

The multiphase circumgalactic medium traced by low metal ions in EAGLE zoom simulations

Benjamin D. Oppenheimer,¹★ Joop Schaye², Robert A. Crain³, Jessica K. Werk⁴ and Alexander J. Richings⁵

¹CASA, Department of Astrophysical and Planetary Sciences, University of Colorado, 389 UCB, Boulder, CO 80309, USA

²Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA, Leiden, the Netherlands

³Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool, L3 5RF, UK

⁴Department of Astronomy, University of Washington, Seattle, WA 98195–1580, USA

⁵Department of Physics and Astronomy and CIERA, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

Accepted 2018 August 17. Received 2018 August 17; in original form 2017 September 20

ABSTRACT

We explore the circumgalactic metal content traced by commonly observed low ion absorbers, including C II, Si II, Si III, Si IV, and Mg II. We use a set of cosmological hydrodynamical zoom simulations run with the EAGLE model and including a non-equilibrium ionization and cooling module that follows 136 ions. The simulations of $z \approx 0.2$ L^* ($M_{200} = 10^{11.7} - 10^{12.3} M_{\odot}$) haloes hosting star-forming galaxies and group-sized ($M_{200} = 10^{12.7} - 10^{13.3} M_{\odot}$) haloes hosting mainly passive galaxies reproduce key trends observed by the COS–Halos survey – low ion column densities show 1) little dependence on galaxy-specific star formation rate, 2) a patchy covering fraction indicative of 10^4 K clumps with a small volume filling factor, and 3) a declining covering fraction as impact parameter increases from 20–160 kpc. Simulated Si II, Si III, Si IV, C II, and C III column densities show good (mostly within 0.3 dex) agreement with observations, while Mg II is under-predicted. Low ions trace a significant metal reservoir, $\approx 10^8 M_{\odot}$, residing primarily at 10–100 kpc from star-forming and passive central galaxies. These clumps preferentially flow inwards and most will accrete onto the central galaxy within the next several Gyr, while a small fraction are entrained in strong outflows. A multiphase structure describes the inner CGM ($< 0.5 R_{200}$) with low-ion metal clumps surrounded by a hot, ambient medium. The dense, cool phase is separate from the O VI observed by COS–Halos, which arises from the outer CGM ($> 0.5 R_{200}$) tracing virial temperature gas around L^* galaxies. Our simulations support previous ionization models indicating that cloud covering factors decline while densities and pressures show little decline with increasing impact parameter (typically < 0.3 dex from 40 to 160 kpc). We find the cool clumps have lower pressures than the ambient medium they are embedded in, and discuss that numerical effects within the hydrodynamic solver likely play a role.

Key words: galaxies: evolution, formation, halo – galaxies: intergalactic medium – cosmology: theory – quasars: absorption lines.

1 INTRODUCTION

The circumgalactic medium (CGM) is thought to contain significant reservoirs of baryons and metals outside of galaxies, extending to the virial radius and beyond (e.g. Tumlinson et al. 2011; Stocke et al. 2013; Johnson, Chen & Mulchaey 2015). Absorption line spectroscopic observations by the Cosmic Origins Spectrograph (COS) on the *Hubble Space Telescope* allow the study of the CGM

around galaxies at redshift $z \lesssim 0.5$, where it is easier to characterize the galaxies’ properties observationally, including their stellar mass, star formation rates (SFRs), and morphologies. The far-ultraviolet (FUV) spectral range of the COS instrument (1100–1700 Å) covers numerous electronic transition lines of metal species, including C II, C III, C IV, Si II, Si III, Si IV, and O VI that can probe the physical state of the gas around galaxies in the evolved Universe.

The COS–Halos survey (Tumlinson et al. 2013) exploited the full FUV spectral range of COS, targeting a series of 44 $z \approx 0.2$ galaxies spanning stellar masses $M_* = 10^{9.6} - 10^{11.3} M_{\odot}$ to explore

* E-mail: benjamin.oppenheimer@colorado.edu

the CGM properties out to an impact parameter, $b = 160$ kpc. The galaxies were selected to be either “blue” or “red,” where the blue sample comprises star-forming galaxies and the red sample comprises passive galaxies with little detectable star formation. Throughout, we define the COS–Halos blue and red samples as galaxies with specific star formation rates ($\text{sSFR} \equiv \text{SFR}/M_*$) greater than or less than 10^{-11} yr^{-1} , respectively. Tumlinson et al. (2011) showed that this division in galaxy properties is reflected in the CGM properties probed by O VI with the blue star-forming sample showing significantly higher O VI column densities ($N_{\text{O VI}}$) than the red passive sample.

However, the low metal ions do not show the same behaviour as O VI. Firstly, unlike O VI, the low ions do not show an obvious dependence on sSFR (Werk et al. 2013, hereafter W13). Secondly, W13 observed a large scatter in the column densities of ions such as C II, C III, Si II, Si III, and Mg II. W13 argued that the large dispersion in low-ion absorption strengths suggests that the cool CGM is patchy in nature and hence spans a large range of densities and/or ionization conditions. This contrasts with O VI, which shows a significantly smaller spread in column densities around the blue star-forming galaxies (Peeples et al. 2014). A third difference is that the covering fractions of low ions decline at larger b when splitting the blue galaxy sample into two impact parameters bins divided at $b = 75$ kpc (W13). O VI shows a much smaller decline in column density and covering fraction with impact parameter (Tumlinson et al. 2011; Johnson et al. 2015). Throughout this paper, we use the general term “low” ion for any metal ion that is not O VI, even though C IV and Si IV are usually considered intermediate ions (e.g. W13).

The binary division of COS–Halo galaxies into active and passive categories based on an sSFR cut belies the complexity of galaxy properties in relation to CGM observables. While we note that the two samples also have far different median stellar masses, $10^{10.1}$ and $10^{10.8} M_\odot$ for the active and passive samples, respectively, Tumlinson et al. (2011) demonstrated that there still exists a statistical division in O VI column densities when using only massive COS–Halos galaxies with $M_* > 10^{10.5} M_\odot$. More recently, surveys have targeted specific types of galaxies representing key evolutionary stages for galaxies in the COS–Halos mass range. Heckman et al. (2017) targeted starbursts in the COS–Burst survey with a median $M_* = 10^{10.4} M_\odot$ and $\text{sSFR} = 10^{-9.2} \text{ yr}^{-1}$ finding stronger Si III and C IV than a control sample without starbursts mainly at larger radii than we explore here. Berg et al. (2018) targeted Active Galactic Nuclei (AGN) with median bolometric luminosities calculated to be $10^{42.9} \text{ erg s}^{-1}$, $M_* = 10^{10.7} M_\odot$, and $\text{sSFR} = 10^{-10.8} \text{ yr}^{-1}$ finding similar Si III strengths compared to a control sample without AGN. Both of these surveys have a majority of their targets at impact parameters larger than 160 kpc, so a direct comparison to COS–Halos preferentially targeting the inner CGM is not possible.

Werk et al. (2014, hereafter W14) performed photoionization modelling using CLOUDY (Ferland et al. 1998) to derive the physical properties traced by H I and low metal ions believed to trace temperatures $T \sim 10^4$ K. W14 found the gas density as a function of impact parameter to decline from a hydrogen number density $n_{\text{H}} \sim 10^{-3} \text{ cm}^{-3}$ inside $b \sim 30$ kpc to $\sim 10^{-4} \text{ cm}^{-3}$ at $b \gtrsim 100$ kpc. A two-phase model based on Maller & Bullock (2004) with cool $T \sim 10^4$ K clouds embedded in a hot $T \sim 10^6$ K halo medium is in tension with these derived physical parameters. This model assumes hydrostatic equilibrium with a Navarro, Frenk & White (1997, NFW) dark matter halo potential within a cooling radius, and predicts cool CGM densities more than $100 \times$ higher than inferred from the COS–Halos observations when combined with single-phase CLOUDY models. However, the Maller & Bullock (2004) model does not ac-

count for mechanical or thermal superwind feedback imparted by star formation (SF)-driven or AGN feedback.

Cosmological hydrodynamic simulations that reproduce the observed properties of galaxies require superwind feedback to eject baryons and metals from galaxies into the CGM and intergalactic medium (IGM) to reduce the efficiency of galactic stellar build-up (e.g. Springel & Hernquist 2003; Oppenheimer et al. 2010; Schaye et al. 2010). A consequence of metal-enriched material leaving galaxies is an enriched IGM/CGM (e.g. Aguirre et al. 2001; Theuns et al. 2002; Cen & Fang 2006; Oppenheimer & Davé 2006; Wiersma et al. 2010; Smith et al. 2011). The Evolution and Assembly of GaLaxies and their Environments (EAGLE) simulation project calibrated the sub-resolution prescriptions for SF and AGN feedback to reproduce the observed $z \approx 0.1$ galactic stellar mass function, as well as the galactic disk size and super-massive black hole– M_* relations (Schaye et al. 2015, hereafter S15; Crain et al. 2015). Because observations of the CGM and IGM were not used to calibrate the EAGLE simulations, the properties of gas outside galaxies are genuine predictions of the model.

Oppenheimer et al. (2016, hereafter O16) integrated the non-equilibrium (NEQ) ionization and dynamical cooling module introduced in Oppenheimer & Schaye (2013a) into the EAGLE simulation code to trace the evolution of 136 ions of 11 elements. O16 ran a set of 20 zoom simulations of individual galactic haloes, 10 of which host blue, star-forming galaxies and have virial masses $\sim 10^{12} M_\odot$ and another 10 haloes at $\sim 10^{13} M_\odot$ most of which host red, passive galaxies. They turned on the NEQ module at low redshift to follow non-equilibrium effects over the redshift range of COS–Halos galaxies. They found that for a COS–Halos-like sample, O VI is strongest around blue galaxies, because the temperatures of the virialized gas in their host haloes overlap with the 3×10^5 K collisional ionization temperature of O VI. O VI is less strong around red COS–Halos galaxies, because their host halo virial temperatures exceed 10^6 K resulting in CGM oxygen being promoted to O VII and above. O16 argued that the correlation between circumgalactic $N_{\text{O VI}}$ and galactic sSFR observed by Tumlinson et al. (2011) is not causal, but reflects the increasing ionization state of oxygen with virial mass and temperature.

The same COS–Halos sight lines that show O VI often also show low metal ions and H I (Thom et al. 2012), implying that the CGM is multiphase. The EAGLE NEQ zoom hydrodynamic simulations are well-suited for a study of the multiphase CGM. Our zooms self-consistently follow the nucleosynthetic production of heavy elements in stars, their propagation out of galaxies due to superwind feedback, and the detailed non-equilibrium atomic processes setting the ionization states in the CGM. Here, we extend the work of O16 to the COS–Halos low metal ions in the same zooms that were used by O16 to explain the $N_{\text{O VI}}$ –sSFR correlation. We mention that the recent work of Oppenheimer et al. (2018) includes fluctuating AGN radiation added to one of our zooms, which results in enhanced O VI column densities around COS–Halos galaxies. We discuss this work throughout our investigation of low ions, but note that these ions, unlike O VI, are not nearly as strongly affected by fluctuating AGN radiation.

Throughout this paper, we compare and contrast our results to other simulation work, although very few of them have considered the low metal ions in comparison to CGM observations. Investigations using grid-based codes, including Hummels et al. (2013, ENZO) and Liang, Kravtsov & Agertz (2016, RAMSES) as well as Smooth Particle Hydrodynamic (SPH) codes (Ford et al. 2016, GADGET-2) find low-ion column densities, and in particular Mg II, are lower than observational results. A general

challenge appears to be that simulations are incapable of matching CGM low ion strengths while making galaxies with realistic properties.

The paper is organized as follows. We describe our simulations and review our non-equilibrium module in Section 2. We describe how the low ion-traced CGM changes as a function of halo mass in Section 3, and then compare directly with COS–Halos observations in Section 4. The physical and evolutionary state of low ions is addressed in Section 5 with discussions of metal masses, physical gas parameters, evolution of low ion-traced gas from $z = 0.2 \rightarrow 0.0$, and ion ratios. We summarize in Section 6. Resolution and non-equilibrium effects are explored in the Appendix, as well as statistical methods.

2 SIMULATIONS

We briefly describe the simulations in this section, and refer the reader to §2 of O16 for further details. We employ the EAGLE hydrodynamic simulation code described in S15, which is an extensively modified version of the N-body+SPH GADGET-3 code last described in Springel (2005). We assume the Planck Collaboration (2014) cosmological parameters adopted in EAGLE simulations: $\Omega_m = 0.307$, $\Omega_\Lambda = 0.693$, $\Omega_b = 0.04825$, $H_0 = 67.77 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\sigma_8 = 0.8288$, and $n_s = 0.9611$. EAGLE uses the Hopkins (2013) pressure–entropy SPH formulation applying a C2 Wendland (1995) 58-neighbour kernel along with several other hydrodynamic modifications collectively referred to as “Anarchy” (Appendix A of S15 and Schaller et al. 2015). The EAGLE code includes subgrid prescriptions for radiative cooling (Wiersma, Schaye & Smith 2009a), star formation (Schaye & Dalla Vecchia 2008), stellar evolution and chemical enrichment (Wiersma et al. 2009b), and superwind feedback associated with star formation (Dalla Vecchia & Schaye 2012) and black hole growth (S15; Rosas-Guevara et al. 2015).

EAGLE provides an ideal testbed for the study of the CGM, because it successfully reproduces an array of galaxy observables (e.g. S15; Furlong et al. 2015, 2017; Trayford et al. 2015; Bahé et al. 2016; Segers et al. 2016) in a model that explicitly follows the hydrodynamics. Even though the EAGLE model was not calibrated on observations of the IGM/CGM, EAGLE simulations show broad but imperfect agreement with absorption line statistics probing H I (Rahmati et al. 2015) and metal ions (Rahmati et al. 2016; Turner et al. 2016, 2017), and the O VI bimodality observed around COS–Halos galaxies (O16). In particular, Rahmati et al. (2016) finds too few strong O VI absorbers with $>10^{14.5} \text{ cm}^{-2}$ when probing IGM sight lines, which agrees with the O16 under-estimate from NEQ zooms that use the same EAGLE prescription. The Oppenheimer et al. (2018) application of flickering AGN was explicitly done to address the shortfall in O VI. Additionally, Turner et al. (2016) finds EAGLE significantly under-estimates C IV pixel optical depths as a function of H I pixel optical depth at $z \approx 3.5$, although we will show that the carbon ions at low- z agree well with our higher resolution simulations. Pixel optical depths are also highly sensitive to small-scale velocities as Turner et al. (2016) shows that adding subgrid turbulence post-processing helps with some of the C IV discrepancy; however, our approach of simulating integrated system column densities would be unaffected by this.

2.1 Non-equilibrium network

The NEQ module (Richings, Schaye & Oppenheimer 2014), integrated into the EAGLE GADGET-3 simulation code by O16, explic-

itly follows the reaction network of 136 ionization states for the 11 elements that significantly contribute to the cooling (H, He, C, N, O, Ne, Si, Mg, S, Ca, & Fe) plus the electron density of the plasma. Our reaction network is described in Oppenheimer & Schaye (2013a). It includes radiative and di-electric recombination, collisional ionization, photo-ionization, Auger ionization, and charge transfer. Cooling is summed ion-by-ion (Gnat & Ferland 2012; Oppenheimer & Schaye 2013a) over all 136 ions. The method has been verified to reproduce results obtained from other codes and is interchangeable with the equilibrium elemental cooling tables of Wiersma et al. (2009a) that were used in other EAGLE runs.

Our EAGLE zooms assume an interstellar medium (ISM), defined as gas having non-zero SFR, with a single phase where we do not follow the NEQ behaviour and instead use equilibrium lookup tables for ionization fractions tabulated as functions of density assuming $T = 10^4 \text{ K}$. This makes little difference for most ions, but it can affect the balance between the lowest ion states such as Si I and Si II or Mg I and Mg II. However, the covering factor of the star-forming ISM in a mock COS–Halos survey using our simulations is insignificant. Metal enrichment from stars onto gas particles releases new metals in their ground-state ions. However, the vast majority of the enrichment occurs in ISM gas where ISM equilibrium tables are used.

We run simulations using the standard “equilibrium” EAGLE code to low redshift and then turn on the NEQ network at $z \leq 0.5$ as described in §2.3 of O16. The only difference with standard EAGLE runs (S15), which use kernel-smoothed metal abundances, is that we use particle-based metal abundances in all EAGLE equilibrium runs and particle-based ion abundances in NEQ runs. Appendix B1 of O16 found that circumgalactic O VI is nearly unchanged when using particle-based instead of kernel-smoothed metallicities, but stellar masses decline by 0.1 dex when using particle-based metallicities. We compare the different metallicity smoothing schemes for L^* haloes in Appendix B.

2.2 Runs

We use the set of zoom simulations listed in table 1 of O16, but we also add a 12.5 Mpc simulation periodic volume described in detail below. Our main resolution is the $M5.3$ resolution of O16 corresponding to an initial SPH particle mass $m_{\text{SPH}} = 2.2 \times 10^5 M_\odot$, using the notation $M[\log(m_{\text{SPH}}/M_\odot)]$. This resolution has a Plummer-equivalent softening length of 350 proper pc at $z < 2.8$, and 1.33 comoving kpc at $z > 2.8$.

Zooms: Twenty zoom simulations centred on haloes with mass $M_{200} = 10^{11.8} - 10^{13.2} M_\odot$, where M_{200} is the mass within a sphere within which the mean internal density is $200 \times$ the critical overdensity. Ten haloes corresponding to “ L^* ” masses ($M_{200} = 10^{11.7} - 10^{12.3} M_\odot$) were selected from the EAGLE Recal-L025N0752 simulation and 10 haloes corresponding to “group” masses ($M_{200} = 10^{12.7} - 10^{13.3} M_\odot$) were selected from the Ref-L100N1504 simulation. Additionally, several zooms contain “bonus” haloes that were verified to reside completely within the region resolved with the high-resolution SPH and dark matter particles.

We only activate the NEQ module at low redshift in order to reduce computational cost, and because the NEQ effects on CGM ionization levels are short-lived compared to the Hubble timescale. L^* (group) zooms are run using the NEQ module beginning at $z = 0.503$ (0.282). We use outputs of zooms at $z = 0.250$, 0.205, and 0.149 in our analysis here. O16 used additional outputs at $z = 0.099 - 0.0$ to obtain a wider range of galaxy properties to simulate COS–Halos, but we decided not to use these additional outputs

since they do not overlap with COS–Halos redshifts and they do not statistically alter the simulation results. Standard equilibrium EAGLE runs are also run to $z = 0$ and we use these runs at $z = 0.20$ for comparison to NEQ runs in Section 4.2.

Periodic Volumes: We add a 12.5 Mpc, 376^3 SPH + DM particle simulation to our analysis here, which is a Recal-L012N0376 simulation using EAGLE terminology and also has $M5.3$ resolution. This simulation was run in NEQ from $z = 0.503 \rightarrow 0.0$, and we use outputs at $z = 0.351, 0.250, 0.205$, and 0.149 . This volume contains several L^*/group halos, which we add to our halo sample, plus a large range of haloes with $M_{200} < 10^{11.7} M_\odot$, which we term “sub- L^* ” haloes. This allows us to simulate the three lower mass COS–Halos galaxies at $M_* < 10^{9.7} M_\odot$, which were removed from the comparison in O16. This periodic simulation was run with the same parameters as the zooms.

2.3 Isolation criteria

For each central galaxy, we test whether it is defined as “isolated” using similar criteria as those used to select the COS–Halos sample. However, there exists some ambiguity in how isolated the COS–Halos galaxies truly are. Tumlinson et al. (2013) reports that COS–Halos galaxies are “the most luminous galaxy within 300 kpc of the QSO sight line at its redshift.” The spectroscopic and photometric galaxy field follow-up of Werk et al. (2012) agreed with this statement, but additionally found many $L > 0.1L^*$ galaxies often within 160 kpc of the targeted COS–Halos galaxy. The initial COS–Halos galaxy selection used only photometric redshifts to select “isolated” galaxies, so it is not surprising that deeper follow-up has resulted in the discovery of lower luminosity neighbouring galaxies at similar redshifts as described in detail in W13.

We therefore make two isolation criteria: 1) the “stringent” criteria that there should not be any galaxies within $b = 300$ kpc having $M_* > 2 \times 10^{10} M_\odot$ applied in O16, and 2) the “loose” isolation criteria that reduces b to 100 kpc but also reduces the minimum stellar mass to $M_* > 10^{10} M_\odot$. We project all central galaxies in three directions (x , y , & z) and test the criteria in each direction. Nearly, all L^* galaxies satisfy both isolation criteria, while over half of group galaxy directions are thrown out using the stringent criteria, which reduces to ≈ 20 per cent using the loose criteria. O16 used the stringent criteria for O VI, but we favour the loose criteria for this work given the follow-up of Werk et al. (2012). This is because the stringent criteria result in the elimination of about half of the prospective passive COS–Halos targets from the COS–Halos sample selection, and such a cut was not applied to that survey (J. Tumlinson, private communication).

We will show in Section 4 that the chosen isolation criteria make a more significant difference for low ions than for O VI around group galaxies. $N_{\text{O VI}}$ increased by ≈ 0.2 dex upon eliminating the stringent isolation criteria with no isolation criteria around group galaxies (O16). The fit to the COS–Halos O VI using the new loose criteria is essentially identical to the stringent criteria, because the passive galaxy sight lines are mostly upper limits. In general, we use the loose isolation criteria in our mock observational samples; however, we will compare to the stringent isolation criteria in certain instances.

3 LOW METAL IONS IN THE CIRCUMGALACTIC MEDIUM

We begin our presentation of results by considering the CGM as traced by different ions within 300 kpc of galaxies as a function

of halo mass. Our purpose is to provide an understanding of how observations of column density as a function of impact parameter depend on host halo mass before we compare directly to COS–Halos data in Section 4 and consider the physical conditions of the cool CGM traced by low ions in Section 5. We mainly concentrate on low silicon ions. W13 showed that the three main differences between low ions and O VI, which we studied in O16, are that the former have 1) little dependence on sSFR, 2) a larger range in column density indicating a patchier covering fraction, and 3) a more strongly declining covering fraction at large impact parameters for blue galaxies.

Fig. 1 shows column density maps of three $z = 0.20$ haloes with masses $10^{11.2}$, $10^{12.2}$, and $10^{13.2} M_\odot$ corresponding to sub- L^* , L^* , and group galaxies, respectively. Projections are 600 kpc across and 2400 kpc deep, and provide the method we use throughout to determine integrated system column densities. Focusing first on the L^* halo in the centre column, which we and O16 argue corresponds to the blue COS–Halos sample, we see that the silicon species (top three rows) have patchier distributions and are much more concentrated around the galaxy compared to the O VI shown in the lower panel.

Moving from lower to higher halo mass for the silicon species, we see dramatic changes. The sub- L^* CGM shows silicon absorption in an extended disk structure out to at most $b \sim 50$ kpc, while the L^* halo is covered with low-ion silicon absorption out to 150–200 kpc. The group halo shows much less low-ion silicon absorption in the central regions, but the patchy distribution extends beyond the radius where it falls off for the L^* galaxy. The O VI by contrast is stronger everywhere within the central 300 kpc for the L^* galaxy than for the sub- L^* and group halo, owing to the O16 explanation of the 3×10^5 K collisionally ionized band overlapping with the virial temperature of a $\approx 10^{12} M_\odot$ halo.

In Fig. 1, the silicon species look stronger in the L^* CGM than in the group CGM, but we present a more quantitative approach in Fig. 2 by plotting the mean column densities, \bar{N}_{ion} , as a function of impact parameter, b , by taking the average N_{ion} in annuli of 15 kpc width. We plot the silicon species (Si II, Si III, & Si IV) along the left column for all central galaxies that appear as isolated in the x , y , and z directions as thin lines. We plot averages as thick lines with black borders for sub- L^* ($M_{200} = 10^{11.0} - 10^{11.3} M_\odot$, dark blue), L^* ($M_{200} = 10^{11.7} - 10^{12.3} M_\odot$, aquamarine), and group ($M_{200} = 10^{12.7} - 10^{13.3} M_\odot$, orange) subsamples. The Si species in the sub- L^* CGM fall off rapidly, while the L^* CGM has stronger Si species inside ≈ 50 kpc than groups, but groups have more extended, shallower distributions of low Si ions.

We also show Mg II in Fig. 2 (top right-hand panel), which shows very similar trends as Si II, although Mg II has a slightly lower ionization potential than Si II meaning that it acts like a slightly lower ion. We also show C IV (right middle panel), which has a higher ionization potential than Si IV, and the high ion O VI (lower right-hand panel). For each species, we show a 5 dex range in column density with the y -axis scaled to the same relative abundance based on the atomic number density within the simulation (e.g. the $\bar{N}_{\text{O VI}}$ range is 1.2 dex higher than for $\bar{N}_{\text{Si II}}$, $\bar{N}_{\text{Si III}}$, and $\bar{N}_{\text{Si IV}}$, because there are $\approx 10^{1.2}$ more oxygen than silicon atoms). This helps visualize the effect of the ion fractions on the strengths of the various ion species. For example, the average group has more Si in Si II than O in O VI, since the bordered orange line is higher for $\bar{N}_{\text{Si II}}$ than for $\bar{N}_{\text{O VI}}$ at all radii. It may be surprising that Si II, which traces $\approx 10^4$ K gas, is relatively more abundant than O VI, which primarily traces $> 10^5$ K gas, in a group whose CGM is dominated by $\gtrsim 10^6$ K gas.

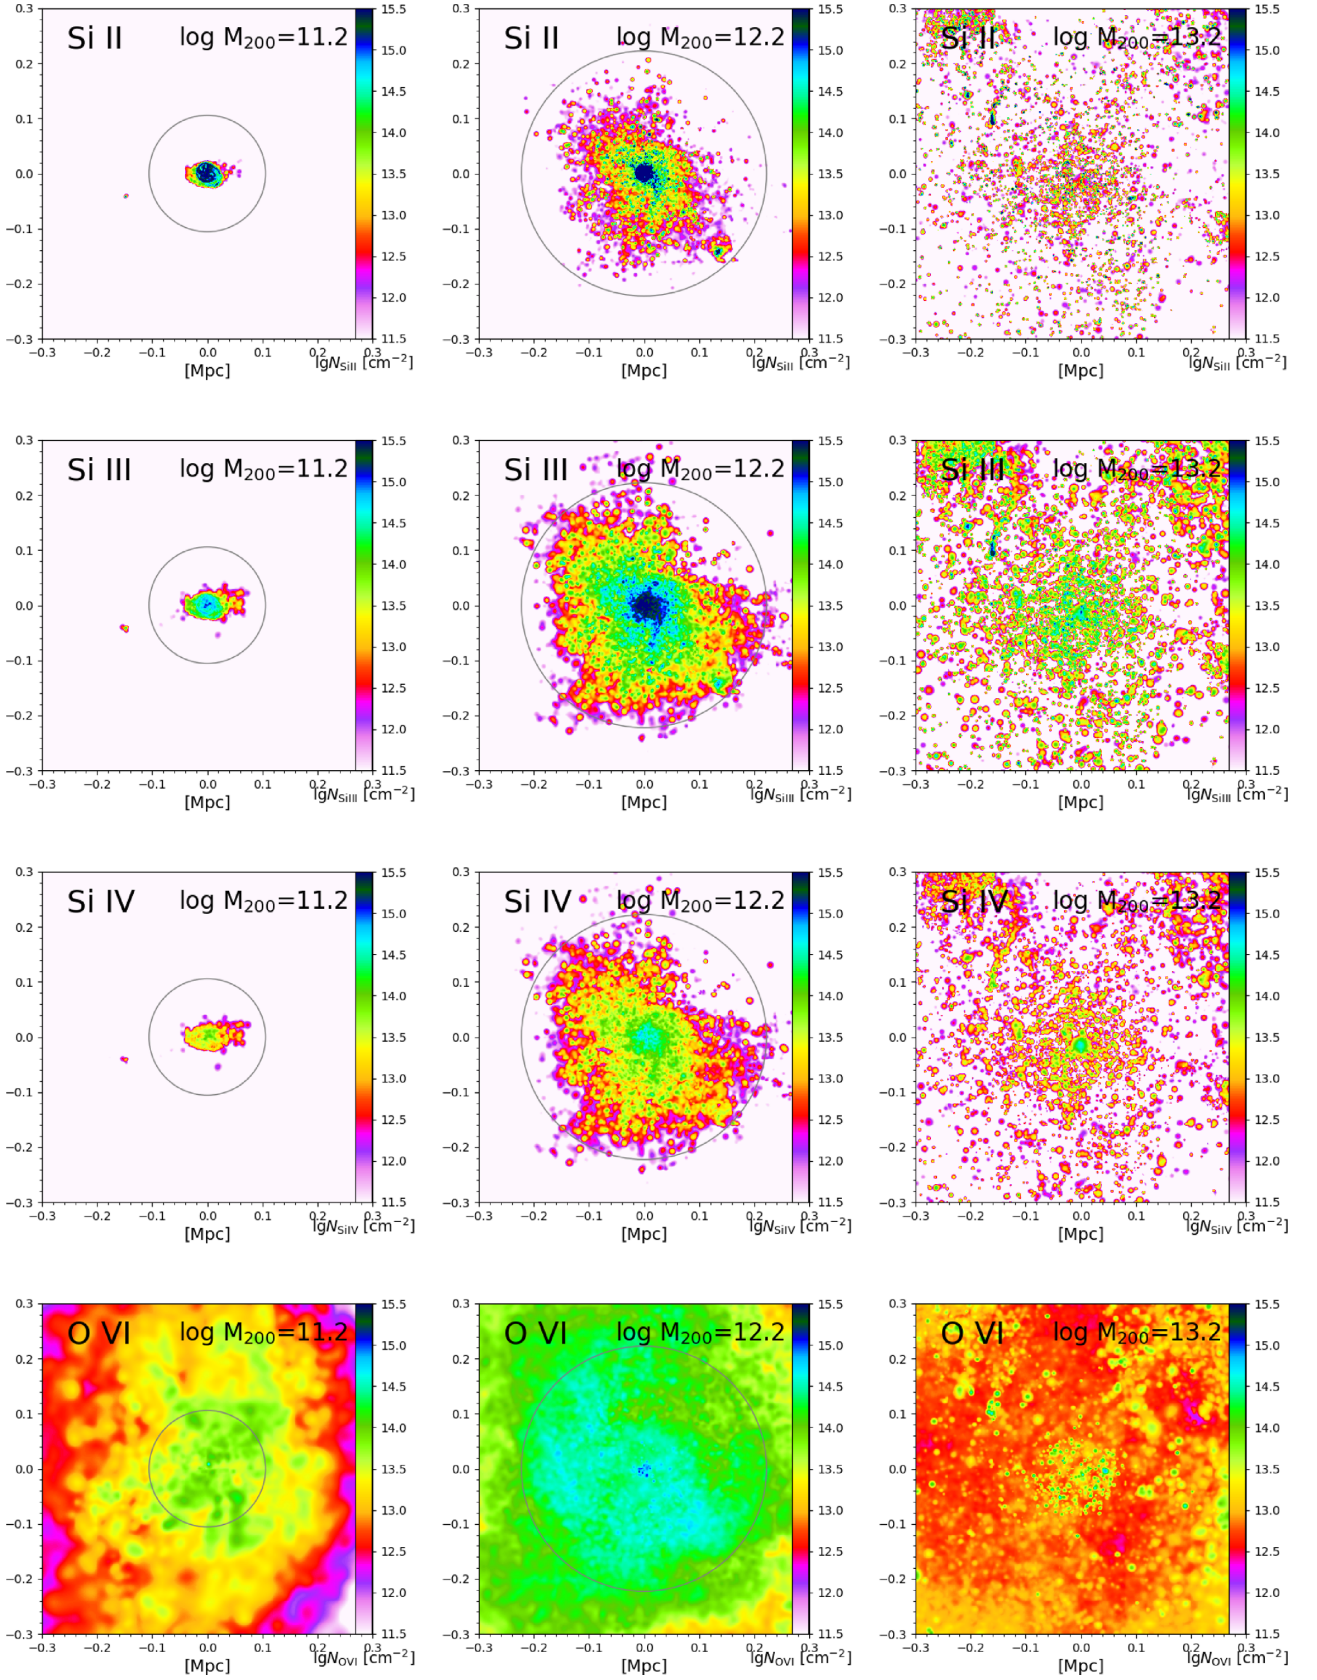


Figure 1. Column density maps for $z = 0.205$ snapshots of three haloes with mass $10^{11.2}$, $10^{12.2}$, and $10^{13.2} M_{\odot}$ representative of sub- L^* , L^* , and group-sized haloes, respectively, from left to right. From top to bottom, the rows show Si II, Si III, Si IV, and O VI column densities on a 600×600 kpc grid. Grey circles indicate R_{200} , which is too large (486 kpc) to appear in the group halo frame.

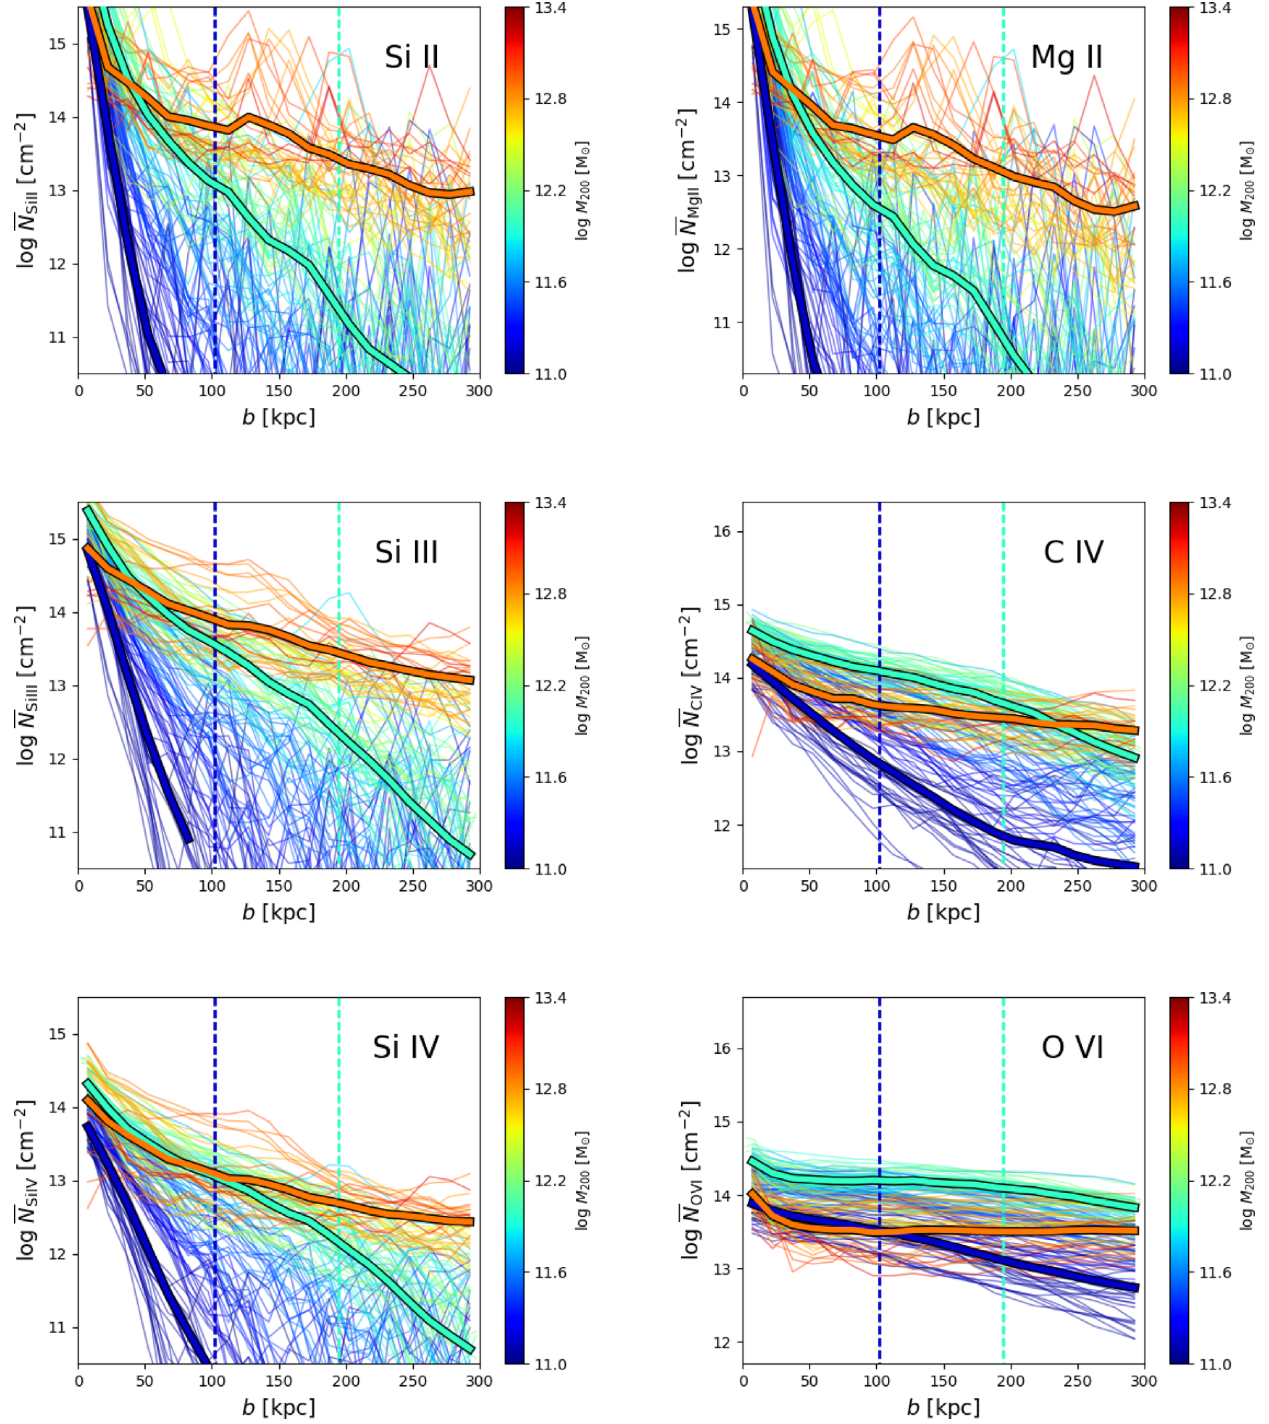


Figure 2. Mean column densities as a function of impact parameter at $z = 0.2$, coloured by halo mass for silicon species on the left (Si II, Si III, & Si IV from top to bottom) and for Mg II, C IV, and O VI on the right. Individual galaxies in isolated projections are shown as thin lines, and averages for sub- L^* , L^* , and group-sized haloes (blue $M_{200} = 10^{11.0} - 10^{11.3} M_{\odot}$, aquamarine $10^{11.7} - 10^{12.3} M_{\odot}$, & orange $10^{12.7} - 10^{13.3} M_{\odot}$) are shown as bordered, thick lines. Dashed vertical lines indicate average R_{200} for the three samples (the group haloes have $R_{200} = 384$ kpc). The column density range for each panel is scaled according to the relative abundance of each element in the simulation. Hence, the relative locations of the curves within their panels reflect the differences in ion fractions. A total of 50 galaxies, 22 from zooms and 28 from the periodic volume, contribute to this plot (11, 15, and 9 contribute to the sub- L^* , L^* , and group averages, respectively).

The progression of ions from lowest (Mg II) to highest (O VI) shows the following trends: 1) lower ions have less extended distributions, 2) lower ions have significantly more scatter, even when plotting quantities binned in 15 kpc-wide annuli, and 3) the lower

the ion, the smaller the impact parameter within which the L^* column densities exceed the group column densities. Several of the trends visible in Fig. 2 have been observed. C IV declines faster around sub- L^* galaxies in COS-Dwarfs (Bordoloi et al. 2014) than

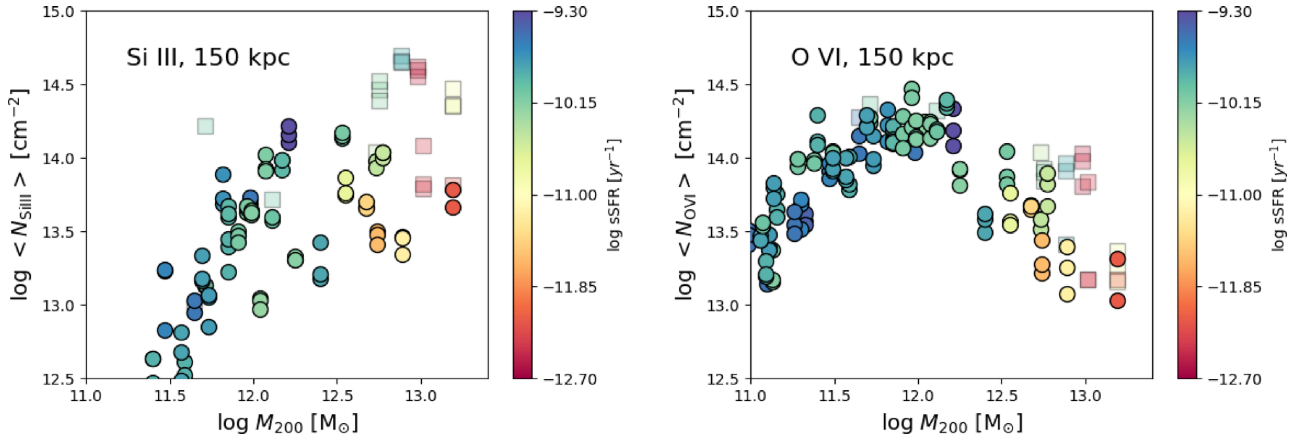


Figure 3. Aperture column densities of Si III (left) and O VI (right) averaged within 150 kpc of the central $z = 0.2$ galaxy and plotted as a function of halo mass. Solid circles indicate stringently isolated galaxies and transparent squares indicate non-isolated galaxies. Colour indicates sSFR.

around more massive galaxies ($M_* > 10^{9.5} M_\odot$) as observed by Burchett et al. (2016) (cf. blue and aquamarine bordered lines in right middle panel). Burchett et al. (2016) also see a decline in C IV detection for higher halo masses ($M_{200} > 10^{12.5} M_\odot$), especially inside 160 kpc (cf. orange & aquamarine bordered lines). Liang & Chen (2014) and Borthakur et al. (2016) observed virtually no Si III beyond the virial radius around their samples dominated by L^* halo mass objects and sensitive down to $N_{\text{Si III}} \sim 10^{12.5} \text{cm}^{-2}$, which agrees with the steep decline seen in Si III in our L^* sample (left middle panel). Local ionizing radiation from galaxies, not included in these simulations, could reduce low-ion column densities preferentially in the inner CGM as we discuss in Section 4.2.

Finally, the O VI averages in the lower right-hand panel of Fig. 2 show remarkably similar impact parameter profiles inside 150 kpc for sub- L^* and group galaxies, but the origins of O VI are very different. As discussed in O16, sub- L^* galaxies have photo-ionized O VI in their $<10^5$ K CGM, while group galaxies have very low O VI fractions in their collisionally ionized $>10^6$ K CGM. However, individual O VI sight line measurements are predicted to be quite different with sub- L^* galaxies showing less scatter in the O VI column densities, and group galaxies showing more scatter with significantly lower median O VI column densities (cf. lower left and right panels of Fig. 1).

3.1 The effect of neighbouring galaxies

It may be counter-intuitive that low ions are more abundant around hotter gas haloes, which mostly host passive galaxies with little star-formation. For C II, C III, Mg II, Si II, Si III, and Si IV, group column densities exceed L^* column densities at every impact parameter >90 kpc (cf. orange and aquamarine bordered lines in Fig. 2 panels). We reconsider the stringent isolation criteria used in O16 to check how the loose isolation criteria we use in this figure differs. The stringent criteria results in a decline of ≈ 0.2 – 0.3 dex between 20 and 160 kpc for low ions, meaning that neighbouring galaxies at $b = 100$ – 300 kpc can increase low-ion column densities by a factor of ≈ 1.5 – 2 .

Fig. 3 shows the 150 kpc “aperture” column densities for Si III and O VI in the x , y , and z directions (there are three data points for each halo), where the aperture column density is defined as

$$\langle N \rangle_b = \frac{\sum_{<b} N(x, y) dx^2}{\pi b^2} \text{cm}^{-2}, \quad (1)$$

where b is the impact parameter and dx is the pixel size ($dx \ll b$). The Si III aperture columns, $\langle N_{\text{Si III}} \rangle_{150}$, increase faster with M_{200} from sub- L^* to L^* haloes than for O VI. However, while $\langle N_{\text{O VI}} \rangle_{150}$ declines from L^* to group haloes, as extensively detailed in O16, $\langle N_{\text{Si III}} \rangle_{150}$ shows no decline and a much larger scatter.

Stringently isolated projections, shown as solid circles, have $\langle N_{\text{Si III}} \rangle_{150}$ values that are similar between groups and L^* haloes, but non-isolated counterparts with neighbouring galaxies within 300 kpc, plotted as transparent squares, indicate a separate branch where $\langle N_{\text{Si III}} \rangle_{150}$ increases from L^* to group haloes. Thus, the typical column density of low ions observed in COS-Halos at $b < 150$ kpc would likely show less dependence on the properties of the central galaxy than is the case for O VI. However, Fig. 2 shows that the $N_{\text{ion}}(b)$ relations for these low ions are most different between the L^* and group samples beyond $b = 150$ kpc, which are not plotted in Fig. 3. Group haloes, even if isolated, still have more low ions beyond 150 kpc than L^* haloes. Finally, the O VI column density shows less dependence on whether the central has neighbouring galaxies than is the case for Si III.

Lastly, in Fig. 3 we colour the $\langle N_{\text{O VI}} \rangle_{150}$ values by sSFR (Fig. 3 right-hand panel) to show the decline in O VI columns going from L^* to groups is driven primarily by halo mass rather than sSFR (O16). $\langle \text{Si III} \rangle$ does not appear to be driven by sSFR either, although there is no obvious decline from L^* to groups and the scatter is more affected by the presence of neighbours.

4 COMPARISON TO COS-HALOS DATA

We now compare our simulation results to the COS-Halos observational survey using the python module named Simulation Mocker Of Hubble Absorption-Line Observational Surveys (SMOHALOS) described in O16. SMOHALOS creates mock COS-Halos surveys using observed impact parameters for galaxies chosen to match the COS-Halos M_* and sSFR. We use the latest spectroscopic galaxy data (Werk et al. 2012) and the most recently published values of the absorption line observations (W13) in our SMOHALOS realizations.

As also described in O16, SMOHALOS applies observational errors from Werk et al. (2012) to simulated galaxy measurements using a random number generator. Gaussian dispersions of 0.2 and 0.3 dex are applied to simulated $\log M_*$ and $\log \text{sSFR}$ values, respectively. The dispersed simulated M_* and sSFR values closest to the observed M_* and sSFR are then selected. Stellar masses assume

a Chabrier (2003) initial mass function (IMF), which requires us to reduce the stellar masses reported by Werk et al. (2012) by 0.2 dex, because they assumed a Salpeter (1955) IMF. sSFR values for passive galaxies are often upper limits, but we treat the value of the upper limit as a detection when matching our simulated galaxies, because these correspond to such low star formation rates that are highly variable in our simulation. The observed b is matched by SMOHALOS through a random number generator picking a pixel at the same b in one of three column density maps (x , y , & z projections) that satisfies the isolation criteria. Like O16, we do not require a simulated galaxy to have the same redshift as the observed galaxy, because we find little evolution over the considered redshift range ($z = 0.15\text{--}0.35$). One hundred SMOHALOS realizations are run to compare with the 44 galaxies from COS-Halos for a total of 4400 measurements, which samples 38 unique simulated galaxies (25 from the zooms and 13 from the periodic volume) at up to three different redshifts.

Fig. 4 shows the COS-Halos observations (W13) for Si II, Si III, Si IV, C II, Mg II, and O VI as a function of impact parameter. W13 reports integrated system column densities that are not broken down by individual component, which allows a comparison with our pixel projection method that does not distinguish components. Blue and red symbols indicate sSFR greater than and less than 10^{-11} yr^{-1} , respectively. Squares indicate detections, upside-down triangles indicate upper limits for non-detections, and upwards pointing triangles indicate lower limits for saturated lines. The median SMOHALOS column density as a function of impact parameter for the blue and red samples, using a division at $\text{sSFR} = 10^{-11} \text{ yr}^{-1}$, are shown as cyan and magenta lines, respectively. One and 2σ dispersions are indicated by thick and thin dashed lines where the simulated absorbers are “perfect” data (i.e. exact column densities, no upper or lower limits).

As in Fig. 2, all y-axis ranges are scaled to the same relative abundances. It is difficult to assess the agreement with the observations from this plot alone given the dominance of lower and upper limits in the data. There are no low ion detections outside the $2\text{-}\sigma$ simulated SMOHALOS bands, unlike the case for O VI, which is $2\text{--}3 \times$ too weak in our simulations (O16). The simulations show the same contrasts between low ions and O VI as observed in COS-Halos: 1) less dependence on the sSFR as indicated by the smaller differences between the medians, 2) a patchier distribution as indicated by larger 1 and $2\text{-}\sigma$ dispersions, and 3) more strongly declining column densities at larger impact parameters particularly around blue galaxies. In the simulations, the dispersion and the dependence on impact parameter decline going from the lowest species (Si II, C II, Mg II) through “intermediate” species (Si III, Si IV) up to O VI. The patchiness of the low ions represents the dispersion within individual haloes as we discuss in Section 5.1 where we find Si III dispersion statistics remain unchanged when combining multiple haloes versus single haloes.

4.1 Survival analysis

4.1.1 Application method

To test the quality of the match between simulations and COS-Halos, we need to account for the observed upper and lower limits, for which we turn to survival analysis. Survival analysis allows a statistical interpretation of incomplete datasets where a portion of the data are “censored.” The Kaplan & Meier (1958, K–M) method provides a general one-variable, non-parametric survival statistic that produces a maximum likelihood distribution using both uncen-

sored (detections) and censored (upper or lower limits) data. While this method has been applied to astronomical datasets including upper limits (e.g. Feigelson & Nelson 1985) and discussed extensively in the context of absorption line surveys in Simcoe, Sargent & Rauch (2004), we apply a two-sided censored K–M estimator that we argue applies more proper treatment as well as limitations compared to a one-sided censoring K–M estimator.

Fig. 5 shows the K–M estimator for six ions in COS-Halos, which plots the cumulative distribution function (CDF) of the fraction of absorbers with a higher column density in black with shading indicating 95 per cent confidence intervals. We apply a two-sided Kaplan–Meier estimator to account for upper limits (i.e. non-detections) and lower limits (i.e. saturated absorption) on the observed column densities. Vertical dotted lines in each panel encompass the range over which the K–M estimator is not required and uncensored detections set the CDF. The K–M estimator applies to the column density ranges where censored and uncensored data overlap, and the K–M method uses the assumption that the censored data distributions follow the uncensored data distributions. This allows us to use censored data points that have different detection limits to estimate the probability distribution according to the detected data points. At the column density below (above) which all observations become upper (lower) limits, the K–M estimator cannot provide a constraint, therefore this assures that the CDF never reaches 1 or 0 for all low ions, because these data are censored on both sides.

The main motivation of the two-sided K–M estimator is to achieve a more statistically appropriate *and limited* CDF to compare with simulated datasets. For example, we limit the CDF at column densities above which absorbers are all saturated, while an analysis like W13 assumes lower limits are uncensored detections, making a statistical comparison with simulations more constrained. Our application remains agnostic about the true distribution of these lower limits, because it is improper to assume saturated absorbers are at that column density and in fact could be much higher as our simulations predict. This is critical for the interpretation of the COS-Halos data, because the ion mass estimates would be under-estimated in such a case, which is a point we further detail in Section 5.1.

To calculate the two-sided K–M estimator, we apply a normal one-sided K–M estimator using the right-censored data (saturated lower limits) and temporarily setting upper limits as detections. We then apply another one-sided K–M estimator using left-censored data (undetected upper limits) setting the lower limits as detections. The first K–M estimator equals one minus the second K–M estimator between the highest upper limit and lowest lower limit (i.e. between the vertical dotted lines). The two-sided K–M estimator uses the first K–M estimator above the lowest lower limit and one minus the second K–M estimator below the highest upper limit. Our two-sided K–M estimator relies on the highest upper limit always being lower than the lowest lower limit. If there is such a violation for a censored measurement, we do not include it in our CDF. However, this only happens for one upper limit in Si II and C II from a very low signal-to-noise (S/N) sight line and so we do not worry about the statistical bias. It is thus rarely expected that upper limit non-detections overlap lower limit saturated absorbers. Our algorithm is similar to the Turnbull (1974) method for doubly censored data, but we do not apply this iterative method, because the requirement for the highest upper limit being lower than the lowest lower limit should prevent this need.

The requirements of the K–M method applied on a dataset as discussed for absorption line measurements in Simcoe et al. (2004) are 1) the censored data points must be independent of one another,

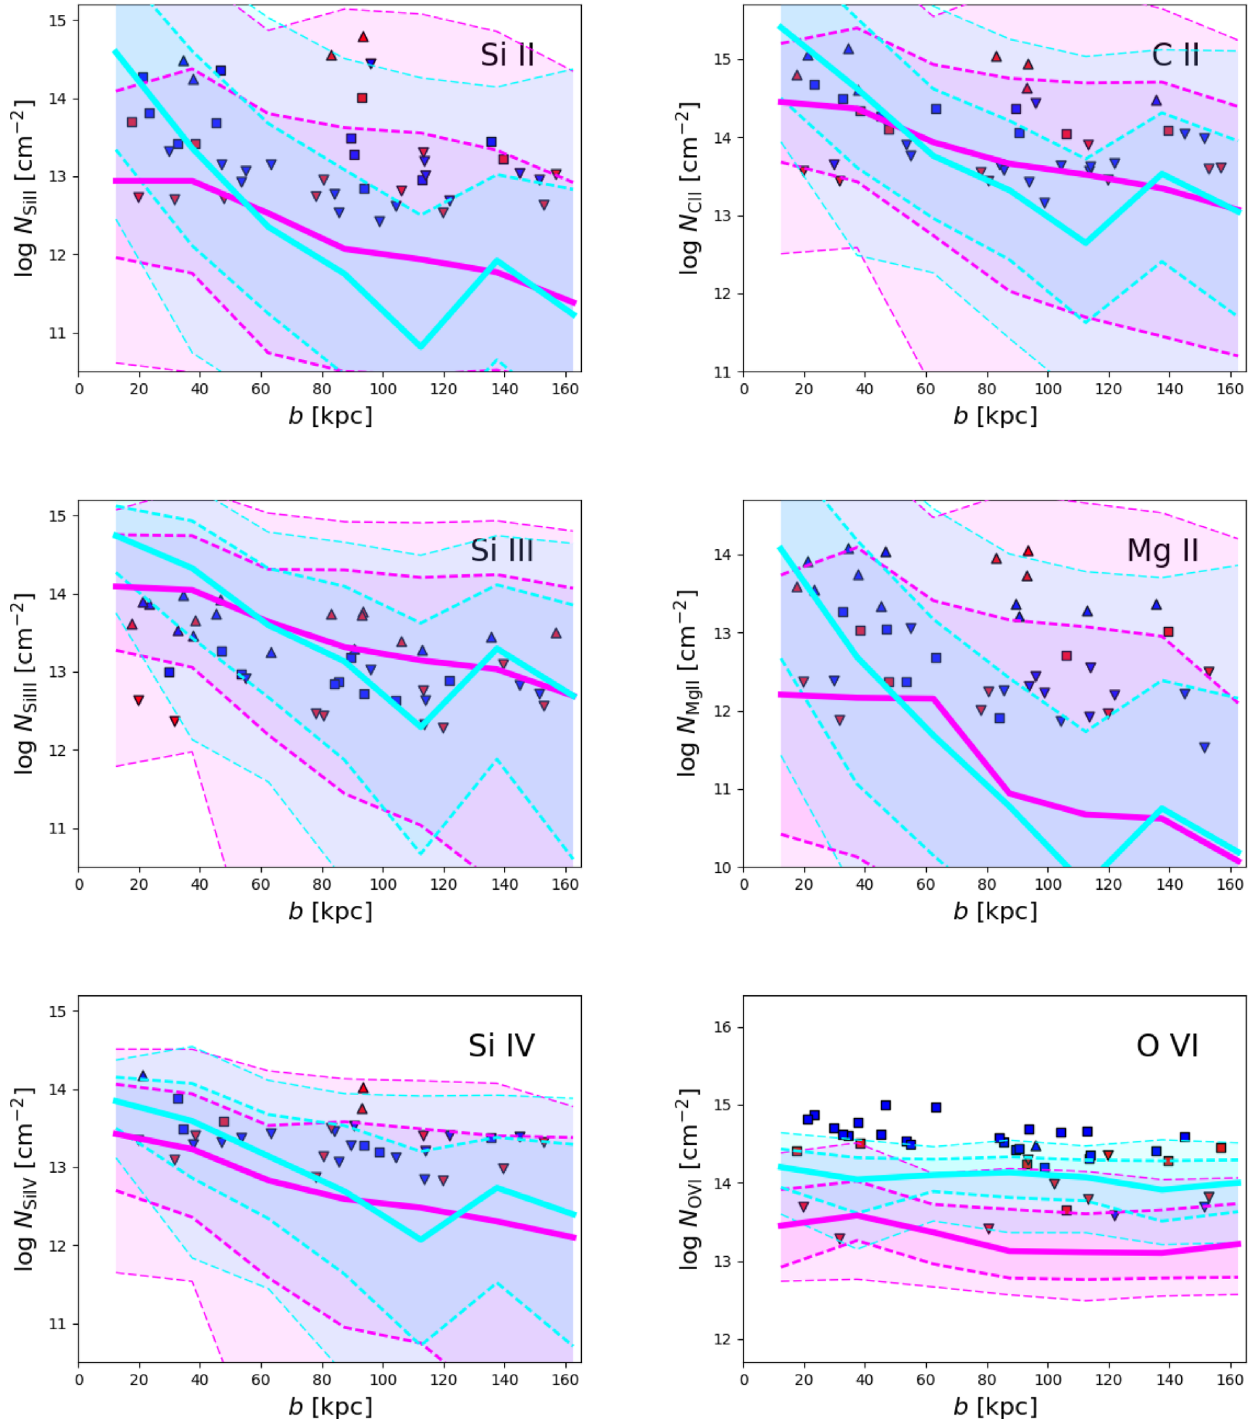


Figure 4. Column density as a function of impact parameter for the COS–Halos blue and red samples (squares are detections, upside-down triangles are upper limits, and upwards pointing triangles are lower limits). Simulated column densities for all 44 galaxies from 100 SMOHALOS realizations are plotted as solid cyan and magenta lines for galaxies with $sSFR$ higher and lower than 10^{-11} yr^{-1} . One and 2σ dispersions are indicated by thick and thin dashed lines, respectively. Three silicon ions (Si II, Si III, & Si IV, left-hand panels) are shown along with C II, Mg II, and O VI (right-hand panels). Compared to O VI, the simulated low ions have less dependence on $sSFR$, as indicated by overlapping cyan and magenta regions, a patchier distribution, as indicated by larger dispersions, and column densities that decline faster at larger impact parameters. A comparison of the simulations to the observations is difficult owing to the dominance of upper and lower limits (see Figs 5 and 6 instead).

and 2) the probability that the data point will be censored should not correlate with its value. The first requirement is true, because the observed data points for a given ion species come from different galaxies with no relation. The second requirement is not true as

upper limits depend on the S/N of the spectrum. Although COS–Halos sight lines aimed to have an S/N of 10, the S/N varies between and within sight lines, which explains in part why there exists a range of values for upper limits. While this second requirement is




Figure 5. Cumulative distribution functions (CDFs) of COS-Halos column densities for various ions (black step function with shading indicating 95 per cent confidence limits) compared to the simulated CDFs (thin green band) generated from 100 SMOHALOS realizations. The Kaplan–Meier method is applied to handle upper and lower limits. The total number of observations (n) is listed on the top along with the average difference in column density between observations and simulations ($\delta \log N$). The three silicon ions (Si II, Si III, & Si IV) are shown on top and C II, Mg II, and O VI are shown on the bottom. The input data points are shown along the bottom: squares for detections, upside-down triangles are upper limits, and upwards pointing triangles are lower limits. Two-sided censoring results in the observed CDFs never reaching 0 or 1. Vertical dotted lines encompass the range over which K–M estimation is not required.

not fulfilled, we argue we can still apply the K–M estimator owing to the implicit assumption that the observed column densities have a much larger range of true uncensored values compared to the range over which detections and censored limits are observed. Hence, we are arguing that the significant dispersion of low-ion columns, as predicted by the simulations (cf. Fig. 4), makes their appearance as censored data points in the COS-Halos sample essentially random. This is not true for O VI around blue galaxies, where the dispersion is smaller than the observed range, but fortunately these data points are almost all uncensored detections.

The Si III panel in Fig. 5 demonstrates our K–M estimator. Upper limits overlap detections between $N_{\text{Si III}} = 10^{12.5} - 10^{13} \text{ cm}^{-2}$, and the K–M method applies the probability distribution of the detections to the range of upper limits where they overlap, as indicated by where the black line and grey limits extend below the left dotted vertical line. Below $10^{12.5} \text{ cm}^{-2}$, there are no detections, and the K–M estimator provides no constraints. Above $10^{13.2} \text{ cm}^{-2}$, all Si III absorbers are saturated, but there exist no detections to guide the K–M prediction of non-detection and again there are no constraints on the CDF. This demonstrates the conservative nature of how we use the two-sided K–M estimator, which limits us to comparing simulations and observations where there exist detections. There is more often significant overlap between detections and upper limits than between detections and saturated lower limits. We urge caution when interpreting the K–M estimator for column densities in the

overlap regions, owing to the assumption of censored data following uncensored detections.

The K–M estimator lacks information on additional parameters beyond the first, column density in our case, but we can sub-divide the sample based on a second parameter, as we do for sSFR, and plot multiple CDFs as in Fig. 6. While we lose information on the dependence on impact parameter for each ion, we are generating mock SMOHALOS surveys with the same impact parameters around similar galaxies as observed. This means that we do not generate mock observations for COS-Halos galaxies when the given ion is not observed, due to instrument coverage or blanketing by unrelated absorbers.

4.1.2 Results

The K–M CDFs from 100 SMOHALOS realizations with perfect column densities (i.e. no upper/lower limits) are shown in green in Fig. 5 and in cyan (blue sample) along with magenta (red sample) in Fig. 6. Also listed for each observation–simulation pair is the number of observations (uncensored and censored, n) and the average column density deviation in dex of the simulation from the observation ($\delta \log N$). $\delta \log N$ is calculated at each step in the observed K–M function corresponding to each uncensored data point. Observed data face complete censorship on one or both ends of the

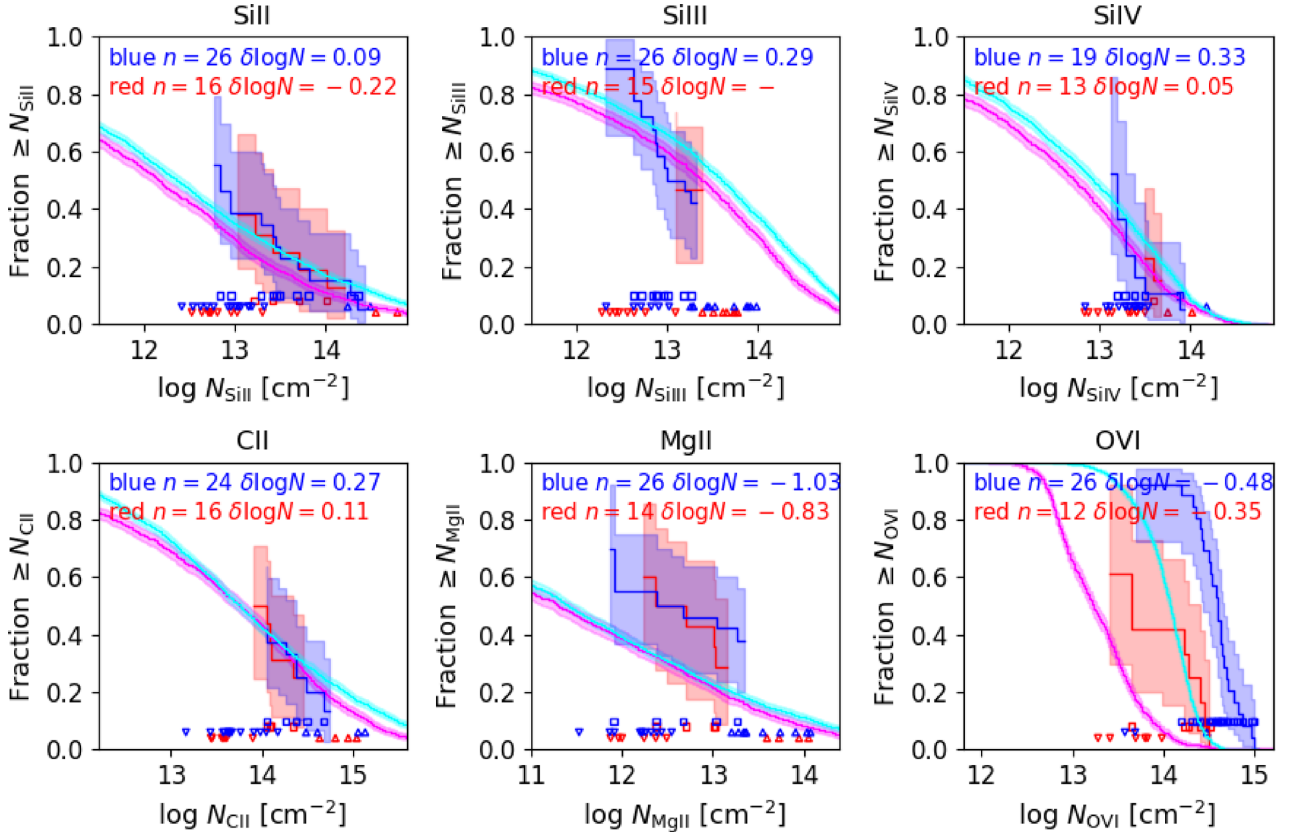


Figure 6. CDFs using the Kaplan–Meier method of COS–Halos and SMOHALOS as in Fig. 5, but divided into blue and red galaxy samples. COS–Halos observations are plotted in blue and red, and SMOHALOS simulations are plotted in cyan and magenta. Sample statistics for blue and red galaxies are listed in each panel as in Fig. 5. No $\delta \log N$ is given for the Si III red sample owing to only having censored data points, but there still is a CDF measurement between the upper and lower limits. We do not plot the vertical dotted lines encompassing the ranges over which the K–M estimation is not required as we do in Fig. 5, but note that these ranges are equal to or larger for the individual blue and red samples as for the entire sample shown in that plot.

column density distribution, where the K–M functions cutoff before reaching 0 or 1.

The level of agreement between mock and real data in Fig. 5 varies from ion to ion. The Si II, Si III, Si IV, and C II total sample distributions usually agree within a factor of ≈ 2 or better (0.3 dex). SMOHALOS Mg II is however 1 dex too low according to this metric. O VI is 0.5 dex too low, which agrees with O16 and indicates that our slightly modified simulated sample at $z = 0.15$ – 0.35 is not different from O16’s sample. The simulations overlap with the 95 per cent confidence limits of the observations (large shaded bands) for the silicon and carbon species, but not for Mg II and O VI. C III and C IV are not shown due to their limited COS–Halos datasets, but the simulations show reasonable agreement with COS–Halos. C III has 25 observations, of which only two are uncensored, while C IV only has three observations since COS–Halos was not designed to cover this ion.

Fig. 6 shows the subdivision into red and blue galaxies. As W13 showed, for the observed low ions the confidence limits of the two galaxy samples always overlap. Simulated red galaxies have slightly lower column densities than blue galaxies with the gap growing towards higher ions, but COS–Halos does not have a sufficiently large sample to probe such small differences except for O VI. We also subdivide the blue sample into small and large impact parameter bins, divided at 75 proper kpc (not shown), and the results show statistically significant increases in column densities at smaller impact parameters for all low ions, which agrees with W13.

Overall, the level of agreement between SMOHALOS and COS–Halos using the K–M estimator is good for Si II, Si III, Si IV, and C II, being within a factor of two for the red and blue subsamples. While there exist some discrepancies – simulated silicon ions are higher than COS–Halos for blue galaxies and simulated Si III has a larger spread – the simulations overlap the 95 per cent confidence limits for these measurements. While simulated Si II and C II column densities are in reasonable agreement with observations, Mg II is too weak for the entire COS–Halos sample and any sub-division by sSFR and impact parameter. No self-shielding is included in these simulations, which we next consider using standard equilibrium simulations.

4.2 Model modifications

Our fiducial simulation results use non-equilibrium ionization assuming a uniform Haardt & Madau (2001) background without self-shielding. Therefore, we now explore NEQ ionization and the expected influence of self-shielding on low ions. We also comment on the effects of simulation resolution (see also Appendix C) and other sources of photo-ionization.

4.2.1 Equilibrium ionization and cooling

Fig. 7 compares the CDFs for our $z = 0.20$ subsample (NEQ in blue) to standard EAGLE equilibrium simulations where we iterate $z = 0.20$ outputs to ionization equilibrium using our NEQ network

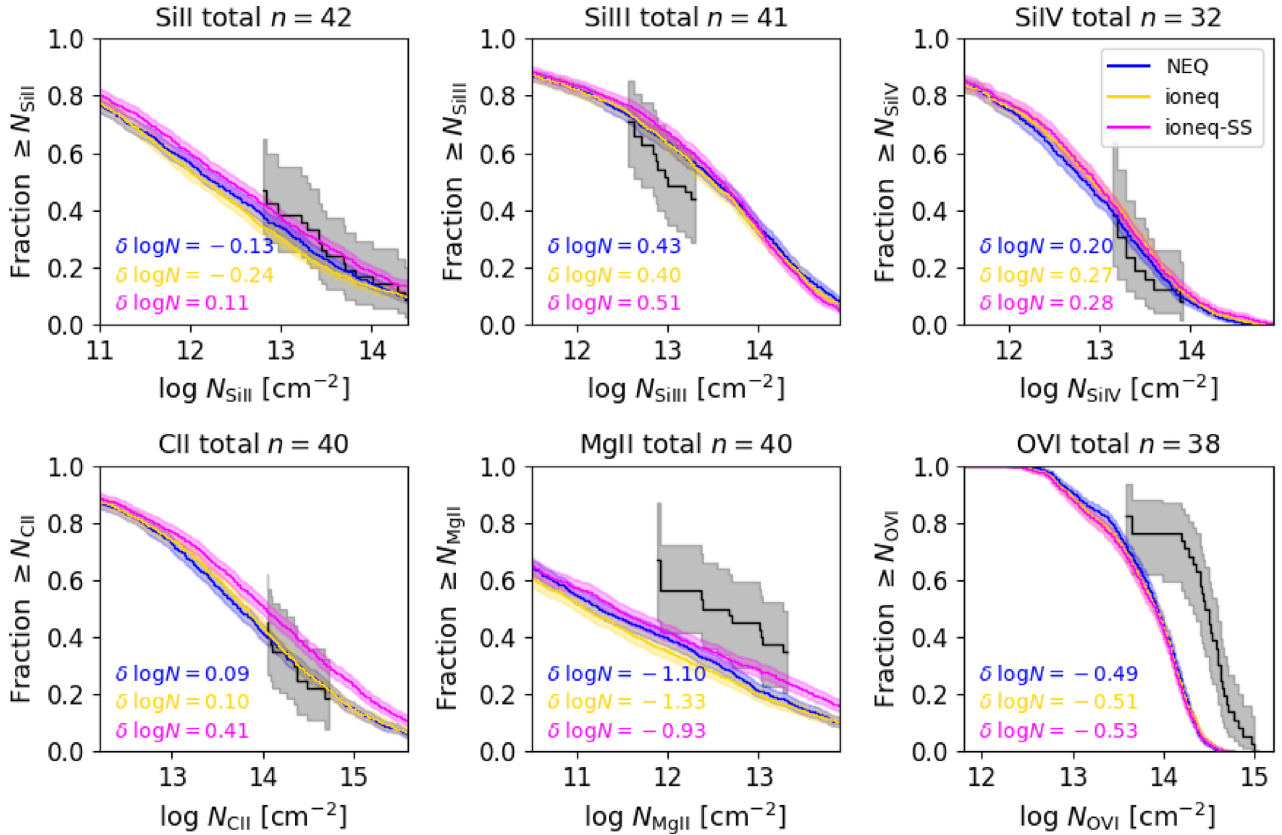


Figure 7. Kaplan–Meier CDFs for the entire COS–Halos sample, as in Fig. 5, compared to different model variations. The $z = 0.20$ NEQ sample in blue is compared to standard equilibrium EAGLE simulations of the same haloes assuming ionization equilibrium for the uniform Haardt & Madau (2001) background (gold), and then applying a self-shielding criterion following the Rahmati et al. (2013) prescription (magenta).

for the same haloes (ioneq in gold), and then apply the Rahmati et al. (2013) self-shielding correction (ioneq-SS in magenta). To simulate self-shielding in post-processing, we modify the regular NEQ network by reducing the ionizing radiation for metal ions with ionization potentials above 1 Ryd according to the density and redshift dependencies derived by Rahmati et al. (2013) from radiative transfer simulations that reproduce the H I column density distribution. The simulation output is then iterated to this new self-shielded ionization equilibrium. This method is only an approximation because it ignores the frequency-dependent attenuation that declines for higher ionization potentials. However, multiply ionized species with higher potentials are not appreciably photo-ionized at the densities where the correction is used. Thirty SMOHALOS realizations of each model are run, and the baseline NEQ model shows essentially identical behaviour as the full sample including all redshifts in Fig. 5.

The NEQ and ioneq runs overlap for the most part, which we further elaborate upon in Appendix A – NEQ ionization does not significantly alter low-ion abundances when assuming a uniform ionization background. Mg II and Si II decline the most, but such differences are expected given that the ioneq and NEQ simulations are separate runs once the NEQ is turned on as described in Section 2.2, and this does not mean there is a significant difference that can be attributed to non-equilibrium ionization.

4.2.2 Self-shielding

Applying the self-shielding criterion increases singly ionized species at higher column densities as indicated by the average $\delta \log N$

value increasing by 0.2–0.4 dex for C II, Si II, and Mg II over the ioneq model. Si III is also boosted by 0.1 dex. This slightly degrades the excellent agreement for C II, yields a similar fit for Si II as the NEQ model, and still leaves a factor of 10-fold too small Mg II column densities. Mg II traces the highest densities of all these species, so it is not unexpected to see the greatest increase due to self-shielding for higher Mg II column densities.

The under-prediction of Mg II is concerning. Part of this discrepancy likely reflects that magnesium nucleosynthetic yields are too low in EAGLE as Segers et al. (2016) demonstrated that Mg in stars is ≈ 0.3 dex too low compared to other elements. We also consider the effect of resolution for a subset of L^* haloes in Appendix C and we find that Mg II column densities increase by 0.2 dex when the mass resolution is increased by a factor of eight over our standard runs, while C II and Si II decrease by 0.1 dex. This could result from these higher resolution simulations better resolving dense substructure. These effects combined could raise Mg II to overlap with the 95 per cent confidence limits of the CDF in Fig. 7. However, they are unlikely to simultaneously solve the O VI discrepancy.

4.2.3 Additional sources of photo-ionization

The last model modifications we consider are additional sources of photo-ionization from the central galaxy due to on-going star-formation (e.g. Stocke et al. 2013, W14) and/or AGN. The latter is explored in Segers et al. (2017) and Oppenheimer et al. (2018) for EAGLE simulations such as these where the addition of fluctuating AGN can enhance the ionization levels even when the AGN is off as appears to be the case for COS–Halos galaxies. These works argue

the proximity zone fossil effect proposed by Oppenheimer & Schaye (2013b) is capable of enhancing O VI levels by ≈ 0.5 dex for typical Seyfert-like AGN episodes in star-forming galaxies. The key is that the timescale to recombine from higher ionization species to O VI is equal to or longer than the typical times between AGN activity, even though the AGN is active for only a small fraction of the time. Proximity zone fossils can solve the under-estimates of O VI in standard NEQ simulations (O16) while not significantly reducing low ions, because even though low ions are ionized to higher levels when the AGN is on, they rapidly recombine to equilibrium after the AGN turns off as shown in Oppenheimer et al. (2018).

The uniform ionization background may also be supplemented by a local ionizing radiation field from emission sources within the galaxy associated with ongoing star formation. Scaling this total local ionizing flux with radius from the galaxy ($\propto r^{-2}$), galaxy star formation rate (\propto SFR) and the escape fraction of ionizing photons ($\propto f_{\text{esc}}$) were explored by W14 in CLOUDY models. The inclusion of stellar radiation from a Starburst99 spectrum (Leitherer et al. 1999) can moderately affect COS-Halos results in the inner CGM for fiducial values of star-forming galaxies in that survey: SFR = $1\text{M}_{\odot}\text{yr}^{-1}$ and an assumed $f_{\text{esc}} = 5$ per cent. At the average impact parameter in the COS-Halos blue sample, $b = 72$ kpc, the ionizing radiation from such a galaxy provides slightly more ionizing radiation ($1.2\times$) than the Haardt & Madau (2001) background. W14 considers these effects from the Starburst99 model and concludes that this emission likely does not play a large role in setting the ionization fractions of Si and O. However, if we post-process our simulation outputs using the physical densities predicted by the model, then we see an increase of intermediate species like Si III and Si IV and a decline in singly ionized species at $b \lesssim 75$ kpc. O VI remains mostly unchanged since it is collisionally ionized in L^* haloes at large radii, which appears to agree with Suresh et al. (2017) who found no difference outside 50 kpc while applying a more extreme stellar radiation field that uses $f_{\text{esc}} = 5$ per cent for lower energy radiation and 100 per cent for soft X-rays.

On the other hand, not included in the Starburst99 models is the soft X-ray emission produced by mechanical energy released into the ISM during a starburst phase, from both supernovae and additional X-ray sources produced by star formation (Cantalupo 2010; Werk et al. 2016), which may have a substantial affect even at large impact parameters. In contrast to Starburst99 models that provide radiation mainly below 4 Ryd, this soft X-ray emission contributes only above 4 Ryd and enhances O VI while not affecting low ions.

Our discussion of model modifications makes clear that there are multiple potential effects that can alter low-ion column densities by factors of two or more. Self-shielding can increase low ions while extra ionization from star formation and AGN can reduce low ions. Thus, the agreement of silicon and carbon species within a factor of ≈ 2 for the standard prescription can be classified as a success of the model given these uncertainties. The Mg II column densities are severely under-estimated, but could be remedied by going to higher resolution, and by increasing the Mg yields, which show evidence for being too low in EAGLE.

5 PHYSICAL PROPERTIES OF LOW METAL IONS

Having explored how simulated observations compare to COS-Halos, we now focus on the physical properties of the gas and metals traced by low ions. We first sum the metal mass budget traced by low ions and follow up by linking observed ions to the




Figure 8. CGM silicon as a function of radius for a sample of 9 $z = 0.2$ L^* haloes with an average mass of $M_{200} = 10^{12.0}\text{M}_{\odot}$, subdivided by ion and summed in 0.1 dex radial bins. The dotted vertical lines indicate R_{200} and $2R_{200}$, and the silicon beyond R_{200} can belong to the CGM of neighbouring sub- L^* galaxies.

physical properties of the gas they trace. The evolutionary state of CGM metals is considered next. Finally, we explore ion ratios used to constrain CLOUDY models, e.g. in W14 and Keeney et al. (2017), and test the validity of single-phase models.

5.1 Low-ion CGM mass estimates

In O16, we used our zoom simulations to explore the circumgalactic oxygen budget, finding that only 0.9 – 1.3 per cent of oxygen at $<R_{200}$ is in the O VI state for L^* haloes spanning $M_{200} = 10^{11.8} - 10^{12.2}\text{M}_{\odot}$. A much larger fraction of oxygen inside the virial radius of those same halos, 27 – 52 per cent, resides in O I – O III. Having followed the NEQ ionization and cooling in our zooms for 11 elements, we can self-consistently trace the 15 silicon ion species in the same way as O16 traced the nine oxygen ion species. Si II, Si III, and Si IV comprise 19 – 42 per cent of L^* haloes’ silicon budget (Si I is negligible), therefore these “low”-ion silicon species provide a good proxy for the low-ion CGM mass estimate. Our simulations show that the Si I – Si IV ion fraction is consistently between 70 and 80 per cent of the O I – O III ion fraction, which O16 plotted in their fig. 10. The median low-ion silicon CGM budget (i.e. Si II, Si III, & Si IV) for L^* haloes is $3.9 - 4.5 \times 10^6\text{M}_{\odot}$, which converts to $7 - 8 \times 10^7\text{M}_{\odot}$ for all metals using the simulation-averaged Si/Z ratio. Relative to solar abundances, our simulation-averaged [Si/Z] and [Si/O] values are within 0.05 dex of Asplund, Grevesse & Sauval (2005).

Fig. 8 illustrates the breakdown of the CGM silicon budget from the average of 9 L^* haloes averaging $M_{200} = 10^{12.0}\text{M}_{\odot}$ at $z = 0.2$ with shading indicating the contribution of various silicon ions as a function of radius. Purple, magenta, and red correspond to Si II, Si III, and Si IV, respectively. These low ions are primarily found inside R_{200} indicated by the left dotted line. Significant silicon at $T < 10^5$ K exists also in Si V (orange) and Si VI (yellow). Green and blue colours correspond to higher Si ions tracing warm-hot CGM. Most silicon (like the majority of the metal mass) resides beyond $0.5R_{200}$, but low ions trace a biased set of interior metals with most of these ions arising in the inner 50 kpc. We do not count silicon from the ISM, defined as gas above the star formation density threshold, which is assumed to be in ionization equilibrium at $T =$




Figure 9. The median low-ion metal surface densities for the 28 star-forming galaxies (thick solid cyan line from 100 SMOHALOS realizations) derived by summing Si II, Si III, and Si IV column densities from Fig. 4 compared to the COS-Halos low-ion metal surface densities derived from W14 CLOUDY modelling as reported in Peebles et al. (2014) (blue dotted lines, two different functional fits shown). One σ dispersions are indicated by dashed lines. The passive SMOHALOS realizations are displayed in magenta for comparison.

10^4 K, but the typical L^* galaxy has an average of $6.0 \times 10^6 M_\odot$ mostly in Si II.

We compare our values to those of Peebles et al. (2014) derived from low-ion CGM budgets traced by these three silicon species and several other low ions with similar ionization energies (C II – C III, N II – N III, Mg I – Mg II), finding an average metal mass of $2.3 \times 10^7 M_\odot$ from 10 to 150 kpc and $\delta v < 600 \text{ km s}^{-1}$ of L^* COS-Halos galaxies, but with values up to $4 \times$ higher, $\approx 9 \times 10^7 M_\odot$, when including systematic uncertainties owing to ionization modelling. We compare the Peebles et al. (2014) fits, derived from the CLOUDY ionization modelling in W14, shown in dotted blue in Fig. 9 for the 28 star-forming COS-Halos galaxies. We calculate low-ion metal surface densities by summing the Si II, Si III, and Si IV SMOHALOS column densities from Fig. 4, taking medians and one-sigma spreads of the summed column densities in 25 kpc radial bins, and converting to a low-ion metal surface density ($M_\odot \text{ kpc}^{-2}$) assuming solar abundances. The comparison between the median SMOHALOS surface density (thick solid cyan line) and the Peebles et al. (2014) fits are promising, except for a dip in the former at 75–125 kpc. Despite the SMOHALOS medians being at or below the Peebles et al. (2014) fits, we derive a $4 \times$ higher low-ion metal mass around L^* galaxies owing to the significant dispersion of column densities at a given impact parameter (dashed cyan lines show the $1-\sigma$ dispersion). The high-end Peebles et al. (2014) mass estimate of $\approx 9 \times 10^7 M_\odot$ owes to their consideration of systematic uncertainties in CLOUDY modelling, and not the dispersion in column densities.

We emphasize the high metal mass value, $\approx 10^8 M_\odot$, indicating a significant reservoir of low ions, and argue that it is consistent with COS-Halos. The biggest difference between Peebles et al. (2014) and our summation is the treatment of the significant scatter at a given impact parameter, which is also seen in COS-Halos (Peebles et al. 2014, their Fig. 7). Fig. 10 shows the residual dispersions of Si III when subtracting out the median Si III binned in $\delta b =$




Figure 10. The residual distributions of Si III column densities after subtracting the median column density profiles plotted in Fig. 2. The distributions of $\delta N_{\text{Si III}}$ are calculated across 9 L^* galaxies and binned into two impact parameter bins (10–75 kpc in the upper panel, and 75–150 kpc in the lower panel). For each bin, we list the median value of $\log N_{\text{Si III}}$ and the $1-\sigma$ range, as well as the mean $\log N_{\text{Si III}}$.

15 kpc bins for a given L^* halo. We divide the figure into two impact parameter ranges, 10–75 and 75–150 kpc, to show that the residuals are well-described by a log-normal distribution that increases in width at larger impact parameters. Fig. 10 shows the dispersion statistics across 9 L^* haloes when subtracting the median Si III halo-by-halo, but using a single L^* halo results in similar dispersions. The main point of Fig. 10 is to show that low-ion species are often well-described by log-normal distributions at a given impact parameter. We list the median logarithmic column density (Si III_{50}) and $1-\sigma$ range in the two panels, and also show that the mean logarithmic column density ($\text{Si III}_{\text{mean}}$) is significantly higher than the median. We suggest using the median, Si III_{50} , and the logarithmic dispersion obtained from observations to calculate $\text{Si III}_{\text{mean}}$ assuming a log-normal distribution.¹

We also wish to contrast our low-ion metal budget with the results of Muratov et al. (2016), who found significantly fewer CGM metals in $\approx 10^{12} M_\odot$ FIRE-zoom-simulated haloes. They found between 0.27 and $1.4 \times 10^8 M_\odot$ of total metals in the CGM, while our low-ion CGM component alone is $\approx 10^8 M_\odot$. Muratov et al. (2016) explained their lower CGM metal content has to do in part with FIRE using lower yields than Peebles et al. (2014), the latter of which is similar to our zooms (O16). The Muratov et al. (2016) zooms additionally having more metals locked in stars (20–70 per cent) compared to our zooms (25–35 per cent, see Fig. 9 of O16). FIRE has not yet divided CGM metals into ionization species, but we would predict that they would find smaller low-ion column densities than observed. However the difference may not be as large with newer FIRE zooms, since their recent m11.9a zoom has a census more similar to ours, although it has a late-time merger that recently enriched the CGM (Muratov et al. 2016).

We also plot the galaxies with $M_{200} = 10^{12.7} - 10^{13.2} M_\odot$ in Fig. 9 and integrate a low-ion metal mass of $6.6 \times 10^7 M_\odot$ using stringently isolated galaxies. Thus the low-ion content of group

¹The mean of a log-normal distribution is $\exp(\mu + \sigma^2/2)$ where μ and σ are natural logarithms and not \log_{10} , therefore the values given for Si III_{50} and the $1-\sigma$ range must be multiplied by $1/\log_{10}(e) = 2.303$.

haloes is $\approx 2/3$ rd the amount of L^* galaxies within 150 kpc. While we make the point that COS–Halos passive galaxies likely have neighbouring galaxies that increase their low-ion column densities in Section 3, this isolated group subsample nonetheless harbours a comparable amount of low-ion metals within 150 kpc as is the case for L^* haloes.

Finally, we tally the amount of cool CGM gas mass, defined as all non-star-forming gas with $T < 10^5$ K. In L^* haloes, the median mass of cool gas is $1.5 \times 10^{10} M_\odot$ for a median halo mass of $8.4 \times 10^{11} M_\odot$, and for group haloes, the cool gas sums to $2.8 \times 10^{10} M_\odot$ for a median halo mass of $7.2 \times 10^{12} M_\odot$. We do not restrict to stringently isolated galaxies for these sums, and there exists more cool gas associated with satellites in group haloes compared to L^* haloes. We discuss mass budgets in future work, but note that the L^* or group sums are lower than the entire COS–Halos sample cool mass sum within 160 kpc from Prochaska et al. (2017) of $(9.2 \pm 4.3) \times 10^{10} M_\odot$.

5.2 Low-ion CGM physical properties

The column density (N)-weighted pixel value of a physical property, p , is calculated according to

$$p_N = \frac{\sum_i p_i \times N_i}{\sum_i N_i}, \quad (2)$$

from column density maps where i is a pixel with a column density greater than N_{\min} . Physical quantities are smoothed using the same SPH kernel as column densities, and we make image projections that allow us to calculate medians and $1 - \sigma$ spreads for density (n_H), temperature (T), and pressure (P) as a function of impact parameter in Fig. 11. We apply a minimum column density based on typical observational column density limits: $10^{12.5} \text{ cm}^{-2}$ for Si species, $10^{13.5} \text{ cm}^{-2}$ for O VI, and $10^{15.0} \text{ cm}^{-2}$ for O VII. Weighted pixel values are not highly sensitive to N_{\min} , although it prevents contributions from pixels below observational capabilities.

The upper panel shows that the density that an ion traces declines with ionization potential with little dependence on impact parameter outside the inner CGM ($b \gtrsim 50$ kpc) for our sample of L^* haloes. Conversely, temperature increases with ionization potential, while also showing little dependence on impact parameter: low silicon ions clearly trace photo-ionized $T = 10^{4-4.5}$ K gas, O VI traces collisionally ionized warm-hot gas, and O VII traces the $\approx 10^6$ K hot halo. The trends of lower densities and higher temperatures with increasing ionization potential and the weak dependence on b were also found in the Ford et al. (2013) simulations and form the basis of the Stern et al. (2016) universal density CGM model. However, in the simulations, those models both have O VI photo-ionized in L^* haloes.

Also shown for density and temperature relations are Mg II and C II, where we use column density limits $10^{12.0}$ and $10^{13.0} \text{ cm}^{-2}$, respectively. Mg II absorbers trace denser and cooler gas than Si II, but still remain within 0.1–0.2 dex of Si II physical values. C II traces gas more like Si II in the very interior, but behaves more like Si III and Si IV throughout most of the CGM, because the ionization potential of C II (24.4 eV) is much higher than those of Si II (16.4 eV) and Mg II (15.0 eV). We overplot the COS–Halos n_H values derived by modelling these low ions (Prochaska et al. 2017),² which clus-

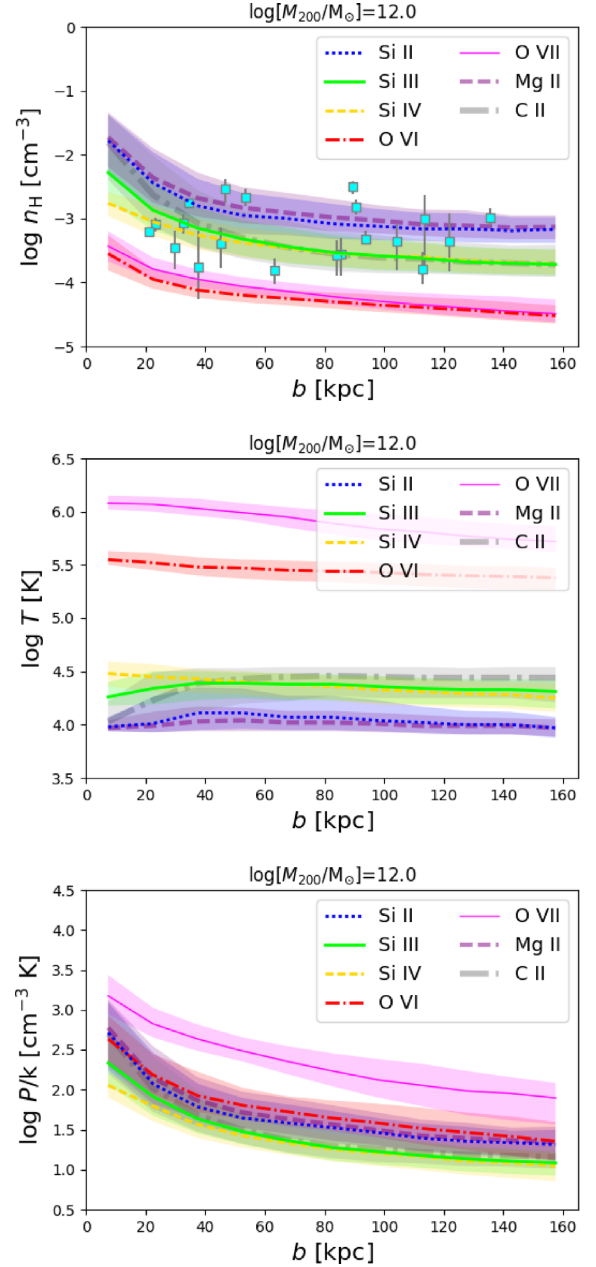


Figure 11. Physical properties weighted by ion column density as a function of impact parameter around 9 L^* haloes averaging $M_{200} = 10^{12.0} M_\odot$. Hydrogen number density and temperature are shown in the upper and middle panels for low species (Si II, Si III, Si IV, C II, Mg II) as well as O VI. Additionally, O VII is shown to indicate the hot halo component, where most CGM metals reside in an L^* halo. Pressure, calculated by taking $n_H \times T/(X_H \mu)$ is also plotted. COS–Halos densities derived from low-ion modelling (Prochaska et al. 2017) around L^* appear as cyan squares.

ter around $n_H \sim 10^{-3} \text{ cm}^{-3}$ and show negligible dependence with impact parameter, in reasonable agreement with the simulation.

The lower panel of Fig. 11 shows the pressure, computed as $P/k = n_H T/(X_H \mu)$, where k is the Boltzmann constant, X_H is the mass fraction of hydrogen, and μ is the mean molecular weight.

²The Prochaska et al. (2017) densities are higher than the CLOUDY-derived densities of W14 (their Figs 10 and 12), owing to a factor of $4 \times$ miscal-

culation in the latter, which did not account for the isotropic nature of the radiation field.

The pressures traced by low silicon ions show a mild decline with impact parameter for $b \gtrsim 100$ kpc reaching pressures $\sim 10 \text{ cm}^{-3} \text{ K}$. This is in agreement with the Stocke et al. (2013, their Fig. 14) CLOUDY-derived pressures determined for their observed $z < 0.2$ warm CGM clouds with impact parameters less than R_{vir} around $L > 0.1L^*$ galaxies, which are comparable to the star-forming COS-Halos galaxies. The pressures derived from W14 and Prochaska et al. (2017) are in a similar range as Stocke et al. (2013); hence, our simulation predictions appear to agree with the densities and pressures derived by CLOUDY modelling of COS low-ion metal column densities.

The density–temperature phase–space diagrams in Fig. 12 show the distribution of metals and ions inside R_{200} for our set of L^* haloes (upper panel) and a set of nine group haloes with average $M_{200} = 10^{12.9} M_{\odot}$ (lower panel). Dotted diagonal lines indicate isobars of $P/k = 1, 10$, and $100 \text{ cm}^{-3} \text{ K}$. The two samples of haloes show similar distributions of low ions, $> 10 \text{ cm}^{-3} \text{ K}$, although the pressures are slightly higher in the group-sized halo. CGM metals peak at a density $n_{\text{H}} \approx 10^{-4.1} \text{ cm}^{-3}$ for both samples, but the high-temperature distribution peaks track the virial temperatures indicated by the dashed horizontal lines, which increase according to $T_{\text{vir}} \sim M_{200}^{2/3}$ or a factor of $4 \times$ across these two samples.

Applying a dividing line between hot and cool CGM metals of $T = 10^5 \text{ K}$, we find that about half the metals are in each phase in the L^* haloes, but only 7 per cent are in the cool phase for the group haloes. Surprisingly, there exists a similar mass of cool CGM metals inside R_{200} for the two of halo mass ranges, $M_{\text{Z, cool}} \sim 1.0\text{--}1.5 \times 10^8 M_{\odot}$, which is borne out by observations of low silicon species being similar between L^* and group halos. This $M_{\text{Z, cool}}$ value range holds among L^* halos, but a $10^{11.8} M_{\odot}$ halo has as much as 75 per cent of its metals in the cool phase while for a $10^{12.2} M_{\odot}$ halo this fraction is as low as 1/3rd. Therefore, the cool phase metal mass of a halo is nearly invariant across a factor of over 20 in halo mass, while the total CGM metal mass increases monotonically with halo mass. Connecting with the mass estimates in Section 5.1 traced by low silicon ions, of the order of half the cool, $T < 10^5 \text{ K}$ CGM metals are traced by the Si II – Si IV phases with the rest in higher Si species tracing densities $n_{\text{H}} \lesssim 10^{-4.5} \text{ cm}^{-3}$.

Returning to radial trends, the upper panel of Fig. 13 shows N -weighted radial distance as a function of impact parameter, indicating that O VI arises at much larger physical radii than its observed impact parameter. This trend was noted by O16, who showed that the typical O VI absorber observed at $b < 150$ kpc traces gas at 200–500 kpc from the galaxy. Conversely, $r_{\text{N}} \approx b$ for the low silicon ions meaning that these ions are tracing gas at a physical distance similar to their observed impact parameter. One conclusion is that the O VI phase is spatially distinct from low ions, so it is not unexpected that O VI shows different kinematics from low ions (Werk et al. 2016). However, we note that Werk et al. (2016) found many aligned components between low ions and O VI, which may be a challenge for our model here and is further explored in Oppenheimer et al. (2018) with the inclusion of the proximity zone fossil effect.

5.2.1 Pressures in the CGM

Finally, we plot pressure as a function of radial distance in the lower panel of Fig. 13. The result is very different from $P(b)$, because the O VI and O VII arise at much greater radial distances than the observed impact parameters. The low silicon ions appear to be out of pressure equilibrium with the hotter phase oxygen ions, although

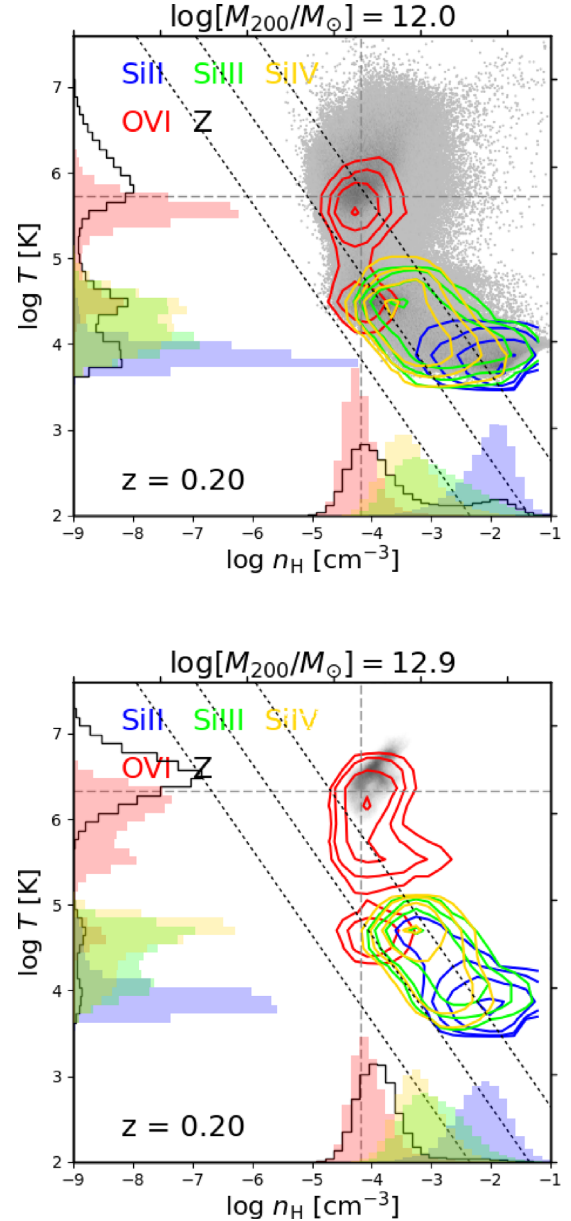


Figure 12. Density–temperature phase–space diagrams of Si II (blue), Si III (green), Si IV (gold), O VI (red), and metals (grey shading) gas inside R_{200} for a sample of L^* haloes (upper panel) and group haloes (lower panel) at $z = 0.2$. Corresponding histograms along the bottom and left indicate the density and temperature distribution for each ion, respectively. The dashed grey horizontal line indicates T_{vir} and the dashed grey vertical line indicates $200 \times$ the critical overdensity. Dotted diagonal lines indicate isobars at 1, 10, and $100 \text{ cm}^{-3} \text{ K}$ from left to right.

the ions are typically not spatially coincident as demonstrated in the upper panel. The silicon (O VI) ions are primarily at $r \lesssim 0.5 R_{200}$ ($\gtrsim 0.5 R_{200}$). Nonetheless, in Appendix D, we examine the lack of pressure equilibrium at fixed radius, finding that low-ion clumps containing 50–100 SPH particles are embedded in a high-pressure ambient medium leading to numerical effects of the SPH formalism at an under-resolved boundary. Therefore, the lack of pressure equilibrium is likely numerical as opposed to physical.

Our result of low ions tracing $P \sim 10\text{--}100 \text{ cm}^{-3} \text{ K}$ gas while O VI also traces pressures at the higher end of this range is surprisingly similar to the pressures calculated by McQuinn & Werk (2017)

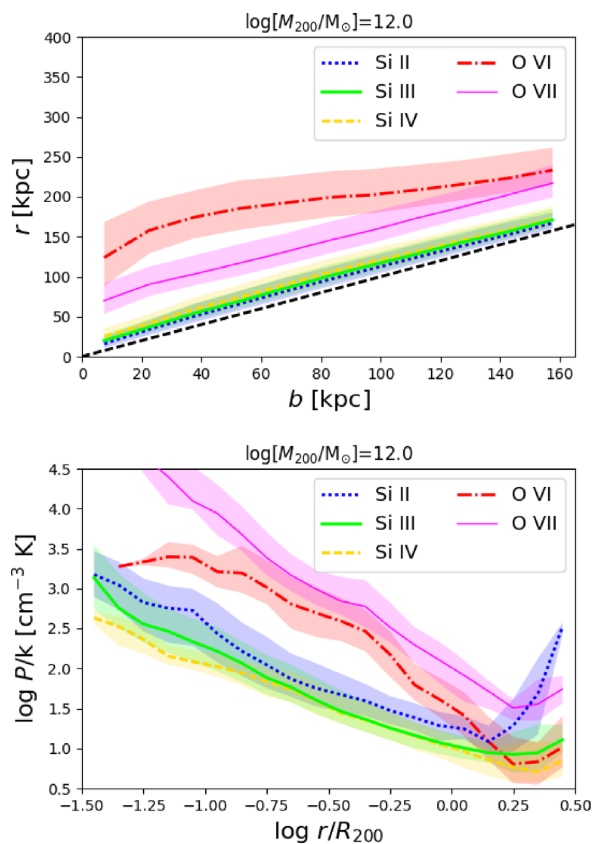


Figure 13. The upper panel shows column density-weighted mean physical radius, r , for a given ion plotted as a function of impact parameter for silicon species, O VI, and O VII for the set of L^* haloes. The 1-to-1 correspondence is plotted as a dashed black line. Cool silicon species show $r \sim b$, while O VI in contrast arises from much larger physical radii at a given impact parameter (O16) and traces a physically and spatially distinct phase at a given impact parameter. The lower panel plots ion-weighted pressure as a function of radius, indicating cool silicon ions and warm-hot oxygen species trace different pressures at the same physical radius.

using their analytic cooling flow model. They argue that low ions are out of pressure equilibrium with O VI, indicating other non-thermal pressure support for the low-ion clouds. Our simulations find that low ions are not spatially coincident with O VI, but arise in the interior CGM and are radially coincident with *even higher* pressures traced by ions like O VII and above. McQuinn & Werk (2017) argued that the velocity alignment between low ions and O VI indicates they are spatially related, suggesting low-ion clouds are directly cooling out of the $10^{5.5}$ K phase. We consider aligned absorbers in Oppenheimer et al. (2018) where we argue that photo-ionized O VI in proximity zone fossils, not included here, increases the amount of O VI components aligned with low-metal-ion components. Without AGN proximity zone fossils, there is still alignment between O VI and low ions, although it is not as easy to line up individual components between the two types of species.

Lastly, we return to the factor of $>100 \times$ higher cool CGM densities predicted by the Maller & Bullock (2004) model versus the observationally constrained results of W14. The updated densities of Prochaska et al. (2017) appear to cluster around $n_H \approx 10^{-3} \text{ cm}^{-3}$ reducing the tension a little bit, but still not close to the cool phase densities predicted by W14 using the Maller & Bullock (2004) model of $\sim 10^{-1.7} - 10^{-1} \text{ cm}^{-3}$. It is worth asking why there is so much tension with this measurement given that the mean density of




Figure 14. Radial velocity relative to the central galaxy around the L^* halo sample for metals (grey shading) and various ions (coloured contours) as a function of radial distance. Negative radial velocities indicate inward motion at $z = 0.20$, and the dashed line delineates the velocity needed to reach the central galaxy by $z = 0$. The dot-dashed line indicates the approximate escape velocity from the haloes, showing very few winds are on trajectories that can escape. Coloured histograms correspond to the radial and velocity distributions of the ions.

halo gas should be of order $100 \times$ the mean density of the Universe at $z = 0.2$, or $n_H \sim 10^{-4.5} \text{ cm}^{-3}$. Using this mean density along with the halo virial temperature of $10^{5.8}$ K for a $10^{12.1} M_\odot$ halo (O16), cool gas at $\sim 10^4$ K would require $n_H \sim 10^{-2.7} \text{ cm}^{-3}$ for pressure equilibrium, which nearly matches the Prochaska et al. (2017) results. However, this implied pressure of $40 \text{ cm}^{-3} \text{ K}$, while agreeing with low-ion pressures in our simulations, is much lower than the O VI and O VII-traced gas in Fig. 13 at all radii except at $\gtrsim R_{200}$. The disagreement with Maller & Bullock (2004) arises because their density and temperature profiles steadily rise towards smaller radii where most low-ion clouds are found. Our density, temperature, and pressure profiles also rise at smaller radii, which is why our low-ion pressures of $10 - 40 \text{ cm}^{-3} \text{ K}$ at $20 - 160 \text{ kpc}$ are out of pressure equilibrium with the hot ambient medium at pressures $> 100 \text{ cm}^{-3} \text{ K}$. There still exists tension between the COS-Halos-derived densities and the densities predicted for pressure equilibrium with the inner, hot CGM, but it is significantly less than the factor $\gtrsim 100$ predicted in W14.

5.3 Low-ion CGM kinematic properties

We begin by considering the radial velocity of the gas particles relative to the central galaxy, determined by calculating $v_{\text{rad}} \equiv \frac{v \cdot r}{r}$ as a function of radial distance in Fig. 14 for our sample of L^* haloes. We weight gas particles by ion or metal mass on a 2D plane in the same manner as the density–temperature diagrams in Fig. 12. Contours show that SPH particles with strong Si II are mainly below the dashed line indicating the average inward velocity required to reach the galaxy in the time elapsed between $z = 0.20$ and 0, which is 2.5 Gyr. This indicates that most Si II has a velocity trajectory consistent with gas accreting onto the galaxy by $z = 0$.

SPH particles with strong Si III and Si IV have radial distributions weighted toward slightly larger radii than Si II (see the bottom histograms), and most have negative, infalling velocities, albeit slightly less than half of these Si ions have velocities below the dashed line. A small fraction of silicon species show signatures of strong outflows ($v_{\text{rad}} > 200 \text{ km s}^{-1}$) in the inner ≈ 30 kpc, but even fewer exceed the average escape velocity from the haloes (dot-dashed line) indicating that low ions rarely escape from low-redshift haloes. O VI, indicated by the red contours, has a much larger radial extent and only a small fraction is on a velocity trajectory that reaches the central galaxy by $z = 0$, being mainly above the dashed line. Hence, the fate of the metals observed in the CGM appears to be very different depending on its ionization state (e.g. Ford et al. 2014). The inflow dividing line does not exactly determine the future of the gas as gas can accelerate or alternatively be reversed by outflows, but we show in Section 5.4 the locations of the different ions in this plane correlate with their $z = 0$ fate.

It should be noted that even though the O VI is collisionally ionized and shows a net inward flow, our simulations do not produce the cooling flow structures theorized by Heckman et al. (2002) and applied to COS-Halos by Bordoloi, Heckman & Norman (2016). Cooling behind a virial shock, at $\sim R_{200}$ in that model, would produce a structured spread of velocities inside the virial shock with O VI occupying a confined post-shock region corresponding to efficient cooling around $T \sim 10^{5.5}$ K. Instead, O VI occupies a range of radii, mainly outside R_{200} corresponding to $\sim 10^{5.5}$ K metal-enriched gas with long cooling times, a large fraction of which also happens to be inflowing.

5.4 Evolution of CGM metals

While the physical conditions of CGM gas can be ascertained by observed metal ions, the evolutionary phase of the gas remains difficult if not impossible to determine directly from observations. Therefore, we use additional information within the simulations to consider the $z = 0$ fate of CGM metals observed at $z = 0.20$. Our simulations allow particle tracking, so we can take subsets of gas and determine if it reaches a galaxy or remains in the CGM. We perform this exercise on a range of halo masses and plot the phase fraction of SPH particles above a specific ion threshold for a given ion species. The fate of $z = 0.2$ CGM SPH particles for nine haloes sorted by halo mass is shown using the bar plots in Fig. 15. Each bar representing particles in the CGM at $z = 0.2$ that at $z = 0$ 1) are converted into stars, 2) are in the ISM, 3) have been recorded to be in a galaxy’s ISM between $z = 0.2$ and 0 (“recycled CGM”), but are in the $z = 0$ CGM, and finally particles that have not been in the ISM between $z = 0.2 \rightarrow 0.0$ and are 4) “cold CGM” ($T < 10^5$ K) or 5) “hot CGM” ($T \geq 10^5$ K) at $z = 0$. The threshold ion fraction to select a $z = 0.2$ SPH particle is 10 per cent $n_{\text{Si}, Z_0}/n_{\text{H}, Z_0} = 10^{-5.46}$ for silicon species and 5 per cent $n_{\text{O}, Z_0}/n_{\text{H}, Z_0} = 10^{-4.61}$ for O VI – this results in $\gtrsim 10^4$ particles selected per L^* halo and corresponds to a level that creates significant absorption in a spectrum. We consider all particles within 500 kpc instead of R_{200} to include the extended O VI around L^* haloes because this barely affects silicon ions. The nine haloes plotted do not undergo major mergers between $z = 0.2$ and 0, and because we do not distinguish between accretion onto a central versus satellite, we verified that most metals concentrate around the central galaxy. The exception is that satellite accretion dominates for Si II for haloes with $M_{200} \geq 10^{13.0} M_\odot$.

We show Si II, Si III, and O VI in haloes ranging from sub- L^* to our largest group ($M_{200} = 10^{11.0} - 10^{13.2} M_\odot$). Over 2.5 Gyr of evolution, more than half of every ion for every halo resides in the z




Figure 15. Tracking the $z = 0$ fate of CGM ion species (Si II, Si III, and O VI from top to bottom) inside 500 kpc at $z = 0.2$ for haloes increasing in mass from left to right ($\log[M_{200}/M_\odot]$ indicated in white). The bars indicate the fate of the gas traced by each ion at $z = 0$, with the possibilities that it joined a galaxy (converted into stars, resides in ISM, or recycled back into the CGM) or remains in the CGM (cold and hot phases divided by $T = 10^{5.0}$ K).

$= 0$ CGM, but there are large differences between species and trends with halo mass. The vast majority of circumgalactic Si II observed at $z = 0.2$ will be accreted onto a galaxy by $z = 0$, defined as stars+ISM+recycled CGM, although more than half of this accretion will be re-ejected into the CGM (cf. green to cyan+orange bars) either by direct outflows (the majority) or by stripping (a minority). More of the $z = 0.2$ circumgalactic Si II around group-sized haloes remains in the CGM (cold+hot) indicating low-ionization metals are more extended and further from massive galaxies relative to L^* haloes (cf. Fig. 2), which results in less accretion onto galaxies. The recycled CGM at $z = 0$ is primarily hot around all galaxies, because of the nature of our thermal feedback.

In contrast, O VI rarely accretes onto any galaxy, which agrees with the radii and radial velocities in Fig. 14. Our focus remains on the low ions, and Si III shows behavior intermediate between Si II and O VI. Most Si III remains within 100 kpc of a central galaxy for haloes L^* and below, and most accretes onto the galaxy. Around group haloes, most Si III and Si IV (not shown) remain in the CGM, with the largest fraction remaining as cold CGM at $z = 0$.

Ford et al. (2014) performed metal tracking in SPH simulations, finding the same general trends we show here. Like them, we find that lower ionization state metals recycle onto galaxies more often, and that lower mass haloes recycle more than higher mass haloes. They also showed that a low-ion-like Mg II will fall onto a galaxy within several Gyr in contrast to O VI, which was injected by outflows many Gyr ago (O16). However, their O VI is primarily photo-ionized (Ford et al. 2013) and a greater fraction of their O VI gas recycles back onto their galaxies, in contrast with our primarily collisionally ionized O VI that almost never recycles. Crain et al. (2017) also performed tracking of H I gas using EAGLE in their Figure 11, but over a shorter time-scale from $z = 0.1 \rightarrow 0$. H I is most comparable with our Si II, and they found that most atomic hydrogen remains associated with its galaxy (59 per cent remains H I, 13 per cent turns into stars) and a majority of the remaining 28 per cent is heated to high CGM temperatures by feedback.

5.5 Ion ratios

Ion ratios have long been used to derive physical properties of absorbers. For low-ion absorbers that are photo-ionized, modelling ion ratios can strongly constrain an absorber’s ionization parameter, which provides a measure of its physical density for a given ionization field and temperature. W14 used ion ratios from a variety of carbon, silicon, nitrogen, magnesium, and oxygen species to derive physical parameters of the metal-enriched CGM. Their method involves using “integrated” column densities where they sum the entire column density along the sight line within $\pm 600 \text{ km s}^{-1}$ of the galaxy systematic velocity. They found a single-phase CLOUDY model in thermal equilibrium (but ignoring adiabatic cooling) by varying the ionization parameter assuming the Haardt & Madau (2001) ionization background and metallicity that matches the hydrogen column density of the sight line.

Our fiducial simulations include NEQ effects, do not assume thermal equilibrium, include adiabatic cooling, and unlike W14 do not include self-shielding. Self-shielding is important for higher densities, especially above $n_H > 10^{-2} \text{ cm}^{-2}$, but our analysis here focuses on the Si II/Si III ratio, which mainly constrains lower densities. We show the COS–Halos $N_{\text{Si II}}/N_{\text{Si III}}$ ratios, nearly all of which are upper limits with Si II detected and Si III a saturated lower limit, in the upper panel of Fig. 16. SMOHALOS ratios with no censoring are shown as grey-gridded shading indicating mostly lower Si II/Si III ratios consistent with the observed upper limits. The shading scales




Figure 16. Ion ratio of Si II over Si III as a function of impact parameter (b). The upper panel shows COS–Halos blue and red subsamples from W13 with arrows indicating upper and lower limits. The SMOHALOS distributions (all “detections”) are shown as gridded shading. The middle panel shows the Si II/Si III ratio as a function of b for all pixels where $N_{\text{Si II}}$ and $N_{\text{Si III}}$ are $> 10^{12.5} \text{ cm}^{-2}$, coloured by Si II-weighted n_H to show that this ratio is primarily sensitive to the physical gas density. The lower panel shows the relationship between $N_{\text{Si II}}/N_{\text{Si III}}$ and n_H in the simulations (data points), in CLOUDY models for various temperatures (thin lines), and variable-temperature CLOUDY models using Si II-weighted n_H and T from the simulations (thick line). Colours indicate the median temperature for each point. The ion-weighted densities for the two ions (Si II & Si III) differ, indicating that these ions probe a multiphase medium. Fortunately, the close agreement between the thick line (CLOUDY models) and squares (simulated values) show that a density derived from an ionization ratio in a simulation gives a similar answer as a single-phase CLOUDY model, at least for Si II-weighted values.

linearly with number density of SMOHALOS ratios. High ratios are observed around two star-forming galaxies at 90 kpc and a high lower limit is observed around a passive galaxy at 140 kpc, which we argue next indicates dense gas.

To understand how the observed ion ratios relate to physical parameters, we plot the integrated ion ratios for sight lines where $N_{\text{Si II}}$ and $N_{\text{Si III}}$ are greater than $10^{12.5} \text{ cm}^{-2}$ as a function of impact parameter for a projection of a $M_{200} = 10^{12.1} M_{\odot}$ L^* halo in the middle panel of Fig. 16. Column density ratios are coloured by Si II-weighted n_{H} . There is a clear relation between this ratio and n_{H} that is mostly independent of impact parameter. High ratios are however seen preferentially in the interior where there exist more dense clumps. The lower envelope of ratios owes to the precipitous decline in the Si II ionization fraction at $n_{\text{H}} \lesssim 10^{-3.7} \text{ cm}^{-3}$. We choose to focus on the Si II/Si III ratio because it is an adjacent ion ratio with high sensitivity to n_{H} .

To compare the physical properties probed by the Si II/Si III ratio in the simulation with those interpreted by the CLOUDY ionization modelling as applied in W14, we show several relations between n_{H} and Si II/Si III in the lower panel. For each of our mock integrated sight lines, we show two physical densities using equation (2), an Si II-weighted (squares) and an Si III-weighted (circles) n_{H} . We show the median densities for several Si II/Si III bins to show that while they differ due to the multiphase nature of the CGM, they remain within 0.5 dex of each other, which matches the expected spread of densities probed by Si II and Si III in Fig. 11 (blue and green lines in the upper panel). We colour data points by the median temperatures to show that these vary as well.

CLOUDY modelling assumes a single-phase density and temperature, and we plot the ratios inferred from CLOUDY models as thin solid lines for five different temperatures. Although there is a significant dependence on temperature as shown by the division between these lines, this mainly owes to Si II becoming collisionally ionized above $\sim 10^{4.2}$ K and the ratio losing sensitivity to physical density. Fortunately, when Si II is observationally detected, it is unlikely to be collisionally ionized³ and we argue that the CLOUDY-derived single-phase model provides a reasonable constraint on n_{H} . For an integrated Si II/Si III ratio, the simulated Si II-weighted n_{H} agrees well with the CLOUDY-derived n_{H} given that one selects the correct temperature, which we show by the colour of the squares for Si II-weighted n_{H} . To compare the squares to the CLOUDY models of the correct temperature, we draw a variable-temperature thick line that shows temperature-dependent CLOUDY models that use the Si II-weighted temperatures (colour of squares). The good agreement between the CLOUDY Si II-weighted models and the EAGLE Si II-weighted densities instills confidence that ionization ratios provide meaningful constraints on gas densities. However, the dependence on temperature is significant, and it is recommended in the future that W14 and other similar works (e.g. Keeney et al. 2017) publish their assumed equilibrium temperatures derived from their ionization parameters and metallicities. The ionization ratio heavily depends on this missing information.

³Although low ions can be collisionally ionized, they almost always are photo-ionized in simulations and CLOUDY models, owing to rapid cooling at $T \gtrsim 10^4$ K to a thermal equilibrium at $T \lesssim 10^4$ K at the high CGM densities where low ions are abundant. We use Si II-weighted quantities in the following analysis, because Si III-weighted quantities predict higher temperatures that would collisionally ionize Si II leading to the inability to use the Si II/Si III ratio to predict density.

The takeaway message is that CLOUDY ion ratio modelling of low ions provides physical constraints that have the correct order of magnitude, but the simulated CGM is a more complex multiphase medium with adjacent low ions probing a spread in physical parameters. Our example here shows that the CLOUDY models best predict the physical conditions traced by the lower ion, Si II, which is why we show Si II-weighted n_{H} in the middle panel of Fig. 16. It is not unexpected that an ion ratio such as Si II/Si III probes a multiphase spread in gas properties given the different ionization potentials to ionize to these states (8.2 eV for Si II, 16.3 eV for Si III), but our exploration of simulated physical properties provides confidence that CLOUDY single-phase modelling yields physically meaningful constraints. Therefore, it is not surprising that our simulations that provide good matches to Si II, Si III, and their ratios find similar densities for low-ion clouds as Prochaska et al. (2017).

We recommend expanding this type of analysis in future work. Generating mock multi-ion observations from simulations, assessing their goodness of fit, and determining the underlying physical parameters of the simulation ion-by-ion is necessary to understand how single-phase models used to constrain physical properties perform. A potential next step would be to perform this analysis on mock spectra with individual components evaluated (e.g. Werk et al. 2016) instead of an integrated system column density.

6 SUMMARY

We explore the low metal ions observed by the COS-Halos survey (Werk et al. 2013, 2014) using the set of EAGLE non-equilibrium (NEQ) zoom simulations that Oppenheimer et al. (2016) found to reproduce the observed O VI bimodality around star-forming and passive galaxies. Our exploration considers C II, C III, C IV, Mg II, Si II, Si III, and Si IV, with the primary focus on the silicon ions. Simulated column densities of silicon and carbon almost always agree with COS-Halos to within a factor of two, although Mg II is nearly a factor of 10 too weak. Our main results regarding L^* haloes ($M_{200} = 10^{11.7} - 10^{12.3} M_{\odot}$) hosting star-forming galaxies, and group-sized haloes ($M_{200} = 10^{12.7} - 10^{13.3} M_{\odot}$) hosting mostly passive galaxies are as follows:

(i) Simulated low metal ion column densities show 1) little dependence on galaxy sSFR, 2) a patchy covering fraction, and 3) a declining covering fraction at larger impact parameters. Low ions observed by COS-Halos trace a phase that is physically and spatially distinct from that traced by O VI. Low ions trace 10^4 K clumps mainly within 100 kpc of both star-forming and passive galaxies, while the O VI traces the ambient $\sim 10^{5.5}$ K medium at radii $\gtrsim 150$ kpc around L^* galaxies. (Section 3, Figs 1 and 2)

(ii) Simulated group galaxies often have neighbouring galaxies with $M_* > 10^{10} M_{\odot}$ inside 300 kpc, which is also true for some COS-Halos passive galaxies. Neighbouring galaxies in and around group haloes increase the average low-ion column densities within 150 kpc. However, when stringently restricting the sample to galaxies without neighbouring galaxies within 300 kpc, group haloes still show more low-ion metals in the outer CGM at > 100 kpc than L^* galaxies do. (Section 3.1, Fig. 3)

(iii) Going from the lowest ionization species (Mg II, Si II, C II) through intermediate species (Si III, Si IV, C IV) to O VI, the dispersions and dependence on impact parameter of the column densities decline. At the same time, the dependence on sSFR increases with O VI showing the strongest (i.e. O VI-sSFR correlation, O16) and C IV showing the second strongest dependence. (Section 4, Fig. 4)

(iv) Cumulative distribution functions of simulated silicon and carbon ion column densities overlap the 95 per cent confidence limits of COS–Halos with survival statistics applied to censored observations. This is true when considering the entire sample, subdividing by sSFR, and sub-dividing star-forming galaxies into two impact parameter bins. The good agreement for C and Si is neither seen for Mg II (≈ 1 dex too weak) nor for O VI (≈ 0.5 dex too weak). (Section 4.1, Figs 5 and 6)

(v) Modifications to our fiducial NEQ simulation model can significantly alter column densities. Self-shielding can raise low-ion column densities by up to 0.4 dex, while ionization from local sources can decrease low-ion column densities but raise high-ion column densities including O VI. However, NEQ ionization does not significantly alter low-ion abundances when assuming a uniform ionization background. The under-prediction of Mg II is likely in part caused by too low-Mg yields in the EAGLE enrichment model. (Section 4.2, Fig. 7)

(vi) The total mass in metals traced by circumgalactic Si II – Si IV is nearly $10^8 M_\odot$ inside L^* haloes, with group haloes having $\approx 2/3$ rd this amount. The mass of cool ($T < 10^5$ K) CGM metals inside R_{200} is nearly invariant from $M_{200} = 10^{11.8} - 10^{13.2} M_\odot$ at $1.0 - 1.5 \times 10^8 M_\odot$, while the fraction of CGM metals that are cool falls from 75 per cent to 3 per cent as haloes increase in mass and virial temperature. (Section 5.1, 5.2, Figs 8 and 9)

(vii) The pressures of low-ion metal clumps agree with Werk et al. (2014) and McQuinn & Werk (2017) with values of $10 - 40 \text{ cm}^{-3} \text{ K}$ for the CGM traced by COS–Halos. Except inside 40 kpc, there is little dependence of n_H , T , and P on impact parameter, which is also supported by Stocke et al. (2013), who found similar pressures independent of galaxy impact parameter. The pressure of the O VI phase at $r \gtrsim 150$ kpc tends to be higher ($30 - 100 \text{ cm}^{-3} \text{ K}$). (Section 5.2, Figs 11, 12 and 13)

(viii) Most silicon ions observed in absorption at $z = 0.2$ will accrete onto a galaxy by $z = 0$, yet most of this accreted gas will be returned to the CGM mainly due to superwind feedback. Higher ionization states and metals in higher mass haloes are less likely to accrete onto the galaxy over this 2.5 Gyr interval. (Section 5.3, 5.4, Figs 14 and 15)

(ix) The silicon ion ratios observed by COS–Halos are broadly reproduced. We show that single-phase CLOUDY models using observed ion ratios of integrated system column densities to calculate physical parameters yield densities of the correct order of magnitude. The agreement is however imperfect because the gas is multiphase. The Si II/Si III ratio in particular sensitively constrains the physical density. (Section 5.5, Fig. 16)

The relatively good match between our simulations and COS–Halos is a genuine prediction of the EAGLE model, because these simulations were not calibrated to reproduce any CGM observations. The low ions are explored as a follow-up to the EAGLE O VI results showing the bimodality of star-forming/passive galaxies in the CGM (O16), and actually exhibit better agreement for silicon and carbon ion column densities here than those O VI results that under-predict COS–Halos column densities by about a factor of three. Our results indicate that low ions trace an almost completely distinct phase of the CGM than O VI – one that traces the metal-enriched re-accretion of gas onto galaxies. It is therefore ironic that low ions do not reflect the star-forming/passive bimodality of O VI, which O16 argues is a result of the O VI fraction peaking for the virial temperatures of haloes hosting star-forming COS–Halos galaxies, while the passive galaxies are predicted to reside in more

massive haloes. This work indicates that even low ions observed in $z \sim 0.2$ $10^{13} M_\odot$ haloes will eventually accrete onto passive galaxies.

Our use of non-equilibrium ionization and cooling does not significantly alter low-ion column densities using the Anarchy formulation of SPH, which is the same finding by O16 for O VI column densities. While we consider only a uniform extra-galactic ionization background here, significant non-equilibrium effects arise for a fluctuating ionization field as Oppenheimer & Schaye (2013b) and Segers et al. (2017) demonstrated exploring AGN proximity zone fossils. High ions, including as C IV and O VI, remain out-of-equilibrium for time-scales $\sim 10^7$ yrs since they are flash photo-ionized at lower CGM densities than the low ions we explore here. While Oppenheimer et al. (2018) argued the O VI may be enhanced by ≈ 0.5 dex around typical star-forming galaxies with recent but mostly inactive AGN, they also showed that low ions are much less affected, owing to their shorter recombination times at the higher densities they trace.

It is concerning that our good match for low ions relies on the low-ion clumps being somewhat out of pressure equilibrium with the higher pressure, hot medium. Although this confirms other calculations (e.g. McQuinn & Werk 2017), our result likely indicates a numerical issue with our implementation of SPH and insufficient resolution at the phase transition (see Appendix D). However, we stress that there is not a >2 order of magnitude pressure imbalance implied in Werk et al. (2014) using the model of Maller & Bullock (2004). New COS–Halos cloud densities are derived to be higher (Prochaska et al. 2017), while EAGLE feedback reduces densities in the CGM probed by COS–Halos by heating and transporting baryons to the outer CGM, often beyond R_{200} . Therefore, the pressure imbalance between the cool clouds and the hot halo medium may be as little as a factor of a few instead of greater than 100.

Finally, we must also consider the under-prediction of Mg II, especially when compared to the similar ions of C II and Si II, which match or slightly over-predict COS–Halos. We identified several modifications that could enhance this ion including self-shielding and more realistic Mg yields than used in EAGLE, but higher resolution may ultimately be needed to resolve the dense sub-structures that Mg II traces. Recent work by McCourt et al. (2016) suggested dense clouds below the parsec scale could be responsible for dense CGM structures, which is far beyond the capabilities of our simulations even though they show good resolution convergence. Depletion onto dust, which is also not included in our simulations, could also alter the relative abundances. Although, there is still a long way to go before we will fully understand the complex, multiphase CGM observed by COS, we have shown that simulations like ours can already guide the interpretation of the observations and confront fundamental questions of how galaxies and their haloes evolve across a Hubble time.

ACKNOWLEDGEMENTS

The authors are grateful to the referee, Dylan Nelson, for detailed and constructive comments that enhanced this work. The authors would like to thank Joe Burchett, Ryan Horton, Ali Rahmati, John Stocke, Todd Thompson, Colin Norman, Ryan O’Leary, and Jason Prochaska for interesting conversations that contributed to this manuscript. We acknowledge discussions with Sean Johnson and Zachary Hafen about the Turnbull statistical method. Support for Oppenheimer was provided through the NASA ATP grant NNX16AB31G. This work was supported by the Netherlands Organisation for Scientific Research (NWO) through VICI grant 639.043.409. This work used the DiRAC Data Centric

system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grants ST/H008519/1 and ST/K00087X/1, STFC DiRAC Operations grant ST/K003267/1, and Durham University. DiRAC is part of the National E-Infrastructure. RAC is a Royal Society University Research Fellow. AJR is supported by the Lindheimer Fellowship at Northwestern University.

REFERENCES

- Agertz O., et al., 2007, *MNRAS*, 380, 963
- Aguirre A., Hernquist L., Schaye J., Weinberg D. H., Katz N., Gardner J., 2001, *ApJ*, 560, 599
- Asplund M., Grevesse N., Sauval A. J., 2005, in Barnes T. G., III, Bash F. N., eds, ASP Conf. Ser. Vol. 336. Cosmic Abundances as Records of Stellar Evolution and Nucleosynthesis. Astron. Soc. Pac., San Francisco, p. 25
- Bahé Y. M., et al., 2016, *MNRAS*, 456, 1115
- Berg T., et al., 2018, *MNRAS*, in press
- Bordoloi R., et al., 2014, *ApJ*, 796, 136
- Bordoloi R., Heckman T. M., Norman C. A., 2016, preprint ([arXiv:1605.07187](https://arxiv.org/abs/1605.07187))
- Borthakur S., et al., 2016, preprint ([arXiv:1609.06308](https://arxiv.org/abs/1609.06308))
- Burchett J. N., et al., 2016, *ApJ*, 832, 124
- Cantalupo S., 2010, *MNRAS*, 403, L16
- Cen R., Fang T., 2006, *ApJ*, 650, 573
- Chabrier G., 2003, *PASP*, 115, 763
- Crain R. A., et al., 2015, *MNRAS*, 450, 1937
- Crain R. A., et al., 2017, *MNRAS*, 464, 4204
- Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
- Danforth C. W., et al., 2016, *ApJ*, 817, 111
- Feigelson E. D., Nelson P. I., 1985, *ApJ*, 293, 192
- Ferland G. J., Korista K. T., Verner D. A., Ferguson J. W., Kingdon J. B., Verner E. M., 1998, *PASP*, 110, 761
- Ford A. B., et al., 2013, *MNRAS*, 432, 89
- Ford A. B., et al., 2014, *MNRAS*, 444, 1260
- Ford A. B., et al., 2016, *MNRAS*, 459, 1745
- Furlong M., et al., 2015, *MNRAS*, 450, 4486
- Furlong M., et al., 2017, *MNRAS*, 465, 722
- Gnat O., Ferland G., 2012, *ApJS*, 199, 20
- Haardt F., Madau P., 2001, in Neumann D. M., Van J. T. T., eds, Clusters of galaxies and the high redshift universe observed in X-rays, Recent results of XMM-Newton and Chandra, XXXVth Rencontres de Moriond
- Heckman T. M., Norman C. A., Strickland D. K., Sembach K. R., 2002, *ApJ*, 577, 691
- Heckman T., Borthakur S., Wild V., Schiminovich D., Bordoloi R., 2017, *ApJ*, 846, 151
- Hobbs A., Read J., Power C., Cole D., 2013, *MNRAS*, 434, 1849
- Hopkins P. F., 2013, *MNRAS*, 428, 2840
- Hummels C. B., Bryan G. L., Smith B. D., Turk M. J., 2013, *MNRAS*, 430, 1548
- Johnson S. D., Chen H.-W., Mulchaey J. S., 2015, *MNRAS*, 449, 3263
- Kaplan E. L., Meier P., 1958, *J. Am. Stat. Assoc.*, 53, 457
- Keeney B. A., et al., 2017, *ApJS*, 230, 6
- Leitherer C., et al., 1999, *ApJS*, 123, 3
- Liang C. J., Chen H.-W., 2014, *MNRAS*, 445, 2061
- Liang C. J., Kravtsov A. V., Agertz O., 2016, *MNRAS*, 458, 1164
- Maller A. H., Bullock J. S., 2004, *MNRAS*, 355, 694
- McCourt M., Oh S. P., O'Leary R. M., Madigan A.-M., 2016, preprint ([arXiv:1610.01164](https://arxiv.org/abs/1610.01164))
- McQuinn M., Werk J. K., 2017, preprint ([arXiv:1703.03422](https://arxiv.org/abs/1703.03422))
- Muratov A. L., et al., 2016, preprint ([arXiv:1606.09252](https://arxiv.org/abs/1606.09252))
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
- Oppenheimer B. D., Davé R. A., 2006, *MNRAS*, 373, 1265
- Oppenheimer B. D., Schaye J., 2013a, *MNRAS*, 434, 1043
- Oppenheimer B. D., Schaye J., 2013b, *MNRAS*, 434, 1063
- Oppenheimer B. D., Davé R., Kereš D., Fardal M., Katz N., Kollmeier J., Weinberg D. H., 2010, *MNRAS*, 406, 2325
- Oppenheimer B. D., et al., 2016, *MNRAS*, 460, 2157 (O16)
- Oppenheimer B. D., Segers M., Schaye J., Richings A. J., Crain R. A., 2018, *MNRAS*, 474, 4740
- Peebles M. S., et al., 2014, *ApJ*, 786, 54 (P14)
- Planck Collaboration, 2014, *A&A*, 571, A1
- Prochaska J. X., et al., 2017, *ApJ*, 837, 169
- Rahmati A., Oppenheimer B. D., 2018, *MNRAS*, 476, 4865
- Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, *MNRAS*, 430, 2427
- Rahmati A., et al., 2015, *MNRAS*, 452, 2034
- Rahmati A., et al., 2016, *MNRAS*, 459, 310
- Richings A. J., Schaye J., Oppenheimer B. D., 2014, *MNRAS*, 440, 3349
- Rosas-Guevara Y. M., et al., 2015, *MNRAS*, 454, 1038
- Saitoh T. R., Makino J., 2013, *ApJ*, 768, 44
- Salpeter E. E., 1955, *ApJ*, 121, 161
- Schaller M., et al., 2015, *MNRAS*, 454, 2277
- Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
- Schaye J., Aguirre A., Kim T.-S., Theuns T., Rauch M., Sargent W. L. W., 2003, *ApJ*, 596, 768
- Schaye J., et al., 2010, *MNRAS*, 402, 1536
- Schaye J., et al., 2015, *MNRAS*, 446, 521 (S15)
- Segers M. C., et al., 2016, *MNRAS*, 456, 1235
- Segers M. C., Oppenheimer B. D., Schaye J., Richings A. J., 2017, *MNRAS*, 471, 1026
- Simcoe R. A., Sargent W. L. W., Rauch M., 2004, *ApJ*, 606, 92
- Smith B. D., Hallman E. J., Shull J. M., O'Shea B. W., 2011, *ApJ*, 731, 6
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 312
- Stern J., Hennawi J. F., Prochaska J. X., Werk J. K., 2016, *ApJ*, 830, 87
- Stoeck J. T., et al., 2013, *ApJ*, 763, 148
- Suresh J., et al., 2017, *MNRAS*, 465, 2966
- Theuns T., Bernardi M., Frieman J., Hewett P., Schaye J., Sheth R. K., Subbarao M., 2002, *ApJ*, 574, L111
- Thom C., et al., 2012, *ApJ*, 758, L41
- Tilton E. M., Danforth C. W., Shull J. M., Ross T. L., 2012, *ApJ*, 759, 112
- Trayford J. W., et al., 2015, *MNRAS*, 452, 2879
- Tumlinson J., et al., 2011, *Science*, 334, 948 (T11)
- Tumlinson J., et al., 2013, *ApJ*, 777, 59
- Turnbull B. W., 1974, *J. Am. Stat. Assoc.*, 69, 169
- Turner M. L., Schaye J., Crain R. A., Theuns T., Wendt M., 2016, *MNRAS*, 462, 2440
- Turner M. L., et al., 2017, *MNRAS*, 471, 690
- Wendland H., 1995, *Adv. Comput. Math.*, 4, 389
- Werk J. K., et al., 2012, *ApJS*, 198, 3
- Werk J. K., et al., 2013, *ApJS*, 204, 17 (W13)
- Werk J. K., et al., 2014, *ApJ*, 792, 8 (W14)
- Werk J. K., et al., 2016, *ApJ*, 833, 54
- Wiersma R. P. C., Schaye J., Smith B. D., 2009a, *MNRAS*, 393, 99
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, *MNRAS*, 399, 574
- Wiersma R. P. C., Schaye J., Dalla Vecchia C., Booth C. M., Theuns T., Aguirre A., 2010, *MNRAS*, 409, 132

APPENDIX A: IONIZATION EQUILIBRIUM TESTS

Non-equilibrium effects are noted to be small when comparing equilibrium and NEQ runs at $z = 0.2$ in Section 4.2. Fig. A1 considers NEQ effects more in-depth by making a different comparison: our standard NEQ runs (blue) are compared to NEQ runs where we iterate our solver to ionization equilibrium in post-processing (termed ioneq(NEQ) runs, green). This differs from the comparison than the one in Fig. 7 where the ioneq runs are separate runs. We use only three zoom simulations in haloes of $10^{12.1-12.3} M_{\odot}$




Figure A1. SMOHALOS cumulative distribution functions simulating the COS–Halos blue L^* galaxies using three zoom simulations hosting star-forming galaxies in haloes of $10^{12.1}–10^{12.3}M_{\odot}$, all at $z = 0.2$. Blue indicates the standard NEQ simulations, green (often overlapping blue curves) indicates the NEQ simulations where ionization equilibrium is assumed, gold indicates the same haloes in the Recal-L025N0752 volume, and magenta shows higher resolution NEQ $M4.4$ resolution simulations.




Figure A2. Median silicon ionization fractions, x_{ion} , of gas particles as a function of density at $z = 0.2$ for three temperatures ($T = 10^{3.7}$ K in blue, $10^{4.0}$ K in black, and $10^{4.3}$ K in red) for the NEQ runs. Particles are selected from L^* zooms to lie within 300 kpc of the central galaxy. We also show the ionization fractions for ioneq(NEQ) runs in cyan, grey, and magenta, and for the most part they lie directly beneath the NEQ curves except for instances where there is little data. The $1-\sigma$ spreads are also shown with thin lines and do not indicate a greater spread for NEQ versus ioneq(NEQ) runs.

to simulate COS–Halos galaxies with $M_* = 10^{9.6}–10^{10.5}M_{\odot}$ and $\text{sSFR} > 10^{-11}\text{yr}^{-1}$. The comparison of the green and blue CDFs shows hardly any difference between assuming ionization equilibrium and using the NEQ runs, which bolsters our point that NEQ effects are generally unimportant for the diagnostics of absorption line column densities.

To further investigate NEQ effects, we show gas particle ionization fractions of Si II, Si III, and Si IV in Fig. A2 as a function of density selected from our L^* zoom simulations. Three different tem-

peratures ($T = 10^{3.7}, 10^{4.0}, 10^{4.3}$ K) show little difference between NEQ and the ioneq(NEQ) ionization fractions, which overlap each other.

Finally, we remake but do not show Fig. 2 plotting the mean column densities as a function of impact parameter for the distinct ioneq simulation runs using 23 L^* and group haloes at $z = 0.20$. We do not see any distinguishing characteristics between the noneq and ioneq runs, which is consistent with Fig. 7.

APPENDIX B: COMPARISON TO EAGLE VOLUMES

We select the same three L^* haloes from the previous section in the Recal-L025N0752 volume and plot them in gold in Fig. A1. As discussed in Section 2.1, these simulations at the same resolution provide a comparison between particle-based (blue) versus kernel-smoothed (gold) metallicities, modulo NEQ effects that we have shown to be small. There is considerable overlap between the two samples indicating that the treatment of metallicity smoothing may affect low ions even less than O VI.

Looking more closely, we find Si II – Si IV columns are ~ 0.1 – 0.2 dex lower inside 100 kpc and ~ 0.1 dex higher outside 100 kpc for the kernel-smoothed EAGLE volume across 9 L^* haloes (not shown). The K–M estimator split by impact parameter for the blue L^* sample shows lower columns inside 75 kpc by ~ 0.2 dex and similar outside 75 kpc, but does not favor an overall better match to the observations (volume is better for Si II and C II, but worse for Mg II, not shown). Cool clump formation does not appear to be inhibited by the metallicity smoothing scheme.

APPENDIX C: RESOLUTION TEST

High-resolution $M4.4$ simulations ($m_{\text{SPH}} = 2.7 \times 10^4 M_\odot$) with $8 \times$ better mass resolution than our standard runs are shown in magenta in Fig. A1. Their softening length is 175 proper pc below $z = 2.8$, which is a factor of two smaller than our fiducial resolution. These runs show 0.1 – 0.3 dex lower column densities for silicon species and C II than standard $M5.3$ runs, but a 0.2 dex higher column density for Mg II. The $M4.4$ runs show a similar increase in O VI of ≈ 0.1 dex as O16 showed in their exploration of $M4.4$ runs compared to $M5.3$ NEQ runs.

The $M4.4$ runs have 0.18 dex lower stellar masses than $M5.3$ runs, which makes their stellar masses $3 \times$ lower than abundance matching constraints (O16). This is the main reason we do not use these high-resolution runs in the main portion of the paper, since we did not recalibrate the feedback prescriptions for this resolution in a similar manner that S15 did for standard (Ref) and high-resolution (Recal) EAGLE prescriptions. The lower stellar masses are likely the result of improved numerical efficiency of the stochastic thermal feedback and indicates $8 \times$ higher resolution simulations need a recalibrated feedback prescription as expected (see S15 for a discussion). Because metal production scales approximately with stellar mass, such a recalibration may boost the column densities. Interestingly, Mg II moves 0.3 dex higher relative to C II and Si II when increasing to $M4.4$ resolution, which may help explain the Mg II under-estimate.

APPENDIX D: PRESSURES IN SPH SIMULATIONS

To explore why our low-ion clumps appear to have lower pressures than the hot, ambient medium, we consider whether these phases are spatially coincident. We display Mollweide projections of the CGM at $0.3R_{200}$ of our $M_{200} = 10^{12.1} M_\odot$ zoom in Figure D1. These projections are made using the SPH smoothing lengths to project particles onto a thin slice at $0.3R_{200}$ (61 kpc) from the galaxy center. The two left-hand panels show n_H and T , indicating cool $< 10^5$ K clumps in an ambient $\gtrsim 10^6$ K medium. Clumps with $T \lesssim 10^5$ K and $n_H \gtrsim 10^{-3} \text{ cm}^{-3}$ have a very small filling factor in this kernel-smoothed slice projection.

In the upper right-hand panel of Fig. D1, we plot P/k by multiplying the particle $n_H/(X_H \mu)$ and T , as calculated in Figs 11 and 13. Most cool clumps appear as blue spots indicating lower pressures, albeit not as low as the individual particle pressures owing to the smoothing method. There is variation in pressure across the projection, albeit less variation than in n_H and T . However, some of the greatest pressure differences occur around the cool clumps, which leads to a discussion of the Anarchy formulation of SPH used in EAGLE (e.g. Schaller et al. 2015). The pressure-entropy or density-independent formulation SPH formulation (Hopkins 2013; Saitoh & Makino 2013) aims to preserve pressure across a boundary to overcome the creation of an artificial surface tension barrier that spuriously prevents mixing between cool clumps embedded in a hot medium (Agertz et al. 2007). However, this pressure is calculated differently than our particle pressures as it uses smoothed entropy and internal energy to determine the “weighted” pressure and the “weighted” density that affects the hydrodynamic equations of motion. Hence, the pressures we calculate throughout this work are inconsistent with the pressure gradients determining the dynamics.

When we calculate the weighted pressure, using a separately tracked smoothed entropy variable, this pressure (lower right-hand panel of Fig. D1) displays much less variation around the cold clumps. For our exploration of low ions occupying the centres of cold clumps, ionization fractions do not differ much between using a particle density versus a smoothed “weighted” density, $n_{H, \text{wd}}$ as introduced in equation (8) of Schaller et al. (2015) that relates to the weighted pressure such that $P_{\text{wd}}/k = n_{H, \text{wd}} T/(X_H \mu)$. The bigger differences are in interface particles in between the hot, ambient and cool, clump phases, which will result in different ionization fractions and cooling rates if we used the weighted density. For our CGM multiphase densities, the Si III and Si IV ionization fractions are more likely to be affected than the Si II ionization fractions. The outer CGM harbouring significant O VI at $r \gtrsim 100$ kpc is not affected by these interface issues and represents a predominantly single-phase CGM.

The difference between weighted and particle pressures at the clump interfaces indicates that the SPH scheme suffers from smoothing where the calculation of cold clump densities (particle and weighted) is affected by lower density, ambient particles with smoothing lengths overlapping even the centers of cold clumps. This is fundamentally a resolution issue, where smoothing lengths based on overlapping 58 neighbours using the C2 Wendland (1995) kernel will lead to a smoothed gradient of physical parameters across a boundary. Unfortunately, while our clumps are resolved with multiple particles (of the order of 50–100), their central densities are affected by low-density ambient particles. Higher resolution could yield a different answer, but our resolution tests here and in O16 show similar pressure differences between phases, indicating the clump masses scale with the SPH particle mass. Hence, the same issues persist at the $M4.4$ resolution, but the clumps are of lower mass. Resolving the multiphase CGM spanning 2–3 dex in density and temperature at a boundary remains a challenging problem, and the numerical behaviour of cold clumps in an ambient medium should also be assessed using other hydro solvers.

The Anarchy formulation of SPH prevents the spurious formation of cool clumps in a hot medium relative to standard SPH as explored in comparative EAGLE simulations (Schaller et al. 2015). This study showed many fewer clumps for Anarchy in a cluster-sized halo. Another modified SPH formulation by Hobbs et al. (2013, SPHS) also showed a reduction in cool clumps versus standard

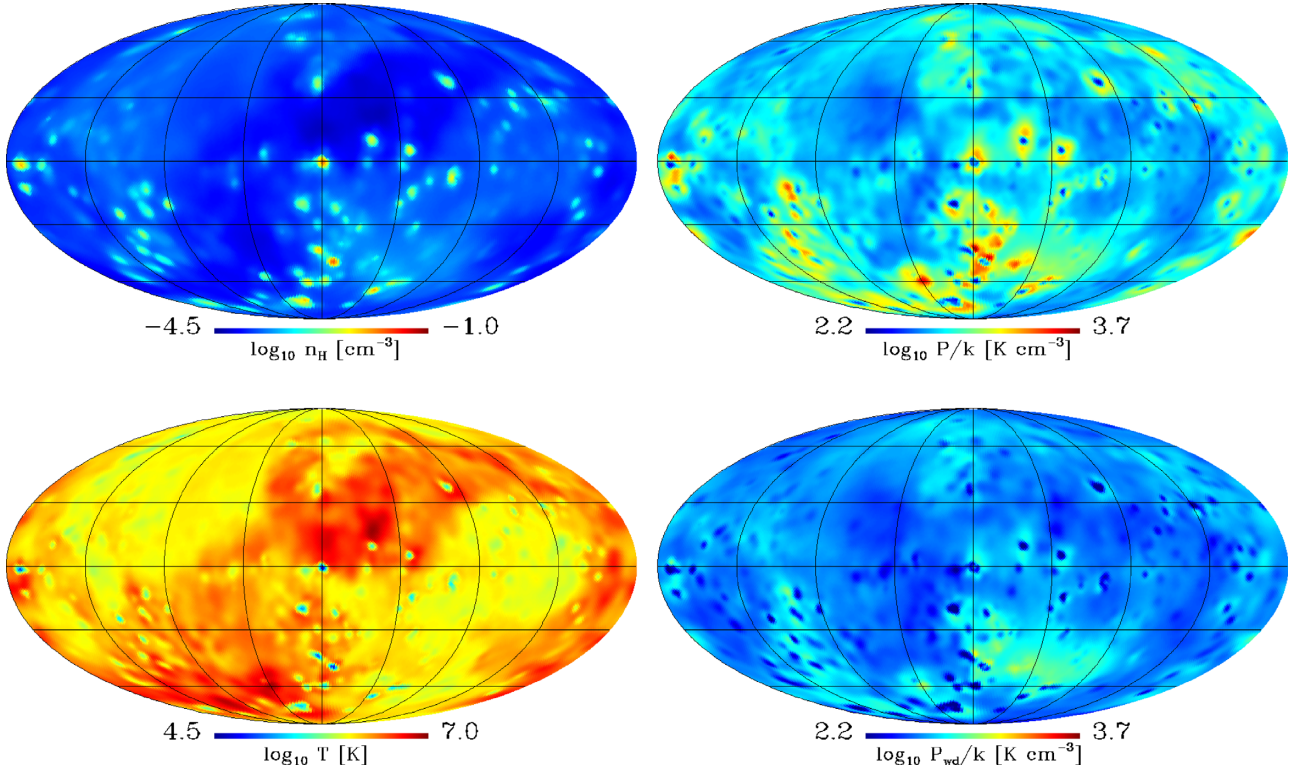


Figure D1. Mollweide projections of SPH particles at $0.3R_{200}$ of a $10^{12.1} M_{\odot}$ halo. Density and temperature (upper and lower left-hand panels) show a multiphase structure of dense, cool CGM embedded in a hot ambient medium. The pressures (upper right-hand panel) show much less variation between the phases, although some of the cool clumps appear to be at lower particle pressures. The weighted pressure used in the hydro equations of motion is shown in the lower right-hand panel and yields much smaller pressure differences.

SPH, but they instead showed the formation of dense filamentary structures in modified SPHS. Their study assumed primordial cooling, which may be a large part of the difference with our zooms (and Schaller et al. (2015) that have enhanced cooling in metal-enriched particles leading to more clumpy structures. Further investigations should explore the effect of metallicity in cool clump formation,

although we have found little difference between the metallicity scheme when comparing kernel-smoothed to particle-based metallicities in Appendix B.

This paper has been typeset from a \LaTeX file prepared by the author.