Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **499,** 2617–2635 (2020)
Advance Access publication 2020 October 1

doi:10.1093/mnras/staa3007

Jet feedback and the photon underproduction crisis in SIMBA

Jacob F. Christiansen,[1]★ Romeel Davé ,[1,2,3] Daniele Sorini [1] and Daniel Anglés-Alcázar[4,5]

[1]*Institute for Astronomy, Royal Observatory, Edinburgh EH9 3HJ, UK*
[2]*Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town 7535, South Africa*
[3]*South African Astronomical Observatories, Observatory, Cape Town 7925, South Africa*
[4]*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA*
[5]*Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269-3046, USA*

Accepted 2020 September 21. Received 2020 August 29; in original form 2019 November 2

ABSTRACT

We examine the impact of black hole jet feedback on the properties of the low-redshift intergalactic medium (IGM) in the SIMBA simulation, with a focus on the Lyα forest mean flux decrement D_A. Without jet feedback, we confirm the photon underproduction crisis (PUC) in which $\Gamma_{\rm H\,I}$ at $z = 0$ must be increased by 6 times over the Haardt & Madau value in order to match the observed D_A. Turning on jet feedback lowers this discrepancy to ∼2.5 times, and additionally using the recent Faucher–Giguère background mostly resolves the PUC, along with producing a flux probability distribution function in accord with observations. The PUC becomes apparent at late epochs ($z \lesssim 1$) where the jet and no-jet simulations diverge; at higher redshifts SIMBA reproduces the observed D_A with no adjustment, with or without jets. The main impact of jet feedback is to lower the cosmic baryon fraction in the diffuse IGM from 39 per cent to 16 per cent at $z = 0$, while increasing the warm-hot intergalactic medium (WHIM) baryon fraction from 30 per cent to 70 per cent; the lowering of the diffuse IGM content directly translates into a lowering of D_A by a similar factor. Comparing to the older MUFASA simulation that employs different quenching feedback but is otherwise similar to SIMBA, MUFASA matches D_A less well than SIMBA, suggesting that low-redshift measurements of D_A and $\Gamma_{\rm H\,I}$ could provide constraints on feedback mechanisms. Our results suggest that widespread IGM heating at late times is a plausible solution to the PUC, and that SIMBA's jet active galactic nucleus feedback model, included to quench massive galaxies, approximately yields this required heating.

Key words: methods: numerical – galaxies: evolution – galaxies: formation – intergalactic medium – quasars: absorption lines .

1 INTRODUCTION

1.1 Background

The intergalactic medium (IGM) contains the vast majority of cosmic baryons at all cosmic epochs (Meiksin 2009). After the epoch of reionization, the IGM is highly ionized by a cosmic background of ultraviolet photons (UVB) emitted by star-forming galaxies and active galactic nuclei (AGNs). The trace neutral component is detectable as H I Lyα absorption in the spectra of background sources such as quasars, which is known as the Lyman alpha forest. The temperature of this gas is set by a balance between local adiabatic expansion and photo-heating from the metagalactic flux, leading to a relatively simple equation of state (Hui & Gnedin 1997). Combined with the fact that absorbing gas mostly tracks gravitationally driven large-scale structure, this has made the Lyα forest useful for a wide range of cosmological applications.

The optical depth τ of Lyman alpha forest absorbing gas along a given line of sight (LOS) depends on the gas density and the neutral fraction. The neutral fraction is itself proportional to the density and inversely proportional to the H I photoionization rate ($\Gamma_{\rm H\,I}$). If we consider the mean optical depth in the Lyα forest, it thus scales as the square of the mean baryonic density (which is $\propto\Omega_b$), and inversely with $\Gamma_{\rm H\,I}$: $\bar{\tau} \propto \Omega_b^2/\Gamma_{\rm H\,I}$, with constants that depend on cosmology (Rauch et al. 1997), and a small correction owing to the temperature dependence of the H I recombination rate. The fluctuations around this mean optical depth can thus be used to measure the matter power spectrum, assuming that the baryons trace matter (e.g. Weinberg, Katz & Hernquist 1998). The mean optical depth, meanwhile, can be used to constrain a combination of Ω_b and $\Gamma_{\rm H\,I}$.

Rauch et al. (1997) applied this approach to measurements of the mean flux decrement in the Lyα forest at $z \sim 2$–3 in order to estimate Ω_b, assuming $\Gamma_{\rm H\,I}$ taken from Haardt & Madau (1996), and obtained $\Omega_b > 0.021h^2$. The Haardt & Madau (1996) background was estimated from the number density of observed quasars and star-forming galaxies plus radiative transfer through a clumpy IGM, assuming that all ionizing photons from quasars and a small fraction of such photons from star-forming galaxies escaped. Despite substantial uncertainties in source count observations at that time, this value for Ω_b turned out to be in good agreement with determinations from the deuterium abundance (Tytler, Fan & Burles 1996) and subsequently the cosmic microwave background (Planck Collaboration XIII 2016).

At lower redshifts, the growth of the cosmic web results in gas shock heating on filamentary structures as it accretes supersonically (Davé et al. 1999). This generates the so-called warm-hot

★ E-mail: jacobfc96@yahoo.com

intergalactic medium (WHIM; Cen & Ostriker 1999; Dave et al. 2001) of gas outside bound haloes in the $T \sim 10^5$–10^7 K temperature range. Owing to the non-linear processes involved, gas dynamical simulations are required to study the growth of the WHIM, and concomitantly, the reduction in Lyα forest baryons. Such simulations broadly predict that roughly one-third of cosmic baryons at the present epoch are in the WHIM (Dave et al. 2001; Davé et al. 2010; Smith et al. 2011). It is very challenging to detect such warm-hot gas observationally since the hydrogen is fully ionized, so metal line absorbers must be used instead, which are weaker and more uncertain. None the less, an observational census primarily from O VI absorption suggests that such predictions are broadly consistent with current data (Tripp, Savage & Jenkins 2000; Shull, Smith & Danforth 2012).

In spite of the increased complexity introduced by the WHIM, it is still possible to use the Lyα forest mean flux decrement to measure $\Gamma_{\rm H\,I}$, given that Ω_b is now well determined from other avenues. Indeed, at $z \sim 0$, this is currently the most robust approach to measuring $\Gamma_{\rm H\,I}$, because it is impossible to directly detect the 912 Å photon background directly given foreground Galactic absorption, and other approaches such as Hα fluorescence are extremely challenging (though see Fumagalli et al. 2017). Dave & Tripp (2001) used this approach on *Hubble Space Telescope* (*HST*) Imaging Spectrograph data to measure $\Gamma_{\rm H\,I}(z=0) = 10^{-13.3\pm0.7}\,{\rm s}^{-1}$. In the meantime, Haardt & Madau (2001) had improved upon their estimate of $\Gamma_{\rm H\,I}$ evolution from source count modelling, and determined $\Gamma_{\rm H\,I}(z=0) = 10^{-13.08}\,{\rm s}^{-1}$, consistent with the Ly$\alpha$ forest measurements. Thus, it appeared that $\Gamma_{\rm H\,I}$ at $z = 0$ was now pinned down to within a factor of a couple.

Measurements of cosmic ionizing photon sources continued to improve. In particular, it became clear that the assumption in Haardt & Madau (2001) of a constant 10 per cent escape fraction of Lyman continuum photons from galaxies was inconsistent with observations; stacked measures of dwarf galaxies at intermediate redshifts suggested instead values below 2 per cent (Rutkowski et al. 2016). Faucher-Giguère et al. (2009) did a new calculation of $\Gamma_{\rm H\,I}(z)$, and estimated $\Gamma_{\rm H\,I}(z=0) = 3.9\times10^{-14}\,{\rm s}^{-1}$. Haardt & Madau (2012) further updated their estimate assuming an evolving escape fraction of $1.8\times10^{-4}(1+z)^{3.4}$ and found an even lower $\Gamma_{\rm H\,I}(z=0) = 2.3\times10^{-14}\,{\rm s}^{-1}$. Hence, as these calculations became more precise, they diverged substantially from the original determination by Haardt & Madau (2001) of $\Gamma_{\rm H\,I}(z=0) = 8.3\times10^{-14}\,{\rm s}^{-1}$, with the latest determinations lower by nearly a factor of 4.

In light of this, Kollmeier et al. (2014) re-investigated constraints on $\Gamma_{\rm H\,I}$ at $z = 0$ from the Lyα forest using new simulations that were substantially improved in dynamic range and input physics compared to those in Dave et al. (2001). This study was also enabled by an improved census of Lyα forest absorbers from *Hubble*'s Cosmic Origins Spectrograph (COS) by Danforth et al. (2016). Kollmeier et al. (2014) found that, in order to match the amplitude of the observed column density distribution or the mean flux decrement, it was necessary to increase the Haardt & Madau (2012, hereafter HM12) value of $\Gamma_{\rm H\,I}(z=0)$ by a factor of $\approx$5, that is, $\Gamma_{\rm H\,I}(z=0) \approx 10^{-13}\,{\rm s}^{-1}$. In other words, if the Lyα forest is robustly predicted in simulations as expected from the simple physics involved, then there was a gross shortfall of observed photon sources relative to that needed to match the observed IGM ionization level. Most of the new found discrepancy owed to the change in the source count estimates of $\Gamma_{\rm H\,I}$. What had initially seemed like a solved problem in 2001 was now, with improved measurements and simulations, yielding a substantial discrepancy. Kollmeier et al. (2014) dubbed this the photon underproduction crisis (PUC) – the Universe did not seem to be producing nearly enough photons to explain the ionization level seen in the Lyα forest.

1.2 Previous investigations into the PUC

The PUC could potentially be solved in a number of ways: (i) it could be that the ionizing background strength in HM12 was underestimated; (ii) it could be that the simulations are simply incorrect, due to numerics; and (iii) it could be that new physics impacts the diffuse IGM, causing it to be more ionized. The subsequent results have been somewhat disparate and controversial, but as we argue below, it is becoming clear that the PUC indeed exists at a level comparable to that presented in Kollmeier et al. (2014).

There has been growing consensus that the HM12 ionizing background may be too weak at $z = 0$. Potential systematic uncertainties include not only the source population emissivity and escape fraction, along with the Lyα column density distribution at the high-$N_{\rm HI}$ end which provides a sink term for ionizing photons. The ionizing background model of Faucher-Giguère et al. (2009) predicts $\Gamma_{\rm H\,I}(z=0)$ about twice the HM12 value, which would partially mitigate the PUC. Khaire & Srianand (2015) did a UV background calculation using updated quasar (QSO) emissivities that were 2 times higher than those in Haardt & Madau (2012), and suggested that this combined with a 4 per cent escape fraction from galaxies could increase the source count estimate of $\Gamma_{\rm H\,I}$ up to the levels required to match Kollmeier et al. (2014). While their assumed QSO emissivity is plausible, the 4 per cent global escape fraction of ionizing photons from galaxies seems less plausible given current measurements (e.g. Rutkowski et al. 2016). A recent update of the Faucher-Giguère et al. (2009) background in Faucher-Giguère (2019, hereafter FG19) similarly found a value of $\Gamma_{\rm H\,I}(z=0)$ about twice that in Haardt & Madau (2012). Furthermore, a recent determination of $\Gamma_{\rm H\,I}(z)$ from Khaire et al. (2019) preferred a higher value for $\Gamma_{\rm H\,I}(z=0)$, but not by more than a factor of 2. Kulkarni, Worseck & Hennawi (2019) found that AGN can only account for half the required photons even though they are expected to greatly dominate the low-z ionizing photon budget. Hence, it appears that a $\sim$2 times systematic difference on the determination of $\Gamma_{\rm H\,I}$ in HM12 is reasonable. However, a factor of $\sim$5 seems difficult to accommodate, as no recent background yields such a large difference (since Haardt & Madau 2001). Thus, the PUC may yet reflect some underlying missing physics in the low-z IGM.

The second option is to appeal to numerics in order to ionize the IGM and lower the required flux. Certainly, the complicated non-linear growth of structure driving the WHIM may be subject to the details of hydrodynamical or other methodology. In general, theoretical studies of the PUC studies can be divided into two classes: those using cosmological hydrodynamical simulations including galaxy formation physics, and those that do not. We recap results from the latter category first.

Shull et al. (2015) compared measurements of the mean flux decrement from COS versus uniform-mesh ENZO simulations, and determined $\Gamma_{\rm H\,I}(z=0) = 4.6\times10^{-14}\,{\rm s}^{-1}$. While still a factor of 2 off from the Haardt & Madau (2012) value, this could be probably accommodated within systematic uncertainties in $\Gamma_{\rm H\,I}$. However, there are two significant caveats. First, uniform-mesh simulations are known to overproduce entropy in low-Mach number shocks and hence increase the amount of numerical heating in the IGM; indeed, in the Dave et al. (2001) comparison of the WHIM in various simulations, the fixed mesh code of Cen & Ostriker (1999) yielded $\gtrsim$50 per cent the baryons in the WHIM, while adaptive resolution codes (both Eulerian and Lagrangian) yielded $\sim$30 per cent. Second, their predicted Lyα absorber column density distribution was sub-

stantially steeper than observed, so while at high column densities ($N_{\rm H\,I} \sim 10^{14}$ cm^{-2}) the amplitude agreed with Haardt & Madau (2012), at low columns ($N_{\rm H\,I} \sim 10^{13}$ cm^{-2}), it agreed better with Haardt & Madau (2001).

Gaikwad et al. (2017a, b) employed a model based on non-radiative hydrodynamic simulations, post-processed to mock up the impact of shock heating and radiative cooling. They also found a substantially reduced PUC, about a factor of ∼2, which they argued could be accommodated by a change in the photoionizing background. While their model enabled efficient exploration of parameter space and could be calibrated to match simulation results in some diagnostics, it is unclear that such a post-processed treatment of heating and cooling can accurately capture the highly non-equilibrium thermodynamics in the IGM.

Viel et al. (2017) examined the PUC in two hydrodynamics simulations: Sherwood (Bolton et al. 2017) and Illustris (Vogelsberger et al. 2014). Sherwood has a large volume but low resolution and did not include star formation or feedback. Illustris included these, and particularly strong AGN feedback. Viel et al. (2017) concluded that in both cases, there was a preference for $\Gamma_{\rm H\,I}(z = 0)$ being ∼1.5–3 times higher than the HM12 value. They further highlighted Lyα linewidths as a potential discriminant between models.

In contrast to these studies, the results from high-resolution simulations with well-constrained galaxy formation physics paint a different picture. These include the original Kollmeier et al. (2014) result, but also studies with different numerical techniques. Tonnesen et al. (2017) confirmed the Kollmeier et al. (2014) result using adaptive mesh refinement simulations with ENZO, which suggested that the PUC is not sensitive to hydrodynamics methodology. Results from Gurvich, Burkhart & Bird (2017) with Illustris using the AREPO moving mesh code also showed a similar PUC if no AGN feedback is included; we discuss these results further below. Hence, the numerics of the hydrodynamic solver does not seem to play a role, so long as one has high resolution and full galaxy formation physics.

If the solution to the PUC cannot be fully solved obtained by appealing to uncertainties in source population modelling or numerical methodology, then the remaining potential solution is that models are missing some widespread IGM heating mechanism that would lower the Lyα absorbing gas. Kollmeier et al. (2014) investigated whether the then-popular blazar heating model of Broderick, Chang & Pfrommer (2012) could accommodate this, and determined that it could go partway, but it produced a column density distribution that was shallower than observed. Since the mean flux tends to be dominated by near-saturated lines ($N_{\rm H\,I} \sim 10^{13.7}$ cm^{-2}) occurring in mildly overdense regions (Davé et al. 1999), it was not possible to solve the PUC by mostly heating void gas. Wakker et al. (2015) strengthened the case for the PUC using the same simulations as Kollmeier et al. (2014) but a different observational measure, by showing that the H I column density as a function of filament impact parameter required the HM12 ionizing background at $z = 0$ to be increased by 4–5 times.

A landmark study was that of Gurvich et al. (2017), who investigated the PUC in Illustris. Unlike most previous simulations studying the PUC, Illustris included strong AGN feedback. This was primarily designed to quench star formation in massive galaxies by heating halo gas, but as a by-product it also deposited energy not only in the vicinity of the AGN, hence decreasing Lyα absorption in quasar environs (Sorini et al. 2018; Sorini 2017), but also into the diffuse IGM gas, thus affecting the Lyα forest. Specifically, by comparing the fiducial Illustris run to a run with no AGN feedback, Gurvich et al. (2017) showed that such injection of energy into the diffuse IGM clearly went towards resolving the PUC in Illustris. Assuming a Faucher-Giguère et al. (2009) UVB (∼×2 higher than HM12), Gurvich et al. (2017) were able to match the observed mean flux decrement, although their column density distribution slope did not match COS data. Such a large impact from feedback was somewhat surprising, since it is commonly believed that galactic feedback does not strongly impact the diffuse IGM far from galaxies.

Although a promising solution to the PUC, Illustris at the same time greatly overevacuates gas from massive haloes (Genel et al. 2014), so it is likely that their AGN feedback model is too strong, or adds energy in the wrong manner. These results also clarify the findings of Viel et al. (2017), which at face value suggested that the inclusion of AGN feedback did not have a large impact owing to the similarity of results between Sherwood and Illustris. Gurvich et al. (2017) instead showed that this was not the case: AGN feedback had a substantial effect, and the agreement owed to comparing a low-resolution simulation with no galaxy formation to one with full galaxy formation physics but including AGN.

In short, hydrodynamic models that include full galaxy formation physics constrained to match a variety of other observations all uniformly show a PUC at the ∼4–6 times level. It is possible that 2 times of this may be explained via a re-evaluation of the low-redshift photoionization rate from HM12. However, this still leaves a factor of ∼2–3 to explain. Surprisingly, AGN feedback may impact the diffuse IGM far from galaxies to mitigate the PUC, but it remains unclear whether a fully successful model can be developed that self-consistently reproduces both galaxies and the low-z Lyα forest. This is the goal of our present study.

1.3 This work: exploring AGN feedback in the IGM with SIMBA

AGN feedback is crucial for reproducing the observed galaxy population (Somerville & Davé 2015), but it is not well understood. Recent years have seen the development of various AGN feedback models within cosmological hydrodynamic simulations, primarily designed to quench massive galaxies as observed (Vogelsberger et al. 2014; Schaye et al. 2015; Weinberger et al. 2018; Henden et al. 2018). One successful recent model is the SIMBA simulation. SIMBA uses an observationally motivated two-mode feedback model, where at high Eddington rates it follows observed ionized or molecular gas outflow scalings, while at low Edington rates it switches to a jet mode with outflow speeds up to ∼8000 km s^{-1}, broadly implementing the physical scenario outlined in Best & Heckman (2012). The two-mode approach is qualitatively similar to the model in IllustrisTNG (Weinberger et al. 2018), although SIMBA uses stably bipolar outflows and significantly less total energy which is more consistent with observations of the kinetic power in radio jets (e.g. Whittam et al. 2018). Such jet feedback can potentially carry matter and hence energy far away from its host galaxy into the diffuse IGM (Borrow, Angles-Alcazar & Dave 2019). SIMBA is able to quench galaxies in good agreement with observations over cosmic time, and more relevantly for this work, yields a hot baryon fraction in massive haloes that is consistent with observations (Davé et al. 2019), so is not over- or underevacuating halo baryons. Moreover, it reproduces observed X-ray scaling relations in groups and clusters (Robson & Davé 2020). Hence, it provides a plausible AGN feedback model that can be used to investigate the PUC.

In this paper, we examine the PUC in the SIMBA simulation. To do so, we generate simulated LOSs in Lyα absorption, and quantify the variation needed in the strength of the assumed photoioniZing background in order to match observations of the mean flux decrement D_A. We focus on D_A and not the column density distribution of absorbers in order to avoid uncertainties associated with line

identification and fitting, which can be quite sensitive to spectral resolution and signal-to-noise ratio (S/N, e.g. Dave et al. 2001). In particular, we investigate the role of the jet mode of AGN feedback in SIMBA. We show that this type of AGN feedback has a large impact on the PUC, while other AGN feedback modes in SIMBA (cf. radiative and X-ray) have minimal impact. We also compare to the MUFASA simulation results, which assumed a different halo-based quenching model that did not employ jets, though still matched massive galaxy properties. We find that SIMBA's AGN jet feedback model is crucial for obtaining agreement between the $\Gamma_{\rm H\,I}$ required to match the D_A observations and modern determinations of $\Gamma_{\rm H\,I}$ from source population modelling, suggesting that widespread IGM heating from AGN is a key factor in helping to solve the PUC.

This paper is organized as follows. In Section 2, we review the SIMBA simulations used in this work. In Section 3, we present some global IGM physical characteristics in the SIMBA runs with and without AGN jets. In Section 4, we present our main results in examining the PUC in SIMBA in runs with and without jets. In Section 6, we discuss the PUC in other AGN feedback tests in SIMBA, and in MUFASA. In Section 7, we discuss various modelling uncertainties. In Section 8, we summarize our results.

2 THE SIMBA SIMULATIONS

2.1 Input physics and cosmology

SIMBA (Davé et al. 2019) is a cosmological hydrodynamic simulation that uses a meshless finite mass (MFM) hydrodynamics solver (Hopkins 2015), which can be classified as an Arbitrary Lagrangian Eulerian (ALE) code. MFM employs a Riemann solver that is able to handle strong shocks and shear flows accurately, without introducing an artificial viscosity (Hopkins 2015). This is particularly beneficial in situations where high Mach number flows and strong shocks are an important physical aspect in the problem, which is the case here in studying the impact of high-velocity jet outflows (described below) on diffuse IGM gas.

SIMBA further employs a number of state of the art subgrid physical processes to form realistic galaxies. Photoionization heating and radiative cooling are implemented using the GRACKLE-3.1 library[1] (Smith et al. 2017) assuming ionization but not thermal equilibrium, with collisional ionization rates for H taken from Abel et al. (1997) and recombination rates from Hui & Gnedin (1997). GRACKLE accounts for a Haardt & Madau (2012) ionizing background modified to account for self-shielding based on the Rahmati et al. (2013) prescription (Emerick, private communication). The strength of the ionizing background has a very weak impact on the gas dynamics during the simulation, hence it is possible to meaningfully vary this assumption in post-processing without introducing significant errors (Katz, Weinberg & Hernquist 1996). The production of 11 different elements (H, He, C, N, O, Ne, Mg, Si, S, Ca, and Fe) are tracked, from Type II and Ia supernovae and stellar evolution. SIMBA tracks dust growth and destruction on the fly, for each individual element (a detailed investigation of the dust model can be found in Li, Narayanan & Davé 2019). Star formation is based on a Kennicutt–Schmidt Law (Kennicutt 1998) scaled by the H_2 fraction, which is calculated for each particle using its local column density and metallicity following Krumholz & Gnedin (2011). Galactic outflows are implemented as kinetic decoupled two-phase winds, as in MUFASA (Davé, Thompson & Hopkins 2016), with an updated mass-loading factor based on particle tracking results from the Feedback in Realistic Environments zoom simulations (Anglés-Alcázar et al. 2017b). For more details on these implementations, see Davé et al. (2019).

[1] https://grackle.readthedocs.io/

2.2 Black hole accretion and feedback

The energy release from black holes,that is, AGN feedback, has a significant impact on the properties of the galaxy and surrounding matter (Fabian 2012). SIMBA is notably unique in its way of modelling black hole processes. Owing to the importance of SIMBA's black hole growth and feedback model for this study, we describe it more detail here; further details are available in Davé et al. (2019).

SIMBA employs a unique two-mode black hole accretion model. Cold gas ($T < 10^5$ K) is accreted via a 'torque-limited' subgrid model that captures how angular momentum loss via dynamical instabilities limits gas inflows into the region near the black hole (Hopkins & Quataert 2011; Anglés-Alcázar et al. 2017a). Meanwhile, hot gas is accreted following the Bondi (1952) formula. The torque-limited mode is appropriate for when black holes are growing in a cold rotationally supported disc, while Bondi mode is more appropriate for hot gas since it models gravitational capture from a dispersion-dominated medium. SIMBA's accretion model thus represents a step up in realism as opposed to simply using Bondi accretion for all forms of gas, as most other current simulations do. This unique black hole accretion model underpins the implementation of AGN feedback in SIMBA.

As material accretes into the central region, SIMBA assumes that 10 per cent of it falls on to the black hole; this accretion efficiency is calibrated to match the amplitude of the black hole mass–galaxy stellar mass relation (Anglés-Alcázar, Özel & Davé 2013; Anglés-Alcázar et al. 2017a) for massive galaxies from Kormendy & Ho (2013). Accreted gas elements are subtracted a fraction of their mass and immediately ejected as AGN feedback such that the desired momentum flux in the wind ($20L/c$, where $L = 0.1\dot{M}c^2$) is achieved. This ejection is purely kinetic, and purely bipolar – i.e. it is ejected in the $\pm\mathbf{L}$ direction where $\boldsymbol{L}$ is the angular momentum vector of the inner disc (i.e. the 256 nearest neighbours to the black hole). The physical motivation and detailed implementation for SIMBA's kinetic AGN feedback are described more extensively in Davé et al. (2019) and Thomas et al. (2019), but we recap the key points below.

There are two modes for this type of feedback: radiative mode feedback, and jet mode feedback. The radiative mode in SIMBA happens when there is a high relative accretion rate around a black hole, above a few percent of the Eddington rate. In this mode, the ejected material is kicked with speeds typically around 1000 km s^{-1}, scaled to follow observations of ionized gas outflows from Perna et al. (2017), and its temperature is not changed in order to represent a multiphase outflows as observed. At lower Eddington ratios, the jet feedback mode begins to switch on, with full jets achieved below 2 per cent. The jet mode ejects gas at much higher velocities than the radiative mode, reaching a maximum of $\sim$8000 km s^{-1}. The jet mode also raises the temperature of the ejected particles, based on observations indicating that jets are mostly made of hot plasma (Fabian 2012). At all times, the amount of matter ejected is mass-loaded from the inner disc in order to have the momentum flux of the outflow be $\approx 20L/c$. This two-mode kinetic feedback broadly follows the physical scenario developed in Best & Heckman (2012) and Heckman & Best (2014).

Besides radiative and jet mode feedback, SIMBA includes also X-ray radiation pressure feedback broadly following Choi et al. (2012). This has the effect of pushing outwards on the gas surrounding

the accretion disc based on the high-energy photon momentum flux generated in the black hole accretion disc. It is only activated in low-cold gas content galaxies and when the jet mode is active, because jets tend to be accompanied by strong X-rays and cold dense gas will tend to absorb X-ray energy and radiate it away quickly.

These three forms of AGN feedback – radiative mode, jet mode, and X-ray – combine to create a quenched massive galaxy population in good agreement with observations (Davé et al. 2019), as well as populating them with black holes as observed (Thomas et al. 2019). The jet mode is primarily responsible for quenching, although the X-ray feedback has a non-negligible impact. Radiative mode, meanwhile, has a minimal effect on the galaxy population.

2.3 SIMBA runs

The SIMBA simulations analysed in this paper are run in a cubic box with length $50\,h^{-1}$ Mpc, with 2×512^3 elements. We employ these runs and not the full-size $100\,h^{-1}$ Mpc run with 2×1024^3 from Davé et al. (2019) because we have variants at this box size that enable direct tests of the impact of assumed input physics, particularly AGN feedback. Owing to computational cost, we do not have such variants for the full SIMBA run. None the less, for all checked properties, the 50 and $100\,h^{-1}$ Mpc SIMBA runs agree very well. SIMBA assumes a cosmology consistent with Planck Collaboration XIII (2016) results: $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$, $\Omega_b = 0.048$, $H_0 = 68\,\mathrm{km\,s^{-1}Mpc^{-1}}$, $\sigma_8 = 0.82$, and $n_s = 0.97$. The resulting mass resolution is $1.82 \times 10^7\,\mathrm{M_\odot}$ for gas elements and $9.6 \times 10^7\,\mathrm{M_\odot}$ for dark matter particles.

We run several variants of AGN feedback, turning off one input physics quantity at a time, denoted as follows:

(i) 'SIMBA' denotes a run with all forms of AGN feebdack on.
(ii) 'No-X' denotes a run turning off only X-ray AGN feedback.
(iii) 'No-jet' denotes a run turning off both jet and X-ray feedback.

We also have a run where all AGN feedback is turned off ('No-AGN'), but it turns out the results are indistinguishable from the No-jet case, hence for simplicity we do not show it here. Apparently, the radiative portion of AGN feedback has little impact on the Lyα forest. The other three runs allow a direct quantification of the effects of the jet and X-ray AGN feedback modes in SIMBA. All these runs are started with identical initial conditions.

We will also compare to the MUFASA simulation, the predecessor to SIMBA which does not contain black holes or an explicit AGN feedback model, but rather utilized a heuristic model in which hot halo gas was prevented to cool in order to quench galaxies as observed (Davé et al. 2016; Davé, Rafieferantsoa & Thompson 2017). This also employed a $50h^{-1}$ Mpc box size with 2×512^3 elements, with identical initial conditions to the SIMBA runs.

2.4 Generating spectra

To generate spectra, we employ PYGAD[2] (Roettgers et al. 2020). PYGAD is a full-featured toolkit for analysing particle-based simulations, including creating mock spectra in any desired ion. To generate H I spectra, PYGAD computes the neutral hydrogen fraction for each gas element based on an input (spatially uniform) UVB via a CLOUDY lookup table (including both collisional and photoionization) interpolated to the redshift of the snapshot, puts that gas element into velocity space, smooths its neutral component into velocity bins along a chosen LOS using the same cubic spline kernel employed in SIMBA, and computes the resulting optical depth in each bin. It further computes the optical depth-weighted density and temperature of H I absorbing gas. For these spectra, we use a velocity-space pixel size of $6\,\mathrm{km\,s^{-1}}$. The procedure closely follows what is done in the SPECEXBIN code presented in Oppenheimer & Davé (2006), and PYGAD has been checked to give essentially identical results.

We generate 1000 spectra for each simulation snapshot through the entire box accounting for periodic boundary conditions. We apply a line spread function (LSF) for the COS G-130M grating interpolated to each redshift, since COS data will provide our main comparison sample. We include Gaussian noise with an S/N = 12 pixel^{-1}, equivalent to the S/N = 20 per resolution element that is typical of the COS Guaranteed Time Observations (GTO) team data presented in Danforth et al. (2016). Finally, we apply a continuum fitting procedure broadly following that described in Danforth et al. (2016): we obtain the median flux value all pixels in a given spectrum, remove those that are $>2\sigma$ below that median (where σ is the inverse of the S/N), and then re-fit the remaining pixels, and iterate until convergence in the median at $<10^{-4}$ relative to the previous iteration.

From these spectra, the mean flux decrement D_A was calculated using

$$D_A = \left\langle \sum_i \left[1 - \exp\left(-\tau_i\right)\right] \right\rangle, \tag{1}$$

where τ_i is the optical depth in velocity bin i of a given spectrum, and the average is taken over all 1000 generated spectra.

Since Lyα forest gas is optically thin, the optical depth of any pixel to good approximation scales as $\tau \propto \frac{1}{\Gamma_{\rm H\,I}}$. This means any adjustment to $\Gamma_{\rm H\,I}$ can be related to an adjustment in τ. This then gives us a way to constrain $\Gamma_{\rm H\,I}$ using the observed value of D_A. To do this, we multiply $\Gamma_{\rm H\,I}$ (e.g. from Haardt & Madau 2012) by a value we denote $F_{\rm UVB}$, which corresponds to multiplying each value of τ_i by $1/F_{\rm UVB}$; in practice, we do the latter, since optically thick absorption is extremely rare and does not contribute significantly to D_A. The value of $F_{\rm UVB}$ was then adjusted iteratively until the value of D_A computed via equation (1) matched the observational determination from the combined data of Danforth et al. (2016) and Kirkman et al. (2007, see equation 4)) to within 0.0001, at each snapshot redshift. $F_{\rm UVB}$ can be regarded as the 'photon underproduction factor' – that is, the amount by which $\Gamma_{\rm H\,I}$ must be increased in the simulations (assuming a given photoionizing background) in order to match the observed D_A. This will be a useful metric for us to quantify the PUC in this work.

2.5 Sample mock spectra

Fig. 1 shows some example $z = 0$ mock spectra generated using PYGAD. These spectra were all generated down the same LOS, from our three SIMBA variants: one from the SIMBA simulation with jet feedback enabled (green), one from the No-jet simulation with jet and X-ray feedback turned off (blue), and one from the No-X simulation with jets enabled but with X-ray feedback disabled (red). The five panels, starting from the top show: (1) flux (here shown directly calculated from optical depths without any post-processing such as noise being added, to better facilitate comparisons between the models); 2) gas density, normalized to the cosmic mean (baryonic overdensity); (3) temperature; (4) peculiar velocity; and (5) density of neutral hydrogen.

At $z = 0$, the flux panel shows that the Lyα forest is quite sparse compared with higher redshifts, but a number of absorption lines are still visible. Not all of these features are strong

[2] https://bitbucket.org/broett/pygad

Figure 1. An example of three spectra generated using PYGAD, down the same LOS at $z = 0$: one from the SIMBA simulation with jets turned on (green line), one from the No-jet run (blue line), and one from the No-X run (red line). The five panels, starting from the top show: (1) flux, directly calculated from the optical depths; (2) gas density, normalized to the cosmic mean (baryonic overdensity); (3) temperature; (4) peculiar velocity; and (5) density of neutral hydrogen. All five quantities are plotted in wavelength space. It can be seen that high-density gas at low temperatures results in absorption.

enough to be detectable with existing instruments, but this gives an impression of what the underlying HI distribution is within the variants of the SIMBA simulation, without any noise or instrumental broadening.

The temperature panel shows that the temperatures are much higher in some parts of the simulations with the jets turned on (SIMBA and No-X) than when they are turned off (No-jet). This illustrates how AGN jet feedback provides an extra source of heating that permeates a significant fraction of the IGM. The additional heating means that the fraction of neutral hydrogen in those regions will be dramatically reduced, and hence that there will be much less Lyα absorption. The densities are also significantly impacted, as the higher temperatures result in smoothing the density distribution.

The panels show that in some regions, the spectra appear to be almost identical for all feedback variants. These regions are probing portions of the simulation that have not been affected by jets. The regions that are affected also usually seem to be relatively denser, which owes to the fact that AGN (and hence AGN feedback) are in galaxies that are biased towards the denser regions. However, the lowest density regions, for example, towards the right of the spectrum are also unaffected, presumably because they are too far away for jet feedback to have reached there.

Comparing the green and red lines that differ by the inclusion of X-ray feedback, we see that this form of feedback has a small but non-negligible impact on IGM gas heating. Turning on X-ray feedback (green line) tends to create a slightly more widespread temperature increase around the densest regions, which are presumably closest to galaxies. The stronger absorption feature around 1219 Å in particular shows an interesting case where the X-ray feedback actually has a bigger impact on the absorption than the jet feedback. This is somewhat unexpected, but it shows that X-ray feedback, despite being explicitly confined to dense ISM gas, still provides an energy input that can somewhat impact larger scales. None the less, it is clear the primary impact on the density and temperature structure, and hence IGM absorption, occurs due to the inclusion of jet feedback. In subsequent sections, we will quantify these trends in our ensemble

of spectra, and use this to understand the implications for the PUC.

3 IGM PHYSICAL PROPERTIES

We begin by examining some global properties of the IGM in SIMBA, particularly related to the evolution of the diffuse IGM gas that predominantly gives rise to the Lyα forest.

3.1 Visualizing IGM jet heating

Fig. 2 shows $50 \times 50\,h^{-1}$ Mpc temperature maps from our simulations at $z = 2$, 1, and 0. The left-hand panels show the full SIMBA run, while the right-hand panels show the No-jet run. The brightest regions represent $T \gtrsim 10^7$ K, and the darkest regions down to temperatures approaching a few times 10^3 K that is set by pure photoionization heating. These images are obtained by computing the mean temperature in each pixel on the $y - z$ plane through the middle of the simulation volume (i.e. at $x = 25\,h^{-1}$ Mpc), using YT's `slice` function. There is also an inset on the $z = 2$ SIMBA panel (top left), showing a $1\,h^{-1}$ Mpc $\times$ $1\,h^{-1}$ Mpc zoom on a massive $z \sim 2$ galaxy with a jet, showing the bipolar features of the feedback.

Large-scale filamentary structures are clearly visible in both simulations. These structures stand out as being somewhat hotter than the voids owing to the density–temperature relation in the diffuse photoionized IGM (Hui & Gnedin 1997). Around denser structures, there is additional shock heating caused by gravitational collapse on to filamentary structures, which raises temperatures to $T \gtrsim 10^5$ K. As the simulations evolve to lower redshifts, many of the smaller filamentary structures drain into the larger ones owing to the hierarchical growth of structure, and the IGM is generally cooler owing to its lower physical density and the lower $\Gamma_{\rm H\,I}$.

Comparing the left- and right-hand panels with and without jets, it can be seen that there is only slightly more heating at $z = 2$ for simulations with the jets included. In the jet run, individual bipolar jets are visible around the largest objects, as these generally have the largest black holes and hence low Eddington ratios that transition into jet mode (Thomas et al. 2019). The No-jet simulation also has some heating owing to gravitational shock heating in large haloes as well as weak feedback. In general, there are not large differences in the large-scale thermal structure at $z = 2$ with the inclusion of jets.

The differences become more drastic at lower redshifts. The No-jet simulation shows heating close to the filamentary structures owing to accretion shocks around large haloes, but this heating does not extend very far out. In contrast, the full SIMBA simulation including jets shows heating at $\gg$Mpc scales away from galaxies, which is consistent with the very high velocities at which these wind particles are ejected. For instance, an unimpeded 8000 km s^{-1} jet will travel $\approx$8 Mpc in a Gyr. While gravity and interactions with surrounding gas will retard this, it is still plausible that such jets will impact gas out to many Mpc over cosmic time (Borrow et al. 2019). At $z = 0$, many of the locations where the No-jet simulation has cold, diffuse IGM, SIMBA has very hot gas typically in the $T \sim 10^6$–10^7 K range. This clearly demonstrates that jet feedback in SIMBA can have widespread impact in the IGM.

3.2 Cosmic phase diagram

An illustrative global diagnostic for understanding IGM evolution is the cosmic phase diagram, that is, gas temperature versus density of all baryons. In phase space, gas broadly divides into four regimes (Davé et al. 2010): *condensed* gas that is cool and dense gas within galaxies and the circumgalactic medium, typically seen neutral and molecular gas; *hot halo* gas that has been shock heated typically to near the halo virial temperature, typically observable via X-ray emission; *diffuse* gas that is mostly photoionization heated in the IGM, which gives rise to the Lyα forest; and *WHIM* gas that has been shock heated to higher temperatures, and which hosts the so-called missing baryons (Dave et al. 2001).

Fig. 3 shows the $z = 0$ cosmic phase diagram for SIMBA (top panel) and the No-jet (bottom) simulations. The density has been scaled by the cosmic mean baryonic density. The black points show a randomly selected 0.1 per cent of the gas (to avoid saturation). Cyan points show gas that is currently star forming. Magenta points show gas elements that have recently been ejected in a galactic outflow, and are currently decoupled from hydrodynamics; note that the temperatures of these particles are arbitrary, as they do not currently experience pressure forces. Finally, the red points show gas elements that have been ejected by either radiative and/or jet AGN feedback at some point in their history.

We divide the phase diagram into four regions, demarcated by the horizontal and vertical dotted lines. The temperature cut is set at $T = 10^5$ K, which is a temperature that cannot be obtained without shock heating or feedback, and the traditional definition of the WHIM (Cen & Ostriker 1999). The density threshold follows Davé et al. (2010) as an estimate of a typical overdensity relative to Ω_m at the virial radius (based on Kitayama & Suto 1996), given by:

$$\delta_{\rm th} = 6\pi^2(1 + 0.4093(1/f_\Omega - 1)^{0.9052}) - 1, \quad (2)$$

where f_Ω is given by

$$f_\Omega = \frac{\Omega_m(1+z)^3}{\Omega_m(1+z)^3 + (1-\Omega_m-\Omega_\Lambda)(1+z)^2 + \Omega_\Lambda}. \quad (3)$$

At $z = 0$, this results in $\delta_{\rm th} \approx 105$. We list the mass fraction of baryons in each of these phases on Fig. 3, along with the baryon fraction in stars that is not included in any of these gas phases but tends to live in dense regions.

The overall phase diagrams in the two cases are generally similar. The condensed phase consists mostly of photoionized gas at $\sim 10^4$ K, along with dense gas forming stars that in SIMBA is forced to lie along a density–temperature relation that explicitly resolves the Jeans mass. The wind particles are artificially set to 10^3 K, but as they do not interact hydrodynamically, their temperature has no impact on their dynamics. The hot halo gas extends up to $T \gtrsim 10^7$ K and generally lies near the virial temperature of its host halo (e.g. Davé, Oppenheimer & Sivanand am 2008). The most massive halo in this box is somewhat anomalously large, giving rise to a distinct clump of high-T gas. The diffuse phase shows the tight density–temperature relation characteristic of photo-heated gas expanding with Hubble flow. Finally, the WHIM phase shows gas that has been shock heated by filamentary accretion as well as feedback processes.

The most notable difference between the SIMBA and No-jet runs is the large decrease in the baryon fraction in the diffuse phase, and a corresponding increase in the baryon fraction contained in the WHIM, when jet feedback is on. The WHIM increase mostly but not entirely comes from the diffuse phase; the baryon fraction of every other phase is at least halved in the jet simulation compared to the simulation without jets.

The No-jet simulation has baryon phase fractions that are broadly similar to the fiducial model at $z = 0$ in Davé et al. (2010), which had stellar feedback but did not have any AGN feedback. Hence non-jet AGN feedback has a fairly minimal impact on the cosmic phase diagram. We have confirmed this for SIMBA by examining the No-AGN simulation, which is not substantially different than No-jet.

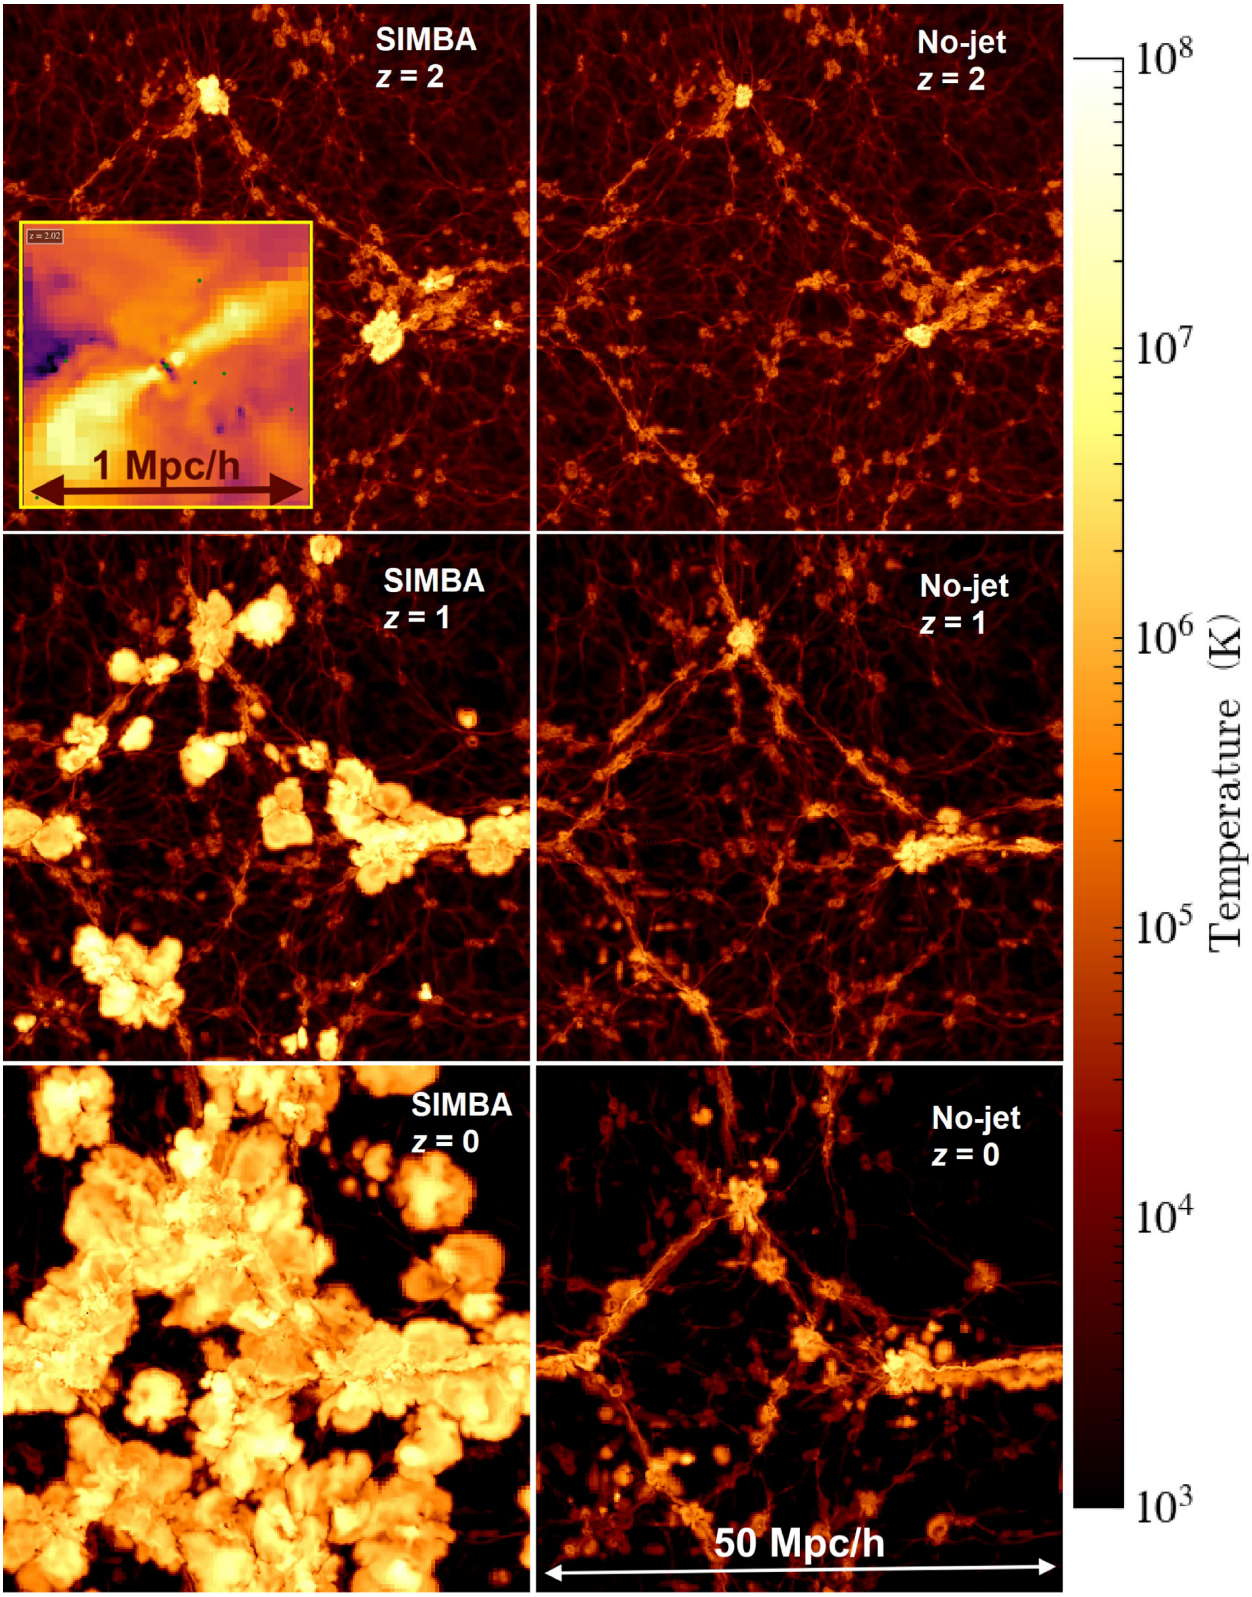

Figure 2. $50\,h^{-1}$ Mpc $\times$ $50\,h^{-1}$ Mpc temperature slices from SIMBA simulations with AGN jet feedback (left-hand column) and from the No-jet run (right-hand column). Top to bottom rows show $z = 2$, 1, and 0, respectively. The inset at $z = 2$ shows a $1\,h^{-1}$ Mpc $\times$ $1\,h^{-1}$ Mpc zoom on a massive $z \sim 2$ galaxy with a jet, showing sustained bipolar feedback. By $z = 0$, the jet feedback clearly has a dramatic effect on the temperature of the IGM by, with many Mpc-scale regions heated by jet energy.

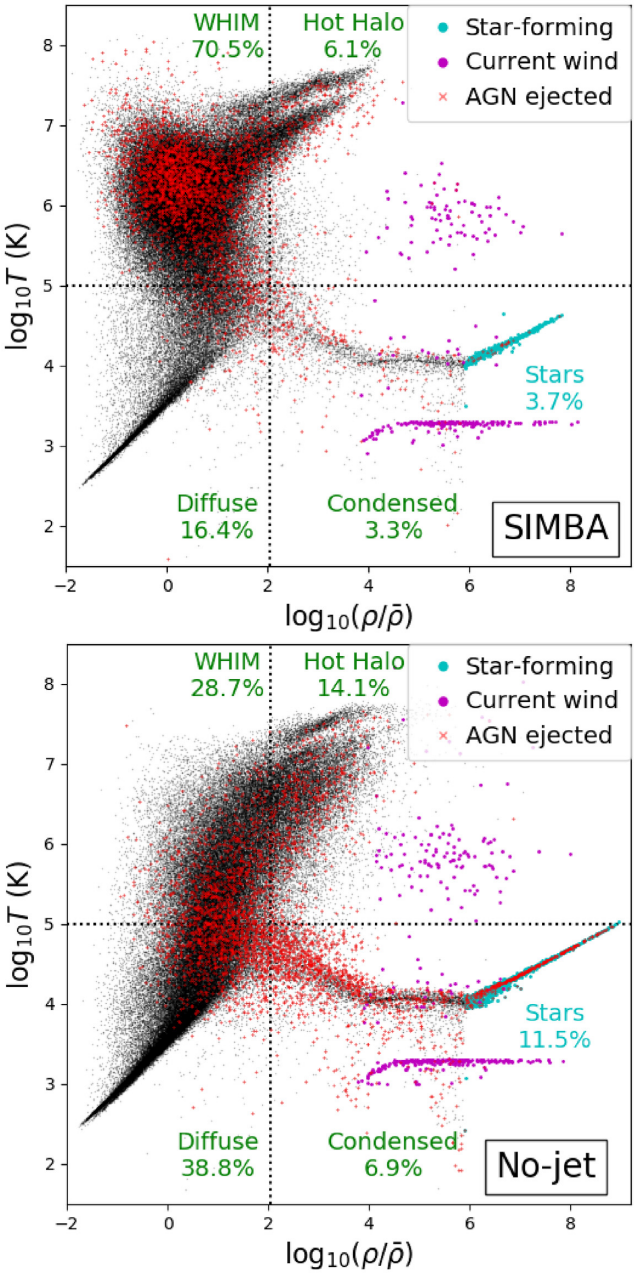

Figure 3. Phase diagrams at $z = 0$ for $50\,h^{-1}$ Mpc SIMBA simulations, for the full SIMBA run including jets (top panel) and for the No-jet run (bottom panel). A randomly selected 0.1 per cent of gas elements are shown for clarity, as black points. Red points are gas elements that have at some point been ejected via AGN feedback; this includes from non-jet (radiative mode) AGN feedback. Magenta points are elements which are currently in a decoupled wind, owing to star formation feedback. Cyan points show star-forming gas. The dotted lines indicate the boundaries between cosmic phases (cf. Fig. 6): The vertical division is the approximate density at the virial radius of dark matter haloes, while the horizontal division at $T = 10^5$ K separates cool from warm/hot phases. Percentages of baryons in each phase are indicated. AGN jet feedback results in AGN-ejected particles reaching much further into voids, while entraining diffuse gas, thus generating substantially more hot gas well outside of galaxy haloes and causing a strong reduction in the amount of cool diffuse IGM gas.

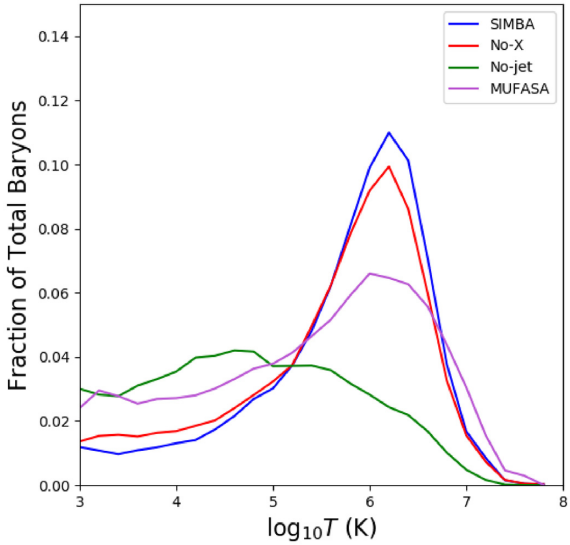

Figure 4. Temperature histograms at $z = 0$ of IGM gas ($\rho/\overline{\rho} < \delta_{\rm th}$; i.e. the WHIM and diffuse phases). Results are shown for $50\,h^{-1}$ Mpc simulations with various runs: the main SIMBA simulation (blue line), the No-jet run (green line), the No-X run (red line), and the MUFASA simulation (purple line). Including jets (either in SIMBA or No-X) strongly shifts the distribution of IGM gas temperatures, producing a peak at $T \sim 10^{6.2}$ K.

Fig. 3 also indicates which gas elements have been ejected by AGN feedback, as red points. In No-jet, we still have radiative AGN feedback up to $\sim$1000 km s^{-1}, which distributes some gas into the diffuse and WHIM phase. However, it does not strongly change the phase of a significant amount of ambient gas; much of it stays at relatively cool temperatures.

In the full SIMBA run with jets, elements touched by AGN feedback can reach well into the diffuse region. In doing so they create a new feature in the cosmic phase diagram at $T \sim 10^6$–7 K near the cosmic mean density, that is not present in the No-jet run. This region is actually populated mostly by particles that have not been directly kicked by jet feedback, but rather have been entrained (and heated) by jet-ejected gas (Borrow et al. 2019). Also, in this simulation, very few particles that are ejected by AGN feedback end up in the condensed star-forming gas phases, unlike in the No-jet case. The reason is that the AGN-touched particles are significantly hotter, so do not have a chance to fall back in to bound systems. This is an important factor for suppressing star formation in massive galaxies having jet feedback, and is a key preventive feedback mechanism that keeps galaxies quenched.

Fig. 4 quantifies the increase in temperature in unbound gas. It shows histograms of the baryon fraction for low-density phases (i.e. the WHIM and diffuse phases), binned in temperature, for various models. The most distinct feature is that the SIMBA runs with AGN jet feedback enabled (SIMBA and No-X) have a large peak in their diffuse baryon fractions at $T \sim 10^{6.2}$ K. This shows that jet feedback strongly increases the overall temperature distribution in WHIM gas, compared to the No-jet run (green). The MUFASA simulations also produce a peak in approximately the same location, but not as sharply; we thus expect that the MUFASA simulation will show results intermediate between the No-jet and jet runs.

When looking at Fig. 3, remember that the diffuse phase gives rise to Lyα absorption; the WHIM is too highly ionized for any H I absorption to occur. This means that a decrease in the diffuse fraction will correspond to a decrease in Lyα absorption. It is therefore clear

that jet feedback will have a significant impact on the amount of H I absorption. This is the primary manner by which AGN jet feedback impacts the Lyα forest. The extra WHIM gas could potentially generate more high ionization metal absorption, such as O VI, O VII, and O VIII. Note however that O VI absorption may not be strongly impacted since O VI absorption is best at tracing the range $T \approx 10^5$–$10^{5.7}$ K, while Fig. 4 shows that most of the jet-heated gas is hotter. Thus O VII which is strong in $T \approx 10^{5.7}$–$10^{6.3}$ K gas (Nicastro et al. 2018) may be a better tracer (e.g. Chen et al. 2003). Ne VIII could also provide a useful tracer of $T \sim 10^6$ K WHIM gas in the extreme UV (Burchett et al. 2019).

Finally, Fig. 5 explores how the distribution of pixel flux decrements changes with and without jet feedback. These show 2D histograms of optical depth versus overdensity, for SIMBA (including jets) in the top panels, and the No-jet case in the bottom panels. Left- and right-hand panels show the results for unweighted pixels and flux decrement-weighted pixels, respectively.

Looking at the unweighted case (left-hand panels), the vast majority of pixels have low flux decrement, reflecting the sparseness of the Lyα forest at low-z. The volume-averaged mean overdensity is ∼0.3. Comparing SIMBA (top) versus No-jet (bottom), we see the emergence of a distinct new cloud of points at moderate overdensities, but low optical depths. These are pixels where jet feedback has heated the gas such that its Lyα flux has been substantially lowered from where it was in the No-jet case. This demonstrates that jet feedback is particularly impacting pixels in the moderate overdensity regime ($\rho/\bar{\rho} \sim 0.5 - 5$). This is consistent with the heating of near mean-density gas in the phase diagrams by AGN jets, as seen in Fig. 3, and will have an impact on the total absorption as discussed in Section 4.

We can see why the impact of jets in this overdensity regime strongly impacts the flux decrements by looking at the right-hand panels of Fig. 5, which shows the result of weighting each pixel by its flux decrement. These panels illustrate where the majority of the flux decrement is coming from in overdensity. Unsurprisingly, the absorption is dominated by flux near $\tau \sim 1$, since at higher τ values the absorption enters into the logarithmic part of the curve of growth where additional optical depth adds little flux decrement. In overdensity, $\tau \sim 1$ regions occur at mild overdensities of $\rho/\bar{\rho} \sim$ few, which from the left plots is exactly the regime that is being impacted by jets. There is also a clear trend of increasing absorption with overdensity, broadly consistent with previous results (e.g. Davé et al. 2010). When including jets, there is a minor but noticeable shift in the peak overdensity, in which the absorption occur over a somewhat broader range in overdensity, shifted to modestly higher values. There is also an increased scatter at overdensities where jet feedback is impacting the pixel optical depths as shown in the left panels. It is not immediately clear how these trends could be tested observationally, but they may be important in interpreting the underlying cosmic densities traced by low-z Lyα absorbers. Overall, this shows that the impact of jets is to heat gas at precisely the range of moderate overdensities that dominate the overall flux decrement, and thus demonstrates why (as we will show later) jets have a significant impact on the flux decrements in the Lyα forest.

3.3 Baryonic phase evolution

Jet feedback clearly has a large impact on the cosmic phase of baryons at $z = 0$. At very high redshifts before jet feedback begins, it should obviously have no impact. The question is then, when do the SIMBA and No-jet diverge in terms of their baryon fractions in the various phases?

Fig. 6 shows the evolution from $z = 3 \rightarrow 0$ of the baryon fraction in each phase as defined in Fig. 3: green is WHIM, cyan is condensed, blue is diffuse, red is hot halo, and magenta is stars. The dashed lines show the predictions for the No-jet simulation, and the solid lines show the results from the SIMBA simulation with jets.

The simulations both with and without jets have identical baryon fractions in each phase at $z \sim 3$, since there are essentially no massive black holes with jets yet at these early epochs. The evolutionary tracks begin to diverge shortly thereafter, with the jet simulation showing more WHIM gas and less in every other phase. By $z = 0$, the jet simulation has almost 2.5 times as many baryons in the WHIM as the simulation without jets, and a corresponding reduction in the diffuse phase. At $z \lesssim 1$, the WHIM phase dominates the baryon fraction in SIMBA, which never happens in the No-jet case.

The late onset of these differences is to be expected, as the jet feedback in SIMBA only activates for black holes with masses $M_{\rm BH} \geq 10^{7.5}$ M$_\odot$ with low Eddington ratios, and black holes in SIMBA only reach the required typical sizes at late epochs (see Thomas et al. 2019). The No-jet case broadly reproduces the same evolution of the baryon fractions as the fiducial model used in Davé et al. (2010), which did not include any AGN feedback.

The SIMBA results with jets show a significantly higher fraction of baryons in the WHIM than previous simulations (Dave et al. 2001). These predicted fractions are also at the high end of current inferences from observations of O VII absorbers at $z \sim 0.4$, which suggest baryon fractions 20–60 per cent (Nicastro et al. 2018) in IGM gas with $T = 10^{5-7}$ K (see their table 1). Our predicted value from the jet simulation is at the top end of this, while from no-jets it is at the bottom end. We will examine predictions for high-ionization metal lines from SIMBA in future work, which could be a key discriminant between these types of models with future X-ray missions such as *Athena* and *Lynx*.

4 THE PUC: MEAN FLUX DECREMENT EVOLUTION

Armed with an understanding of the physical properties of the IGM, we now examine how AGN feedback impacts H I absorption in the IGM, and thereby investigate the PUC. To study this, we will use the metric of D_A, the mean flux decrement in the Lyα forest. This avoids the uncertain and non-unique process associated with line identification and fitting, which can depend fairly sensitively on S/N, spectral resolution, and other specific aspects that would need to be more closely reproduced in the mock spectra when comparing to observations, and impart greater uncertainties. For our purposes, D_A provides a robust and well-defined measure that accurately quantifies the PUC.

Fig. 7 encapsulates our main results. Here, we show D_A as a function of redshift in the top panels, and the inferred $F_{\rm UVB}$ versus redshift in the bottom panels. In the left-hand panels, D_A and $F_{\rm UVB}$ have been computed from spectra assuming a Haardt & Madau (2012) UVB (henceforth referred to as HM12), while in the right-hand panels a Faucher-Giguère (2019) UVB (henceforth referred to as FG19) has been assumed. We choose these two backgrounds since the former is the one in which the PUC was originally found, and the latter is a recent state of the art UVB model. The dashed green line represents values measured from spectra from the full SIMBA simulation with jets, and the dashed blue line represents the No-jet results. Dotted lines indicate uncertainty due to cosmic variance, which was estimated by splitting the spectra into four quadrants based on their LOS down the simulation box, and computing the standard deviation on the value of D_A found in each of the four quadrants.

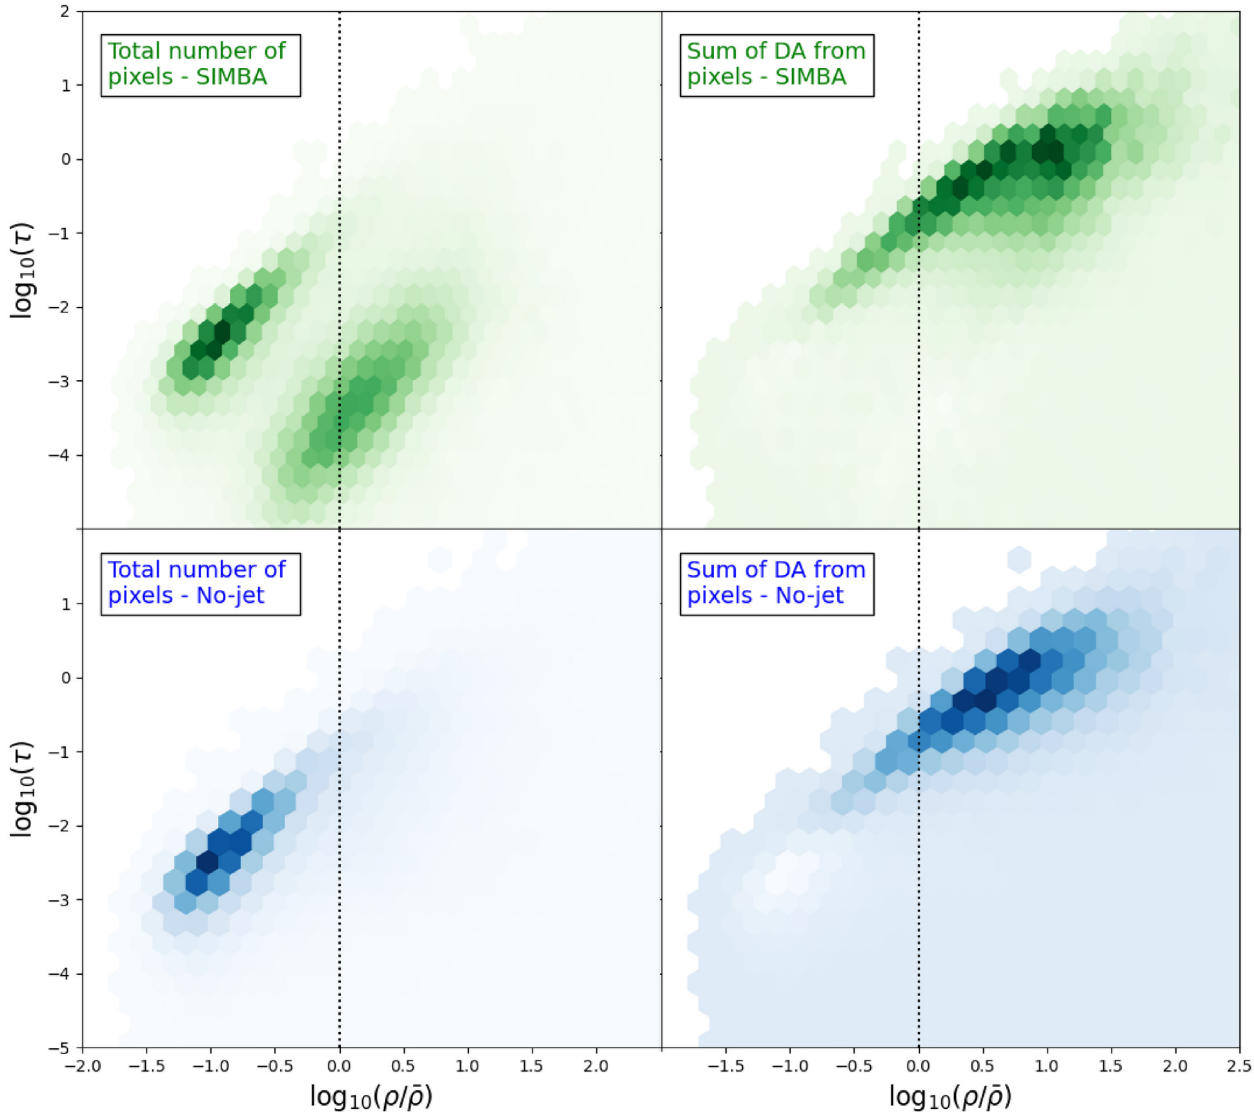

Figure 5. Optical depth τ versus gas overdensity $\rho/\overline{\rho}$ at $z = 0$ for SIMBA (upper panels) and the No-jet Model (lower panels), showing aggregate properties of all the pixels in a sample of 1000 spectra. Left-hand panels: the total number of pixels with given values of $\rho/\overline{\rho}$ and τ. A large concentration of pixels can be found in the same region for both SIMBA and No-jet, which lies along the optical depth–overdensity relation set by photoionization as canonically seen (e.g. Davé et al. 1999). But jet feedback in SIMBA has taken a substantial amount of the absorption that occurs around the mean density in the No-jet case, and heated the gas to yield low optical depths. Right-hand panels: the sum of D_A from all pixels, at given values of $\rho/\overline{\rho}$ and τ. As expected, in both cases pixels from overdense gas ($\rho/\overline{\rho} > 1$) are the primary contributors to absorption; however, there is a slight change in the distribution of these overdense pixels when jets are included.

This cosmic variance uncertainty is typically ∼10 per cent for the full SIMBA results, and ∼20 per cent for the No-jet results. The effect of cosmic variance appears to be somewhat larger in the No-jet case, which may owe to the fact that without jets, Lyα absorbing gas is present in highly overdense regions where the variance in absorption is higher, whereas jet feedback removes this. The estimated effect of cosmic variance is in all cases greater than the statistical uncertainty on D_A, which is $\lesssim$1.0 per centfor all samples.

In the top panels, the black data points show the Danforth et al. (2016) measurements from *HST*/COS data, while the brown data points show the compilation from Kirkman et al. (2007, from their table 5 covering $z \approx 0$–3). To get the observed evolution of D_A, we fit a single power law to the combined data sets from $z = 0$–3:

$$\log D_A = -1.848 + 1.982 \log (1 + z) \quad (4)$$

We show this as the black dashed line, with a shaded variance computed from the deviations to the individual data points (σ = 0.136 dex). We will use this line as our baseline observations for comparison to our predicted D_A values.

The bottom panels of Fig. 7 show $F_{\rm UVB}$ for the SIMBA simulation variants as a function of redshift. The calculation of $F_{\rm UVB}$ is described in Section 2.4, and can be regarded as the 'photon underproduction factor', by which $\Gamma_{\rm HI}$ must be adjusted for simulations to match the observed value of D_A. As with the top panels, the bottom left panel shows the results when using the HM12 background, and the bottom right panel shows the results when using the FG19 background. As $F_{\rm UVB}$ is a rather theoretical quantity, it has been calculated directly from the optical depths, without any continuum fitting being performed.

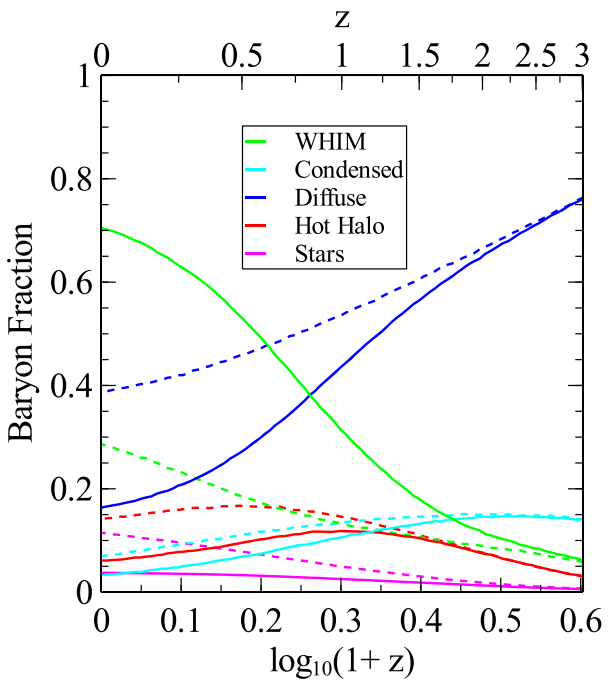

Figure 6. Baryon fraction evolution from $z = 3 \rightarrow 0$ in the various phases shown in Fig. 3, for the $50\,h^{-1}$ Mpc SIMBA simulations. The solid lines are from the full SIMBA simulation with jet feedback turned on; the dashed lines are from the No-jet run. The impact of jets appears at $z \lesssim 2$ and becomes more evident as redshift decreases, with increases in the WHIM phase, and decreases in the diffuse, condensed, hot halo gas, and stellar phases.

Fig. 7 (top panels) clearly illustrates the PUC. The No-jet simulation (blue line) shows significantly higher absorption than *HST* observations, more so for HM12. Meanwhile, the absorption in SIMBA simulation with jets is significantly closer to matching the *HST* data at low redshifts ($z \lesssim 0.5$), though the HM12 case is still mildly discrepant. This illustrates our primary result, that including jet feedback and employing a modern determination of the UVB from FG19 essentially solves the PUC in SIMBA, and allows consistency between source count determined UVB estimates and the estimate obtained from the Lyα forest.

As we have shown that the jets are a source of additional heating, and heating should reduce the amount of Lyα absorption, the reduction in D_A with jets is expected. The discrepancy between jet and no-jet results is also expected to be greater going towards lower redshifts, as this is when the jets have had more time to affect the IGM gas in the simulation. At $z \gtrsim 1$, the jet and No-jet simulations do not show strong differences in D_A for either HM12 or FG19, which is expected because there are only minor differences in the diffuse baryon fraction above this redshift (cf. Fig. 6). Thus the PUC is only present at $z \lesssim 1$, and increases strongly to lower redshift.

The predicted D_A generally follows a power-law slope in $(1 + z)$, but that slope is different depending on whether the HM12 or FG19 UVB is adopted. The slopes when using the FG19 background match the slope of the Danforth et al. (2016) data better than they match the Kirkman et al. (2007) data, while the converse is the case when using the HM12 background. The contrast between the two backgrounds is particularly stark at $z \sim 0$. Fitting a power law with $D_A \propto (1 + z)^{\alpha}$, we obtain slopes of $\alpha = [1.5, 0.9]$ for the jet and No-jet cases respectively for HM12, and $\alpha = [1.8, 1.3]$ for FG19. This can be compared to the Danforth et al. (2016) slope of $\alpha = 2.2 \pm 0.2$, showing that the full SIMBA case with FG19 produces a $D_A(1 + z)$ slope in very good agreement with observations, and as a result a non-evolving $F_{\rm UVB}$. At higher redshifts ($z \gtrsim 1$), the predicted values of D_A match fairly well with expectations from either HM12 or FG19, which nicely demonstrates that prior to the impact of AGN jet feedback, the UVB amplitude determined from source count modelling is in good agreement with that inferred from the Lyα forest.

To more precisely quantify this excess of absorption shown in the top panels, we show $F_{\rm UVB}$ as a function of redshift in the bottom panels of Fig. 7. For the No-jet case and the HM12 UVB, the photon underproduction factor reaches $\sim$6 at $z = 0$, and is already $\sim$3 at $z = 0.5$. This confirms the PUC found by Kollmeier et al. (2014) in the case with no AGN feedback and HM12. In fact, even though the No-jet run has some AGN feedback, the underproduction factor is higher compared to the 5 times discrepancy found by Kollmeier et al. (2014). This may owe to the lower star formation-driven wind speeds in SIMBA relative to the Davé et al. (2013) simulations used in Kollmeier et al. (2014), and/or the use of MFM rather than smoothed particle hydrodynamics for the hydrodynamics. In any case, the overall results are very similar, and confirm that the PUC is present in state of the art simulations when no feedback is included that heats the IGM.

With jets on, the green line shows that the PUC is not completely eradicated – at $z = 0$, with HM12, the photon underproduction factor is still 2.5 (lower left panel). However, this is clearly much closer to unity, which would be the value if the predicted D_A exactly matched the Danforth et al. (2016) measurements. Given that there are $\sim$2 times uncertainties in the source count modelling determinations of $\Gamma_{\rm HI}$ (Khaire & Srianand 2015), such a discrepancy may not be considered severe.

Looking at the lower right panel which assumed FG19 instead of HM12, the PUC is essentially gone. The No-jet case still has a factor of 3 discrepancy in $F_{\rm UVB}$, while the jet simulations reduces this to $\sim$1.2, which is now likely well within current uncertainties. Interestingly, the evolution of $D_A(z)$ predicted in the SIMBA simulation is in very good agreement when assuming FG19, but with HM12 we predict a fairly strongly increasing PUC to lower redshifts. Thus, the SIMBA simulations with jet feedback and using the FG19 background are in quite good agreement with the low-redshift Lyα forest data from Danforth et al. (2016).

It is interesting to note that $F_{\rm UVB}$ is actually somewhat larger than the discrepancy in D_A from the top panel. For instance, at $z = 0$ for the No-jet case and HM12 background, the ratio of the predicted D_A (blue line) and the Danforth et al. (2016) value is about a factor of 3. However, when one goes through the exercise of iteratively adjusting the ionizing background to match D_A, this indicates that a factor of 6 is needed to match the observations. The reason is that saturated lines provide a subdominant but non-negligible contribution to D_A. Saturated lines move into the logarithmic portion of the curve of growth, so their flux decrement no longer scales linearly with optical depth and $\Gamma_{\rm HI}^{-1}$. Hence, it is important to do the exercise of iteratively fitting to the observed D_A as we have done, since the PUC is actually worse that it appears simply by examining the discrepancy in D_A.

The impact on D_A at lower redshifts owes not only to the increasing filling factor of hot gas as evident from Fig. 2, but also to the fact that the largest contribution to D_A comes from marginally saturated lines ($N_{\rm H\,I} \approx 10^{13.5\text{–}14}$ cm^{-2}), since below this the column density distribution has a slope shallower than -2 (Danforth et al. 2016), and above this the increase in absorbers' column densities no longer

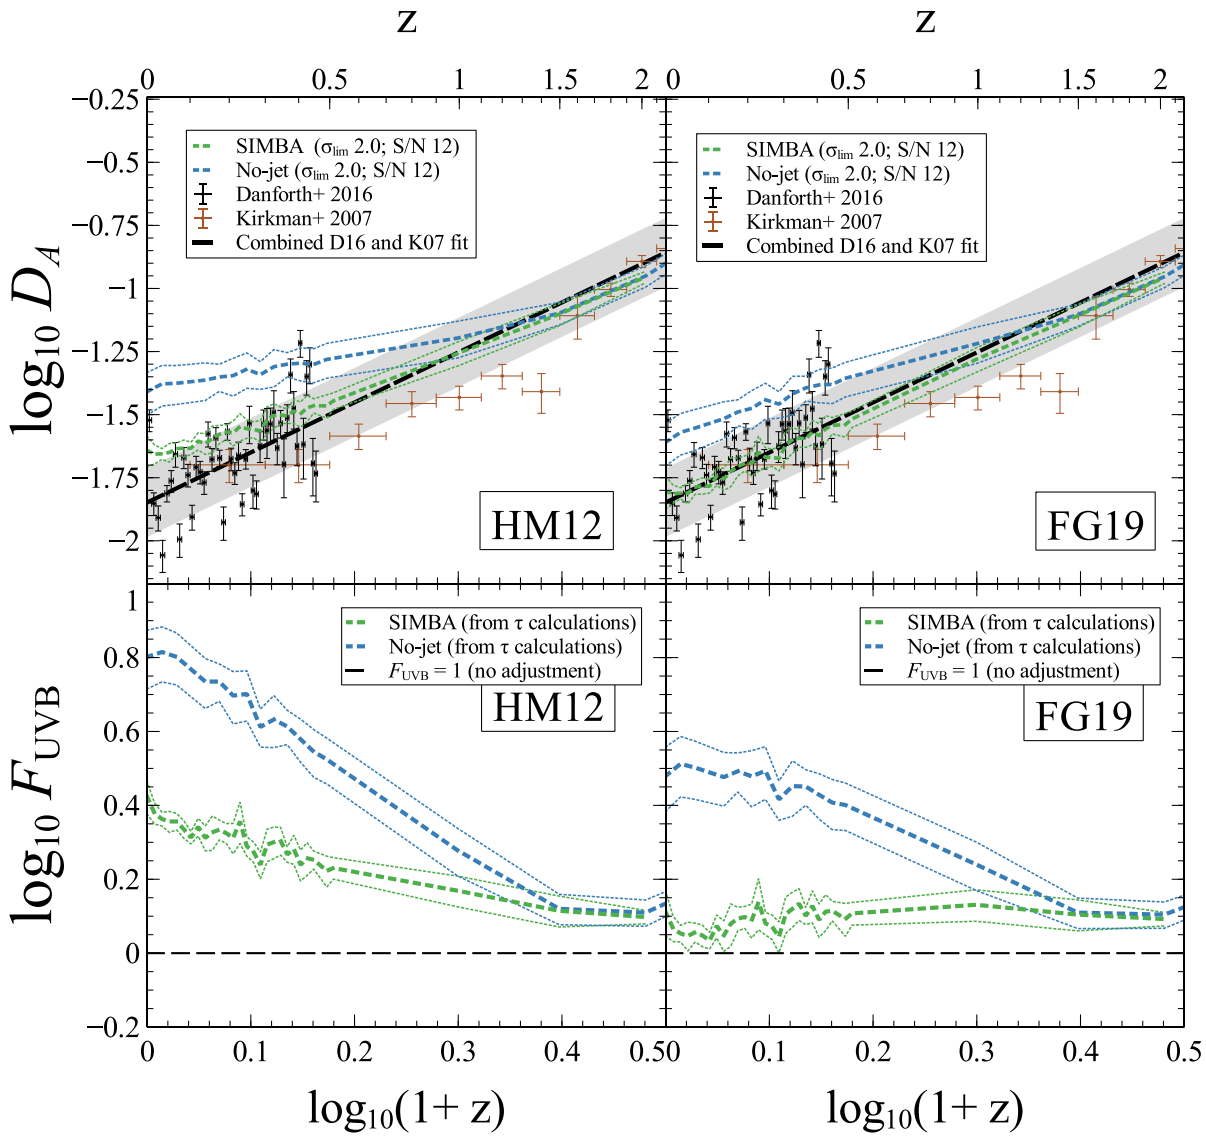

Figure 7. Top panels: D_A versus redshift for simulations versus observations, with the left-hand panels showing the results when simulated spectra are generated using the HM12 background, while right-hand panels show the results when assuming an FG19 background. The dashed green line is from the full SIMBA simulation with jets, while the dashed blue line is from the No-jet run; dotted lines indicate estimates of the uncertainty due to cosmic variance. Black points with error bars are the binned observations from COS data by Danforth et al. (2016) , and brown points with error bars show the observational results from Kirkman et al. (2007) from Faint Objects Spectrograph data, as well as data from Kast and Keck High Resolution Sepctrometer (HIRES). The dashed black line shows a combined fit of both observational data sets, with the grey shaded region indicating the uncertainty on the fit. Bottom panels: photon underproduction factor $F_{\rm UVB}$, that is, the factor by which $\Gamma_{\rm H\,I}$ must be multiplied in order for simulated predictions of D_A (calculated directly from the optical depths, foregoing the continuum fitting process done for D_A in the top panels) to match D_A given by the fit to observational data shown in the top panels. The dashed black line shows the value which indicates no adjustment ($F_{\rm UVB} = 1$). It can clearly be seen in the top panels that the SIMBA simulations including jets are much closer to matching observed values of D_A than the No-jet runs, regardless of the background used. The FG19 background provides a closer match to observation than the HM12 background. By $z = 0$ and using the HM12 background, $F_{\rm UVB} \approx 1.5 - 2.5$ for SIMBA, while $F_{\rm UVB} \approx 4 - 6$ for the No-jet run, showing that jets strongly mitigate the PUC.

contribute linearly to D_A. At $z \sim 2$–3, marginally saturated lines correspond to gas at moderate overdensities of a few, but by $z = 0$, these lines arise in diffuse gas of overdensities of ∼20–50 (Davé et al. 1999). As a result, they move into the regions nearer to galaxies that are most dramatically impacted by the jet heating. This exacerbates the effect on $D_A(z)$.

As mentioned in Section 3, the diffuse phase of matter at low densities and temperatures is responsible for Lyα absorption. The 2.5 times reduction in D_A when jets are turned on as seen in Fig. 7 is consistent with the 2.5 times reduction in the fraction of baryons in the diffuse Lyα-absorbing phase, as seen in Figs 3 and 6. In light of this, a straightforward physical interpretation of the impact of AGN feedback on the low-redshift IGM is that it serves to heat a sufficient fraction of diffuse gas into the WHIM phase in order to provide a potential resolution to the PUC.

5 FLUX PROBABILITY DISTRIBUTION FUNCTION

While D_A measures the mean absorption, the distribution of pixel fluxes provides a more detailed test of whether SIMBA yields an accurate description of the low-redshift Lyα forest (Gaikwad et al. 2017b). Therefore, in this section, we examine the flux probability distribution function (FPDF), which is the density histogram of normalized flux values from a set of quasar spectra. The FPDF tests whether the *distribution* of pixel flux decrements is in accord with observations, which is a higher order constraint as compared to D_A.

To compare to observations, we download the reduced quasar spectra obtained by the COS GTO team (Danforth et al. 2016) from the Mikulski Archive for Space Telescopes.[3] We mask out pixels that have significant foreground contamination, identified as having $>2\sigma$ absorption in the composite foreground spectrum. For the remaining pixels, we collate the fluxes for all the pixels with rest-frame Lyα wavelengths of >1040 Å (to avoid Lyβ and O VI absorption), are not within 5000 km s^{-1} of the quasar systemic redshift, and are >5000 km s^{-1} redwards of $z = 0$ in order to avoid Galactic absorption. For these fluxes, we subtract off any remaining foreground absorption, and normalize the flux using the continuum provided by the COS GTO team. We then compute the FPDF from this continuum-normalized flux, in two bins of $\Delta\log(1+z) = 0.0385$, which correspond to $z = 0$–0.093 and 0.194–0.305. We have checked that the intermediate-redshift bin gives very similar answers, and while the sample contains some higher z data, it is fairly sparse and noisy so does not give useful constraints. The mean redshifts for the pixels in these bins are $\bar{z} = 0.06$ and 0.26.

Fig. 8 shows the FPDF (i.e. dP/dF where P is the probability density of a pixel having a flux F) in SIMBA (green line) and No-jet (blue), versus the COS GTO data (cyan) from Danforth et al. (2016). Errors on the COS GTO data are computed by propagating the errors from the individual pixels and continuum level, but do not include cosmic variance, so may be underestimated. In the simulations, we compute the FPDF at the snapshot whose redshift is closest to the mean redshift of COS GTO pixels in that bin, as indicated in the legend. For clarity, we only show the FG19 results for the simulations, to focus on how jet feedback impacts the FPDF.

SIMBA provides a very good match to the shape and amplitude of the observed FPDF. The only significant deviation is seen for near-saturated fluxes ($\lesssim 0.1$) in the lower z bin, but given that such saturated absorption is rare, it is likely that cosmic variance dominates the uncertainty here. Otherwise, over most of the fluxes, SIMBA is a much better match to the observations than the No-jet case, which yields an FPDF that is $\sim$2 times higher for fluxes with significant absorption.

This demonstrates that SIMBA's jet feedback not only suppresses the mean absorption in accord with observations, but it does so in a way that further yields a pixel flux distribution in good agreement with observations. This confirms that SIMBA's jet feedback, together with the FG19 background, provides a much better representation of the low-z Lyα forest as compared to a simulation that does not include such widespread AGN feedback.

6 AGN FEEDBACK VARIANTS

In the previous section, we focused on comparing the full SIMBA simulation with the No-jet run, because these provide the greatest differences illustrating the impact of AGN feedback. In this section, we further consider two additional model variants, to gain insights into how well these PUC measurements might be able to discriminate between AGN feedback models. In the No-X case, we have left jets on but turned off X-ray AGN feedback; if jets are the dominant mechanism impacting the IGM, we expect this model to be similar to the full SIMBA run, as opposed to the No-jet run which turns off both jet and X-ray feedback. We will also consider MUFASA, which used a completely different method for quenching galaxies in which hot gas in haloes above an (evolving) mass threshold was prevented to cool (Davé et al. 2016). The No-X model produces mostly quenched galaxies but with insufficiently low specific star formation rates compared to observations (Davé et al. 2019), while MUFASA produces a quenched population in very good agreement with observations (Davé et al. 2017), in some ways even better than SIMBA, but it uses a less physical approach that does not directly model black holes. Here, we examine $D_A(z)$ and $F_{\rm UVB}(z)$ in these two variants.

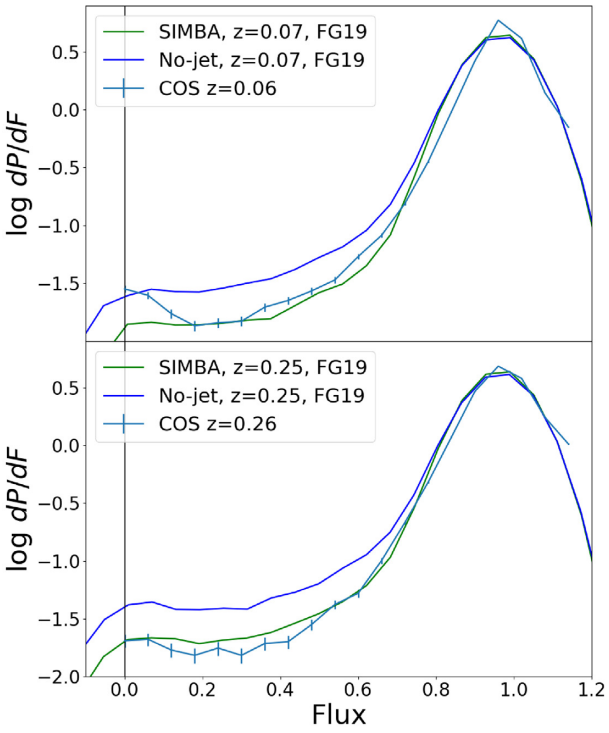

Figure 8. The FPDF in SIMBA (green) versus No-jet (blue), compared to observations from the COS GTO team (cyan; Danforth et al. 2016), compiled as described in the text. SIMBA provides a very good match to the observed FPDF, much better than the No-jet case, demonstrating that SIMBA's jet feedback suppresses absorption within flux bins in accord with data.

Fig. 9 shows D_A (top panel) and $F_{\rm UVB}$ (bottom) as a function of $\log(1+z)$, as in Fig. 7. Here we focus on just the FG19 background, as this one is overall more successful for SIMBA. We show the results from No-X and MUFASA as the red and purple lines, respectively. For comparison, we continue to show the SIMBA and No-jet lines in green and blue, respectively. The observations are also shown as presented in Fig. 7. For clarity, the uncertainties due to cosmic variance are omitted from the graph, but are typically $\sim$10 per cent for the No-X case (similar to the full SIMBA run) and up to $\sim$30 per cent for the MUFASA results (somewhat higher than the No-jet case). It is not immediately evident why MUFASA would exhibit such large cosmic

[3] https://archive.stsci.edu/prepds/igm/

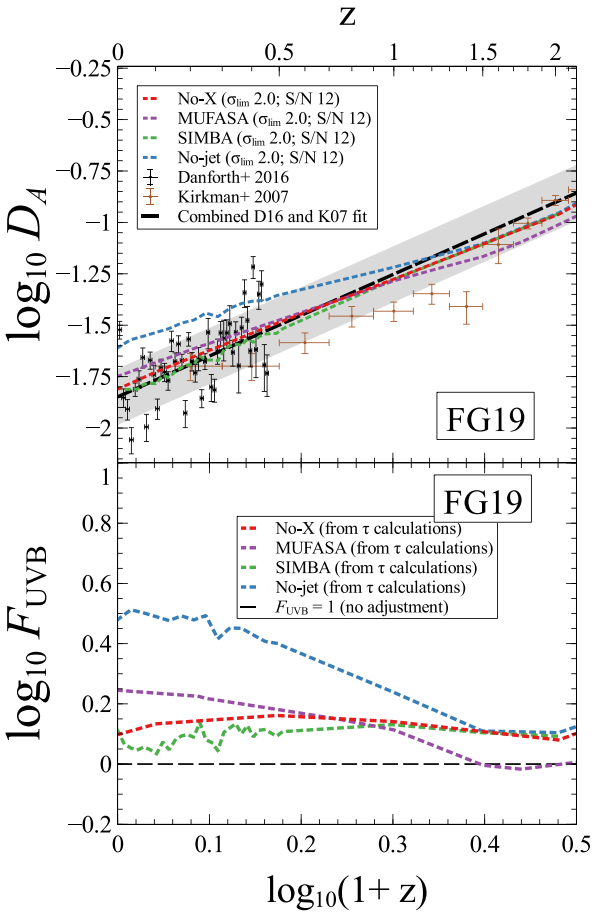

Figure 9. D_A (top panel) and $F_{\rm UVB}$ (bottom) as a function of redshift for various SIMBA simulations using the FG19 background, similar to Fig. 7. The red line shows results from the No-X simulation with jets but without X-ray feedback, and the purple line shows results from the MUFASA simulation. Green and blue lines are reproduced from Fig. 7 showing the SIMBA and No-jet runs for comparison. Observations are also reproduced from Fig. 7, as indicated. The No-X simulation is quite similar to the SIMBA run with X-ray feedback, showing that X-ray feedback has negligible impact on the diffuse IGM, and thus the impact comes from the jets. MUFASA matches observational data better than the No-jet case, but not as well as with jets, indicating that these observations could potentially discriminate between otherwise successful AGN feedback models.

variance, none the less this is still relatively small compared to the values needed to solve the PUC.

Fig. 9 demonstrates that X-ray feedback has a negligible impact on Lyα absorption in the IGM; the values in red and green are nearly overlapping at all redshifts, though the D_A values are very slightly higher than SIMBA without the additional feedback from X-rays. This is expected, as the X-ray feedback primarily acts within the inner disc of galaxies close to the AGN, and thus is not expected to directly impact IGM gas. This conclusively demonstrates that it is in particular the AGN jet feedback that is responsible for lowering $F_{\rm UVB}$ in SIMBA.

For MUFASA, it is interesting to note that there is still a substantial reduction in $F_{\rm UVB}$, moving D_A closer to the observed values, though not as strongly as in SIMBA. This was anticipated from Fig. 4, which showed that MUFASA generates a substantial shift in the IGM temperature distribution from that expected with no or weak AGN feedback. This is somewhat surprising because the direct impact of the feedback is confined to halo gas (by adding heat to offset cooling), yet it appears to have a wider impact on IGM gas. None the less, by $z = 0$, the photon underproduction factor is still $\approx$2, so significantly higher than in SIMBA, though well lower than in the No-jet case. We do not show the HM12 results here, but the corresponding factor for MUFASA in this case is $\approx$4. Hence one might envision, with improved measurements of $\Gamma_{\rm H\,I}$ in the local universe such as from flourescence (Fumagalli et al. 2017), it may be possible to discriminate between variants of AGN feedback based on their impact on the diffuse IGM.

7 MODELLING UNCERTAINTIES

As can be seen in the description of spectra generation in Section 2.4, there are many factors which may in principle have an influence on the calculation of D_A. These include: the continuum fitting process, and the choice of the free parameter involved; the application of an LSF to the flux to imitate the effect of light passing through an instrument; the S/N of the noise added to the spectra; the boxs ize of the simulation; and the mass element resolution of the simulation, among others.

Fig. 10 shows D_A for variations in different steps of the spectra generation process. Each of the four lines labelled 'SIMBA' is from the primary SIMBA simulation including jets. The SIMBA line labelled 'τ calculations' shows the value of D_A when calculated directly from the optical depths, with no effect from the LSF, added noise, or continuum fitting process. The other three SIMBA lines all involve the calculation of D_A using flux that has undergone the application of the COS LSF, and a continuum fitting process; they vary in either the S/N, or in the value of $\sigma_{\rm lim}$ used in continuum fitting. Specifically, pixels which are $\sigma_{\rm lim}$/S/N below the median pixel flux are removed during each iteration of the continuum fitting process. The application of the COS instrument LSF was found not to make any discernible difference in the overall D_A of the sample, and thus we only show values calculated with the LSF applied in Fig. 10.

It can be seen from Fig. 10 that using a value of $\sigma_{\rm lim} = 2.0$ produces a significantly lower value of D_A compared to using $\sigma_{\rm lim} = 1.5$, regardless of the S/N used. A higher S/N also lowers D_A, but not as drastically as the effect of $\sigma_{\rm lim}$. This trend with $\sigma_{\rm lim}$ is to be expected, as using a larger value of $\sigma_{\rm lim}$ will result in a lower continuum level. The proper value to set $\sigma_{\rm lim}$ can be difficult to determine, hence why we provide this comparison. One way to set this is to see which value results in the closest match to the true value (i.e. 'tau calculations'). At $z \lesssim 0.5$, we find that $\sigma_{\rm lim} = 2.0$ matches reasonably well. This is also a typical value taken in observations. However, Danforth et al. (2016) actually used $\sigma_{\rm lim} = 1.5$, which would tend to raise D_A above the true value. Note that they were fitting to an intrinsically non-uniform continuum and generally over longer stretches than our spectra using a higher order fitter, hence their choice of $\sigma_{\rm lim}$ is not directly comparable to ours, but it is illustrative of the changes that can occur owing to this choice. In our procedure with SIMBA, such a choice clearly overfits the continuum at $z = 0$.

While we have shown the full SIMBA run here, we have checked that the same trends generally hold in the No-jet simulation. This means had we used $\sigma_{\rm lim} = 1.5$, the PUC effect noted in the No-jet simulation would be even more severe than we presented. It would also mean that even including SIMBA's jet feedback would be unable to fully solve the PUC.

Fig. 11 shows D_A for the primary SIMBA simulation used in this report (labelled 'SIMBA', which as noted previously has a box size of $50\,h^{-1}$ Mpc and 512^3 mass elements), compared with two variants

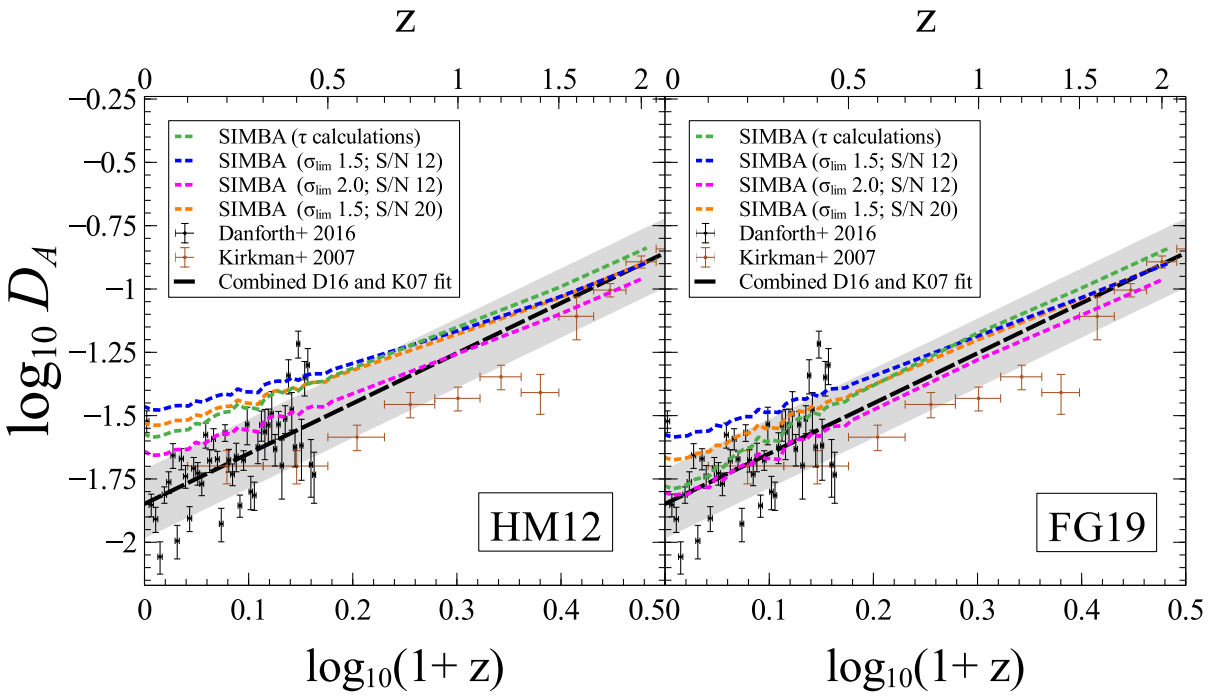

Figure 10. Graph showing the changes in D_A from spectra in SIMBA resulting from changes in parameters used in the continuum fitting process, assuming the HM12 background (left) and FG19 background (right). The dashed green line shows the SIMBA results when calculated directly from the optical depths, with no continuum fitting process being done. The dashed blue, magenta, and orange lines show the SIMBA results when a continuum fitting process is performed on the spectra, using a S/N of 12 and a value of $\sigma_{\rm lim} = 1.5$, S/N = 12 and $\sigma_{\rm lim} = 2.0$, and S/N = 20 and $\sigma_{\rm lim} = 1.5$, respectively. The magenta line is our default choice in presenting our earlier results (e.g. Fig. 7). The observational results from Danforth et al. (2016) and Kirkman et al. (2007), along with a combined fit of their data are shown identically to previous graphs of D_A.

of the SIMBA simulation. One of these variants uses a box size of $25\,h^{-1}$ Mpc with 512^3 mass elements (labelled 'm25n512'), and the other uses a box size of $25\,h^{-1}$ Mpc with 256^3 mass elements (labelled 'm25n256'). Aside from the box size and number of mass elements, the simulation physics for all three are identical.

It can be seen in Fig. 11 that both the m25n512 and m25n256 simulations have slightly more absorption than the main SIMBA results; however, both of the simulations using a box size of $25\,h^{-1}$ Mpc have almost identical absorption, with the slight exception of right around $z = 0$. The variation in parameters between the three simulations sheds light on the effects of box size and mass resolution on our results for D_A. There appears to be a minor increase in absorption at close to $z = 0$ when the resolution is increased, but we note that this is insignificant in magnitude compared with the differences between the SIMBA and No-jet simulations seen in Fig. 7.

The fact that both m25n512 and m256n256 are more similar to each other than to regular SIMBA suggests that box size has a greater effect on D_A than resolution does, at least among the values under consideration. The SIMBA simulation with the larger box size of $50\,h^{-1}$ Mpc but same resolution appears to have somewhat less absorption than the $25\,h^{-1}$ Mpc simulations at $z < 1$. This occurs because the $25\,h^{-1}$ Mpc simulations has fewer large galaxies owing to its non-representative volume, and therefore will have fewer large black holes producing jet feedback. This means that less of the IGM will have been heated, and these simulations will therefore have somewhat higher absorption.

In short, choices in continuum fitting can have an impact of up to ∼0.2 dex on the predicted values of D_A at $z = 0$. While this is substantial, it is roughly independent of the feedback model, so it does not lower the differences between the Jet and No-Jet. It also cannot fully explain the PUC, without rather extreme choices. If for instance we were to match the Danforth et al. (2016) choice of $\sigma_{\rm lim} = 1.5$, we actually produce an even larger PUC that is not fully mitigated even with the inclusion of jets. While using a very large value of $\sigma_{\rm lim}$ could mitigate the PUC, it is a non-standard choice that would give rise to significant a mismatch between the fitted continuum and the true continuum. In addition, we have checked that our results are unlikely to be significantly influenced by the numerical effects of simulation resolution. However, they are significantly impacted by box size when jets are included, because jets preferentially occur in massive red and dead galaxies that are underrepresented in small volumes. We conclude that our main results in Fig. 7 are not driven by numerical or analysis artifacts.

8 SUMMARY AND DISCUSSION

We have examined the evolution of the mean flux decrement in the Lyα forest D_A predicted in various simulations from the SIMBA suite, and used this to infer the H I photoionization rate as a function of redshift from $z = 2 \rightarrow 0$ by iteratively matching it to observations of D_A. We consider the full SIMBA simulation that includes various forms of AGN feedback (jet, X-ray, and radiative), and compare it to identical simulations with either X-ray feedback or X-ray and jet feedback turned off. We find greatest sensitivity to the inclusion of jet feedback: With jet feedback turned off, we recover the so-called PUC (Kollmeier et al. 2014) in which the Lyα forest observations require $\Gamma_{\rm H\,I}$ values at $z = 0$ that are ≈6 times higher than inferred from source count modelling by Haardt & Madau (2012). Including jets (regardless of X-rays or radiative feedback), reduces this discrepancy to ≈2.5 times, and further using an updated ionizing background from Faucher-Giguère (2019) now results in a reasonable match to

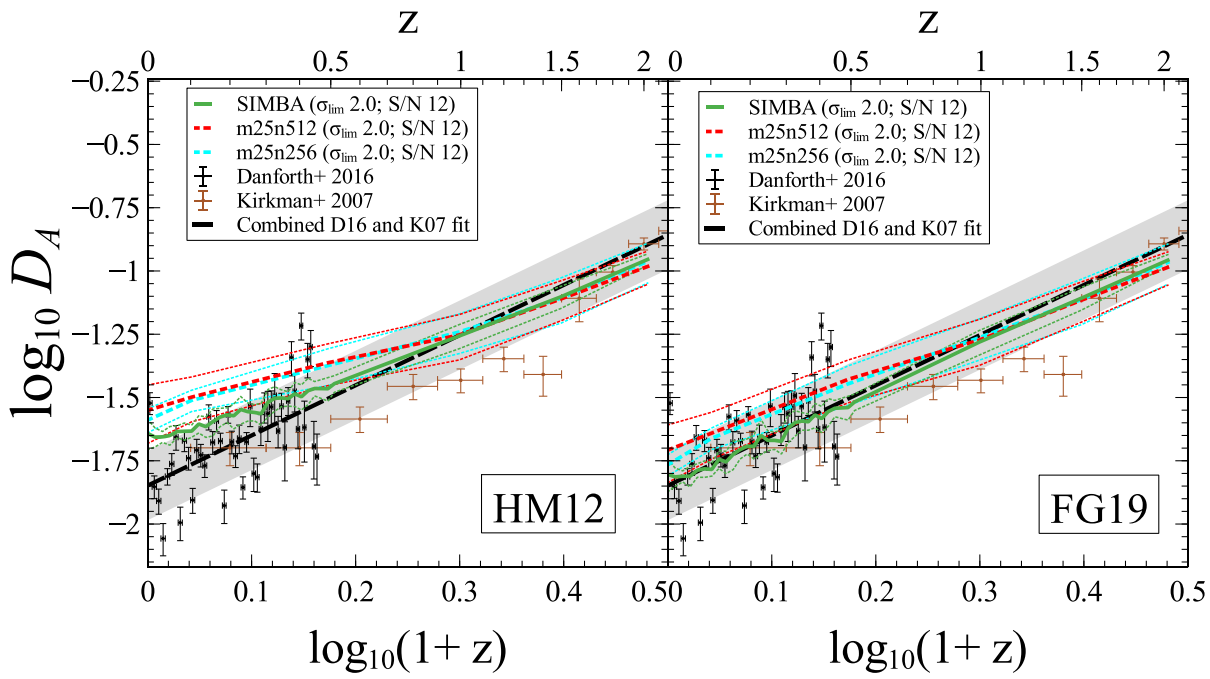

Figure 11. Graph showing effects of box size and resolution for three simulation runs of SIMBA varying these parameters, shown for both the HM12 (left) and FG19 (right) backgrounds. The green line shows the SIMBA run fiducial to this work, with a box size of $50h^{-1}$ Mpc and 512^3 gas elements, labelled 'SIMBA'; the red dashed line shows a run with a box size of $25\,h^{-1}$ Mpc and 512^3 mass elements, labelled 'm25n512'; and the cyan dashed line shows a run with a box size of $25h^{-1}$ Mpc and 256^3 mass elements, labelled 'm25n256'. The observational results from Danforth et al. (2016) and Kirkman et al. (2007), along with a combined fit of their data are shown identically to previous graphs of D_A. Among these tests, the change due to simulation volume is more important than numerical resolution, which has little impact. This owes to jet feedback being more prominent in massive galaxies that are underrepresented in a $25\,h^{-1}$ Mpc volume.

D_A both at $z = 0$ and and its evolution since $z \sim 2$. Hence in SIMBA, it appears that AGN jet feedback strongly mitigates the PUC.

To understand the physical origin of the impact of jets, we examine the physical and evolutionary properties of the IGM, and their impact on D_A. Our main findings are as follows:

(i) Heating from AGN jets leads to a significantly increased fraction of baryons in the WHIM in SIMBA – up to 70 per cent at $z = 0$, versus 30 per cent when jets are excluded. This increase in the WHIM fraction comes primarily from a decrease in the baryon fraction in the diffuse IGM, from 39 per cent to 16 per cent at $z = 0$, along with reductions in other cool phases. With jets, the IGM baryon temperature distribution strongly peaks at $T \gtrsim 10^6$ K, instead of being predominantly at $T \lesssim 10^5$ K.

(ii) The baryon fractions in various phases most strongly diverge between the full SIMBA and No-jet cases at $z \lesssim 1$, when large black holes form that are responsible for quenching galaxies via jet heating. At $z \gtrsim 1$, there is no PUC, as the predicted $D_A(z)$ in all SIMBA variants matches observations fairly well for either the HM12 or FG19 case.

(iii) The decrease in the diffuse baryon fraction by 2.5 times leads to a decrease in D_A by a commensurate factor at $z = 0$ in SIMBA with jet feedback. Hence, the main impact on D_A of jet feedback is to remove IGM baryons from the Lyα-absorbing phase via heating. This is corroborated by examining the pixel counts as a function of density, where SIMBA shows a new population of low-absorption fluxes near and above the cosmic mean density compared to No-jet, owing to jet heating.

(iv) Assuming an FG19 background rather than HM12 results in the predicted D_A matching observations over the full redshift range probed here ($z = 0$–2), thus solving the PUC in SIMBA. Quantitatively, the photon underproduction factor $F_{\rm UVB}$ at $z = 0$ is reduced by a factor of ∼2.5 times owing to the inclusion of jets, and by ∼2 times by using FG19 instead of HM12, thereby reducing $F_{\rm UVB} \approx 6 \rightarrow 1.2$.

(v) The agreement in the redshift evolution of D_A when including jets is a crucial success, as it highlights the importance of having a solution to the PUC that only impacts D_A at late cosmic epochs, primarily at $z \lesssim 1$. This coincides with jet feedback becoming increasingly commonplace in order to yield today's red and dead galaxy population in SIMBA.

(vi) SIMBA further produces a good match to the FPDF as a function of redshift compared to COS GTO data, while No-jet significantly overproduces the FPDF. This shows that SIMBA not only suppresses the mean absorption, but also the absorption as a function of flux in accord with data.

(vii) Examining a simulation with jets on but X-ray feedback off shows very similar results to the full SIMBA simulation with all AGN feedback modes on. This demonstrates that it is the AGN jets that are responsible for heating the IGM and mitigating the PUC.

(viii) Comparing SIMBA to MUFASA which used a different halo mass-based thermal quenching mechanism shows that MUFASA goes partways towards solving the PUC, but does not have as dramatic an effect as SIMBA's jets. This suggests that careful measurements of $\Gamma_{\rm H\,I}$ and D_A could together provide constraints on AGN feedback mechanisms.

(ix) We have tested the sensitivity of these results to various numerical choices. A key choice occurs in the continuum fitting process, and reasonable choices can lead to variations of up to ∼0.2 dex in $D_A(z = 0)$. While this cannot fully explain the PUC,

there is some uncertainty on the specific numbers quoted in above. Generally, our choice of the key parameter $\sigma_{\rm lim} = 2.0$ is conservative, in the sense that it results in less of a PUC than, for example, $\sigma_{\rm lim} = 1.5$ used in Danforth et al. (2016). We further find that our results are insensitive to simulation resolution, but in SIMBA we do require a sufficient volume to representatively produce massive galaxies that drive jets.

While some work has claimed that revising models of the ionizing background is solely sufficient to solve the PUC, we find that it is also necessary for AGN jets to be modelled in order for the crisis to be fully resolved in SIMBA. AGN jets are phenomena that are known to exist, and it is heartening that their inclusion in state of the art simulations can also play a role in addressing other discrepancies between observation and theoretical predictions.

Our results broadly echo those presented in Gurvich et al. (2017), who showed using the Illustris simulation that AGN feedback can have a strong impact on the diffuse Lyα forest. They likewise found that such heating, plus assuming an Faucher-Giguère et al. (2009) background (which is slightly lower than FG19), essentially solved the PUC in Illustris. SIMBA has the advantage of a more plausible AGN feedback model that does a better job of quenching galaxies and does not overevacuate hot haloes, but the resulting impact on the IGM appears broadly comparable. Without jets, our results confirm every other fully hydrodynamical galaxy formation simulation's results that also find a PUC at the ∼4–6 times level. We also find that without jets, the redshift evolution of D_A does not match observations. Any model that aims to solve the $z = 0$ PUC must also account for the fact that it is a late-time cosmic effect, essentially disappearing at $z \gtrsim 1$.

The large increase in WHIM baryon fraction should be testable with future observations, such as with high-ionization oxygen absorption lines. The impact on O VI absorption may be modest because in SIMBA the jet heating does not strongly increase the amount of $\sim 10^{5.5}$ K gas (Fig. 4) where such absorption is strong, but rather moves gas to higher temperatures that would give rise to, for example, O VII absorption in the soft X-rays, or Ne VIII in the extreme UV. Current constraints are insufficient to discriminate between our jet versus no-jet predictions, but upcoming facilities such as *Athena* and *Lynx* would be ideal for this. Another potential avenue for constraints is examining Sunyaev–Zel'dovich integrated IGM pressure measurements (e.g. Lim et al. 2018; de Graaff et al. 2019), which could provide constraints on the phase space distribution of IGM baryons. We plan to investigate whether SIMBA satisfies these constraints in future work.

The shape of the H I column density distribution is also an important constraint for solving the PUC. We have sidestepped this issue here, even though it was an important consideration in previous works (Kollmeier et al. 2014; Shull et al. 2015; Gurvich et al. 2017). Any solution to the PUC must also impact the column density distribution in a way that remains concordant with observations. A proper comparison of this, however, requires carefully mimicking the observational S/N, LSF, wavelength coverage, and profile fitting algorithm used for the data. It is worth noting that in Dave & Tripp (2001), the observed column density distribution using high-resolution *HST*/Space Telescope Imaging Spectrograph (STIS) data was found to be significantly steeper than that found by Danforth et al. (2016) using lower resolution *HST*/COS data, illustrating this sensitivity. We plan to conduct side-by-side Voigt profile fitting comparisons of absorber statistics in the future, but the PUC is already evident even when considering the stacked statistic of the mean flux decrement.

Broadly, our conclusions highlight the perhaps surprising point that the ionization level of the low-redshift IGM as traced by Lyα absorption can potentially be strongly impacted by AGN feedback originating deep within massive galaxies. While current uncertainties around determining the low-z metagalactic photoionization rate complicate the interpretation, this none the less provides new avenues to constrain AGN feedback models in a regime far removed from where it is typically constrained via the properties of quenched galaxies and their black holes.

ACKNOWLEDGEMENTS

The authors thank Sarah Appleby, Blakesley Burkhart, Neal Katz, and Juna Kollmeier for helpful discussions. We thank the YT team for development and support of YT, and Bernhard Röttgers and Miha Cernetic for development of PYGAD. RD acknowledges support from the Wolfson Research Merit Award program of the U.K. Royal Society. DS was supported by the European Research Council, under grant no. 670193. DAA acknowledges support by the Flatiron Institute, which is supported by the Simons Foundation. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility. The equipment was funded by BEIS capital funding via STFC capital grants ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The simulation data underlying this article are available at `simba.roe.ac.uk`. The derived data generated in this research will be shared on reasonable request to the corresponding author.

REFERENCES

Abel T., Anninos P., Zhang Y., Norman M. L., 1997, New Astron., 2, 181
Anglés-Alcázar D., Davé R., Faucher-Giguère C.-A., Özel F., Hopkins P. F., 2017a, MNRAS, 464, 2840
Anglés-Alcázar D., Faucher-Giguère C.-A., Kereš D., Hopkins P. F., Quataert E., Murray N., 2017b, MNRAS, 470, 4698
Anglés-Alcázar D., Özel F., Davé R., 2013, ApJ, 770, 5
Best P. N., Heckman T. M., 2012, MNRAS, 421, 1569
Bolton J. S., Puchwein E., Sijacki D., Haehnelt M. G., Kim T.-S., Meiksin A., Regan J. A., Viel M., 2017, MNRAS, 464, 897
Bondi H., 1952, MNRAS, 112, 195
Borrow J., Angles-Alcazar D., Dave R., 2019, MNRAS, 491, 6102
Broderick A. E., Chang P., Pfrommer C., 2012, ApJ, 752, 22
Burchett J. N. et al., 2019, ApJ, 877, L20
Cen R., Ostriker J. P., 1999, ApJ, 514, 1
Chen X., Weinberg D. H., Katz N., Davé R., 2003, ApJ, 594, 42
Choi E., Ostriker J. P., Naab T., Johansson P. H., 2012, ApJ, 754, 125
Danforth C. W. et al., 2016, ApJ, 817, 111
Dave R., Tripp T. M., 2001, ApJ, 553, 528
Dave R. et al., 2001, ApJ, 552, 473
Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, MNRAS, 486, 2827
Davé R., Hernquist L., Katz N., Weinberg D. H., 1999, ApJ, 511, 521
Davé R., Katz N., Oppenheimer B. D., Kollmeier J. A., Weinberg D. H., 2013, MNRAS, 434, 2645
Davé R., Oppenheimer B. D., Katz N., Kollmeier J. A., Weinberg D. H., 2010, MNRAS, 408, 2051
Davé R., Oppenheimer B. D., Sivanand am S., 2008, MNRAS, 391, 110
Davé R., Rafieferantsoa M. H., Thompson R. J., 2017, MNRAS, 471, 1671
Davé R., Thompson R., Hopkins P. F., 2016, MNRAS, 462, 3265

de Graaff A., Cai Y.-C., Heymans C., Peacock J. A., 2019, A&A, 624, A48
Fabian A. C., 2012, ARA&A, 50, 455
Faucher-Giguère C.-A., 2019, MNRAS, 493, 1614 (FG19)
Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, ApJ, 703, 1416
Fumagalli M., Haardt F., Theuns T., Morris S. L., Cantalupo S., Madau P., Fossati M., 2017, MNRAS, 467, 4802
Gaikwad P., Khaire V., Choudhury T. R., Srianand R., 2017a, MNRAS, 466, 838
Gaikwad P., Srianand R., Choudhury T. R., Khaire V., 2017b, MNRAS, 467, 3172
Genel S. et al., 2014, MNRAS, 445, 175
Gurvich A., Burkhart B., Bird S., 2017, ApJ, 835, 175
Haardt F., Madau P., 1996, ApJ, 461, 20
Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, Clusters of Galaxies and the High Redshift Universe Observed in X-rays, XXXVI Rencotres de Moriond, Savoie, France. p. 64, preprint (astro-ph/0106018)
Haardt F., Madau P., 2012, ApJ, 746, 125 (HM12)
Heckman T. M., Best P. N., 2014, ARA&A, 52, 589
Henden N. A., Puchwein E., Shen S., Sijacki D., 2018, MNRAS, 479, 5385
Hopkins P. F., 2015, MNRAS, 450, 53
Hopkins P. F., Quataert E., 2011, MNRAS, 415, 1027
Hui L., Gnedin N. Y., 1997, MNRAS, 292, 27
Katz N., Weinberg D. H., Hernquist L., 1996, ApJS, 105, 19
Kennicutt Robert C. J., 1998, ARA&A, 36, 189
Khaire V., Srianand R., 2015, MNRAS, 451, L30
Khaire V. et al., 2019, MNRAS, 486, 769
Kirkman D., Tytler D., Lubin D., Charlton J., 2007, MNRAS, 376, 1227
Kitayama T., Suto Y., 1996, ApJ, 469, 480
Kollmeier J. A. et al., 2014, ApJ, 789, L32
Kormendy J., Ho L. C., 2013, ARA&A, 51, 511
Krumholz M. R., Gnedin N. Y., 2011, ApJ, 729, 36
Kulkarni G., Worseck G., Hennawi J. F., 2019, MNRAS, 488, 1035
Lim S. H., Mo H. J., Wang H., Yang X., 2018, MNRAS, 480, 4017
Li Q., Narayanan D., Davé R., 2019, MNRAS, 490, 1425
Meiksin A. A., 2009, Rev. Mod. Phys., 81, 1405
Nicastro F. et al., 2018, Nature, 558, 406
Oppenheimer B. D., Davé R., 2006, MNRAS, 373, 1265
Perna M., Lanzuisi G., Brusa M., Mignoli M., Cresci G., 2017, A&A, 603, A99
Planck CollaborationXIII, 2016, A&A, 594, A13
Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, MNRAS, 430, 2427
Rauch M. et al., 1997, ApJ, 489, 7
Robson D., Davé R., 2020, MNRAS, 498, 3061
Rutkowski M. J. et al., 2016, ApJ, 819, 81
Röttgers B., Naab T., Cernetic M., Davé R., Kauffmann G., Borthakur S., Foidl H., 2020, MNRAS, 496, 152
Schaye J. et al., 2015, MNRAS, 446, 521
Shull J. M., Moloney J., Danforth C. W., Tilton E. M., 2015, ApJ, 811, 3
Shull J. M., Smith B. D., Danforth C. W., 2012, ApJ, 759, 23
Smith B. D., Hallman E. J., Shull J. M., O'Shea B. W., 2011, ApJ, 731, 6
Smith B. D. et al., 2017, MNRAS, 466, 2217
Somerville R. S., Davé R., 2015, ARA&A, 53, 51
Sorini D., 2017, PhD thesis, International Max Planck Research School for Astronomy and Cosmic Physics at the University of Heidelberg (IMPRS-HD)
Sorini D., Oñorbe J., Hennawi J. F., Lukić Z., 2018, ApJ, 859, 125
Thomas N., Davé R., Anglés-Alcázar D., Jarvis M., 2019, MNRAS, 487, 5764
Tonnesen S., Smith B. D., Kollmeier J. A., Cen R., 2017, ApJ, 845, 47
Tripp T. M., Savage B. D., Jenkins E. B., 2000, ApJ, 534, L1
Tytler D., Fan X.-M., Burles S., 1996, Nature, 381, 207
Viel M., Haehnelt M. G., Bolton J. S., Kim T.-S., Puchwein E., Nasir F., Wakker B. P., 2017, MNRAS, 467, L86
Vogelsberger M. et al., 2014, MNRAS, 444, 1518
Wakker B. P., Hernandez A. K., French D. M., Kim T.-S., Oppenheimer B. D., Savage B. D., 2015, ApJ, 814, 40
Weinberg D. H., Katz N., Hernquist L., 1998, in Woodward C. E., Shull J. M., Thronson Harley A. J., eds, ASP Conf. Ser. Vol. 148, Origins, Astron. Soc. Pac., San Francisco, CA. p. 21 preprint (arXiv:astro-ph/9708213)
Weinberger R. et al., 2018, MNRAS, 479, 4056
Whittam I. H., Prescott M., McAlpine K., Jarvis M. J., Heywood I., 2018, MNRAS, 480, 358

This paper has been typeset from a TEX/LATEX file prepared by the author.