
 Bahé Y. M. et al., 2016, MNRAS, 456, 1115  
 Bahé Y. M., Schaye J., Crain R. A., McCarthy I. G., Bower R. G., Theuns T., McGee S. L., Trayford J. W., 2017, MNRAS, 464, 508  
 Behroozi P. S., Wechsler R. H., Conroy C., 2013, ApJ, 770, 57  
 Benson A. J., 2014, MNRAS, 444, 2599  
 Boselli A., Cortese L., Boquien M., Boissier S., Catinella B., Lagos C., Saintonge A., 2014, A&A, 564, 66  
 Bothwell M. S., Kenicutt R. C., Johnson B. D., Wu Y., Lee J. C., Dale D., Engelbracht C., Calzetti D., Skillman E., 2011, MNRAS, 415, 1815  
 Bothwell M. S., Maiolino R., Kennicutt R., Cresci G., Mannucci F., Marconi A., Ciccone C., 2013, MNRAS, 433, 1425  
 Bouché N., Gardner J. P., Katz N., Weinberg D. H., Davé R., Lowenthal J. D., 2005, ApJ, 628, 89  
 Bryan G. L. et al., 2014, ApJS, 211, 19  
 Brown J. S., Martini P., Andrews B. H., 2016, MNRAS, 458, 1529  
 Bruzual G., Charlot S., 2003, MNRAS, 344, 1000  
 Catinella B. et al., 2010, MNRAS, 403, 683  
 Catinella B. et al., 2013, MNRAS, 436, 34  
 Chabrier G., 2003, PASP, 115, 763  
 Colbert J. W. et al., 2013, ApJ, 779, 34  
 Crain R. A. et al., 2015, MNRAS, 450, 1937  
 Crain R. A. et al., 2017, MNRAS, 464, 4204  
 Croton D. J. et al., 2006, MNRAS, 365, 11

- Davé R., 2008, *MNRAS*, 385, 147
- Davé R., Finlator K. M., Oppenheimer B. D., 2011, *MNRAS*, 416, 1354
- Davé R., Finlator K. M., 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. J., Hopkins P. F., 2016, *MNRAS*, 462, 3265 (Paper I)
- Decarli R. et al., 2016, *ApJ*, 833, 70
- Dekel A. et al., 2009, *Nature*, 457, 451
- Delhaize J., Meyer M. J., Staveley-Smith L., Boyle B. J., 2013, *MNRAS*, 433, 1398
- Duffy A. R., Schaye J., Kay S. T., Dalla Vecchia C., 2008, *MNRAS*, 390, L64
- Ellison S. L., Patton D. R., Simard L., McConnachie A. W., 2008, *ApJ*, 672, L107
- Erb D. K., Shapley A. E., Pettini M., Steidel C. C., Reddy N. A., Adelberger K. L., 2006, *ApJ*, 644, 813
- Faucher-Giguere C. A., Lidz A., Zaldarriaga M., Hernquist L., 2009, *ApJ*, 703, 1416
- Feldmann R., Gnedin N. Y., Kravtsov A. V., 2012, *ApJ*, 747, 124
- Fernández X. et al., 2016, *ApJ*, 824, 1
- Finlator K., Davé R., 2008, *MNRAS*, 385, 2181
- Font-Ribera A. et al., 2012, *J. Cosmol. Astropart. Phys.*, 11, 59
- Forbes J. C., Krumholz M. R., Burkert A., Dekel A., 2014, *MNRAS*, 443, 168
- Furlong M., Bower R. G., Theuns T. et al., 2015, *MNRAS*, 450, 4486
- Gabor J. M., Davé R., 2012, *MNRAS*, 427, 1816
- Gabor J. M., Davé R., 2015, *MNRAS*, 447, 374
- Geach J. E., Smail I., Moran S. M., MacArthur L. A., Lagos C. P., Edge A. C., 2011, *ApJ*, 730, 19
- Genel S., 2016, *ApJ*, 822, 107
- Genel S. et al., 2014, *MNRAS*, 445, 175
- Gnedin N. Y., Kravtsov A. V., 2011, *ApJ*, 728, 88
- Grasshorn Gebhardt H. S., Zeimann G. R., Ciardullo R., Gronwall C., Hagen A., Bridge J. S., Schneider D. P., Trump J. R., 2016, *ApJ*, 817, 10
- Gunawardhana M. L. P. et al., 2013, *MNRAS*, 433, 2764
- Hahn O., Abel T., 2011, *MNRAS*, 415, 2101
- Haynes M. P. et al., 2011, *AJ*, 142, 170
- Holwerda B. W., Blyth S.-L., Baker A. J., 2012, *IAUS*, 284, 496
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hughes T. M., Cortese L., Boselli A., Gavazzi G., Davies J. I., 2013, *A&A*, 550, A115
- Ibert O. et al., 2015, *A&A*, 579, 2
- Iwamoto K., Brachwitz F., Nomoto K., Kishimoto N., Umeda H., Hix W. R., Thielemann F.-K., 1999, *ApJS*, 125, 439
- Kaviraj S. et al., 2016, *MNRAS*, preprint ([arXiv:1605.09379](https://arxiv.org/abs/1605.09379))
- Kennicutt R. C., 1998, *ApJ*, 498, 541
- Kereš D., Yun M. S., Young J. S., 2003, *ApJ*, 582, 659
- Kereš D., Katz N., Weinberg D. H., Davé R., 2005, *MNRAS*, 363, 2
- Kewley L. J., Ellison S. L., 2008, *ApJ*, 681, 1183
- Kim J.-h., Abel T., Agertz O. et al., 2014, *ApJS*, 210, 14
- Kriek M. et al., 2015, *ApJS*, 218, 15
- Krumholz M. R., McKee C. F., Tumlinson J. T., 2008, *ApJ*, 693, 216
- Krumholz M. R., McKee C. F., Tumlinson J. T., 2011, *ApJ*, 729, 36
- Lagos C. Del P., Baugh C. M., Lacey C. G., Benson A. J., Kim H.-S., Power C., 2011, *MNRAS*, 418, 1649
- Lagos C. Del P. et al., 2015, *MNRAS*, 452, 3815
- Lagos C. Del P. et al., 2016, *MNRAS*, 459, 2632
- Lara-López M. A. et al., 2010, *A&A*, 521, L53
- Lara-López M. A. et al., 2013, *MNRAS*, 433, L35
- Lilly S. J., Carollo C. M., Pipino A., Renzini A., Peng Y., 2013, *ApJ*, 772, 119
- McNamara B. R., Nulsen P. E. J., 2007, *ARA&A*, 45, 117
- Ma X., Hopkins P. F., Faucher-Giguere C.-A., Zolman N., Muratov A. L., Kereš D., Quataert E., 2016, *MNRAS*, 456, 2140
- Madau P., Dickinson M., 2014, *ARA&A*, 52, 415
- Mannucci F., Cresci G., Maiolino R., Marconi A., Gnerucci A., 2010, *MNRAS*, 408, 2115
- Mehta V. et al., 2016, *ApJ*, preprint ([arXiv:1505.07843](https://arxiv.org/abs/1505.07843))
- Mitra S., Davé R., Finlator K., 2015, *MNRAS*, 452, 1184
- Mitra S., Davé R., Simha V., Finlator K., 2017, *MNRAS*, 464, 2766
- Mo H. J., Mao S., White S. D. M., 1998, *MNRAS*, 295, 319
- Moran S. M. et al., 2012, *ApJ*, 745, 66
- Muratov A. L., Kereš D., Faucher-Giguere C. A., Hopkins P. F., Quataert E., Murray N., 2015, *MNRAS*, 454, 2691
- Narayanan D., Krumholz M. R., Ostriker E. C., Hernquist L., 2012, *MNRAS*, 421, 3127
- Narayanan D., Bothwell M., Davé R., 2012, *MNRAS*, 426, 1178
- Nomoto K., Tominaga N., Umeda H., Kobayashi C., Maeda K., 2006, *Nucl. Phys. A*, 777, 424
- Noterdaeme P. et al., 2012, *A&A*, 547, L1
- Obreschkow D., Croton D., De Lucia G., Khochfar S., Rawlings S., 2009, *ApJ*, 698, 1467
- Oppenheimer B. D., Davé R., 2008, *MNRAS*, 387, 577
- Pasquali A., Gallazzi A., Fontanot F., van den Bosch F. C., De Lucia G., Mo H. J., Yang X., 2010, *MNRAS*, 407, 937
- Peebles M. S., Shankar F., 2011, *MNRAS*, 417, 2962
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Popping A., Davé R., Braun R., Oppenheimer B. D., 2009, *A&A*, 504, 15
- Popping G., Somerville R. S., Trager S. C., 2014, *MNRAS*, 442, 2398
- Popping G., Behroozi P. S., Peebles M. S., 2015a, *MNRAS*, 449, 477
- Popping G. et al., 2015b, *MNRAS*, 454, 2258
- Prochaska J. X., Wolfe A. M., 2009, *ApJ*, 696, 1543
- Rafieferantsoa M., Davé R., Anglés-Alcázar D., Katz N., Kollmeier J. A., Oppenheimer B. D., 2015, *MNRAS*, 453, 3980
- Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, *MNRAS*, 430, 2427
- Rao S. M., Turnshek D. A., Nestor D. B., 2006, *ApJ*, 636, 610
- Robertson P., Shields G. A., Davé R., Blanc G. A., Wright A., 2013, *ApJ*, 773, 4
- Saintonge R. L. et al., 2011, *MNRAS*, 415, 32
- Saintonge R. L. et al., 2016, *MNRAS*, 462, 1749
- Salim S., Lee J. C., Ly C., Brinchmann J., Davé R., Dickinson M., Salzer J. J., Charlot S., 2014, *ApJ*, 797, 126
- Salim S., Lee J. C., Davé R., Dickinson M., 2015, *ApJ*, 808, 25
- Sanders R. L. et al., 2015, *ApJ*, 799, 138
- Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Scoville N. et al., 2016, *ApJ*, 820, 83
- Somerville R. S., Davé R., 2015, *ARA&A*, 53, 51
- Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, *MNRAS*, 391, 481
- Sparre M. et al., 2015, *MNRAS*, 447, 3548
- Springel V., 2005, *MNRAS*, 364, 1105
- Steidel C. C. et al., 2014, *ApJ*, 795, 165
- Steidel C. C., Strom A. L., Pettini M., Rudie G. C., Reddy N. A., Trainor R. F., 2016, *ApJ*, 826, 159
- Tacchella L. J., Dekel A., Carollo C. M., Ceverino D., DeGraf C., Lapiner S., Mandelker N., Primack J. R., 2016, *MNRAS*, 458, 242
- Tacconi L. J. et al., 2013, *ApJ*, 768, 74
- Telford O. G., Dalcanton J. J., Skillman E. D., Conroy C., 2016, *ApJ*, 827, 35
- Thompson R., 2015, *Astrophysics Source Code Library*, record ascl:1502.012
- Thompson R., Nagamine K., Jaacks J., Choi J.-H., 2014, *ApJ*, 780, 145
- Tremonti C. A. et al., 2004, *ApJ*, 613, 898 (T04)
- Vagnozzi S., Freese K., Zurbuchen T. H., 2016, preprint ([arXiv:1603.05960](https://arxiv.org/abs/1603.05960))
- Vallini L., Gruppioni C., Pozzi F., Vignali C., Zamorani G., 2016, *MNRAS*, 456, L40
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
- Weinmann S. M., Pasquali A., Oppenheimer B. D., Finlator K., Mendel J. T., Crain R. A., Maccio A. V., 2012, *MNRAS*, 426, 2797
- White C. E., Somerville R. S., Ferguson H. C., 2015, *ApJ*, 799, 201
- Zahid H. J. et al., 2014, *ApJ*, 792, 75