Seeing the invisible: tracing cosmic structure across time and space with hydrogen lines

CALVIN VERRIER SYKES

How to cite:

SYKES, CALVIN VERRIER (2021) Seeing the invisible: tracing cosmic structure across time and space with hydrogen lines. Doctoral thesis, Durham University.

Seeing the invisible: tracing cosmic structure across time and space with hydrogen lines

Calvin Sykes

A Thesis presented for the degree of
Doctor of Philosophy

Institute for Computational Cosmology
and
Centre for Extragalactic Astronomy

Ogden Centre for Fundamental Physics
Department of Physics
Durham University
United Kingdom

October 2021

Seeing the invisible: tracing cosmic structure across time and space with hydrogen lines

Calvin Sykes

Submitted for the degree of Doctor of Philosophy

October 2021

Abstract:

In this thesis, we explore astrophysical frontiers along two lines of enquiry, on very different physical scales and separated by ten billion years of cosmic time, yet linked by a focus on the properties of diffuse gas modelled by radiation-hydrodynamic simulations and revealed by spectroscopic observations.

Firstly, we investigate RELHICs, a hypothetical population of low-mass dark matter haloes, predicted by hydrodynamical simulations to remain gas-rich yet star-free. We show that RELHICs illuminated by the extragalactic ultraviolet background (UVB) should produce distinctive "fluorescent" emission, the characterisation of which has the potential to yield constraints on the intensity and spectral shape of the UVB, on the nature of dark matter, and on the primordial abundance of helium. We address the substantial practical challenges inherent in detecting this extremely faint emission.

Secondly, we focus on the damped-Lyman alpha (DLA) absorbers originally detected in the spectra of background quasars, which we study using a novel set of cosmological simulations in conjunction with accurate radiative transfer modelling. We make a detailed investigation of the statistical properties of $z \sim 3$ DLAs, and relate the properties of individual absorbers to the haloes that host them. We present evidence furthering the view that the neutral gas which gives rise to DLAs is primarily composed of circumgalactic material undergoing cosmological accretion onto galaxies, with only the sightlines containing very strong absorbers arising in the galaxies themselves.

Declaration

The work described in this thesis was undertaken at the Institute for Computational Cosmology and Centre for Extragalactic Astronomy, both at Durham University, England. It was carried out during the period October 2017–October 2021, while the author was a research student under the supervision of Prof. Michele Fumagalli and Prof. Ryan Cooke. No part of this thesis has been submitted for any degree or qualification at Durham University or elsewhere. It is the sole work of the author, with the exception of certain figures which are given an explicit attribution in the figure caption.

Chapters 2 and 3 are adapted from papers published in the journal Monthly Notices of the Royal Astronomical Society (MNRAS), as follows:

Sykes C., Fumagalli M., Cooke R., Theuns T., Alejandro B. L., 2019, MNRAS, 487, 609
Sykes C., Fumagalli M., Cooke R., Theuns T., 2020, MNRAS, 492, 2151

Chapter 4 is adapted from material in preparation for submission to MNRAS. The author was primarily responsible for all aspects of these publications.

Acknowledgements

So here we are, at the end of a long and occasionally bumpy road. There are many people without whom I would never have made it, and to all of you I offer my heartfelt thanks.

First and foremost thanks must go to my family, for their endless support. It spans the gamut from the most fundamental matters – for instilling and encouraging a desire to ask and to learn – to the more mundane, yet no less appreciated – in particular, for acting as my personal removals service over the last eight years!

To Michele and Ryan, I could not have wished for better supervisors and I am very grateful to have been your student. Thank you for your inexhaustible patience and enthusiasm, and for your support in matters academic and otherwise. Thanks also to Tom, for all your invaluable advice and wisdom. And to Alastair and the other members of cosma-support, without all your effort keeping the machine whirring away, most of the work in this thesis would have been impossible.

I count myself very lucky to have spent my time in Durham with such a fantastic group of colleagues and friends, who made the department such an enjoyable place to work. Particular thanks go to my office-mates through the years, for welcoming me to the department and helping me find my footing; for companionship in the wilds of the Rochester building (aka ‘shanty-town’); and for support of the moral, technical, and caffeinated varieties. To the pub crew, thanks for all the good times we shared. I suspect I still owe most of you a pint or two, which I very much hope I will get an opportunity to repay some day! And finally, thanks to my housemates, Tom C & Tom R, for preserving my sanity through lockdown.

During this thesis, the author was supported by an STFC studentship, project code ST/R504725/1.

The work in this thesis used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (`www.dirac.ac.uk`). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, 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.

The work in this thesis made use of the following public software packages:

- CLOUDY (Ferland et al. 2017)
- GADGET (Springel 2005)

Additionally, the `python` programming language (Van Rossum & Drake Jr 1995) was used, in conjunction with the following libraries:

- `Cython` (Behnel et al. 2011)
- `h5py` (Collette 2013)
- `IPython/jupyter` (Perez & Granger 2007)
- `matplotlib` (Hunter 2007)
- `numba` (Lam et al. 2021)
- `numpy` (Harris et al. 2020)
- `py-sphviewer` (Benitez-Llambay 2015)
- `scipy` (Virtanen et al. 2020)

“Hydrogen is a light, odourless gas, which, given enough time, turns into people.”

— Edward R. Harrison

Contents

List of Figures

List of Tables

CHAPTER 1

Introduction to the thesis

1.1 Cosmology

1.1.1 The beginnings of cosmology

The modern science of cosmology has its origins a little over a century ago, with the development of the theory of general relativity (Einstein 1914, 1915a). GR radically redefined our understanding of the gravitational interactions between massive bodies, replacing the Newtonian interpretation of gravitational forces with a more abstract conception in which gravity is viewed as a geometric effect due to the warping of spacetime in the presence of matter and energy. This description is embodied by the Einstein field equations

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \kappa T_{\mu\nu} \tag{1.1.1}$$

in which the terms on the left-hand side, comprising the Ricci curvature $R_{\mu\nu}$; the scalar curvature R; and the metric $g_{\mu\nu}$, describe the curvature of spacetime, while the right-hand side contains the stress-energy tensor $T_{\mu\nu}$ which represents its mass-energy content. The constant $\kappa = 8\pi G/c^4$, where G is the familiar gravitational constant of Newtonian mechanics, and c is the speed of light. Initially applied to local astronomical phenomena, GR found early success in predicting the observed orbital precession of Mercury (Einstein 1915b) and the gravitational lensing of background starlight by the Sun (Dyson et al.

1920). To build on these successes, attention soon turned to the task of obtaining a general-relativistic description of the Universe as a whole.

An early attempt by Einstein (1917; see also O'Raifeartaigh et al. 2017) started from the assumption that on the largest scales the Universe is homogeneous and isotropic, later referred to as the "cosmological principle", which remains foundational to this day. Einstein struggled to reconcile this principle with the prevailing wisdom of a static, temporally-infinite universe: the dynamics predicted by GR implied collapse under gravity for a spatially-finite universe, or unstable equilibrium for an infinite one. His proposed resolution to this conflict involved augmenting the field equations (Eq. 1.1.1) with an additional term, the cosmological constant Λ, which served to provide a source of negative pressure to counteract the inward pull of gravity. This solution remained an unsatisfactory one, requiring a finely-tuned value of Λ, and again displaying instability against local density perturbations.

The assumption of a static universe was relaxed by Friedmann (1922), leading to a model in which a homogeneous, isotropic universe could undergo uniform metric expansion or contraction. The same results were independently re-derived by Lemaître (1927), and generalised by Robertson (1935, 1936a,b) and Walker (1937). In this model, today known as the Friedmann-Lemaître-Robertson-Walker (FLRW) model, the metric $g_{\mu\nu}$ expressed in spherical coordinates takes the form

$$\mathrm{d}s^2 = -c^2\mathrm{d}t^2 + a(t)^2\left(\mathrm{d}r^2 + S_k(r)^2(\mathrm{d}\theta^2 + \sin^2(\theta)\mathrm{d}\phi^2)\right) , \tag{1.1.2}$$

where the functional form of S_k differs for universes with positive, zero, or negative intrinsic spatial curvature k. The time evolution of the metric is encapsulated by the "scale factor" $a(t)$, whose behaviour is itself determined by the content of the spacetime, as expressed by the Friedmann equation

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho + \frac{\Lambda c^2}{3} - \frac{kc^2}{a^2} \tag{1.1.3}$$

which is obtained from solving the field equations with the metric of Eq. (1.1.2), in conjunction with the continuity equation

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = -3\frac{\dot{a}}{a}\left(\rho + \frac{p}{c^2}\right) \tag{1.1.4}$$

which arises from thermodynamic considerations. In Eqs. (1.1.3) and (1.1.4), dot symbols indicate time derivatives, ρ is the total density, and p is the thermodynamical pressure. The third and final ingredient is an equation of state, which relates the density of matter with the pressure it exerts. In a cosmological context, it is typical to assume a so-called barotropic equation of state, $p = w\rho c^2$, where the “barotropic parameter” w is constant for a given substance (e. g. $w = 0$ for non-relativistic matter). With this specified, it is possible to solve Eq. (1.1.3) to obtain the time evolution of the scale factor for a universe containing a specific type of substance (or more realistically, a mixture of multiple types, as we will discuss in Section 1.1.2). As the detail afforded to this model here would suggest, it remains central to modern cosmology. When initially introduced, its reception was however lukewarm, with Einstein in particular dismissing it (Einstein 1922).

Contemporaneously with these theoretical developments, observational astronomy was also making rapid progress, which would eventually vindicate the expansionary model. In parallel with Einstein’s introduction of GR, vigorous discussion was ongoing regarding the overall scale of the Universe, with two factions forming: one which believed that the Milky Way constituted the entirety of the Universe, and that the observed “spiral nebulae” were relatively small structures located within our galaxy; the other held that these structures were instead distinct “island universes” located far beyond the Milky Way. This dispute culminated in the “Great Debate” of 1920 (Shapley & Curtis 1921), from which the latter view emerged victorious. This was in part thanks to spectral measurements of the spiral nebulae, which appeared to indicate recession velocities of hundreds of kilometres per second, far larger than expected for nearby structures (Slipher 1915, 1917). Truly conclusive evidence arrived a decade later, with Edwin Hubble’s use of Cepheid variable stars (Hubble 1925) to obtain the first reliable indications of the distance to the spiral nebulae, firmly demonstrating that they were in fact distinct galaxies lying far beyond the Milky Way. But Hubble’s most impactful contribution would follow shortly after, when he presented evidence for a linear correlation between the distances to other galaxies and their recession velocities (Hubble 1929; see also Bahcall 2015, review), exactly as predicted by expansionary models (Lemaître 1927). This work provided the necessary impetus for a paradigm shift, with the expansionary model becoming widely accepted – including by the newly-converted Einstein – within the following few years. With this came the abandonment of the cosmological constant, no longer required to ensure a static universe, and later (apocryphally) referred to by its original proponent as his "greatest blunder". However, we will revisit the cosmological constant in Section 1.1.2.2.

With the acceptance of the FLRW model came a new problem - that of the origins of the Universe. Following the evolution of an expanding universe in reverse led to the inevitable prediction of a zero point in cosmic time, at which the universe was condensed into a singular state. The nature of this "primeval atom", as it was named by Lemaître when presaging this issue in his 1927 work, would provide the source of the next major debate in astronomy. Initially on philosophical grounds, many astronomers rejected the idea of a cosmic beginning (e. g. Zwicky 1929, 1933; Einstein 1931; Milne 1935). Another notable dissenter was Fred Hoyle, who first coined the modern term "Big Bang", using it in a gently-pejorative sense. With time, most alternative models were ruled out observationally, such that by the 1950s there remained two viable cosmological paradigms: the temporally-finite Big Bang; and the "steady state" solution favoured by Hoyle, in which new matter was constantly generated during the expansion of an infinitely-old universe, such that the average density remained constant (Bondi & Gold 1948; Hoyle 1948). However, evidence against this theory began to mount as well, with the discovery of a population of bright radio sources dubbed quasi-stellar objects (QSOs; "quasars") exclusively at large distances, implying an evolution in galaxy properties with light travel time that was in conflict with the predictions of the steady-state model (Ryle & Scheuer 1955). Spectacular confirmation of the Big Bang model arrived a decade later, with the serendipitous discovery of a spatially-uniform microwave radiation field, corresponding to a black body temperature of ~3 K (Penzias & Wilson 1965). This "cosmic microwave background" (CMB) was a clear prediction of the Big Bang theory, being a natural result of a Universe that was much hotter and denser in the distant past (Alpher & Herman 1948; Dicke et al. 1965), and its detection ruled out the steady state model for all but its most committed adherents.

1.1.2 The modern cosmological paradigm

To summarise the narrative of the previous section, our current cosmological understanding is built upon the application of general relativity to a Universe which is homogeneous, isotropic, and expanding from an initial hot, dense state. As discussed on page 3, the precise nature of the expansion depends on the mass-energy content of the Universe; this therefore forms the next subject of interest.

1.1.2.1 Dark matter

As early as the 1920s, it was suggested that in addition to the luminous mass inferred from starlight, galaxies required an additional non-luminous component, dubbed "dark matter" (DM), to explain their kinematics (Kapteyn 1922; Oort 1932). Early quantitative measurements supporting this hypothesis were obtained from studying the velocity dispersions of galaxies in the Coma cluster (Zwicky 1937), and the "rotation curve" $v(r)$ of the Andromeda galaxy (Babcock 1939; Mayall 1951). As time passed, the base of evidence in favour of DM continued to grow, encompassing both dynamical arguments (Kahn & Woltjer 1959) and further rotation curve measurements, with the first systematic results obtained by Burbidge & Burbidge (1960; and subsequent works). With improved instrumentation, it became possible to extend rotation curve measurements beyond galaxies' central regions, leading to the puzzling observation that many showed a 'flat' velocity profile, $v(r) \sim$ const, at large radii (Rubin & Ford 1970; Rogstad & Shostak 1972; Roberts & Rots 1973). It was soon recognised that these findings implied a matter distribution with a substantial, yet unseen, extended component (Ostriker et al. 1974; Einasto et al. 1974a,b; Rubin et al. 1980); meanwhile, early numerical work by Ostriker & Peebles (1973) demonstrated that such a "halo" was a necessary condition for the stability of galaxy disks.

However, throughout this period there was little speculation as to the identity of the DM, with the default assumption being that it should be some macroscopic form of baryonic matter e. g. stellar remnants. Only with the growth of interest in particle physics in the 1980s did the idea that DM was instead composed of microscopic particles come to prominence, first motivated by an (ultimately erroneous) claim of a measured 30 eV mass for the electron neutrino (Lubimov et al. 1980). Presently, a variety of particle identities are hypothesised for DM (for a review, see Feng 2010), but the precise nature of the DM particle has relatively little relevance for its cosmological implications.

One factor which does prove significant is the initial velocity distribution of the particles, which distinguishes ‘hot’ dark matter (HDM) models, in which the kinetic energies of DM particles in the early Universe are large compared with their rest mass, from ‘cold’ dark matter (CDM) for which the opposite is true. Many of the early particle candidates, such as the massive neutrinos of Lubimov et al., were HDM. A crucial prediction of these candidates is the existence of a length scale below which structure is erased, as a result of the DM freely streaming out of smaller overdensities before gravitational collapse begins. This length scale is typically of Mpc size, corresponding to galactic superclusters. In an HDM model, galaxies form in a ‘top-down’ manner, through the fragmentation of structures of this size into smaller ones (Bond et al. 1980; Peebles 1982a; Bond & Szalay 1983). As demonstrated numerically by White et al. (1983), this gives rise to a galaxy distribution exhibiting strong clustering on the free-streaming scale and little below it. This result was in direct contradiction with contemporary results from the first large-scale galaxy survey (Fig. 1.1; Davis et al. 1982), effectively ruling out HDM. By contrast, simulations assuming CDM predicted a hierarchical, ‘bottom-up’ method of structure formation, in which small objects formed first and merged to form larger ones. This scenario proved consistent with the data (Blumenthal et al. 1984, 1985; Davis et al. 1985), resulting in CDM rapidly becoming the prevailing hypothesis.

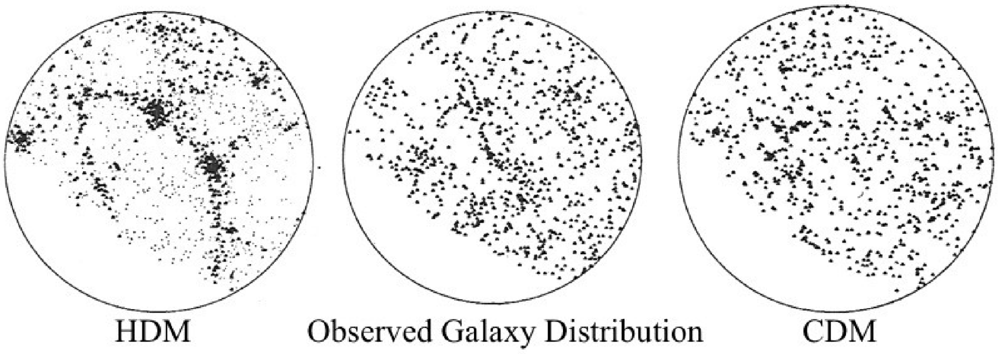

Figure 1.1
Comparison of simulations of CDM and HDM cosmologies with the observed distribution of galaxies. The HDM model predicts much stronger clustering on Mpc scales than observed, whereas the CDM model is generally consistent with the data. However, the observed peculiar velocities of galaxies were higher than expected, and galaxies appeared to exhibit strong biasing (i.e. they are found only at large matter overdensities), providing early indications of the presence of dark energy. Reproduced from Primack (2009).

1.1.2.2 Dark energy

Per Eq. (1.1.2), the overall geometry of an FLRW universe depends on the intrinsic spatial curvature k. There are three possibilities: closed, flat, or open universes, corresponding to $k = \{+1, 0, -1\}$ respectively. For a matter-dominated universe, such as the CDM universe discussed in the previous section, the sign of the curvature also dictates the universe's long-term evolution. A positively-curved, closed universe reaches a maximum value of the scale factor before eventually recollapsing, whereas flat and open universes continue to expand forever. Given a measured value for the Hubble parameter H, exactly one value of the average density corresponds to a flat universe, the 'critical' density

$$\rho_{\text{crit}} = \frac{3H^2}{8\pi G} \sim 10^{-26}\,\text{kg}\,\text{m}^{-3} \,. \tag{1.1.5}$$

Expressing the total density as a ratio to the critical value gives the "density parameter" $\Omega \equiv \rho/\rho_{\text{crit}}$. For closed universes, $\rho > \rho_{\text{crit}}$ and $\Omega > 1$, while for open universes $\Omega < 1$.

In early work, philosophical considerations – specifically, the fact that a closed universe is spatially finite – led to a preference for assuming $\Omega \geq 1$ (e. g. Einstein 1917). As discussed in Section 1.1.1, a "cosmological constant" was briefly introduced to counteract such a universe's tendency for collapse, before being removed once the observational fact that the Universe was expanding became clear. From this point, little thought was paid to the Universe's spatial curvature until the early 1980s, when an early period of accelerated expansion, dubbed "inflation", was hypothesised (Guth 1981; Linde 1982). Inflation was theoretically motivated by the development of then-popular grand unified theories (e. g. Georgi & Glashow 1974), and also provided an explanation for the "horizon problem" – the unexplained uniformity of the CMB across the sky, which required the Universe's temperature to be tightly correlated across regions of space which were causally disconnected at the time of the CMB's emission. During inflation, the scale factor would have grown exponentially, such that any preexisting spatial curvature would be strongly diluted, generating a flat $\Omega = 1$ universe instead.

The resulting model universe, spatially flat due to inflation and with a present-day energy density dominated by CDM, would soon face observational scrutiny. Measurements of galaxy clustering, similar to those which effectively distinguished between HDM and CDM, indicated clustering on larger scales than predicted (Geller & Huchra 1989; Maddox et al. 1990; Efstathiou et al. 1990); while constraints on the baryon abundance (White et al. 1993) and the spectrum of anisotropies in the CMB, then newly-discovered by the COBE satellite (Smoot et al. 1992; Wright et al. 1992; Efstathiou et al. 1992), both pointed to a total matter density well below critical, $\Omega_{\rm M} \sim 0.2$. These discrepancies prompted investigation of alternative dark matter models, as well as more radical changes in assumptions: in particular, they could be alleviated by the reintroduction of the cosmological constant, now interpreted as an additional contribution to the Universe's total density arising from negative-pressure "dark energy" (DE). Just before the new millennium, measurements of type 1a supernovae (SNe1a) led to the striking conclusion that rather than the expansion rate of the Universe slowing, it was instead accelerating (Riess et al. 1998; Perlmutter et al. 1999)[1], exactly as expected if the Universe was presently DE-dominated.

---

[1] Specifically, the "deceleration parameter" $q_0 \equiv -\ddot{a}a/\dot{a}^2$ was determined to be <0.

1.1.3 Current status

Since 2000, additional evidence from a variety of sources, including further SNe1a data; more precise measurements of the CMB anisotropies (Spergel et al. 2003; Komatsu et al. 2011; Aghanim et al. 2020, 2021); baryonic acoustic oscillations (Eisenstein et al. 2005; Alam et al. 2017; des Bourboux et al. 2020); weak lensing (Van Waerbeke et al. 2000; Bacon et al. 2000; Kilbinger 2015); and galaxy cluster statistics (Borgani et al. 2001; Allen et al. 2011), have established the cosmological constant as a permanent addition to the cosmological model. Much of this evidence also reinforces support for CDM, as do additional probes of cluster dynamics (Clowe et al. 2006; Thompson et al. 2015). So convincing is the level of agreement between modern observations that a unified 'concordance' model of cosmology, referred to as 'ΛCDM', has arisen (e. g. Bahcall et al. 1999). This represents the latest stage of cosmology's maturation as a science, having progressed from establishing the fundamental laws which govern the Universe's evolution; to determining the basic facts of its composition; and to the present situation, where most attention is now focused on making increasingly precise determinations of the ΛCDM model's free parameters, as shown in Fig. 1.2.

Despite the many successes of the ΛCDM model, it should be noted that some disagreement is ongoing. On small scales, a number of potential challenges to CDM exist, including the missing satellites problem (Klypin et al. 1999; Moore et al. 1999) and the core-cusp problem (Flores & Primack 1994; Moore 1994); see also Del Popolo & Le Delliou (2017). On larger scales, some discrepancies in measurements of the cosmological parameters have grown increasingly significant; most well-known is the so-called 'H_0 crisis', the 5σ disagreement between measurements of the present-day expansion rate derived from the CMB and from SNe1a (Riess 2020). Meanwhile, in the realm of particle physics, theoretical and practical limitations render the nature of CDM and DE – which together constitute fully 95% of the Universe's mass-energy budget – largely unknown. Numerous explanations have been proposed for these and other issues (Bull et al. 2016; Perivolaropoulos & Skara 2021), and the coming years are hoped to prove fruitful in throwing light upon many of them. In the

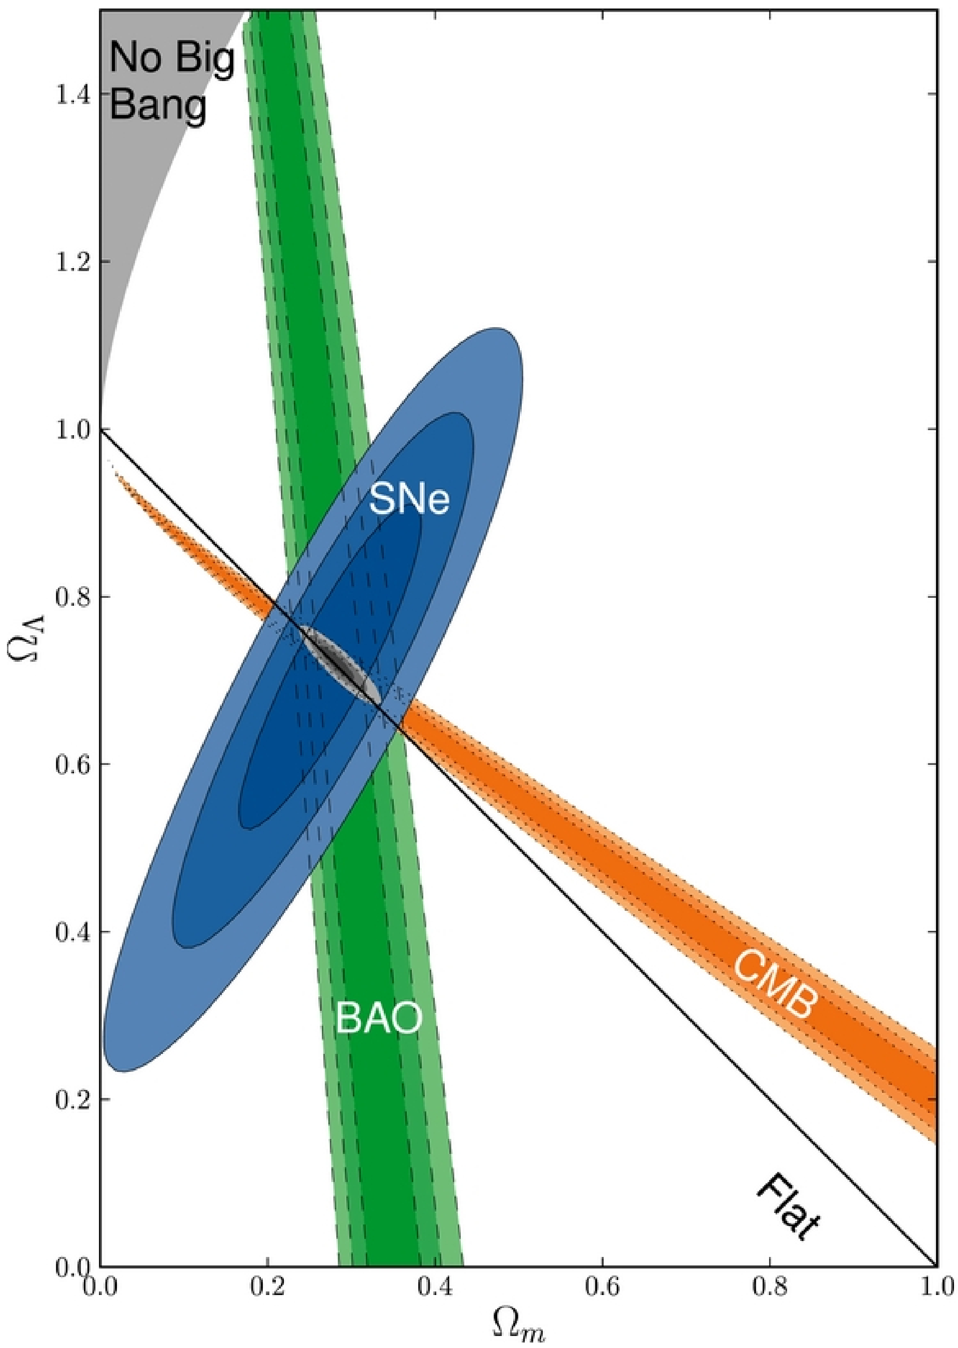

Figure 1.2

Combined constraints on the ΛCDM model in the $\Omega_M - \Omega_\Lambda$ plane, obtained from CMB anisotropies (orange), SNe1a (blue), and the baryon acoustic oscillation (BAO) feature in large scale structure (green). Shading indicates 68.3%, 95.4% and 99.7% confidence regions. For this combination of data, the best-fitting set of parameters are $\Omega_M = 0.272 \pm 0.014$, $\Omega_\Lambda = 0.729 \pm 0.014$, $\Omega_k = 0.002 \pm 0.005$. Reproduced from Suzuki et al. (2012).

context of this thesis though, we shall make the pragmatic choice to assume the standard narrative for most purposes.

1.2 Structure formation

In the FLRW model of cosmology discussed in the previous section, the Universe is modelled as perfectly smooth and homogeneous. While effective in describing the large-scale evolution of the cosmos, this approximation is clearly violated on small scales, where there exists a wealth of non-linear structure. The modern understanding of this structure's formation begins with the discovery of anisotropies in the CMB by the COBE satellite (Smoot et al. 1992). This discovery provided validation of extensive theoretical work, which theorised anisotropies to be the product of slight variations in the Universe's density at the time of CMB emission (Sachs & Wolfe 1967); identified the formation sites of galaxies with local maxima in the density field (Press & Schechter 1974; Bardeen et al. 1986); and predicted their further growth through gravitational collapse (Gunn & Gott 1972).

A second crucial development was the acceptance of the CDM paradigm, which fixes the expected spectrum of anisotropies (Peebles 1982b), and implies a scenario in which, as we have previously remarked, structure formation proceeds via the hierarchical merging of small structures to form larger ones. Tests of the resulting distribution of DM, especially on small scales where nonlinear dynamics dominate, are most readily achieved using N-body simulations. In these calculations, an ensemble of representative 'particles' are evolved from an initial state encapsulating the primordial density perturbations according to gravitational dynamics (e. g. Efstathiou et al. 1985). From these simulations, detailed predictions may be made, including of the mass distribution of DM haloes (Reed et al. 2007; Tinker et al. 2008) and their internal structure (Navarro et al. 1996b, 1997). They also predict that on large scales, the DM distribution forms a filamentary 'cosmic web' (Fig. 1.3; Bond et al. 1996; Springel et al. 2005), with underdense voids containing little structure separating the filaments, and clusters of collapsed haloes found where they intersect.

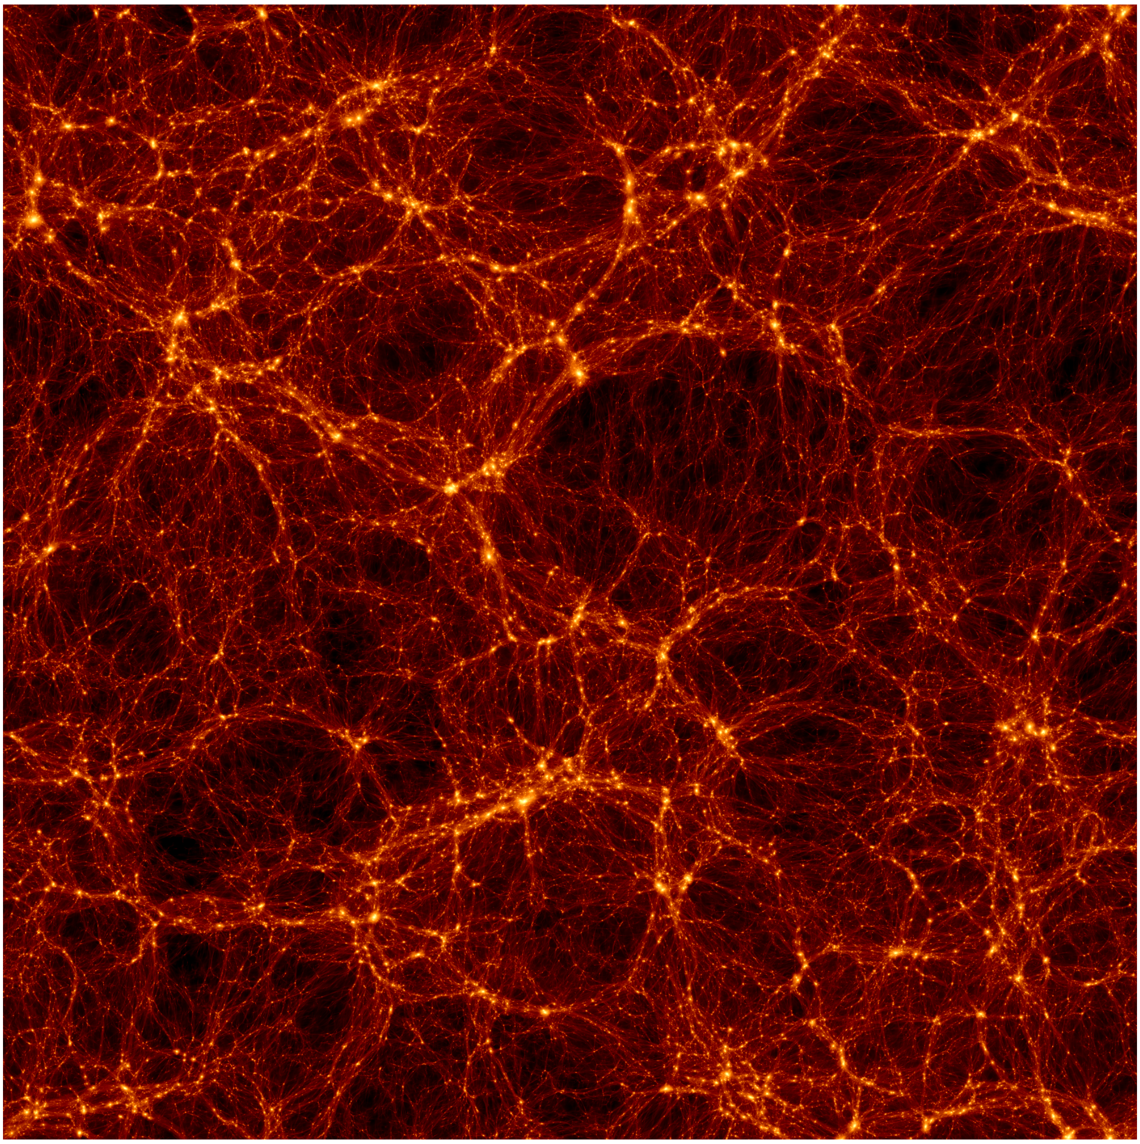

Figure 1.3
Visualisation of the cosmic web produced from an N-body simulation, showing a $(250 \times 250 \times 40)\,\mathrm{Mpc}/h$ slice at $z = 0$; colour encodes the dark matter density. Figure produced by Mansfield & Diemer[2] using simulation data from Diemer & Kravtsov (2014).

---

[2] See: `http://www.benediktdiemer.com/visualization/images/`

Extending numerical modelling from the formation and statistics of dark matter haloes, to that of the galaxies found within them, requires incorporating the dissipative dynamics of gas. Early attempts at accounting for these phenomena (Rees & Ostriker 1977; White & Rees 1978) focused on motivating the observed sizes and angular momentum distributions of galaxies. This work led to the development of semi-analytic models (Cole 1991; Lacey & Silk 1991; White & Frenk 1991), in which a dark matter distribution obtained from an N-body simulation is combined with models for baryonic physics derived analytically or extrapolated from small-scale hydrodynamical simulations (see below). Modern semi-analytic models (e. g. GALFORM: Cole et al. 2000; Lacey et al. 2016; L-GALAXIES: Springel et al. 2001; Guo et al. 2011; Henriques et al. 2020) now include numerous additional astrophysically-relevant processes, and successfully reproduce many observed properties of the galaxy population.

However, the phenomenological nature of semi-analytic models means they cannot capture baryonic dynamics explicitly. This has driven the parallel development of simulations which model hydrodynamical forces in addition to gravitational ones, allowing the evolution of gas to be followed self-consistently. These simulations have employed several numerical techniques, including both particle-based methods e. g. smoothed particle hydrodynamics (SPH; Monaghan 1992; Springel 2010a), and techniques which employ either adaptively-refined regular grids (Kravtsov et al. 1997; Teyssier 2002; Bryan et al. 2014) or irregular meshes (Pen 1998; Springel 2010b). Regardless of the method used, early simulations struggled to produce realistic galaxies (e. g. Katz & Gunn 1991; Katz 1992; Navarro & White 1993, 1994), with the formation of excessively compact, fast-rotating disks being a particular issue (Navarro et al. 1995). The primary cause of these discrepancies was identified as being the failure of these simulations to incorporate the effects of stellar feedback (e. g. Pearce et al. 1999; Sommer-Larsen et al. 1999; Thacker & Couchman 2001), which in the real Universe is implicated in driving galactic winds that act to regulate gas inflow (Veilleux et al. 2005; review).

In subsequent work, these shortcomings have been dealt with through the inclusion of phenomenological "subgrid models", which permit the modelling of behaviour occurring below the resolution limit through methods similar to those used by semi-analytic models. In the case of supernovae feedback, a variety of such models have been developed, which involve injecting energy from supernovae into the surrounding gas in either kinetic (Navarro & White 1993; Springel & Hernquist 2003) or thermal (Kay et al. 2003; Dalla Vecchia & Schaye 2012) form, often accompanied by the temporary decoupling of particles affected by feedback from hydrodynamical forces or radiative cooling (e. g. Springel & Hernquist 2003; Okamoto et al. 2005; Stinson et al. 2006). Many other phenomena are now also included in hydrodynamical simulations via subgrid models, e. g. metal enrichment (Oppenheimer & Davé 2008; Wiersma et al. 2009b), cosmic rays (Enßlin et al. 2007; Jubelgas et al. 2008; Wadepuhl & Springel 2011), and feedback from active galactic nuclei (AGN; Booth & Schaye 2009; Rosas-Guevara et al. 2015; Weinberger et al. 2017).

The approximate nature of these implementations means that they always require calibration against observed galaxy properties, and as such, the predictive power of simulations which use them must be carefully considered. Nevertheless, recent simulations have proven capable of evolving highly realistic galaxies, both in a cosmological context (e. g. Schaye et al. 2015; Pillepich et al. 2018), and in smaller-scale ‘zoom’ simulations which follow individual haloes at higher resolution (Hopkins et al. 2014; Sawala et al. 2016; Grand et al. 2017).

1.3 Gas around galaxies: the inter- and circum-galactic media

While study of the formation and evolution of galaxies is important, they contain only a small fraction of the Universe’s baryonic matter. Most baryons are instead found in the “intergalactic medium” (IGM), a low-density gas phase filling the space between galaxies (Cen & Ostriker 1999; Davé et al. 2001; Nicastro et al. 2018). The IGM was first discovered by Gunn & Peterson (1965), via a broad absorption feature in the spectrum of a distant quasar due to scattering of Ly α photons (see Section 1.4) by diffuse intergalactic

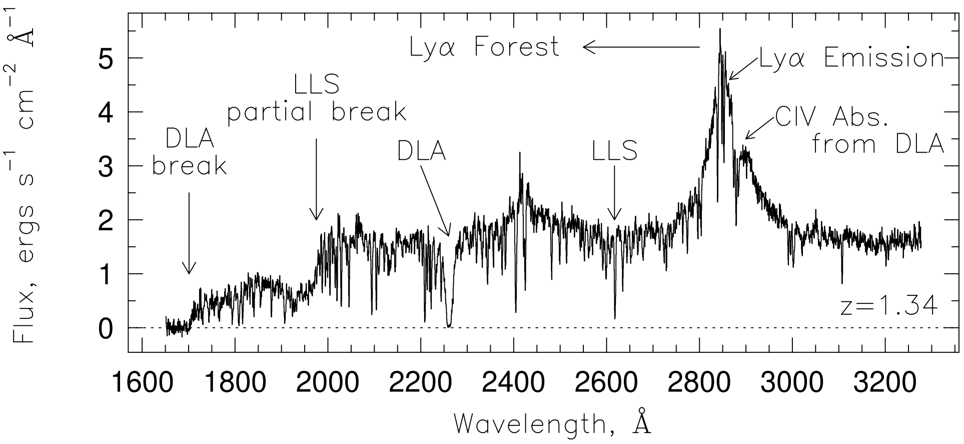

Figure 1.4
Absorption spectrum along sightline to a quasar at $z = 1.34$ with principal features annotated. Reproduced from Charlton & Churchill (2001).

gas along the line of sight. The relatively weak nature of this absorption provided early indications of the physical state of the IGM, namely that it is rarefied ($n_{\rm H} \sim 10^{-6}\,{\rm cm}^{-3}$) and highly ionised. It was soon recognised that individual structures should also give rise to discrete absorption features in QSO spectra (Bahcall & Salpeter 1965; Wagoner 1967). In time these were duly detected in the form of a dense 'forest' of Ly α absorption lines (Fig. 1.4; Lynds 1971), associated with moderately-overdense intergalactic structures located at different redshifts from the background quasar (Sargent et al. 1980; for a review see Rauch 1998).

The precise nature of these structures was initially unclear, initially being conjectured to be the result of compact clouds confined either by the pressure of the surrounding IGM (Sargent et al. 1980), or the gravity of isolated, low-mass dark matter haloes (Rees 1986). The advent of hydrodynamical simulations of cosmological volumes revised this view, by demonstrating that baryons associated with the filamentary 'cosmic web' could effectively reproduce the statistics of the Ly α forest (Cen et al. 1994; Hernquist et al. 1996; Theuns et al. 1998). Subsequent work has found continuing agreement between simulations and observational probes of the IGM, such as the Ly α forest power spectrum (Tytler et al. 2004; Viel & Haehnelt 2006; Rollinde et al. 2013) and the hydrogen column density distribution function (Pontzen et al. 2008; Altay et al. 2011, 2013).

At the typical density of the IGM, gas follows a power-law effective equation of state

(Miralda-Escudé & Rees 1994; Hui & Gnedin 1997; McQuinn & Upton Sanderbeck 2016), whose parameters are sensitive to the thermal history of the gas, in particular the timeline of its reionisation (Hui & Gnedin 1997; Furlanetto & Oh 2009; see also Section 1.5), and their measurement therefore forms another topic of interest. The IGM temperature also sets a threshold halo mass for galaxy formation, below which halo potential wells are insufficiently deep to confine gas with random thermal motion (Navarro & Steinmetz 1997; Okamoto et al. 2008). This cutoff presents a possible resolution to the 'missing satellites' problem faced by ΛCDM (Efstathiou 1992; Quinn et al. 1996; Bullock et al. 2000), and plays a crucial role in the material presented in Chapters 2 and 3.

Some absorption features in QSO spectra have origins in denser gas, shown (see Fig. 1.4) through evidence of optically thick absorption at the 912 Å Lyman limit for the eponymously-named Lyman limit systems (LLSs; Tytler 1982), or by Lorentzian absorption wings in damped Ly α systems (DLAs; Wolfe et al. 1986). Others are not associated with absorption by hydrogen atoms, but by heavier elements ("metals", e. g. Bergeron 1986; Bergeron & Boissé 1991). In both cases, these features are now identified as originating from a separate "circumgalactic medium" (CGM), composed of gas more strongly associated with galaxies. In contrast to the relatively smooth IGM, the CGM has been revealed to be a complex environment, containing both an extended hot phase traced by highly-ionised species e. g. O VI (Tripp et al. 2008; Tumlinson et al. 2011; Werk et al. 2016), and cool clouds associated with absorption by ions such as Mg II (Steidel & Sargent 1992; Churchill & Charlton 1999; Chen et al. 2010). It is interpreted as the region where where gas inflowing from the IGM, outflows of enriched gas driven by feedback, and recirculating gas 'fountains' all interact (Fig. 1.5; see also Tumlinson et al. 2017; review).

Compounding this complexity, the interpretation of observational studies of the CGM is hampered by the intrinsic rarity of fortuitous alignments between galaxies and background quasars, and by the fundamental limitation of sampling only one sightline per absorber (but see Lopez et al. 2018, 2020). Imaging the CGM directly would present an ideal alternative, but is made a daunting challenge by the intrinsic faintness of the emission. Nevertheless, some efforts have already found success (Cantalupo et al. 2014; Martin et al.

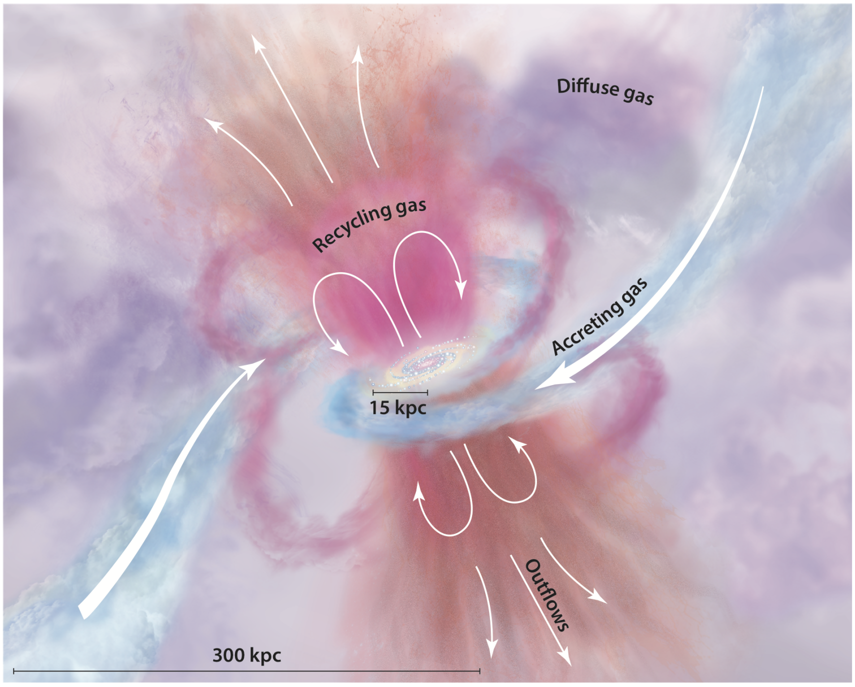

Figure 1.5
Schematic depiction of the circumgalactic medium, illustrating filamentary accretion onto galaxy disks (blue), outflows into the IGM (red), and gas recycling (pink). The surrounding gas (purple) contains a mix of material from all these sources. Reproduced from Tumlinson et al. (2017).

2015; Hayes et al. 2016). Improving modelling of the CGM has also proved challenging, with very high resolutions proving necessary in order to resolve the multi-phase nature of the CGM (Hummels et al. 2019; Peeples et al. 2019; van de Voort et al. 2019), yet precluding their use in cosmologically-representative simulation volumes at present. In the near future, the likely direction for progress instead involves developing idealised models of these phenomena (e. g. Fielding et al. 2017; Gronke & Oh 2018; Abruzzo et al. 2021), to enable their inclusion as subgrid processes in the next generation of simulations.

1.4 Interactions between radiation and matter

Interpreting many observational results, such as those of the gas in and around galaxies, requires knowledge of the physical processes affecting emission from and/or absorption in that gas. As the most abundant element in the Universe, hydrogen has particular significance. In common with all atoms, its electrons occupy quantised energy levels, whereby the nth level corresponds to a binding energy

$$\begin{aligned} E_n &= \frac{m_\mathrm{p} e^4}{8 h^2 \varepsilon_0^2} \frac{1}{n^2} \\ &= 13.6 \frac{1}{n^2}\,\mathrm{eV}\,, \end{aligned} \tag{1.4.1}$$

where m_p and e are the proton mass and electron charge respectively, h is the Planck constant, and ε_0 is the vacuum permittivity. When electrons transition between energy levels E_1 and E_2, a photon with energy $|E_1 - E_2|$ is emitted (if $E_2 < E_1$), or must be absorbed before the transition may occur (if $E_2 > E_1$). Transitions important for the work in this thesis include the previously-introduced Ly α line, which corresponds to a transition between the $n = 1$ and 2 levels; and the H α transition between $n = 2$ and 3. The rates at which these transitions occur depend on the distribution of electrons between energy levels, which are in turn set by the physical conditions within the gas, as well as any radiation fields to which the gas is exposed.

The interaction between radiation and matter is encapsulated by the equation of radiative transfer,

$$\frac{\mathrm{d}I_\nu}{\mathrm{d}l} = \frac{1}{c}\frac{\mathrm{d}I_\nu}{\mathrm{d}t} = -\kappa_\nu I_\nu + j_\nu\,. \tag{1.4.2}$$

The quantity I_ν is the specific intensity, defined such that a beam of photons in the frequency range $\mathrm{d}\nu$ which pass through a surface $\mathrm{d}A$, along directions within a solid angle $\mathrm{d}\omega$ of the normal to the surface, during a time interval $\mathrm{d}t$, carry energy equal to $I_\nu\,\mathrm{d}\nu\,\mathrm{d}\omega\,\mathrm{d}A\,\mathrm{d}t$. Similarly, the absorption coefficient κ_ν is defined such that in the volume element $\mathrm{d}V \equiv \mathrm{d}A\,\mathrm{d}\ell$, energy $\kappa_\nu I_\nu\,\mathrm{d}\nu\,\mathrm{d}\omega\,\mathrm{d}V\,\mathrm{d}t$ is absorbed from the beam, whereas the emissivity j_ν gives the energy emitted from matter into the beam as $j_\nu\,\mathrm{d}\nu\,\mathrm{d}\omega\,\mathrm{d}V\,\mathrm{d}t$. A final important quantity is the "optical depth" τ_ν, defined by

$$\mathrm{d}\tau_\nu \equiv -\kappa_\nu \mathrm{d}l\,, \tag{1.4.3}$$

which appears in the integral form of Eq. (1.4.2),

$$I_\nu(\tau_\nu(l)) = I_\nu(0) e^{-\tau_\nu(l)} + \int_0^{\tau_\nu(l)} \left.\frac{j_\nu}{\kappa_\nu}\right|_{\tau_\nu} e^{-\tau_\nu}\,\mathrm{d}\tau_\nu\,. \tag{1.4.4}$$

Here, $I_\nu(0)$ is the mean intensity incident on the absorbing region, and $\tau_\nu(l)$ is the total optical depth along the sightline to the observer (Spitzer 1978; Rybicki & Lightman 1991). Modelling the transport of radiation is always reducible to solving one of Eqs. (1.4.2) and (1.4.4), but the appearance of I_ν on both sides of the equation necessitates an iterative solution. The functional forms of κ_ν and j_ν may also be very complex: for any given transition, they will depend on the particle densities associated with the initial and final states; while the allowable transitions and their relative probabilities are set by the underlying atomic physics.

To compute the density n_X for particles of species X, the reaction rates for all processes that produce and destroy the species must first be determined. Repeating this for all species yields a system of equations in $\mathrm{d}n_X/\mathrm{d}t$, which may be solved to calculate the individual densities. The set of processes that must be included depends on the problem being solved, but in general terms, they will include collisional processes involving direct interactions between atoms and/or electrons; and radiative ones involving interactions with a photon. Within each category, multiple types of process exist, including excitations and de-excitations from one bound state to another; ionisations, which liberate an electron from a bound state; and recombinations, in which a charged ion captures a free electron.

In the majority of astrophysical situations, including those relevant to this thesis, radiative processes are dominant, so we will discuss these briefly, but neglect their collisional equivalents. Photoionisation occurs when a photon with sufficiently high energy interacts with an atom or ion, liberating an electron. The rate at which photoionisations of species X may occur, Γ_X, is set by the incident radiation spectrum, and is given by

$$\Gamma_X = \int_{\nu_{\mathrm{th}}}^{\infty} \frac{4\pi J_\nu}{h\nu} \sigma_X(\nu)\, \mathrm{d}\nu\,, \tag{1.4.5}$$

where ν_{th} is the threshold frequency for ionisations of X, σ_X is the corresponding photoionisation cross-section, and the specific intensity I_ν is replaced by its angle average, the mean intensity

$$J_\nu \equiv \frac{1}{4\pi} \int I_\nu\, \mathrm{d}\Omega\,. \tag{1.4.6}$$

When X is neutral hydrogen, $\nu_{\mathrm{th}} = 13.6\,\mathrm{eV}/h$, and the value of the cross-section at this frequency is $\sigma_{\nu_{\rm th}} \approx 6 \times 10^{-18}\,{\rm cm}^2$.

In the inverse process, radiative recombination, free electrons and X^+ ions combine to regenerate the species X. The rate for this process is specified by the recombination coefficient, α, with units $\rm cm^3\,s^{-1}$, such that the rate of recombinations per unit volume is given by the product $n_{X^+} n_e \alpha$. Following a recombination occurring, the atom is typically left in an excited state, resulting in a cascade of de-excitation transitions. For hydrogen recombination, an important consideration is the treatment of so-called resonant transitions (those ending at $n = 1$), because of the typically large population of the ground state. The optical depth (Eq. 1.4.3) can be used to establish two limiting cases: when $\tau_\nu \ll 1$, the absorbing material is said to be optically thin and all resonant line photons are assumed to escape the absorbing region; this scenario is referred to as "Case A" (e. g. Osterbrock & Ferland 2006). The opposite situation, "Case B", occurs when $\tau_\nu \gg 1$: in this case the absorber is optically thick to resonance photons, which are assumed to be absorbed on the spot. Under Case B conditions, recombinations to the ground state are effectively ignored, since the photon they produce will immediately cause a compensating ionisation.

In addition to the particle densities, the gas temperature also plays an important role in determining transition rates. In turn, the temperature is set by a balance between heating and cooling processes. The primary contribution to the heating rate is again a radiative process, specifically the gain of thermal energy by gas due to ionisations by photons with energy greater than the ionisation threshold. For the species X, this photoheating occurs at a rate given by

$$\mathcal{H}(X) = n_X \int_{\nu_{\rm th}}^{\infty} \frac{4\pi J_\nu}{h\nu} h(\nu - \nu_{\rm th}) \sigma_X(\nu)\,{\rm d}\nu\,. \tag{1.4.7}$$

Finally, the cooling rate Λ may have significant contributions from both collisional and radiative processes, with their relative importance being a function of temperature and the species for which the cooling rate is being calculated. For hydrogen under typical conditions ($T \sim 10^4\,{\rm K}$, $n \lesssim 0.1\,{\rm cm}^{-3}$), the gas will be mostly ionised, and collisional processes dominate. The cooling rate is additionally sensitive to metal enrichment, as several relatively abundant elements readily undergo collisional excitation in typical conditions, leading to a net energy loss via subsequent radiative de-excitation (Spitzer 1978, see also e. g. Wiersma et al. 2009a).

By way of summary, the solution of a radiative transfer problem can be broken down into the following schematic steps:

(i) For all species under consideration, compute transition rates for all relevant physical processes, given the current temperature, density, and specific intensity.

(ii) Form system of equations in $\mathrm{d}n_X/\mathrm{d}t$ from transition rates, and solve to determine new level populations.

(iii) Use level populations to determine j_ν and κ_ν, and hence solve the transfer equation Eq. (1.4.2) to obtain a new estimate for I_ν.

This procedure must then be iterated until a convergence criterion, e. g. in the predicted level populations, is met. Frequently, simplifications can be made by assuming that the system is in thermal and/or ionisation equilibrium, implying that the heating and cooling rates, and transition rates into and out of each state, are respectively equal.

1.5 The ultraviolet background

An important application of radiative processes is in computing the diffuse ionising ultraviolet background (UVB) which pervades the Universe. This is a necessary input for simulations employing radiative cooling prescriptions (Cen & Ostriker 1992; Wiersma et al. 2009a), and is also used to determine the ionisation correction factors necessary for interpreting certain observations (e. g. Fumagalli et al. 2016; Prochaska et al. 2017). In this case, we consider an expanding FLRW spacetime such that the transfer equation (Eq. 1.4.2) requires modification, becoming (Peebles 1993)

$$\left(\frac{\partial}{\partial t} - \nu H \frac{\partial}{\partial \nu}\right) J_\nu + 3HJ_\nu = -c\kappa_\nu J_\nu + cj_\nu\,, \qquad (1.5.1)$$

where the additional terms represent cosmological redshifting of photons and their dilution by expansion. An isotropic radiation field has also been assumed, such that $I_\nu = J_\nu$.

The integral version of the transfer equation likewise has a different form (Haardt & Madau 1996; Faucher-Giguère et al. 2009), giving the mean intensity measured by an observer with redshift z_0 at a received frequency ν_0 as

$$J_{\nu_0}(z_0) = c \int_{z_0}^{\infty} \frac{\mathrm{d}t}{\mathrm{d}z} \frac{(1+z_0)^3}{(1+z)^3} \langle j_\nu(z) \rangle \exp\left[-\bar{\tau}(\nu_0, z_0, z)\right] \, \mathrm{d}z \,, \tag{1.5.2}$$

where $\nu = \nu_0(1+z)/(1+z_0)$ and $\mathrm{d}t/\mathrm{d}z = 1/[(1+z)H(z)]$. The "effective optical depth" $\bar{\tau}$ is defined by $e^{-\bar{\tau}} \equiv \langle e^{-\tau} \rangle$, such that it represents an average of the optical depths along all lines of sight, as does the proper emissivity $\langle j_\nu(z) \rangle$. These averages are required because we are no longer considering absorption and emission in a continuous medium, but by individual macroscopic absorbers and point source emitters.

Solution of Eq. (1.5.2) therefore requires the distributions of sources and absorbers to be specified. The absorbers are the structures in the IGM and CGM discussed in Section 1.3, whose abundance is given by the column density distribution function

$$f(N_{\mathrm{HI}}, z) \equiv \frac{\partial^2 N_{\mathrm{abs}}}{\partial N_{\mathrm{HI}} \partial z} \tag{1.5.3}$$

which specifies the number of absorbers N_{abs} per unit column density N_{HI} and redshift. For the purposes of computing the UVB, $f(N_{\mathrm{HI}}, z)$ is typically parameterised as a piecewise power law fit to observations. Meanwhile, the sources considered include AGN, star-forming galaxies, and diffuse recombination emission along the line of sight, for each of which the variation of emissivity with redshift must be defined.

With prescriptions for these in place, Eq. (1.5.2) may be evaluated for a range of frequencies and redshifts to produce a model UVB spectrum $J(\nu, z)$. A number of authors have constructed such models (Miralda-Escude & Ostriker 1990; Giroux & Shapiro 1996; Fardal et al. 1998; Faucher-Giguère et al. 2009), with a series of models by Haardt & Madau (1996, 2001, 2012; see also Madau & Haardt 2015; Puchwein et al. 2019) seeing especially wide use.

These models are typically in good agreement, especially for the redshift range $2 < z < 5$, where observations provide reliable constraints on the state of the IGM. At higher z, the major open questions relate to the process of reionisation (Barkana & Loeb 2001, 2007), whereby the first ionising sources transformed the cold, neutral IGM that resulted from recombination at $z \sim 1100$ into a hot, ionised one by $z \sim 6$ (Becker et al. 2015; McGreer et al. 2015). The general outline of this process is well-understood: these sources produce intense UV radiation, which excavates ionised "bubbles" in the surrounding neutral medium. As reionisation proceeds, bubbles expand outward from neighbouring overdensities until they begin to overlap, ultimately leading to the production of a volume-filling ionised medium (e. g. Arons & Wingert 1972; Shapiro & Giroux 1987; Meiksin & Madau 1993). This phase transition marks the point at which the Universe became transparent to ionising radiation on large scales, and hence at which the UVB first reached a homogeneous equilibrium state. Furthermore, it is the event that enabled the use of distant background sources to study intervening structures in absorption, as was discussed in Section 1.3.

Uncertainties around reionisation are centred on the sources responsible, with Population III (metal-free; Bromm & Larson 2004) star formation in the first galaxies being the primary candidate. Steeply-declining high-z luminosity functions imply the majority of these galaxies are low-mass objects that remain too dim to observe today (Bouwens et al. 2015; Finkelstein et al. 2015). Meanwhile, the escape fraction f_{esc}, which quantifies the proportion of ionising photons able to escape from the galaxies in which they are produced, remains highly uncertain despite substantial observational and theoretical efforts (see Dayal & Ferrara 2018, §7.1).

At lower redshift, determining the UVB is instead challenging because widely-studied spectral lines, such as the hydrogen Lyman series, shift into the ultraviolet, necessitating space-based observations. Meanwhile, a greater fraction of the IGM is expected to lie in a hot phase, which remains challenging to detect (Bregman 2007). These difficulties have resulted in uncertainty in the amplitude and spectral shape of the $z = 0$ UVB. In particular, the Haardt & Madau (2012) prediction for $J_\nu(z = 0)$ has been found to represent an underestimate by a factor of five relative to the mean intensity required to obtain simulated

Ly α forest statistics in agreement with observations (Kollmeier et al. 2014). This "photon underproduction crisis" has been confirmed in subsequent work (Viel et al. 2017), although its severity has been reduced by appealing to increased galaxy escape fractions (Shull et al. 2015) or a greater contribution to the UVB from quasars (Khaire & Srianand 2015; Gurvich et al. 2017). Indeed, in more recent UVB models which incorporate updated prescriptions for source emissivities, the crisis has been shown to be alleviated (Puchwein et al. 2019; Faucher-Giguère 2020).

A valuable independent test of the $z = 0$ H I photoionisation rate, $\Gamma_{\rm HI}$, can be provided by observing so-called "fluorescent" emission, the result of radiative de-excitation of atoms previously ionised by the UVB. H α emission provides an ideal candidate for detection, with the non-resonant nature of this line simplifying modelling relative to Ly α, while its rest-frame optical wavelength $\lambda_{\rm H\alpha} = 6563$Å makes it accessible to observations from the ground. Detections of H α fluorescence in local H I clouds have yielded a series of upper limits on $J_\nu(z = 0)$ and/or $\Gamma_{\rm HI}$ (Vogel et al. 1995; Donahue et al. 1995; Weymann et al. 2001), as have equivalent studies of the edges of galaxies' H I disks (Maloney 1993; Dove & Shull 1994; Bland-Hawthorn et al. 1997; Ćirković et al. 1999; Madsen et al. 2001).

More recently, Adams et al. (2011) identified the nearby spiral galaxy UGC7321 as a suitable target for the detection of H α fluorescence. Subsequently, Fumagalli et al. (2017a) conducted this measurement using the MUSE (Multi-Unit Spectrographic Explorer; Bacon et al. 2010) instrument at the VLT (Very Large Telescope), obtaining a constraint on the H α surface brightness $\Sigma_{\rm H\alpha} = (1.2 \pm 0.1 \pm 0.5) \times 10^{-19}$ erg s^{-1} cm^{-2} arcsec^{-2}, where the two errors indicate statistical and systematic uncertainties respectively.

In common with previous studies, substantial uncertainties are incurred when translating this constraint to a measurement of $\Gamma_{\rm HI}$. Photoionisation modelling is required to obtain the run of $\Sigma_{\rm H\alpha}$ with $\Gamma_{\rm HI}$, which requires assumptions regarding the orientation and density distribution of the galaxy's H I disk to be made. Furthermore, the presence of nearby sources of ionising radiation has the potential to bias the measurement of $\Gamma_{\rm HI}$ high due to a "proximity effect", whereby the effective UVB is boosted above the background level by the local sources. Investigating the potential for conducting H α fluorescence imaging on a new class of objects whose physical properties circumvent these issues forms a major focus of this thesis, and is the subject of the following two chapters.

1.6 Thesis outline

In this introduction, we have reviewed the development of modern cosmology, and our understanding of the formation of cosmic structure within the Universe it describes. We have discussed the extended, diffuse structures spanning the voids between galaxies, which are home to the majority of the Universe's baryonic matter. We have presented the fundamentals of the theory underpinning interactions between matter and radiation, and highlighted the computation of the ionising ultraviolet background across cosmic time as an important application of it. The work presented in this thesis brings together elements from all these topics.

In Chapter 2, we introduce Reionisation-Limited H I Clouds (RELHICs), low-mass dark matter haloes that are an inevitable prediction of structure formation in an ΛCDM cosmology, as demonstrated by numerical simulations. If detected, these objects would verify the extension of structure formation into the baryonic domain on an as-yet untested scale. Through the use of radiative transfer modelling, we will show that they also form an ideal environment to conduct measurements of the UVB via H α fluorescence, albeit one that presents substantial practical challenges.

In Chapter 3, we extend our discussion of RELHICs to also consider the detection of fluorescent emission in helium lines, from which a valuable independent constraint on the primordial abundance of helium may be made. Furthermore, we will show that combining measurements of H α and He emission from RELHICs presents the potential to constrain not just the overall amplitude of the $z = 0$ UVB, but its spectral shape as well.

In Chapter 4 we shift our focus to the damped Ly α absorbers at $z \sim 3$. We exploit the large volume and high spatial resolution offered by a novel suite of hydrodynamical simulations, combined with further radiative transfer modelling, to make detailed predictions for their distribution and its sensitivity to modelling assumptions. We also explore how the global absorption statistics are built up from the properties of the individual absorbers.

Finally, in Chapter 5 we provide a closing summary of the work presented here, and discuss potential directions for future work.

CHAPTER 2

Fluorescent rings in star-free dark matter haloes

2.1 Introduction

A fundamental prediction of the cold dark matter (CDM) paradigm is the existence of a large number of dwarf haloes (those with mass $\lesssim 10^{10}\,\mathrm{M}_{\odot}$) resulting from the hierarchical nature of structure formation. To achieve consistency with observational constraints, it is necessary that only a small fraction of these low-mass haloes host luminous galaxies (e.g. Klypin et al. 1999; Moore et al. 1999). This requirement can be fulfilled by the effects of baryonic feedback processes, as demonstrated by hydrodynamical simulations (e.g. Babul & Rees 1992; Efstathiou 1992). Specifically, cosmic reionisation heats baryonic matter to $\sim 10^4$ K (see Meiksin 2009 for a review), inhibiting baryons from condensing into low-mass haloes to form stars. Furthermore, gas which does accrete can later be expelled by external ram-pressure stripping as the halo moves through the intergalactic medium (Benítez-Llambay et al. 2013).

Discovery of these star-free 'dark' haloes would therefore provide convincing evidence in support of CDM, and a number of methods for detecting them have been proposed. The presence of dark haloes may be inferred gravitationally, either via their dynamical influence on other objects (Erkal & Belokurov 2015; Feldmann & Spolyar 2015), or via gravitational lensing (Vegetti et al. 2010; Hezaveh et al. 2016). Alternatively, one may exploit the fact that the haloes do contain a reservoir of almost pristine gas originating from the limited accretion that occurred before reionisation. Observing this gas either in absorption against a luminous background source (Rees 1986), or directly in emission, could allow the presence of the dark haloes to be inferred.

Previously, Sternberg et al. (2002; hereafter S02) considered the emission properties of the gas bound to dark haloes, in the context of comparing them with observed high-velocity clouds (HVCs). This work concluded that for HVCs to be dark halo candidates, they must be 'circumgalactic' objects located relatively close to the Milky Way, and pressure-confined by a hot galactic corona. More recently, Benítez-Llambay et al. (2017; hereafter BL17) used the APOSTLE suite of Local Group hydrodynamical simulations (Sawala et al. 2016) to model the state of gas within dark haloes. This work identified two populations: 'COSWEBs', haloes which are found to lose almost all their residual gas after reionisation due to ram pressure stripping; and 'RELHICs', generally more isolated systems which are not subject to this additional effect, and thus have relatively constant baryonic content since reionisation.[1] RELHICs with sufficiently high halo masses were found to develop an approximately kiloparsec-sized neutral core, with the surrounding gas remaining ionised by the diffuse ultraviolet background (UVB).

Whereas BL17 considered the prospect of detecting H I 21 cm emission from these objects, in this work we follow S02 by investigating a second potential source of emission: the fluorescent H α line, in order to reassess its expected properties in light of the 'extragalactic' environment for RELHICs favoured by APOSTLE. Fluorescent emission occurs when atoms ionised by the UVB recombine to produce neutral atoms in excited states, which subsequently undergo radiative cascades to return to the ground state. Although the $n = 2 \rightarrow 1$ Ly α transition is intrinsically brighter, we consider the H α transition ($n = 3 \rightarrow 2$) due to its non-resonant nature and rest-frame optical wavelength, which are especially advantageous for the ground detection of low-redshift sources.

In this work, we perform photoionisation simulations of the gas within dark haloes in the local Universe. We calculate radial profiles of the H α surface brightness, and investigate its

---

[1] These acronyms stand for "COSmic WEB Stripped systems" and "Reionisation-Limited H I Clouds" respectively.

dependence on physical properties of the haloes and the UVB. We use an adapted version of the code described by Cooke & Pettini (2016; hereafter CP16), which follows the procedure outlined in S02. Our calculations are qualitatively similar to those performed by the CLOUDY spectral synthesis code (Ferland et al. 2017), but with two important additional features. Firstly, we model gas in hydrostatic equilibrium with a dark matter halo to obtain realistic gas density profiles. Secondly, we perform the calculations in (projected) spherical coordinates, in order to capture the geometric dependence of optical depths within the cloud.

This chapter is organised as follows: in Section 2.2, we describe our numerical method. We use this method in Section 2.3 to compute grids of models with a variety of input parameters to determine the resulting radial surface brightness profiles. We then combine our results with the APOSTLE simulations in Section 2.4 to predict the abundance of H α-fluorescent RELHICs on the sky, and the prospects for their detection. Lastly, in Section 2.5 we present a discussion of our results and their significance for constraining the structure and environment of a fluorescent halo, based on its observed H α emission. Throughout, we assume a set of cosmological parameters ($H_0 = 67.3\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$, $\Omega_\Lambda = 0.685$, $\Omega_\mathrm{M} = 0.315$, $\Omega_\mathrm{B} = 0.0491$) consistent with *Planck* measurements (Planck Collaboration et al. 2014).

2.2 Simulation design

We use a modified version of the spherically-symmetric ionisation balance code described in CP16. A full description of this code is given therein; here we will provide a brief summary and describe the modifications we have made.

2.2.1 Code outline

We consider clouds of primordial gas, consisting only of hydrogen and helium with abundance ratio by number $n_\mathrm{He}/n_\mathrm{H} = 0.083$ (corresponding to a primordial helium mass fraction $Y_{\rm P} = 0.24$). This gas is embedded within a spherically-symmetric Navarro-Frenk-White (NFW; Navarro et al. 1996b) dark matter halo, with radial density profile given by

$$\rho(r) = \frac{\rho_s}{x(1+x)^2}\,; \quad x \equiv \frac{r}{r_s}\,, \tag{2.2.1}$$

where r_s and ρ_s are characteristic length and density scales of the halo. The gas is assumed to be in hydrostatic equilibrium with a potential $\varphi(r)$, such that $\mathrm{d}P(r) = -\rho_{\rm g}(r)\mathrm{d}\varphi(r)$, where $P(r)$ and $\rho_{\rm g}(r)$ are the gas pressure and density profiles respectively. For simplicity, we neglect the self-gravity of the gas such that $\varphi(r)$ is defined solely by the dark matter halo, which is a reasonable assumption for the DM-dominated systems we consider.

Haloes are constructed by choosing a virial mass M_{200}, defined as the mass contained in a sphere with average density $200\rho_{\rm crit}$, where $\rho_{\rm crit}$ is the critical density of the Universe. The radius of this sphere is the virial radius R_{200}, such that

$$M_{200} = \frac{4\pi}{3} {R_{200}}^3 \times 200\rho_{\rm crit}\,. \tag{2.2.2}$$

NFW haloes may be parameterised by a concentration parameter $c_{200} \equiv R_{200}/r_s$. In our fiducial models, we determine values of c_{200} as a function of M_{200} using the mass-concentration relation of Ludlow et al. (2016). We investigate alternative mass-concentration relations and density profiles in Section 2.3.2.2 in order to determine the dependence of our results on these assumptions. A total gas mass $\mathrm{M_g}$ is then assigned to the halo using the analytic model employed by BL17, which determines the gas mass required for hydrostatic equilibrium as a function of M_{200}.

Given the assumed absence of local star formation, the gas is illuminated solely by photons from the diffuse UVB. We adopt the Madau & Haardt (2015; hereafter MH15) UVB in our fiducial model, although we explore the effect of varying the shape of this spectrum in Section 2.3.2.1. Under the assumption of spherical symmetry, the UVB irradiates the gas isotropically. However, the gas column density has an angular dependence which varies with position, meaning that the local background radiation intensity, and hence the local photoionisation rate, is a function of both depth within the cloud and direction of the incident radiation.

In addition to photoionisation, we consider primary and secondary collisional ionisation, ionisation resulting from helium recombination radiation, and charge transfer ionisation. We include radiative, dielectronic and charge transfer recombinations. Ionisation equilibrium for each species is enforced by equating the appropriate ionisation and recombination rates in each case. The rate coefficients assumed for these reactions are unmodified from those detailed in CP16, with the exception of the rates for collisional ionisation, which we take from the CHIANTI atomic database (Dere et al. 1997; Zanna et al. 2015) in order to enable calculation of the rates at temperatures $\sim 10^3$ K, which may be reached in the neutral core (the lower limit in the previously-used Dere (2007) rates was instead $\simeq 10^4$ K). The version of the code described in CP16 additionally assumed strict thermodynamic equilibrium, such that the temperature at each radial coordinate is that for which the computed heating and cooling rates are equal. The heating rate includes both photoionisation heating and secondary heating by primary photoelectrons, while the cooling rate incorporates contributions from collisional excitation/ionisation cooling, recombination cooling, Brehmsstrahlung cooling and Compton cooling/heating. We take the formulae used to determine each of these contributions from Cen (1992). At low gas densities ($n_\mathrm{H} \lesssim 10^{-4.8}\,\mathrm{cm}^{-3}$), cooling becomes inefficient and the gas temperature will instead be that which results in a heating timescale equal to the Hubble time. We incorporate this additional regime by switching from calculating the equilibrium temperature to using the tabulated temperature-density relations given in Wiersma et al. (2009a) when the gas density falls below this transition point.

The form of the calculations follows that described by S02: assuming the gas is initially isothermal and fully-ionised, we compute the gas pressure profile required to maintain hydrostatic equilibrium. A boundary condition is required to determine the pressure profile; in common with BL17, we use the assumption that at large distances from the centre of the halo ($\simeq 100\,R_{200}$), the density should approach the cosmic mean baryon (number) density $\bar{n} \simeq 10^{-6.7}\mathrm{cm}^{-3}$ at $z = 0$.

From the computed pressure profile, we determine the radial density distribution, and hence the intensity of the radiation field within the cloud. We then solve for ionisation equilibrium to calculate the neutral fractions of each species as a function of radius. The temperature profile is then derived as described above, and used to recompute the pressure profile. This process is iterated until a convergence criterion is met, requiring that the relative change in the computed neutral fractions between successive iterations is no more than 0.1% at any radial coordinate.

Once a converged ionisation structure has been found, the H α surface brightness $\Sigma_{\mathrm{H}\alpha}$ may be computed. At a radial position within the halo r, H α photons are produced at a rate given by the volume emissivity $\varepsilon_{\mathrm{H}\alpha} = n_{\mathrm{HII}}(r) n_e(r) \alpha_{\mathrm{H}\alpha}(r)$, where n_{HII} and n_e are the densities of ionised hydrogen and electrons respectively. The effective recombination coefficient $\alpha_{\mathrm{H}\alpha}$ expresses the rate per unit electron and ion densities at which the $n = 3 \longrightarrow 2$ transition occurs. We obtain values for this coefficient, as a function of temperature, from Osterbrock & Ferland (2006).

For a line of sight intersecting the halo with impact parameter b, the projected surface brightness is found as the line integral of the emissivity along the line of sight

$$\Sigma_{\mathrm{H}\alpha}(b) = \frac{h\nu_{\mathrm{H}\alpha}}{2\pi} \int_b^{R_{200}} \frac{r}{\sqrt{r^2 - b^2}} n_{\mathrm{HII}}(r) n_e(r) \alpha_{\mathrm{H}\alpha}(r) \, \mathrm{d}r \, , \tag{2.2.3}$$

where h is Planck's constant and $\nu_{\mathrm{H}\alpha} = 4.57 \times 10^{14}$ Hz is the frequency of an H α photon.

2.2.2 Tests of numerical accuracy

To validate our code, we model a series of plane-parallel gas slabs, varying the volume and total column densities of the slab. For each slab, we initially specify a constant temperature $T_{\mathrm{gas}} = 10^4$ K. In this case the ionisation structure of the slab may be predicted analytically; in particular, the surface brightness is expected to approach a constant value, in a similar manner to the result for Ly α fluorescence reported by Gould & Weinberg (1996; see Appendix A.1). The use of a plane-parallel geometry also enables us to compare our results with those obtained from the spectral synthesis code CLOUDY (Ferland et al. 2017).

Fig. 2.1 demonstrates that in the high-column density limit, the code presented here reproduces the analytic prediction well, with the only deviation occurring at high volume

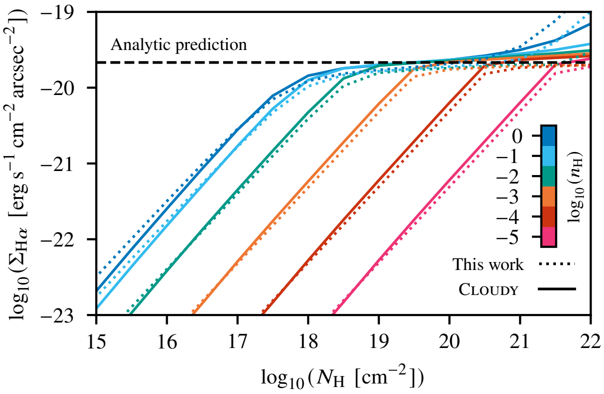

Figure 2.1

H α surface brightness as a function of total H column density for isothermal, H-only plane parallel models. Dotted curves are obtained from the code we describe here, solid ones from CLOUDY; each same-coloured pair corresponds to a particular constant gas density n_H. The horizontal dashed line indicates the analytic result for H α surface brightness in the photoionised, optically-thick limit.

densities where collisional ionisation becomes significant. Note that we make a number of simplifications as compared to CLOUDY: we assume case B recombination, rather than self-consistently computing the diffuse radiation field produced by recombinations within the cloud, and our treatment of collisional processes is less sophisticated than the one in CLOUDY. We further explore the effects of these approximations in Appendix A.2.

In addition to verifying the ionisation balance portion of our code, we compare the physical properties of our model RELHICs with the analytic model of BL17. In Fig. 2.2, we compare radial density and temperature profiles obtained from this model with corresponding ionisation balance calculations, for RELHICs with a range of halo masses M_{200}. We obtain good agreement for low-mass haloes and the outer regions of more massive ones, where the gas remains fully ionised. For haloes which are sufficiently massive to develop an ionisation front, our code predicts higher densities and lower temperatures within the neutral region. This occurs because in the BL17 model, a constant value of the mean mass per particle $\mu = 0.62$ is used, as appropriate for fully-ionised primordial gas. In this work, we explicitly calculate the ionisation state of the gas as a function of radius, and so can determine μ as

$$\mu(r) = \frac{\sum_i X_i(r) A_i(r)}{\sum_i X_i(r) + X_e(r)}, \tag{2.2.4}$$

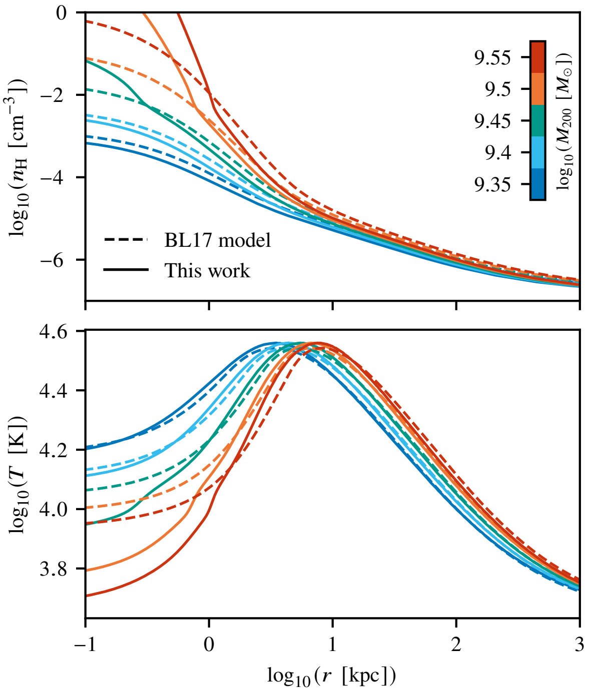

Figure 2.2
Radial density (top) and temperature (bottom) profiles from BL17 model (dashed lines) and our code (solid lines). The models agree well, except for the inner regions of high-mass haloes, where the transition to neutral gas changes the mean mass per particle μ, an effect which is not included in the BL17 model.

where X_i is the number abundance of species i relative to H, A_i is the mass number of species i, and the e subscript denotes electrons. For fully neutral gas, Eq. 2.2.4 yields $\mu = 1.23$, resulting in higher pressures which require higher gas densities to maintain hydrostatic equilibrium. Since the cooling rate is an increasing function of density, the equilibrium temperature for neutral gas is consequently lower.

2.2.3 Star formation threshold

In the following section, we will show that more massive dark matter haloes will retain a more massive reservoir of gas; correspondingly, they will have larger neutral cores and higher peak values of $\Sigma_{\mathrm{H}\alpha}$. However, these cores will be denser and colder, increasing the likelihood of star formation occurring. UV radiation emitted locally from young stars within the halo will readily exceed the intensity of the UVB, while their supernovae would expel gas from the halo. Hence, we require that the total amount of star formation is small, which imposes an upper limit on the halo masses for which fluorescent RELHICs are expected to form. Above this threshold, the objects that form will be more accurately categorised as ultra-faint dwarf galaxies.

Star formation primarily occurs within dense molecular clouds, and so the presence of H_2 may be used as a proxy for star formation. Ideally, we would self-consistently determine the molecular fraction within our code; however, the complexity of doing this puts it beyond the scope of this work. Consequently, we instead consider the formation of H_2 as a post-processing step. Analytic models, such as that described by Krumholz et al. (2009b), permit this by prescribing the surface molecular fraction $f_{H_2} \equiv \Sigma_{H_2}/(\Sigma_{HI} + \Sigma_{H_2})$ as a function of the column density of H I and an estimate of the H_2-dissociating Lyman-Werner flux.[2] We note that this model is primarily intended for disk galaxies with active star formation and higher metallicity than the gas we consider (which we assume to have primordial composition, although for the purposes of computing f_{H_2} we must assume a non-zero metallicity $Z = 0.001Z_\odot$). While the exact values of f_{H_2} we obtain may not be fully accurate, we expect the column density threshold for H_2 formation to be reliable. This threshold is predicted by the model to occur at $N_{HI} \sim 10^{23}\,\mathrm{cm}^{-2}$ which, when the absence of metals or dust is taken into consideration, is in good agreement with the observed paucity of damped Ly α (DLA) systems above a similar column density threshold, believed to be due to the onset of H_2 formation (e.g. Schaye 2001; Krumholz et al. 2009c). Additionally, the upper bound on halo mass of $M_{200} \lesssim 10^{9.6}\,\mathrm{M}_\odot$ implied by this threshold is consistent with the fraction of 'dark' haloes found in the APOSTLE hydrodynamical simulations, which exceeds 50% for halo masses below this value (Benítez-Llambay et al. 2017). Accordingly, excluding all haloes for which $N_{HI} \geq 10^{23}\,\mathrm{cm}^{-2}$ provides a threshold below which we may discount star formation.

An additional upper limit can arise from the constraint that the total gas fraction M_{gas}/M_{200} remains small, such that we are justified in neglecting the gas self-gravity. At the halo masses required for this constraint to become relevant, we find that column densities already

---

[2] More advanced models for f_{H_2} exist, e.g. Krumholz (2013). These model the formation of H_2 at low densities more accurately, but give the same results in the high-density regime that is relevant to our purposes.

exceed the limit for molecular gas formation defined above. However, the steep increase in gas density within the ionisation front raises the possibility of gas becoming locally self-gravitating within the neutral core. Adopting the stricter condition $M_{\rm gas}(\leq r)/M_{\rm DM}(\leq r) \ll 1\, \forall\, r < R_{200}$ to avoid this possibility would result in a small ($\lesssim 0.005$ dex) reduction in our upper threshold on halo mass. Nevertheless, in the absence of further work to model the chemical state of this gas, especially during early times when the greater mean density of the Universe would have yielded even steeper gas density profiles, our upper thresholds on halo mass should themselves be interpreted as upper limits.

2.3 Properties of fluorescent haloes

2.3.1 Fiducial model

We first present results from our fiducial model, in which the UVB is given by the MH15 spectrum and an NFW density profile is assumed for the dark matter halo, along with the Ludlow et al. (2016) mass-concentration relation and a redshift $z = 0$.

We compute models for RELHICs with halo masses $10^{9.38}\,{\rm M_\odot}$ and above, and show the resulting hydrogen neutral fractions $X_{\rm HI} \equiv n_{\rm HI}/n_{\rm H}$ and Hα surface brightnesses $\Sigma_{\rm H\alpha}$ in Fig. 2.3. At the lower end of this range, hydrostatic equilibrium is achieved for a bound gas mass within R_{200} of ${\rm M_{gas}} \sim 10^7\,{\rm M_\odot}$. With projected column densities $N_{\rm HI} \lesssim 10^{15}\,{\rm cm^{-2}}$, this gas is optically thin to the UVB, and hence is highly ionised throughout. The gas remains rarefied, with peak densities $n_{\rm H} \sim 10^{-3}\,{\rm cm^{-3}}$, and therefore $\Sigma_{\rm H\alpha}$ is low at all impact parameters. As the halo mass increases, the central gas density also increases. Consequently, at small impact parameters the projected column density and surface brightness gradually increase, while remaining largely unchanged along lines of sight avoiding the high-density central region. For $M_{200} = 10^{9.47}\,{\rm M_\odot}$, the central column density reaches $10^{18}\,{\rm cm^{-2}}$, corresponding to an optical depth at the Lyman limit $\tau_{912} \sim 1$. A sharp transition in the gas properties occurs as this threshold is exceeded: the gas becomes self-shielded from the UVB and develops a neutral core. By $M_{200} = 10^{9.49}\,{\rm M_\odot}$

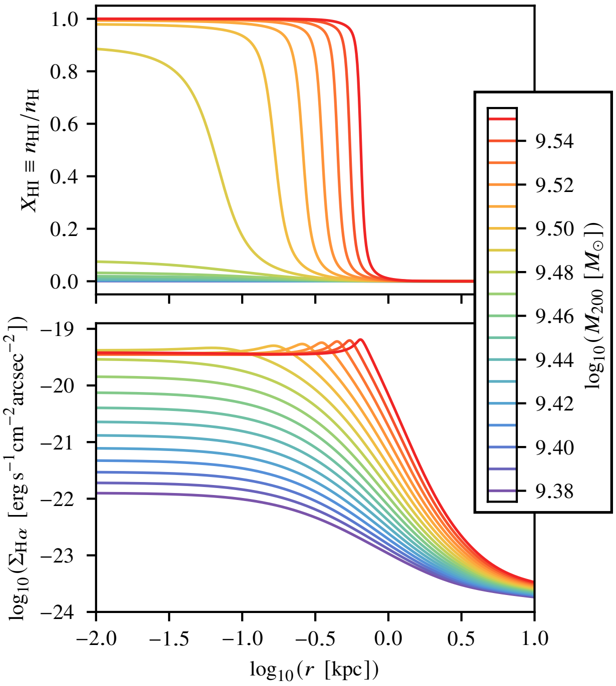

Figure 2.3

Radial profiles for hydrogen neutral fraction (top) and H α surface brightness (bottom) for haloes in our fiducial model, as a function of halo mass M_{200}. Note the sudden transition in the ionisation state of the gas at $M_{200} = 10^{9.5}\,\mathrm{M}_\odot$. This corresponds to the formation of the limb-brightened ring in emission, shown here as a peak in the $\Sigma_{\rm H\alpha}$ profile.

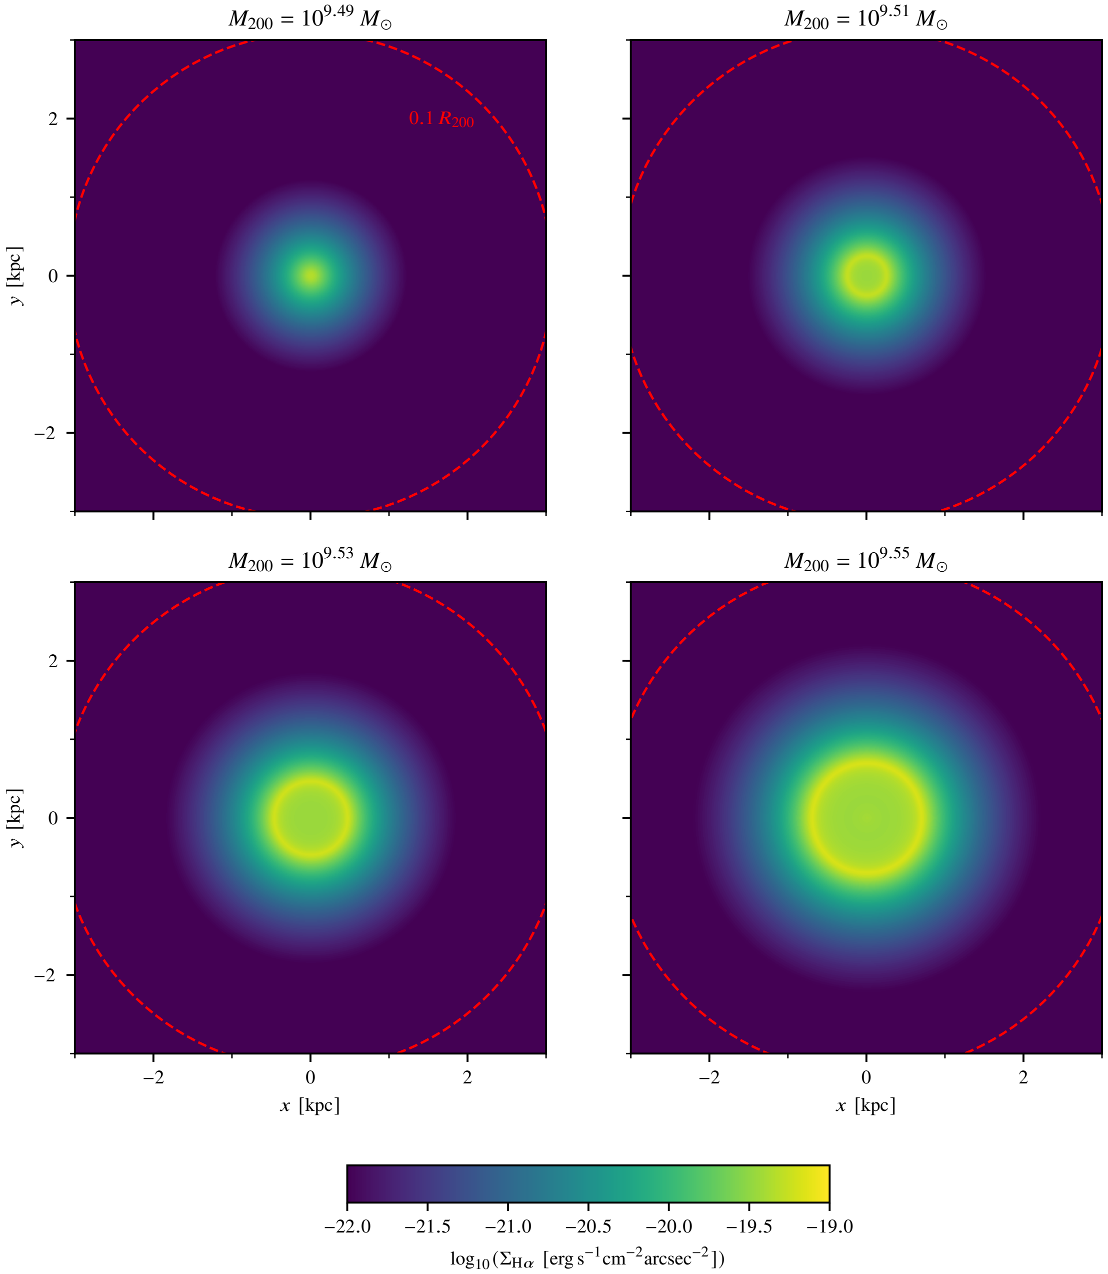

Figure 2.4

Projected H α surface brightness for haloes in our fiducial model. The range of halo masses shown is chosen such that in the top-left panel, a neutral core first begins to form. As the halo mass increases, the neutral core enlarges and the limb-brightened peak in $\Sigma_{\mathrm{H}\alpha}$ becomes more prominent until we reach the most massive halo allowed by our star formation criterion (bottom-right). Red dashed lines indicate 10% of the virial radius.

the central neutral fraction approaches 1.

Fig. 2.4 shows the projected surface brightness distributions for RELHICs at this mass and above, clearly illustrating their characteristic ring-shaped surface brightness morphology. This occurs because the H α emissivity is sharply peaked at the ionisation front: at larger radii, the gas density is reduced, while inside the ionisation front few H II recombinations occur. Hence, lines of sight tangent to the surface of the neutral core pass through a greater fraction of high-emissivity gas than those passing through the centre of the neutral region. This 'limb brightening' effect produces surface brightness profiles which peak at the projected radius of the ionisation front. As the halo mass increases further, the core grows in size (see also Fig. 2.3), and rapidly becomes cooler and denser, leading to the formation of molecular H_2, and our previously-discussed star formation criterion is exceeded for haloes above $10^{9.55}\,M_\odot$. For a halo of this mass, the surface brightness profile peaks at $r = 0.70\,\mathrm{kpc}$ with a value of $\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}} = 6.58 \times 10^{-20}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$.

Ring-shaped H α surface brightness morphologies were also predicted by S02, with a comparable peak value of $\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}} \simeq 4 \times 10^{-20}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$. However, this is attained for a UVB spectrum with hydrogen photoionisation rate ≃20% lower than the MH15 spectrum (see Section 2.3.2.1 below). Additionally, S02 adopt an external bounding pressure, which acts in conjunction with the DM potential to confine the gas toward the centre of the halo, allowing a dense, self-shielding core to form at a much lower halo mass $M_{200} \simeq 10^8\,M_\odot$.

2.3.2 Model variations

2.3.2.1 UVB spectrum

Discrete sources of ionising radiation, such as a nearby AGN or bright galaxy, can produce local inhomogeneities in the UVB. Furthermore, uncertainties present in the modelling of the UVB result in variations in the intensity and spectral shape of the background computed by different models at a fixed redshift. Accordingly, we explore the sensitivity

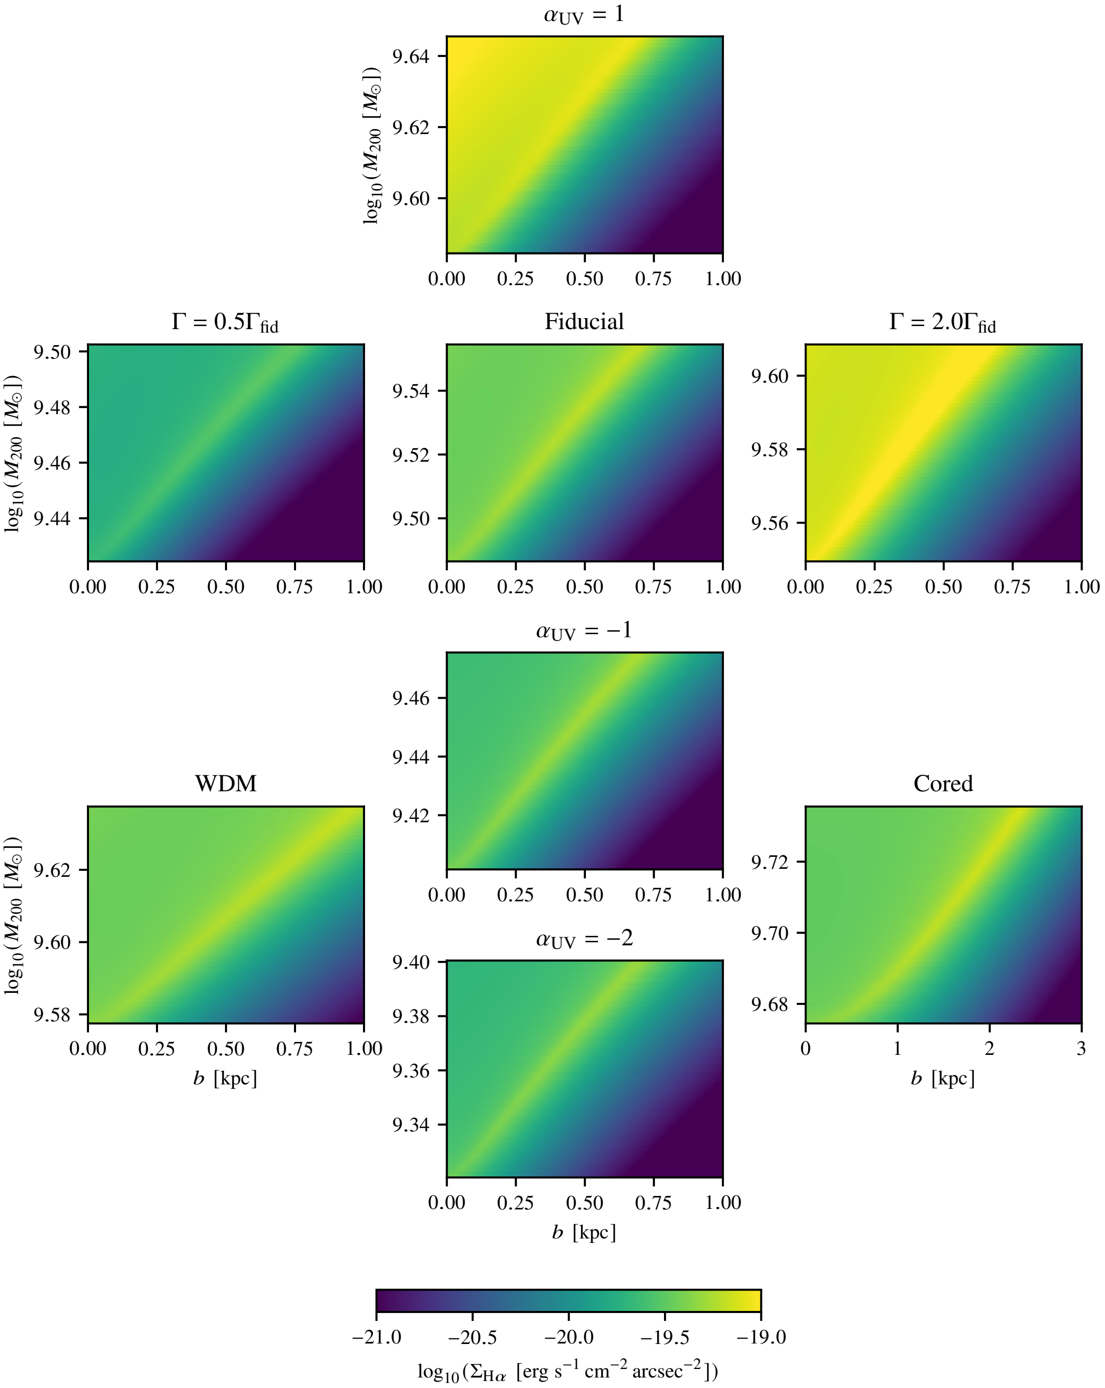

Figure 2.5

$\Sigma_{\rm H\alpha}$ as a function of impact parameter b and M_{200} for different model variations. In each panel, the y-axis limits are set by the minimum mass required for the formation of a neutral core, and the maximum mass allowed by the star formation criterion. Results from our fiducial model are shown on the upper-central panel, with panels to the left and right illustrating the effect of varying the photoionisation rate, and panels above and below that of changing the spectral index $\alpha_{\rm UV}$. Panels in the bottom left and right corners show results for models with a different halo mass-concentration relation and halo density profile (note the wider x-axis limits in the latter).

of our results to such changes in the UVB. The intensity of the UVB may be parameterised by the hydrogen photoionisation rate, which is given by

$$\Gamma = 4\pi \int_{\nu_{\rm th}}^{\infty} \frac{J(\nu)\sigma(\nu)}{h\nu}\,{\rm d}\nu\,, \tag{2.3.1}$$

where $J(\nu)$ is the mean intensity, $\sigma(\nu)$ is the cross-section for photoionisation of $\rm H^0$, and $\nu_{\rm th} = 13.6\,{\rm eV}/h$. In Fig. 2.5, we show stacked radial surface brightness profiles as a function of halo mass, with colour encoding the value of $\Sigma_{\rm H\alpha}$. The panel in the centre of the upper horizontal row corresponds to our fiducial model, in which the photoionisation rate is $\Gamma_{\rm fid} = 5.47 \times 10^{-14}\,{\rm s}^{-1}$. Panels to the left and right show results obtained with Γ respectively decreased and increased by a factor of two from its fiducial value. We find that increasing the photoionisation rate shifts the halo mass at which RELHICs are predicted to possess fluorescent rings upward, while the surface brightness of these rings also increases. Conversely, a reduced photoionisation rate allows dimmer rings to form at lower halo masses.

We interpret these trends as being due to the ability of a more intense UVB to penetrate a greater total column density of gas (i. e. $N_{\rm H}$) before becoming attenuated. In the situation we consider, the gas volume density is a function of radius, which itself becomes steeper with increasing halo mass. This results in the ionisation front forming at a larger characteristic gas number density, which is only reached for more massive haloes. In turn, the greater density of the gas at the ionisation front boosts its emissivity, and hence the peak surface brightness. We further explore the reasons for surface brightness variations between models in Section 2.5.

To consider variations in the shape of the UVB spectrum, we adopt the parameterisation from Crighton et al. (2015a), which introduces the UVB shape parameter $\alpha_{\rm UV}$. This parameter is used to modify the ionising spectrum as follows:

$$J_\nu(E) = \begin{cases} N_\Gamma \times H(E) & E \leq E_0 \\ N_\Gamma \times H(E) \times (E/E_0)^{\alpha_{\rm UV}} & E_0 < E \leq E_1 \\ N_\Gamma \times H(E) \times (E_1/E_0)^{\alpha_{\rm UV}} & E > E_1 \end{cases}\,, \tag{2.3.2}$$

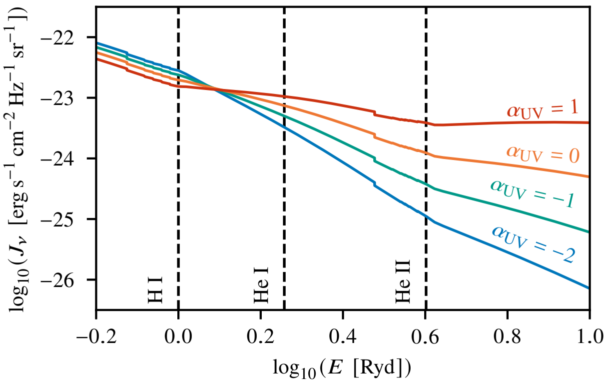

Figure 2.6

UVB spectra in the range 1 Ryd $< E <$ 10 Ryd for different values of the slope parameter α_{UV}, where $\alpha_{\mathrm{UV}} = 0$ corresponds to the fiducial MH15 spectrum. Vertical dashed lines indicate the H I, He I and He II ionisation edges.

where $H(E) \equiv J_{\nu,\mathrm{MH15}}(E)$, $E_0 = 1\,\mathrm{Ryd}$ and $E_1 = 10\,\mathrm{Ryd}$. We additionally scale these spectra by a factor $N_\Gamma \equiv \Gamma_{\mathrm{fid}}/\Gamma$, such that the photoionisation rate is unchanged from its fiducial value, in order to examine the effect of modifying the spectral shape in isolation. The resulting spectra for $\alpha_{\mathrm{UV}} = 0$ (corresponding to our fiducial model), +1, −1 and −2 are plotted in Fig. 2.6.

In Fig. 2.5, the central column shows H α surface brightness profiles as a function of M_{200} for RELHICs illuminated by the same four spectra. We find that the halo mass required to form a fluorescent ring is reduced for more negative values of α_{UV}, corresponding to softer spectra with a greater proportion of ionising photons having energies close to the Lyman limit (since the spectra have been renormalised to maintain a constant value of Γ). These photons are more likely to be absorbed, due to the $\sim(\nu/\nu_{\mathrm{th}})^{-3}$ dependence of the photoionisation cross-section, and so will on average be absorbed within a lower total gas column density. As described previously, the coupling between column and volume densities means that this corresponds to the ionisation front forming in lower-density gas, which maintains hydrostatic equilibrium with a shallower dark matter potential. In addition to the reduction in emissivity due to the lower gas density, an additional effect is caused by the reduced flux of photons with energies close to 24.6 eV, at which helium is singly ionised. This produces a larger He I fraction at the H ionisation front, further reducing the emissivity by lowering the electron density. Opposite arguments apply to the harder

$\alpha_{\rm UV} = 1$ spectrum, which produces larger peak surface brightnesses but inhibits neutral cores from forming until higher halo masses are reached.

2.3.2.2 Halo density profile

In our fiducial model, we have used an NFW density profile to model dark matter haloes, with the concentration parameter c_{200} determined using the Ludlow et al. (2016) relation, which is obtained from haloes in the Copernicus Complexio (CoCo) N-body simulations (Bose et al. 2016; Hellwing et al. 2016). The CoCo suite also includes runs in which a warm dark matter (WDM) cosmology is used. In contrast to CDM cosmologies, where c_{200} is a monotonically increasing function of M_{200}, the WDM mass-concentration relation turns over at a specific mass scale (Bose et al. 2017). H α surface brightness profiles obtained for RELHICs whose haloes follow this relation are shown in the bottom-left corner of Fig. 2.5. At the halo masses at which fluorescent rings are predicted, the WDM mass-concentration relation results in haloes which are approximately 50% less concentrated at fixed M_{200} than the equivalent halo following the L16 relation, and so have systematically lower central dark matter densities. Consequently, hydrostatic equilibrium is maintained with a shallower gas density profile, which produces lower central gas densities. Hence, the column density required to form a neutral core is first reached at higher halo mass, and once neutral cores form they extend to larger projected radii.

We also consider the possibility of varying the density profile assumed for the dark matter, by computing surface brightness profiles for haloes with ‘cored’ density profiles, where the central dark matter density approaches a finite value rather than diverging to infinity, as occurs for ‘cuspy’ profiles such as NFW. Observations of nearby dwarf galaxies, which imply central density gradients shallow than an NFW profile, have been interpreted either as a manifestation of the self-interacting nature of dark matter (e.g. Spergel & Steinhardt 2000), or as a result of baryonic interactions (Navarro et al. 1996a; Pontzen & Governato 2012; and references therein). The latter scenario generally requires supernova feedback from stars forming within the halo (e.g. Benítez-Llambay et al. 2019), and so evidence of a core in a halo hosting a fluorescent ring would be suggestive of self-interacting dark matter (SIDM). We choose to model Burkert density profiles, for which (Burkert 1995)

$$\rho(r) = \frac{\rho_0 {r_0}^3}{(r + r_0)(r^2 + {r_0}^2)}, \tag{2.3.3}$$

where r_0 is the core radius, which we choose to be a function of M_{200} such that $R_{200}/r_0 = c_{200}$, in analogy to the NFW scale radius r_s. This choice results in core radii $r_0 \sim 2-3$ kpc, somewhat larger than the typical size $r_0 \sim 1$ kpc predicted for SIDM haloes of the mass we consider, but within the range allowed by constraints on the self-interaction (see e.g. Tulin & Yu 2018). We caution that these assumptions will influence the resulting fluorescent emission; accordingly, the results we present here, shown in the bottom-right panel of Fig. 2.5, should be treated as illustrative only, rather than a specific prediction of a particular SIDM model.

As with the Bose et al. (2017) mass-concentration relation, the reduction in central dark matter density caused by the dark matter core results in larger, more extended neutral gas cores developing at higher halo mass. However, the more extreme modification to the dark matter density profile results in a much more substantial change in the size of the neutral gas core. In both these models and those using the WDM mass-concentration relation, we find the magnitude of $\Sigma_{\mathrm{H}\alpha,\mathrm{max}}$ to be slightly increased with respect to our fiducial model. This occurs because of the larger physical size of the emitting region results in a longer path length over which $\varepsilon_{\mathrm{H}\alpha}$ is integrated to obtain the surface brightness, but is a relatively minor effect when compared to the impact of varying the properties of the UVB.

2.4 Detectability of fluorescent haloes

We now make predictions for the abundance with which dark haloes exhibiting Hα fluorescence appear on the sky, and hence consider the prospects for their detection, using results from the APOSTLE suite of high-resolution hydrodynamical simulations (Sawala et al. 2016). These simulations follow twelve galaxy groups selected to resemble the Local Group (LG) on the basis of their dynamical properties; we consider the five volumes simulated at the highest ('L1') resolution corresponding to a gas (DM) particle mass $1.0\ (5.0) \times 10^4\ \mathrm{M}_{\odot}$.

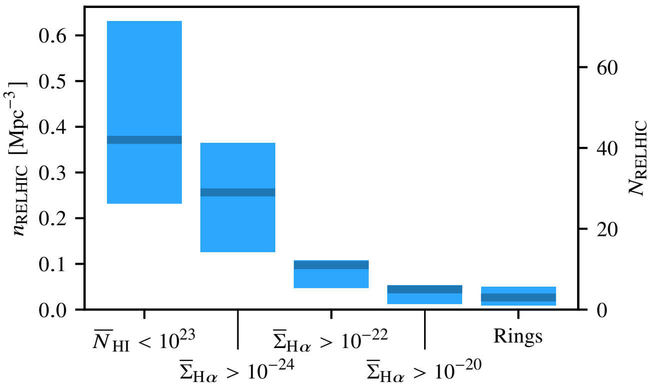

Figure 2.7

Number density distribution of H α-fluorescent RELHICs in APOSTLE. The median values are shown by thick solid lines; shaded regions indicate 10th – 90th percentiles. From left to right, we show the number density of all RELHICs with H I column densities below our star formation threshold, those with surface brightnesses exceeding three different limiting values, and those with discernible ring-shaped surface brightness profiles. Here, $\overline{N}_{\mathrm{HI}} \equiv N_{\mathrm{HI}}/[\mathrm{cm}^2]$ and $\overline{\Sigma}_{\mathrm{H}\alpha} \equiv \Sigma_{\mathrm{H}\alpha}/[\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}]$. The right-hand axis indicates the corresponding number counts per APOSTLE volume.

From the $z = 0$ outputs of these simulations, we select all haloes with $M_{200} > 10^8\,\mathrm{M}_\odot$ located within 3.5 Mpc of the simulated LG's barycentre, before discarding all 'luminous' haloes (i. e. those containing one or more star particles). The remaining 'dark' haloes are then divided into the two populations (COSWEBs and RELHICs) introduced in Section 2.1, where haloes with negligible gas content at $z = 0$ are placed into the former category and discounted.[3] From this classification, we obtain a total of 263 RELHICs across the five volumes, with median halo mass $M_{200} = 10^{9.12}\,\mathrm{M}_\odot$. We match each of these haloes with a corresponding calculation from Section 2.3, deselecting 23 objects due to their exceeding our star formation criterion (Section 2.2.3); for the remaining 240 we obtain a prediction of their peak H α surface brightness under our fiducial model. In Fig. 2.7, we show the resulting number density of all fluorescent RELHICs, and of subsets with surface brightnesses exceeding chosen threshold values.

Fig. 2.7 indicates that high surface brightness fluorescence is rare, as are ring-shaped structures, with a median value of $3^{+2.6}_{-2.0}$ objects per volume, where the sub- and superscripts indicate the 10th and 90th percentile limits. This is not unexpected, due to the very narrow

---

[3] The group finder SUBFIND (Springel et al. 2001) is used to quantify the bound gas mass.

range in halo mass for which we expect bright ring fluorescence.[4] Consequently, detecting fluorescent rings in the surroundings of the LG is impractical, since a prohibitively deep, all-sky H α survey would be required. However, searching for fluorescent RELHICs around nearby galaxies, albeit still challenging, becomes more feasible, since surface brightness remains independent of distance for cosmologically-nearby sources, but the angular extent of a volume equivalent to that simulated in Apostle decreases substantially. In this situation, ultra-deep observations along single pointings have the potential to uncover one or more florescent rings.

Additionally, the requirement for an expensive optical search may be circumvented by exploiting the large H I column densities associated with strongly-fluorescent systems. These imply that the relatively massive RELHICs which host fluorescent rings should also be bright 21 cm H I emitters, and so existing H I sources may be used to provide a set of targets for subsequent H α observations, once deep (but affordable) broad-band imaging has been used to rule out associations with galaxies. Of particular interest is the catalogue of ultra-compact high velocity clouds (UCHVCs) detected in surveys like ALFALFA (Adams et al. 2013), which BL17 demonstrate to have H I emission properties which overlap at least partially with corresponding predictions for Apostle RELHICs. This is in contrast to the (non 'ultra-compact') HVCs, given the 'extragalactic' environment for RELHICs that Apostle favours. S02 found that the latter scenario would require HVCs to be underconcentrated and unrealistically spatially extended in order to reproduce their observed angular sizes.

Throughout the preceding discussion, we have considered fluorescence from haloes at redshift $z = 0$. At earlier times, the UVB has a greater intensity, resulting in an increased photoionisation rate Γ. As detailed in Section 2.3.2.1, this will correspond to greater intrinsic surface brightnesses. However, the observed surface brightness is dimmed by a factor $1/(1+z)^4$ relative to its intrinsic value. In Fig. 2.8, we illustrate that the

---

[4] We caution that 'bright' is used here in a relative sense only; in absolute terms a surface brightness $\sim 10^{-20}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$ is still observationally challenging. Nevertheless, a combination of deep observations, high spectral resolution, and pixel-based stacking techniques has been successfully used to image surface brightnesses approaching this threshold (e.g. Fumagalli et al. 2017a; Wisotzki et al. 2018).

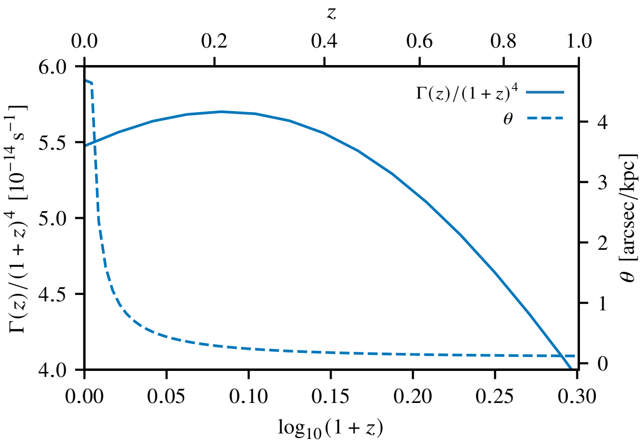

Figure 2.8
Redshift evolution of the apparent surface brightness (solid curve; left scale) and the angular separation corresponding to a comoving kpc (dashed curve; right scale).

quantity $\Gamma(z)/(1+z)^4$, which is directly proportional to the observed surface brightness, is maximised at $z \sim 0.2$. Also shown in Fig. 2.8 is the variation with redshift of the angular separation corresponding to a proper kpc, which decreases rapidly with redshift. In addition, halo concentration increases with redshift, which we expect to result in more compact fluorescent rings. Hence, resolving fluorescent objects at $z \gg 0$ is likely to be challenging. As we noted in Section 2.3.2.1, a boost in the UVB photoionisation rate can be provided even at $z \simeq 0$ by proximity to a highly-luminous source, such as an AGN (Cantalupo et al. 2005). To obtain detailed predictions of the expected surface brightness in this case, further modelling is required to determine the precise effects of the resulting change in spectral shape on the properties of fluorescent rings. Additionally, hydrodynamic modelling is needed to establish the conditions (e.g. separation from the source) needed to avoid stripping of gas from a candidate RELHIC.

2.5 Physics governing the properties of fluorescent haloes

In Section 2.3, we presented surface brightness profiles resulting from varying input parameters to our photoionisation code. In the event of a successful detection of a

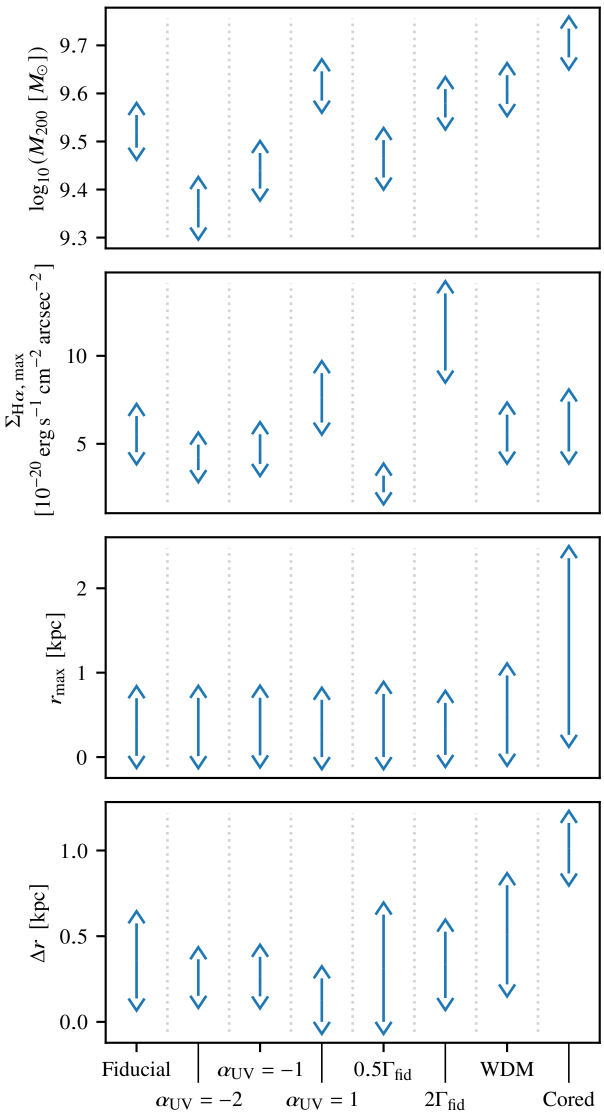

Figure 2.9

Summary of properties of ring fluorescence. The top panel shows the halo mass range for which ring fluorescence is predicted, for all models considered in Section 2.3. Our fiducial model is on the far left, followed by the following variations: three different values of the UVB spectral slope α_{UV}, UVB photoionisation rates Γ a factor of 2 below and above the fiducial value, the WDM (Bose et al. 2017) mass-concentration relation, and the cored dark matter density profile given by Burkert (1995). Subsequent panels show, for the same set of haloes in each model, values of the peak surface brightness, the radius at which this peak is reached, and the full width at half maximum of the radial surface brightness.

candidate RELHIC with ring-shaped fluorescent H α emission, we would like to determine the properties of the object, along with those of the incident radiation field. Hence, in this section we consider the causes of the trends displayed in our results in more detail, and discuss how they may be used to infer the properties of a fluorescent RELHIC and its environment from its appearance.

We summarise our results in Fig. 2.9, which shows the range of halo masses for which ring fluorescence is predicted under the different model variations we consider, along with the corresponding ranges in the peak value surface brightness, the projected radius at which this peak occurs, and the full width at half maximum of the surface brightness profile. Depending on its properties, purely detecting a fluorescent ring may provide constraints on its environment. For example, detection of a ring with $\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}} \gtrsim 10^{-19}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$ would strongly suggest a UVB intensity enhanced with respect to the MH15 spectrum. From the bottom two panels of Fig. 2.9, we also see that with 'non-standard' dark matter (i. e. WDM or cored density profiles), fluorescent rings are expected to reach their peak surface brightness at larger projected radii than in the fiducial cuspy CDM case; meanwhile, the surface brightness profiles for the cored haloes we model are also broader. This implies that combined measurements of r_{max} and Δr could potentially be used to distinguish WDM from CDM, and cored from cuspy density profiles, although we again note that these properties depend on the exact parameterisation of the density profile. For example, choosing a smaller core radius, corresponding to a more concentrated halo, would translate to reductions in r_{max} and Δr.

We now turn to considering why different input parameters affect the properties of fluorescent rings in the ways shown in Fig. 2.9. An important trait universal to all our models is that, despite significant variation in the sizes and peak surface brightnesses of the fluorescent rings they predict, the column density of the gas at r_{max} has a characteristic value $N_{\mathrm{HI}} = N_{\mathrm{crit}} \simeq 10^{19}\,\mathrm{cm}^{-2}$. We illustrate this in Fig 2.10, where we also show that this corresponds to the column density at which the neutral fraction is $X_{\mathrm{HI}} = 0.5$. While the association of the ionisation front with a column density of this magnitude is a standard result of radiative transfer calculations, it is nevertheless significant that we obtain self-

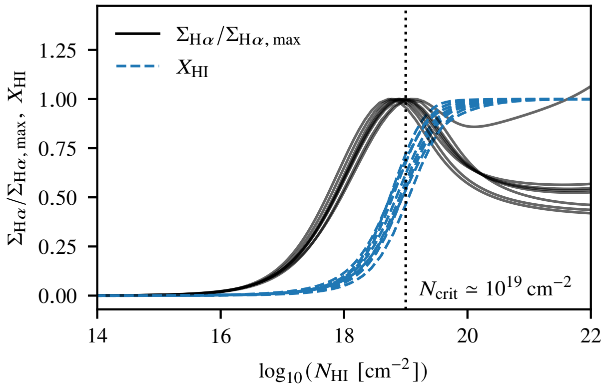

Figure 2.10

Normalised H α surface brightness profiles for the most massive halo which satisfies our star formation criterion in each model we consider (black), shown with corresponding H neutral fractions X_{HI} (dashed blue).

similar X_{HI} and $\Sigma_{\mathrm{H}\alpha}$ profiles despite substantial variation in the shapes of physical density profiles between our different models.

One notable exception to this self-similarity is the $\alpha_{\mathrm{UV}} = 1$ model, in which the surface brightness profile dips after the peak at N_{crit}, but then rises again toward higher N_{HI}. This occurs due to the increased flux of high-energy photons in the harder UVB spectrum assumed by this model, which typically penetrate to smaller radii (corresponding to larger column densities), due to their reduced photoionisation cross-section. Additionally, the energetic photoelectrons they liberate can go on to ionise further atoms via secondary collisional processes. This extra source of ionisation generates a greater degree of residual ionisation within the predominantly-neutral region. In combination with the increase in n_{H} toward smaller radii, this produces the observed doubly-peaked surface brightness profile.

With the 'fixed point' N_{crit} identified, the changes in ring brightness and size between models may be described in terms of differences in the $N_{\mathrm{HI}} - n_{\mathrm{H}}$ and $N_{\mathrm{HI}} - r$ relations respectively. The former follows since the peak emissivity is given by $\varepsilon_{\mathrm{H}\alpha,\,\mathrm{max}} = X_{\mathrm{HII}} n_{\mathrm{H}} n_{\mathrm{e}} \alpha_{\mathrm{H}\alpha}(T)|_{r_{\mathrm{max}}}$. Since r_{max} corresponds to the ionisation front where $X_{\mathrm{HII}} = 0.5$, and so $n_{\mathrm{e}} \approx n_{\mathrm{H}}$ while also $T \simeq 10^4$ K, the value of $\varepsilon_{\mathrm{H}\alpha,\,\mathrm{max}}$ is set by $n_{\mathrm{H}}(r_{\mathrm{max}})$. Hence, a change in the local density of gas where N_{HI} reaches N_{crit} will produce a change in the peak surface brightness. By the same token, a change in r_{max} implies a change in the position of the ionisation front, and therefore a change in the radius at which $N_{\mathrm{HI}} = N_{\mathrm{crit}}$. With these 'scaling relations'

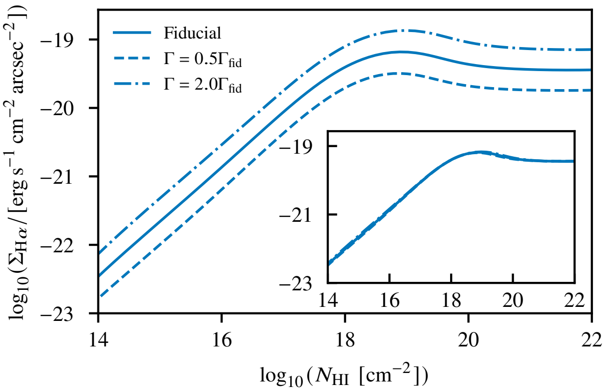

Figure 2.11

Surface brightness as a function of H I column density for different photoionisation rates. In the inset panel, we show that the three N_{HI}–$\Sigma_{\mathrm{H}\alpha}$ relations are self-similar when the surface brightness is scaled by the ratio $f_\Gamma = \Gamma/\Gamma_{\mathrm{fid}}$.

identified, it is then possible to take the properties of an observed fluorescent ring and work backwards to determine the physical properties of the host RELHIC. By comparing these inferred values with those predicted by our simulations, constraints can be placed on which models are applicable to observed objects, as we discuss next.

We first consider trends in the maximum surface brightness $\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}}$; as described previously, this is primarily set by the gas number density at the ionisation front. The physical size of the neutral region provides an additional influence, by altering the path length of sightlines passing through gas with emissivity close to maximal. However, the most extreme change in ring geometry with respect to our fiducial model occurs when a cored density profile is assumed, but the resulting values of $\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}}$ differ by a much smaller margin. Hence, we infer that the ring brightness is set primarily by properties of the UVB. In particular, the models in which we consider a doubling or halving of Γ predict changes in surface brightness by the same amount. This is shown explicitly by plotting $\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}}$ as a function of N_{HI} (Fig. 2.11), where we demonstrate in the inset panel that $\Sigma_{\mathrm{H}\alpha}/f_\Gamma$ curves, where f_Γ is the multiplicative factor applied to the photoionisation rate, are self-similar. Hence, assuming the other properties of the UVB (i. e. α_{UV}) are similar to the MH15 spectrum, measuring the peak surface brightness of an observed ring and inferring a value for f_Γ would allow a direct determination of the UVB photoionisation rate (see also Fumagalli et al. 2017a).

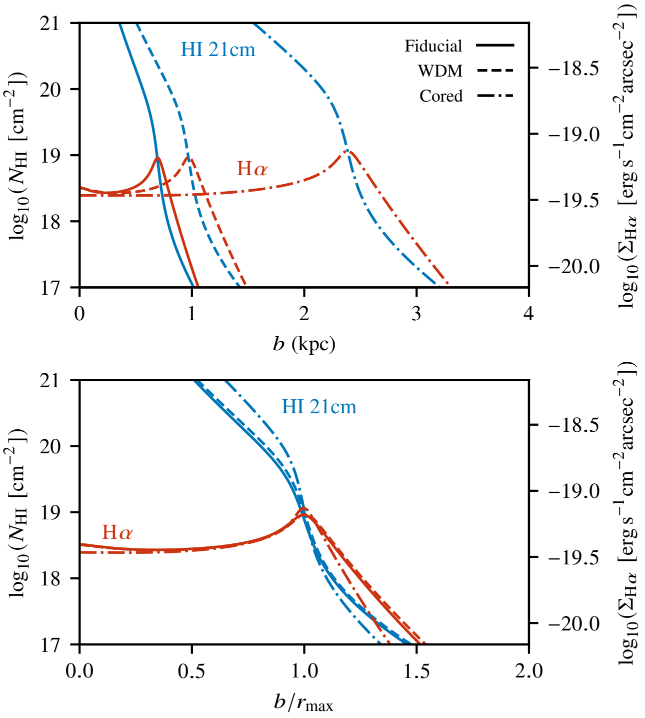

Figure 2.12

H i column density profiles (blue) for models with different dark matter density profiles, with corresponding H α surface brightness profiles (red). Profiles are shown as a function of impact parameter b, both in physical units (top panel) and normalised by $r_{\rm max}$ (bottom panel). For each model, profiles are shown for the most massive halo which satisfies our star formation criterion.

In contrast, the projected radius $r_{\rm max}$ at which fluorescent rings are maximally bright is essentially independent of the properties of the UVB, changing only when the mass-concentration relation or density profile of the underlying dark matter halo is altered. As the column density at $r_{\rm max}$ remains equal to $N_{\rm crit}$, it is the $N_{\rm HI}$–r relation that is changing in this case, as we illustrate in the upper panel of Fig. 2.12. Identifying this trend reinforces our previous observation that measuring $r_{\rm max}$ in combination with Δr permits cuspy CDM, cored DM (e.g. SIDM) and WDM to be differentiated. However, doing so requires knowledge of the distance to the ring in order to determine physical values of $r_{\rm max}$ and Δr. Without this, the H α and H i profiles are only known as a function of $b/r_{\rm max}$, a situation which results in them becoming highly self-similar. We illustrate this in the lower panel of Fig. 2.12. However, if we instead assume a fixed family of models and vary the halo mass, we find a systematic trend in the shape of the H i profile, which becomes steeper for $N_{\rm HI} > N_{\rm crit}$ with increasing halo mass (Fig. 2.13). Changes also occur in the H α

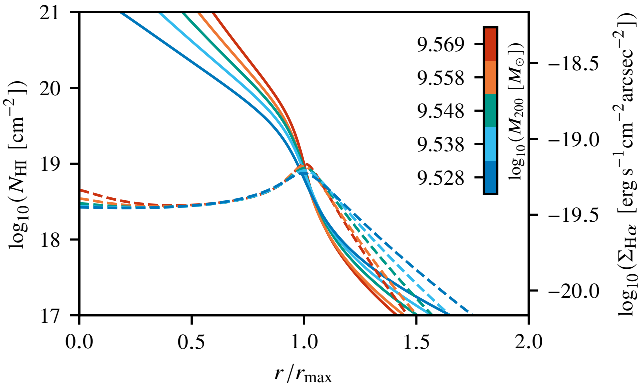

Figure 2.13
H I column density profiles (solid lines), with corresponding H α surface brightness profiles (red). Profiles are shown as a function of $b/r_{\max}$ for haloes of different mass in our fiducial model.

profiles, which become more sharply peaked at $b = r_{\max}$ for higher halo masses. Using complementary observations of H α and H I in this way may permit better constraints to be made, although the relatively subtle change in e.g. $r_{\max}$ produced by the WDM models, which correspond to a significant 50% reduction in halo concentration c_{200}, imply that obtaining the necessary accuracy from observations would be challenging.

2.6 Summary and conclusions

We have reported the results from radiative transfer calculations of gas bound to a gravitational potential such as that due to a dark matter halo and illuminated by an isotropic ionising background. By comparison with the well-known CLOUDY software, as well as previous analytic modelling, we have established that for the intended usage, despite making a number of simplifications, the code we use is robust and its results accurate. We have used this code to model the gas clouds bound to RELHICs, and obtained surface brightness profiles for the H α line, which is produced via fluorescent emission due to the ionising background. Our key findings are:

(i) We confirm the results of previous studies, finding that haloes which develop a self-shielded neutral core possess surface brightness profiles with a characteristic ring-shaped morphology.

(ii) This occurs for only a narrow range of halo masses, bounded from below by the requirement of a sufficiently deep potential, and from above by the requirement that star formation remains suppressed.

(iii) Using numerical simulations, we predict a small population of $3^{+2.6}_{-2.0}$ (10-90% c. i.) fluorescent rings in a spherical 3 Mpc volume centred on the Local Group. Detecting these objects via a blind H α survey would be impractical, but pre-selection techniques, e.g. via 21 cm H I imaging, can provide catalogues of promising target sources, particularly beyond the LG where the sky coverage required to map a region of equivalent volume is reduced.

(iv) We quantify how variations of the intensity (i. e. photoionisation rate) and spectral shape of the ionising background map directly to changes in the intensity of the fluorescent emission. Modifications to the dark matter potential instead alter the apparent size and shape of the fluorescent ring.

(v) We show that reverse-engineering these relations would enable a direct determination of the UVB photoionisation rate from the intrinsic brightness of an observed ring; equally, constraints on the nature and distribution of dark matter may be inferred from the ring's geometry.

The existence of a large number of low-mass dark matter halos which do not host a luminous galaxy is an unavoidable prediction of the CDM paradigm. While challenging, observing the distinctive ring-shaped fluorescent H α emission from the gas bound to these haloes would provide a probe of this cosmological model on an as-yet untested scale. In the longer term, a detailed analysis of this emission would unlock further insights, both into the nature of dark matter and the spectrum of the extragalactic ionising background.

CHAPTER 3

Y_P and the UVB slope from fluorescent RELHICs

3.1 Introduction

Almost all helium atoms in the Universe were synthesised in the first few minutes after the Big Bang, during the period of Big Bang nucleosynthesis (BBN; Alpher et al. 1948; Hoyle & Tayler 1964). The primordial helium mass fraction Y_P, or equivalently the abundance by number y_P[1], which results from this brief period of nucleosynthesis is influenced by the early-time expansion history of the Universe, and by the abundance of free neutrons at the onset of nucleosynthesis, which in turn depends on the neutron half-life. The primordial helium abundance is therefore sensitive to both cosmology and particle physics, making accurate measurements of this quantity highly informative. The precise measurements of the baryon-to-photon ratio obtained from cosmic microwave background satellites such as *Planck* remove the final free parameter from BBN, meaning that the Standard Model prediction of $Y_\mathrm{P} = 0.24672 \pm 0.00017$ ($y_\mathrm{P} = 0.08188 \pm 0.00008$) is reliable (Pitrou et al. 2018).

Comparing observational measures of Y_P with the BBN prediction allows the presence of any new physics beyond the Standard Model to be investigated. To date, the leading

---

[1] These two quantities are related by $Y_\mathrm{P} = 4y_\mathrm{P}/(1+4y_\mathrm{P})$. We note that Y_P is defined as $Y_\mathrm{P} \equiv 4\,n(^4\mathrm{He})/n_\mathrm{b}$, where n_b is the baryon density. Y_P is therefore somewhat of a misnomer; it does not represent the mass fraction of ^{4}He. Since BBN codes naturally calculate a number abundance ratio, and observations also measure the primordial helium abundance in this form, we will predominantly use the number abundance $y_\mathrm{P} \equiv n_\mathrm{He}/n_\mathrm{H}$ in this chapter.

method for determining $Y_{\rm P}$ involves comparing the relative intensity of hydrogen and helium emission lines measured in H II regions, ionized bubbles of gas surrounding regions of active star formation (e. g. Izotov et al. 2014; Aver et al. 2015; Peimbert et al. 2016; Valerdi et al. 2019; Fernández et al. 2018, 2019). These studies select H II regions in metal-poor galaxies ($Z/Z_{\odot} \lesssim 0.1$), to minimise the enrichment by stars where emission is detected. However, the level of contamination remains necessarily non-zero, so the observed ratios of hydrogen to helium emission must be extrapolated down to zero metallicity in order to recover the primordial abundance ratio. This limitation introduces the possibility of systematic errors (see e. g. Izotov et al. 2007; Porter et al. 2009), the characterisation of which becomes increasingly important as statistical errors on the measurements improve. Consequently, it is beneficial to consider independent techniques for determining $Y_{\rm P}$. One such alternative involves studying intergalactic absorption lines arising in almost-primordial clouds located between us and a background quasar. This approach has been demonstrated to yield a primordial value of $Y_{\rm P} = 0.250^{+0.033}_{-0.025}$ (Cooke & Fumagalli 2018) consistent with the Standard Model prediction, although the constraint obtained is not yet as tight as that resulting from H II region measurements, for which a weighted average of recent determinations (see references above) gives $Y_{\rm P} = 0.248 \pm 0.001$.

In this chapter, we discuss a novel method for determining $Y_{\rm P}$. We focus on low-mass dark matter haloes, the existence of which is a robust prediction of the cold dark matter (CDM) model for hierarchical structure formation. Below a mass scale of approximately $10^{10}\,{\rm M}_{\odot}$, observational constraints indicate that many haloes fail to host luminous galaxies (Klypin et al. 1999; Moore et al. 1999). This requirement can be met by appealing to baryonic feedback processes; most prominently, cosmic reionization heats intergalactic gas to $\sim10^{4}$ K, inhibiting star formation in haloes with potential wells too shallow to confine the heated gas (Miralda-Escudé & Rees 1994; Okamoto et al. 2008; Meiksin 2009). Using the APOSTLE suite of Local Group hydrodynamical simulations (Sawala et al. 2016), Benítez-Llambay et al. (2017; hereafter BL17) identified a population of haloes with masses $10^{8} < M_{\rm halo}/{\rm M}_{\odot} < 10^{9.6}$, which additionally experience negligible star formation prior to reionization. Hence, these haloes remain essentially star-free down to redshift $z = 0$, and by avoiding mechanisms such as ram pressure stripping from interactions with the cosmic web, can retain a small reservoir of essentially-pristine gas. This gas consists of an approximately kiloparsec-sized neutral core surrounded by an envelope kept ionized by the diffuse ultraviolet background (UVB), motivating the naming of this population as "REionization-Limited HI Clouds" (RELHICs).

In previous work (Chapter 2, published as Sykes et al. 2019; hereafter S19), we performed radiative transfer simulations to model RELHICs and examine the effects of the UVB on the properties of their gas. The UVB ionizes atoms in the gas, which later recombine to produce hydrogen emission lines such as H α. We found that for RELHICs with masses in the narrow range $10^{9.4} < M_{\rm halo}/{\rm M}_{\odot} < 10^{9.6}$, this fluorescent emission displays a distinctive ring-shaped morphology when seen in projection on the sky. The narrow mass range for which we predict these fluorescent rings, in combination with their intrinsically low surface brightness of the emission, makes fluorescent rings rare and their detection challenging. This intrinsic brightness increases at higher z, due to the greater amplitude of the UVB. However, the need to resolve the ring-shaped emission that distinguishes a fluorescent RELHIC, in combination with the rapid onset of cosmological surface brightness dimming, means that only relatively local RELHICs ($z \lesssim 0.2$) are realistic candidates for detection. Nevertheless, they remain a firm prediction of the CDM paradigm, and their detection would provide a probe of this cosmological model on an as-yet untested scale. Furthermore, we have shown that observable properties of the rings, such as their projected size and peak brightness, are sensitive to the properties of the UVB and the mass of the underlying dark matter halo.

Fluorescent RELHICs will also produce emission in helium recombination lines, which will exhibit a similar ring-like appearance. As a result of their star-free nature, the gas they contain should be almost pristine in composition, and so they have the potential to yield a direct constraint on $y_{\rm P}$, albeit one with substantial observational challenges given current instrumentation, as we will show. In this chapter, we explore this possibility, and find that in addition to being able to measure the helium abundance, a comparison of the fluorescent hydrogen and helium emission lines from RELHICs could provide the first observational constraint on the shape of the ionizing UVB spectrum.

The chapter is organised as follows: in Section 3.2, we provide a brief description of our numerical method and describe how we have extended the calculations presented in S19 to additionally predict surface brightnesses for He I and He II emission lines. We then present our results in Section 3.3, considering constraints on the primordial helium abundance $y_{\rm P}$ and on the UVB slope in turn (Sections 3.3.1 and 3.3.2), and then combined constraints on both parameters (Section 3.3.3). We conclude by discussing our results and their implications in Section 3.4. Throughout, we assume a set of cosmological parameters ($H_0 = 67.3 \ {\rm km\,s^{-1}\,Mpc^{-1}}$, $\Omega_\Lambda = 0.685$, $\Omega_{\rm M} = 0.315$, $\Omega_{\rm B} = 0.0491$) consistent with *Planck* measurements (Planck Collaboration et al. 2014).

3.2 Modelling helium emissivities

As was demonstrated in BL17, the RELHICs identified in APOSTLE are well-described by a simple analytic model in which the gas they contain is in hydrostatic equilibrium with a gravitational potential due to the host dark matter halo, and in thermal equilibrium with the UVB. To predict their emission properties, we implement this analytic model using an ionization balance code originally described in Cooke & Pettini (2016) and with additional modifications introduced by S19. Similar in functionality to photoionization codes such as CLOUDY (Ferland et al. 2017), the code we present includes two important additional features. Firstly, it applies the condition of hydrostatic equilibrium to determine the gas density profile, using the background gravitational of a Navarro-Frenk-White (NFW; Navarro et al. 1996b) dark matter halo. Secondly, although we assume spherical symmetry (implying that the UVB irradiates the gas isotropically), the gas column density has an angular dependence which also varies with radius, meaning that the local attenuated radiation field, and hence the local photoionization and photoheating rates, are functions of both depth within the cloud and direction of incident radiation.

Our calculations take the following form, which is similar to that described by Sternberg et al. (2002). A dark matter potential is defined by choosing a virial mass M_{200} and obtaining the concentration parameter $c_{200} \equiv R_{200}/r_s$ from the Ludlow et al. (2016) mass-concentration relation. Here, r_s is the NFW scale radius, and virial quantities are defined such that within a sphere of radius R_{200}, the average density is $200\rho_{\rm crit}$, where $\rho_{\rm crit}$ is the critical density of the Universe, and the total mass enclosed by this sphere is M_{200}. A total baryonic gas mass $M_{\rm g}$ is then assigned to the halo using the analytic model employed by BL17; the gas is split into $N_r = 1000$ radial and $N_\mu = 30$ azimuthal cells and initialised to be fully ionized and isothermal with temperature $T = 10^4$ K. We assume the gas to have primordial composition, with the helium abundance given by $n_{\rm He} = y_{\rm P} n_{\rm H}$.

We then determine the pressure profile required for hydrostatic equilibrium, using as a boundary condition the assumption that at $r \gg R_{200}$, the gas density approaches the cosmic mean baryon number density, $\bar{n}_{\rm H} \simeq 10^{-6.7} {\rm cm}^{-3}$ at $z = 0$. From the resulting gas density profile, we next determine the intensity of the radiation field within the cloud, using the Madau & Haardt (2015; hereafter MH15) $z = 0$ UVB as the initial, unattenuated spectrum. We calculate ionization rates for photoionization, primary and secondary collisional ionizations, hydrogen ionization resulting from helium recombination radiation, and charge transfer ionization. We further calculate recombination rates for radiative, dielectronic and charge transfer processes, assuming Case B conditions hold throughout the cloud. This means we ignore recombinations directly to the ground state, because at the typical densities associated with the ionization front, the ionizing photons produced by these recombinations will invariably ionize a nearby neutral atom. Thus, they will have no effect on the overall ionization state of the gas and produce no detectable emission.[2] By equating the relevant processes for each atomic species, we enforce ionization equilibrium and thus determine the fractional ionization $X_{A^{i+}} \equiv n_{A^{i+}}/n_A$ of all species A (here H and He) and ionization stages i.

Finally, we determine the temperature profile by assuming thermal equilibrium for higher-density gas ($n_{\rm H} > 10^{-4.8}\,{\rm cm}^{-3}$), where this threshold is set by the condition that the timescale for equilibrium must be shorter than the Hubble time. For gas below this threshold density, we instead set the gas temperature to that resulting from a heating

---

[2] We further justify this assumption in Appendix B of S19.

timescale equal to the Hubble time, interpolating between the two regimes to ensure the temperature profile remains smooth. In the equilibrium case, we compute the heating rate by considering primary photoheating and secondary heating by primary photoelectrons, while the cooling rate includes contributions from collisional excitation/ionization cooling, recombination cooling, Brehmsstrahlung cooling and Compton cooling/heating. For details of the rate coefficients and other atomic data that we use, see S19 and Cooke & Pettini (2016).

We proceed iteratively, using the temperature profile to recompute the pressure profile and repeating the above procedure until a convergence criterion is met: namely, that the fractional ionizations in every radial cell change by less than 0.1% between successive iterations. With a converged ionization structure found, we then compute the volume emissivity ε_ν for an emission line with frequency ν as:

$$\varepsilon_\nu(r) = h\nu\, n_{\rm ion}(r)\, n_e(r)\, \alpha_{\rm eff}(T(r))\,; \quad [\varepsilon] = {\rm erg\,s^{-1}\,cm^{-3}}\,, \tag{3.2.1}$$

where h is Planck's constant, $n_{\rm ion}$ is the density of H II, He II or He III as appropriate, n_e is the electron density, and $\alpha_{\rm eff}$ is a temperature-dependent coefficient expressing the rate per unit ion and electron densities at which the relevant transitions occur. For H α and the analogous He II 4686Å line, we obtain $\alpha_{\rm eff}$ values from Osterbrock & Ferland (2006), while for the He I lines we use the emissivities compiled by Porter et al. (2012, 2013), which are tabulated as functions of n_e and T. Finally, we calculate the projected surface brightness Σ_ν as the integral of $\varepsilon_\nu(r)$ along lines of sight corresponding to an impact parameter b:

$$\Sigma_\nu(b) = \frac{1}{2\pi}\int_b^{R_{200}} \frac{r}{\sqrt{r^2-b^2}}\varepsilon_\nu(r)\,{\rm d}r\,; \quad [\Sigma] = {\rm erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}\,. \tag{3.2.2}$$

3.3 Results

We first discuss the qualitative properties of the emission lines using a fiducial model, which takes the primordial helium abundance to be $y_{\rm P} = 0.083$, corresponding to a primordial

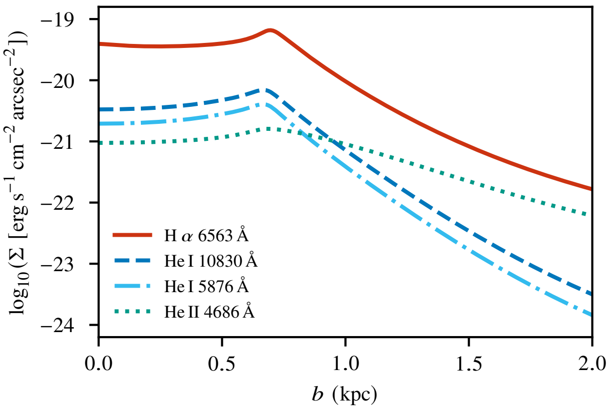

Figure 3.1

Surface brightness Σ as a function of impact parameter b, for the 4686Å He II line and two He I lines, in addition to the hydrogen H α line.

mass fraction $Y_P = 0.249$. In Fig. 3.1, we show surface brightness profiles as a function of impact parameter for a halo with $M_{200} = 10^{9.55}\,\mathrm{M}_\odot$. This is the maximum halo mass such that upper bounds on the column density and gas mass fraction, motivated by the need to avoid star formation, are not exceeded (see S19; §2.3). In addition to the hydrogen H α line, we plot the surface brightness of the two brightest He I lines (at 10830Å and 5876Å), and the He II 4686Å line. For this halo, the H α surface brightness reaches a peak intensity of $\Sigma_{\mathrm{H}\,\alpha,\,\max} = 6.58 \times 10^{-20}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$. The helium line surface brightnesses are significantly dimmer, reaching maximum values of 4.06×10^{-21} and $1.61 \times 10^{-21}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$ for the He I and He II lines, respectively.

$\Sigma_{\mathrm{H}\,\alpha,\,\max}$ occurs at $b = 0.70$ kpc, which corresponds to the projected radius of the fluorescent ring. This location is set by the position of the hydrogen ionization front, at which the H α volume emissivity reaches a maximum. Outside the ionization front, $\varepsilon_{\mathrm{H}\alpha}$ falls rapidly with increasing radius as the gas density drops. Conversely, at radii within the ionization front the emissivity is suppressed by the exponentially-decreasing ionized fraction.[3] We find that the peak helium surface brightnesses occur at a similar radial position to $\Sigma_{\mathrm{H}\,\alpha,\,\max}$, despite the respective ionization fronts being located at different radii. This occurs because the requirement of hydrostatic equilibrium produces gas densities which decrease rapidly with radius, such that the helium emissivities are affected more strongly by the falling

---

[3] At even smaller radii, $\varepsilon_{\mathrm{H}\alpha}$ begins to rise again due to an increased contribution from secondary collisional ionizations.

electron density than by the helium ion densities.

From Eq. 3.2.1, we would expect that the ratio of helium to hydrogen emissivity is set by the product of the ratios of ion densities, line frequencies and rate coefficients. However, recovering the helium abundance $y_{\rm P}$ from this ratio involves some additional considerations. Firstly, the densities in Eq. 3.2.1 refer to single ionization stages, whereas $y_{\rm P}$ is set by the overall atomic abundances. Secondly, the rate coefficients $\alpha_{\rm eff}$ are temperature-dependent. Since the temperature and ionization structure of the gas (which influence $\alpha_{\rm eff}$ and the fractional ionizations respectively) depend both on $y_{\rm P}$ and each other, we proceed by performing the iterative procedure outlined in Section 3.2 for a number of different $y_{\rm P}$ values. By comparing the results obtained in each case, we may determine the effects of changing $y_{\rm P}$ in relative terms.

3.3.1 Determining $y_{\rm P}$

To quantify how varying $y_{\rm P}$ changes the predicted emission (and hence the sensitivity for determining $y_{\rm P}$), we first consider the ratio of H α to helium emissivity R_ε, defined as

$$R_\varepsilon \equiv \frac{\varepsilon(\mathrm{H}\,\alpha)}{\varepsilon(\mathrm{He\,II}\,4686\text{Å}) + \varepsilon(\mathrm{He\,I}\,10830\text{Å}) + \varepsilon(\mathrm{He\,I}\,5876\text{Å})}\,, \tag{3.3.1}$$

such that the net He I emissivity is the sum of the brightest two He I lines. These comprise the near-infrared 10830Å line and the 5876Å line, which is intrinsically fainter but lies in the optical part of the spectrum along with the H α and He II lines, and so may be more convenient to detect. In addition to our fiducial model, we consider variations in which the assumed value of $y_{\rm P}$ is altered by a factor

$$f_y \equiv \frac{y_{\rm P}}{y_{\rm P,\,fid}}\,, \tag{3.3.2}$$

which we allow to take the values $f_y = (1.01, 0.99, 1.10, 0.90)$, corresponding to ±1% and ±10% changes in $y_{\rm P}$.

Calculations for a fixed halo mass are not directly comparable between these variations, since the assumed helium abundance affects the thermal and ionization state of the gas, altering the conditions for hydrostatic equilibrium and leading to gas distributions with different emission properties. We instead require that our results should be "self-similar" across model variations, in the sense that a change in the input parameter which distinguishes the variations should cause the predicted emission properties to change in a systematic way. In S19 we found that this condition is satisfied if we compare haloes whose gas content reaches the same peak column density of neutral hydrogen $N_{\mathrm{HI,\,max}}$. This condition is appropriate because $N_{\mathrm{HI,\,max}}$ provides a proxy for τ_∞, the total optical depth of the gas. This sets the intensity of the radiation field near the centre of the cloud, and thus its ionization state, in a way that is largely independent of the overall density and temperature structure, which does differ between model variations. We employ the same approach here. However, the relation between M_{200} and $N_{\mathrm{HI,\,max}}$ is not known *a priori*, so for each model variation we use an iterative procedure to determine the halo mass which results in the desired value of $N_{\mathrm{HI,max}}$. While we must choose a value of $N_{\mathrm{HI,max}}$ to enable comparisons between our model variations, the results we will present are insensitive to the threshold chosen, provided that it is sufficiently high that a well-defined ionization front is formed ($N_{\mathrm{HI,max}} > 10^{19}\ \mathrm{cm}^{-2}$). We adopt $N_{\mathrm{HI,max}} = 10^{23}\,\mathrm{cm}^{-2}$ which, although large, is unlikely to result in self-shielded gas becoming star-forming due to the extremely metal-poor nature of the gas, which results in inefficient cooling. Furthermore, the steep decline of N_{HI} with radius means that this column density is achieved only for sightlines passing through the very centre of the RELHIC. [4]

We plot the resulting emissivity ratios, as a function of the hydrogen number density n_{H}, in the upper panel of Fig. 3.2. As expected, R_ε decreases (i. e. $\varepsilon_{\mathrm{He}}$ is larger relative to $\varepsilon_{\mathrm{H}\alpha}$) for models with a larger value of y_{P}, and conversely for smaller y_{P}. In the middle panel of Fig. 3.2, we additionally scale the obtained R_ε curves by f_y to demonstrate that these profiles remain almost self-similar, particularly near the ionization front where the emissivities are maximised. We quantify this property in the lower panel, where the

---

[4] See Section 3.4 of this chapter and Section 2.3 of S19 for further discussion.

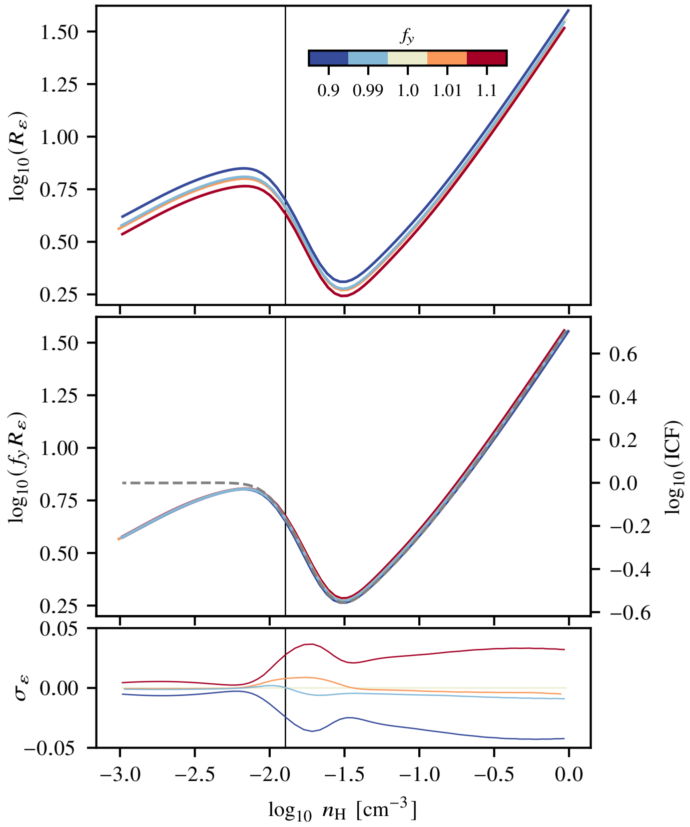

Figure 3.2

Upper panel: Hydrogen to helium emissivity ratio R_ε (Eq. 3.3.1) for models with different f_y (Eq. 3.3.2), as a function of n_H. *Middle panel*: Emissivity ratios, scaled additionally by f_y, for the same models. The ionization correction factor for $f_y = 1$, defined in the text, is shown with a dashed grey line using the right-hand scale. *Lower panel*: Deviation from self similarity σ_ε (Eq. 3.3.3), i. e. the error incurred in assuming $f_y R_\varepsilon$ is independent of y_P. The solid vertical lines indicate the density at the location of the ionization front.

fractional deviation from self-similarity σ_{ε}, defined as

$$\sigma_{\varepsilon} \equiv \frac{f_y R_{\varepsilon} - R_{\varepsilon,\,\mathrm{fid}}}{R_{\varepsilon,\,\mathrm{fid}}} \tag{3.3.3}$$

is plotted for each value of f_y. We find larger values of σ_{ε} at higher densities, corresponding to gas located inside the ionization front. Identifying the causes of these deviations is made difficult by the coupled nature of our calculations, but we expect differences in the thermal state and density profile of the gas (in particular, the electron density n_{e}) between models to be important. Experiments where the variation in y_{P} was ignored when determining these quantities reduced, but did not eliminate, the residual discrepancies. In any case, they represent deviations at the position of the peak emissivity of <1% for 1% changes in y_{P} and ~2% for 10% changes. Hence, they may be safely discounted, and we are justified in interpreting the variations in R_{ε} as being solely caused by changes in y_{P}. This direct relationship between the observable values of R_{ε} and the underlying abundance is the fundamental property that makes RELHICs appealing tools for determining y_{P}.

The emissivities and their ratio depend on the densities of a specific ionization stage, whereas determining y_{P} requires the total densities of H and He. This scenario is commonly encountered in absorption and emission line studies, and is circumvented by introducing an ionization correction factor (ICF) to account for unobserved ionization stages, allowing the total density to be inferred. We define the ICF to be

$$\mathrm{ICF} \equiv \frac{n_{\mathrm{HII}}}{n_{\mathrm{H}}} \frac{n_{\mathrm{He}}}{n_{\mathrm{HeII}} + n_{\mathrm{HeIII}}} = \frac{1 - X_{\mathrm{HI}}}{1 - X_{\mathrm{HeI}}}\,, \tag{3.3.4}$$

noting that this expression has a similar functional form to the emissivity ratio (by using Eq. 3.2.1 to expand each of the terms in Eq. 3.3.1). We illustrate this correspondence in the middle panel of Fig. 3.2, where the ICF is shown by the grey dashed line and plotted on the right-hand axis. The equilibrium temperature of the gas falls with decreasing density, which in combination with the different temperature dependencies of the recombination coefficients, cause R_{ε} to also fall with decreasing density. Conversely, at low densities both H and He become fully ionized, leading to the ICF tending to 1 and breaking the correspondence with R_{ε}. Nevertheless, at densities of $n_{\mathrm{H}} \sim 10^{-2}\,\mathrm{cm}^{-3}$ associated with the peak emissivity, assuming a direct proportionality between R_ε and the ICF allows a measurement of the former to be translated to an ionization correction, and hence a value of y_P, to a precision of within 5%.

However, the volume emissivity is not the directly observable quantity; that is instead its integral along the line of sight, the surface brightness. For situations in which the gas distribution can be described by a plane-parallel model, this distinction is not problematic. In the spherical geometry that we consider here, projection effects are significant since the surface brightness peak occurs due to limb brightening along lines of sight passing through more strongly-emitting gas. To determine the impact of these projection effects, we define the surface brightness ratio, analogously to R_ε, as

$$R_\Sigma \equiv \frac{\Sigma(\mathrm{H}\,\alpha)}{\Sigma(\mathrm{He\,II}\,4686\text{Å}) + \Sigma(\mathrm{He\,I}\,10830\text{Å}) + \Sigma(\mathrm{He\,I}\,5876\text{Å})}\,. \tag{3.3.5}$$

We also define the deviation from surface brightness self-similarity as

$$\sigma_\Sigma \equiv \frac{f_y R_\Sigma - R_{\Sigma,\,\mathrm{fid}}}{R_{\Sigma,\,\mathrm{fid}}}\,. \tag{3.3.6}$$

We plot R_Σ, $f_y R_\Sigma$, and σ_Σ for the same models shown previously in Fig. 3.3, where the x-axis now shows projected neutral hydrogen column densities N_HI (rather than n_H as used in Fig. 3.2). We find that despite the projection effects, self-similarity is closely preserved when moving to surface brightnesses, with values of σ_Σ at the position of the peak surface brightness in fact being smaller than the equivalent quantity for R_ε. This is likely due to the fact that the surface brightness at any radius is calculated by integrating over the entire emissivity profile, allowing for a degree of fortuitous cancellation between errors in opposite directions. However, determining y_P from R_Σ requires accurate measurements of the peak surface brightness Σ_max in H α and the three helium lines we consider. In absolute terms this emission is still extremely faint, particularly for the He II 4686Å line for which $\Sigma_\mathrm{max} \sim 10^{-21}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$. Consequently, obtaining a measurement of the surface brightness with the precision needed to produce a competitive measurement of y_P would be very challenging using current instrumentation. The integrated nature of the surface brightness may also impact our results, since it means that our predicted values

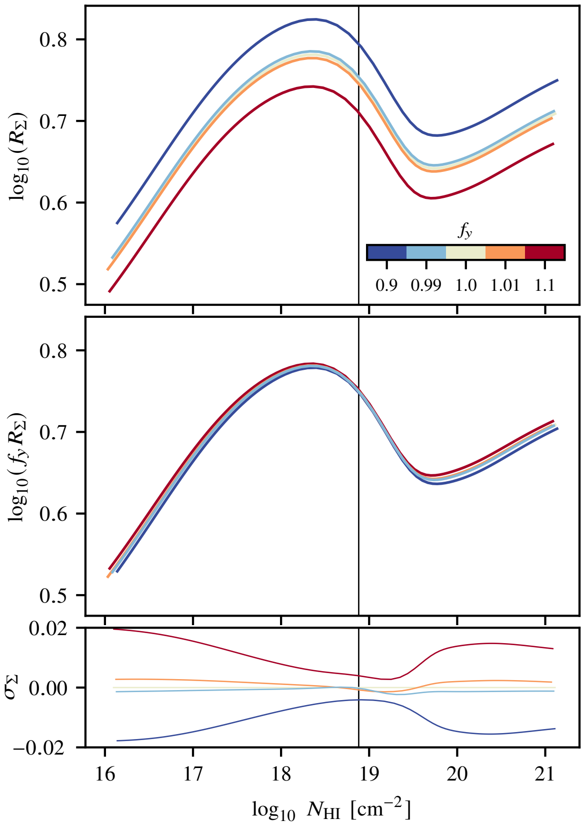

Figure 3.3

Upper panel: Hydrogen to helium surface brightness ratio R_Σ (Eq. 3.3.6) for the same models shown in Fig. 3.2, as a function of $N_{\rm H\,I}$. *Middle panel:* Surface brightness ratios scaled by f_y. *Lower panel*: Deviation from self similarity σ_Σ (Eq. 3.3.6), i. e. the error incurred in assuming $f_y R_\Sigma$ is independent of $y_{\rm P}$. The solid vertical lines indicate the H I column density at the location of the ionization front.

of R_Σ are sensitive to our modelling of the complete temperature and density structure of the gas, whereas R_ε depends only on local values of T and n_e. However, we do not expect this to be a significant disadvantage, given that temperature and density structures for RELHICs are well-specified.

We now consider the possibility of inferring y_P from the total line flux across the projected area of the fluorescent ring. This is observationally more feasible since measuring fluxes does not depend on making a highly-precise measurement of the peak surface brightness. Hence the fluorescent ring itself need not be spatially resolved, and the measurement precision attainable depends solely on the precision with which the flux can be determined [the signal-to-noise ratio (SNR) of the observations]. Moreover, the flux will depend only on the well-specified total gas content of the halo, given that the bulk of the emission originates in optically-thick gas for which the intensity of emission may be predicted analytically (Gould & Weinberg 1996; see also App. A, S19). We integrate the surface brightness over impact parameter to calculate total line luminosities, finding typical values of $L_i \approx (6 \times 10^{35}, 9 \times 10^{34}, 3 \times 10^{34})\,\mathrm{erg\,s^{-1}}$ for the H α, total He I, and He II 4686Å luminosities respectively. Converting these luminosities to line fluxes would require assuming a distance from the observer to the RELHIC that we model. However, we wish to consider the ratio of hydrogen and helium fluxes, which remains distance-independent and will be equal to the ratio of the total line luminosities. Hence, we define the H α to helium flux ratio as

$$R_\mathrm{F} \equiv \frac{F(\mathrm{H}\,\alpha)}{F(\mathrm{He\,II}\,4686\text{Å}) + F(\mathrm{He\,I}\,10830\text{Å}) + F(\mathrm{He\,I}\,5876\text{Å})} \tag{3.3.7}$$

but in practice, we compute the luminosity ratio R_L instead. In the main panel of Fig. 3.4 we show the values of R_F we obtain for different values of y_p, finding the expected trend of decreasing R_F with increasing y_P. More quantitatively, we expect that $R_\mathrm{F} \propto 1/y_\mathrm{P}$, with the normalisation of this relation being set by the relative intrinsic emissivities of the three lines we consider. We use standard non-linear least squares regression to fit a curve of this functional form to the predicted values of R_F, shown by the dashed curve in Fig. 3.4. To improve the fit, we additionally compute values of R_F for $f_y = \pm 20\%$ (not shown in

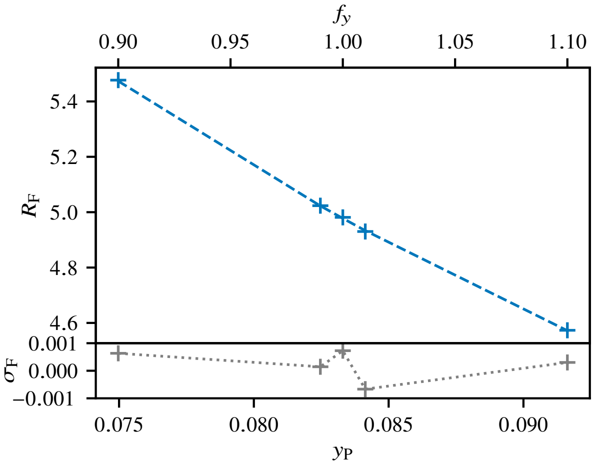

Figure 3.4

Upper panel: Flux ratio R_F (Eq. 3.3.7) as a function of y_P. The dashed line indicates a fit of the form $R_\mathrm{F} = A/y_\mathrm{P} + b$, where $A = 0.372$ and $b = 0.516$ are arbitrary scaling constants. *Lower panel*: The relative error σ_F (Eq. 3.3.8) between the calculated flux ratios and the fitted line.

Fig. 3.4). We define σ_F as the normalised residual of the data with respect to this fit:

$$\sigma_F \equiv \frac{R_\mathrm{F} - R_\mathrm{F,\,fit}}{R_\mathrm{F,\,fit}}, \tag{3.3.8}$$

and plot this quantity as a function of y_P in the lower panel of Fig. 3.4, finding $\sigma_\mathrm{F} \ll 1\%$ over the range of y_P values we consider.

We next use this fit to determine the precision with which y_P may be inferred from measuring R_F, given that the latter is calculated from measurements of the individual fluxes, which will have associated uncertainties. We assume that this uncertainty is described by a single relative error value δ_F for each of the four fluxes that must be measured, and use standard error propagation to determine the resulting error in R_F. By inversion of the fit in Fig. 3.4, we obtain a range of values of y_P consistent with the imprecise value for R_F, the extrema of which we report as δ_{y_P}, the error in y_P. In Fig. 3.5, we plot normalised values of δ_{y_P} as a function of the flux SNR, defined as SNR $= 1/\delta_\mathrm{F}$. We find that for SNRs of 10 and 100, corresponding to $\delta_\mathrm{F} = 10\%$ and 1%, y_P may be inferred to a precision of $^{+13}_{-10}\%$ and $^{+1.2}_{-1.0}\%$ respectively, where we have assumed that the underlying 'true' helium abundance is the fiducial value $y_\mathrm{P} = 0.083$; repeating these calculations assuming different values of y_P does not significantly affect the obtained values of δ_{y_P}. The asymmetry in these limits results from the non-linearity of the function $R_\mathrm{F}(y_\mathrm{P})$, which means that for a flux

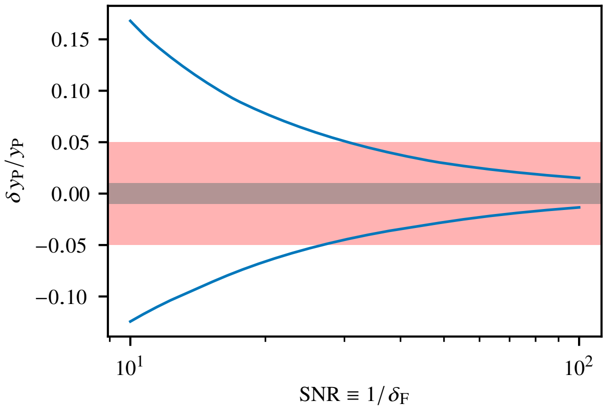

Figure 3.5

The relative precision to which y_{P} can be inferred (blue curves) as a function of δ_{F}, the precision to which the individual line fluxes are measured. The true helium abundance is assumed to be $y_{\mathrm{P}} = 0.083$. Red and grey shaded bands show 5% and 1% errors on y_{P} respectively.

error of fixed magnitude $|\delta_{\mathrm{F}}|$, the magnitude of $\delta_{y_{\mathrm{P}}}$ will vary depending on the sign of δ_{F}. Fig. 3.5 may also be used to determine the SNR required to achieve a constraint on y_{P} of a given precision. The 5% constraint indicated by the red shaded region corresponds to the range of reported values of y_{P} (see references in Section 3.1); in the absence of a uniform systematic offset in these measurements, this is the minimum level of precision which must be reached for an independent measurement to provide additional information. Conversely, the grey region indicates a 1% constraint, as obtained by the most precise determinations of y_{P} currently available (e. g. Valerdi et al. 2019). We find that satisfying these two constraints requires a flux SNR of ~30 and ~140 respectively. Hence, measurements of R_{F} from RELHICs have the potential to provide competitive constraints on y_{P}, provided that the individual emission line fluxes can be determined to a precision of $\delta_{\mathrm{F}} \approx 3\%$ or better. As discussed, we expect this conclusion to be insensitive to the details of the gas distribution provided that it remains optically thick.

3.3.2 Determining the UVB spectral slope

In Section 3.3.1, we assumed that the UVB is known (and is given by the MH15 spectrum) in order to identify the effects of varying y_{P} in isolation. In reality, the UVB spectral shape is poorly constrained at $z \sim 0$, as demonstrated by the variance between different UVB synthesis models (see e. g. Faucher-Giguère et al. 2009; Madau & Haardt 2015; Puchwein et al. 2019). These discrepancies may be further compounded by the uncertain impact of inhomogeneities in the UVB resulting from local sources. In S19, we explored the effects of varying the UVB spectral slope on the properties of H α rings, finding that a harder UVB produced brighter rings at higher characteristic halo masses, and vice versa for a softer UVB. The properties of helium rings will also be affected by the UVB slope, since a harder spectrum contains a greater proportion of helium-ionizing photons and will therefore produce brighter helium emission at fixed $y_{\rm P}$. As discussed in Section 3.1, existing measurements of $y_{\rm P}$ approach a precision of 1%, and so it is reasonable to take this value as exact, and instead use the observable properties of helium rings to infer the UVB slope.

As in S19, we parameterise the UVB slope using the shape parameter $\alpha_{\rm UV}$ introduced by Crighton et al. (2015a), which modifies the slope of a given reference spectrum as follows:

$$J_\nu(E) = \begin{cases} N_\Gamma \times J_{\nu,\rm ref}(E) & E \leq E_0 \\ N_\Gamma \times J_{\nu,\rm ref}(E) \times (E/E_0)^{\alpha_{\rm UV}} & E_0 < E \leq E_1 \\ N_\Gamma \times J_{\nu,\rm ref}(E) \times (E_1/E_0)^{\alpha_{\rm UV}} & E > E_1 \end{cases}, \tag{3.3.9}$$

where $J_{\nu,\rm ref}(E)$ is the mean intensity of the reference spectrum at energy E, and E_0 and E_1 are pivot points between which we modulate the mean intensity by an additional power law with exponent $\alpha_{\rm UV}$. We set $E_0 = 1$ Ryd and $E_1 = 10$ Ryd, and introduce an additional factor $N_\Gamma \equiv \Gamma_{\rm HI,fid}/\Gamma_{\rm HI}$. This acts to renormalise the spectra such that the H I photoionization rate of the modified UVB is the same as that of the reference spectrum, which for our purposes is the MH15 UVB. In addition to our fiducial model which corresponds to $\alpha_{\rm UV} = 0$, we compute models for RELHICs illuminated by UVBs with $\alpha_{\rm UV} = (-2, -1, 1)$. These values are chosen to cover all realistic UVB spectra between the extremes of a soft, starburst-driven spectrum ($\alpha_{\rm UV} = -2$) and a hard, AGN-dominated spectrum ($\alpha_{\rm UV} = 1$). As described at the start of Section 3.3, for each value of $\alpha_{\rm UV}$ we iteratively perform calculations for haloes of different masses to obtain a model RELHIC with H I column density equal to the adopted threshold $N_{\rm HI,max} = 10^{23}$ cm^{-2}. We calculate values for $R_{\rm F}$

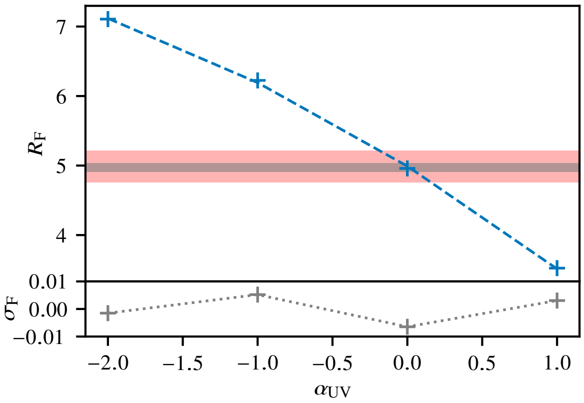

Figure 3.6

Upper panel: Flux ratio R_F as Fig. 3.4, but for models with fiducial value $y_P = 0.833$ and different UVB slopes α_{UV}. The dashed line indicates a quadratic fit $R_F = A\alpha_{UV}{}^2 + B\alpha_{UV} + C$, where $A = -0.138$, $B = -1.34$ and $C = 4.99$. Red and grey shading indicates the range of R_F values consistent with ±5% and ±1% variations of y_P respectively. *Lower panel*: The relative error σ_F between the calculated flux ratios and the fitted line.

as described previously, and plot these as a function of α_{UV} in Fig. 3.6. In red (grey) shading, we show the range of values of R_F resulting from ±5% (1%) variations in y_P, as plotted in Fig. 3.5. We see that changing α_{UV} results in a much wider range of R_F values than changing y_P. Thus, if the UVB slope deviates significantly from that of the MH15 spectrum, R_F will change from its fiducial value by a greater margin than could be caused by any reasonable uncertainty in y_P, allowing the two effects to be distinguished.

We use a quadratic fit to describe the variation of R_F with α_{UV}, which reproduces the data to a precision of 1% or better, as shown in the bottom panel of Fig. 3.6. In the same manner as was done for y_P, we use this fit to compute the precision with which α_{UV} may be reconstructed from uncertain measurements of R_F. This is shown in Fig. 3.7, where we plot $\delta_{\alpha_{UV}}$, the absolute measurement uncertainty in α_{UV}, as a function of δ_F. We find that SNRs of 10 and 100 yield constraints on α_{UV} of ${}^{+0.46}_{-0.51}$ and ${}^{+0.048}_{-0.048}$ respectively. The red and grey shaded regions now show the range of values of $\delta_{\alpha_{UV}}$ for which the corresponding values of R_F could also be obtained in models with $\alpha_{UV} = 0$ and $f_y \neq 1$. Thus, if the underlying UVB slope deviates from the fiducial value by an amount $\delta_{\alpha_{UV}} \lesssim 0.2$, the expected change in R_F is degenerate with that attributable to ±5% changes in y_P.

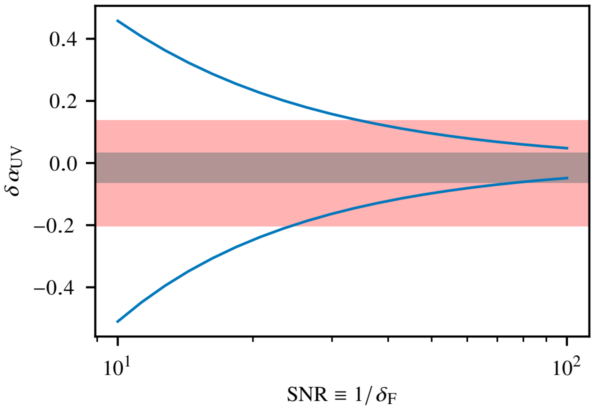

Figure 3.7
The absolute error with which $\alpha_{\rm UV}$ can be inferred as a function of $\delta_{\rm F}$. The true UVB slope parameter is assumed to be $\alpha_{\rm UV} = 0$. Red and grey shading indicates the values of $\delta_{\alpha_{\rm UV}}$ for which the expected change in $R_{\rm F}$ is degenerate with that caused by changes in $y_{\rm P}$ of 5% and 1% respectively.

3.3.3 Combined constraints on $y_{\rm P}$ and $\alpha_{\rm UV}$

We have demonstrated that it is feasible to determine either $y_{\rm P}$ or $\alpha_{\rm UV}$ using the hydrogen-to-helium flux ratio $R_{\rm F}$, assuming perfect knowledge of the other property. However, this is not representative of the real-world scenario in which both $y_{\rm P}$ and $\alpha_{\rm UV}$ are uncertain, as illustrated by the degeneracy visible in Fig. 3.7 and discussed above. In this section we investigate the possibility of simultaneously constraining $y_{\rm P}$ and $\alpha_{\rm UV}$.

In order to do this, it is necessary to break the degeneracy between $y_{\rm P}$ and $\alpha_{\rm UV}$, which both influence the value of $R_{\rm F}$. Since harder UVB spectra will contain more He II-ionizing photons, we expect the ratio

$$R_{\rm He} \equiv \frac{F({\rm He\,II}\,4686{\rm \AA})}{F({\rm He\,I}\,10830{\rm \AA}) + F({\rm He\,I}\,5876{\rm \AA})} \tag{3.3.10}$$

to increase with increasing $\alpha_{\rm UV}$. In contrast, changing $y_{\rm P}$ scales all the helium ion abundances equally, and so will only have a minor, indirect effect on $R_{\rm He}$ arising from the slight change to the abundance of free electrons, which affects the He I and He II emissivities differently. We supplement our existing models, which vary either $y_{\rm P}$ or $\alpha_{\rm UV}$ while keeping the other parameter constant, with additional runs of our photoionization code in which both $y_{\rm P}$ and $\alpha_{\rm UV}$ are varied. We calculate values of $R_{\rm F}$ and $R_{\rm He}$ for this

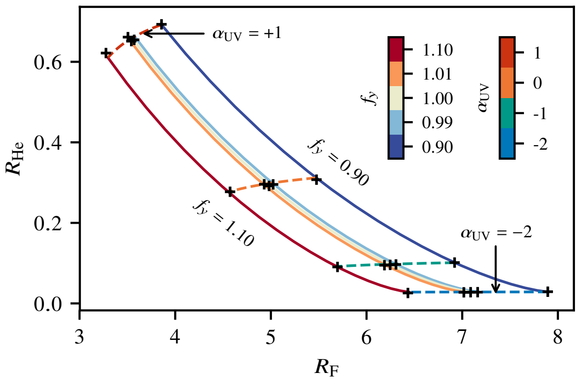

Figure 3.8

$R_{\rm F}$ (Eq. 3.3.7) vs. $R_{\rm He}$ (Eq. 3.3.10) for models with different $y_{\rm P}$ and $\alpha_{\rm UV}$. Black points indicate values calculated from runs of our ionization balance code; curves show best fits to this data using Eqs. 3.3.11 and 3.3.12. Curves at constant $y_{\rm P}$ and different values of $\alpha_{\rm UV}$ are shown with solid lines, coloured according to the left-hand colourbar. Conversely, curves at constant $\alpha_{\rm UV}$ and varying $y_{\rm P}$ are shown as dashed lines and coloured according to the right-hand colourbar.

grid of models, which we show as the black points in Fig. 3.8. As in the one-dimensional cases presented previously, we next fit a 2D surface to the calculated flux ratios, in order to allow interpolation of the ratios for arbitrary values of $y_{\rm P}$ and $\alpha_{\rm UV}$. We define these fits as follows:

$$R_{\rm F} = f_1(y_{\rm P})\, g_1(\alpha_{\rm UV}) \tag{3.3.11}$$

$$R_{\rm He} = f_2(y_{\rm P})\, g_2(\alpha_{\rm UV})\,, \tag{3.3.12}$$

where g_1, g_2 and f_2 are cubic polynomials, and $f_1 \propto 1/y_{\rm P}$ as in Section 3.3.1. We emphasise that these fits are intended to be empirical only, and are chosen for their simplicity. Nevertheless, they are able to reproduce the flux ratios obtained from our simulations to an accuracy of 3% or better across the range of $y_{\rm P}$ and $\alpha_{\rm UV}$ values we consider. We evaluate these fits and plot the resulting curves in Fig. 3.8 in order to show the degree to which they reproduce the data.

As a result of the non-linear mapping between the $R_{\rm F}$–$R_{\rm He}$ and $y_{\rm P}$–$\alpha_{\rm UV}$ axes, we employ a Monte Carlo technique to estimate the precision with which the latter parameters may be recovered. We begin by choosing the magnitudes of the uncertainties with which $R_{\rm F}$ and $R_{\rm He}$ are measured, which we derive by choosing a single flux error $\delta_{\rm F}$ and propagating this

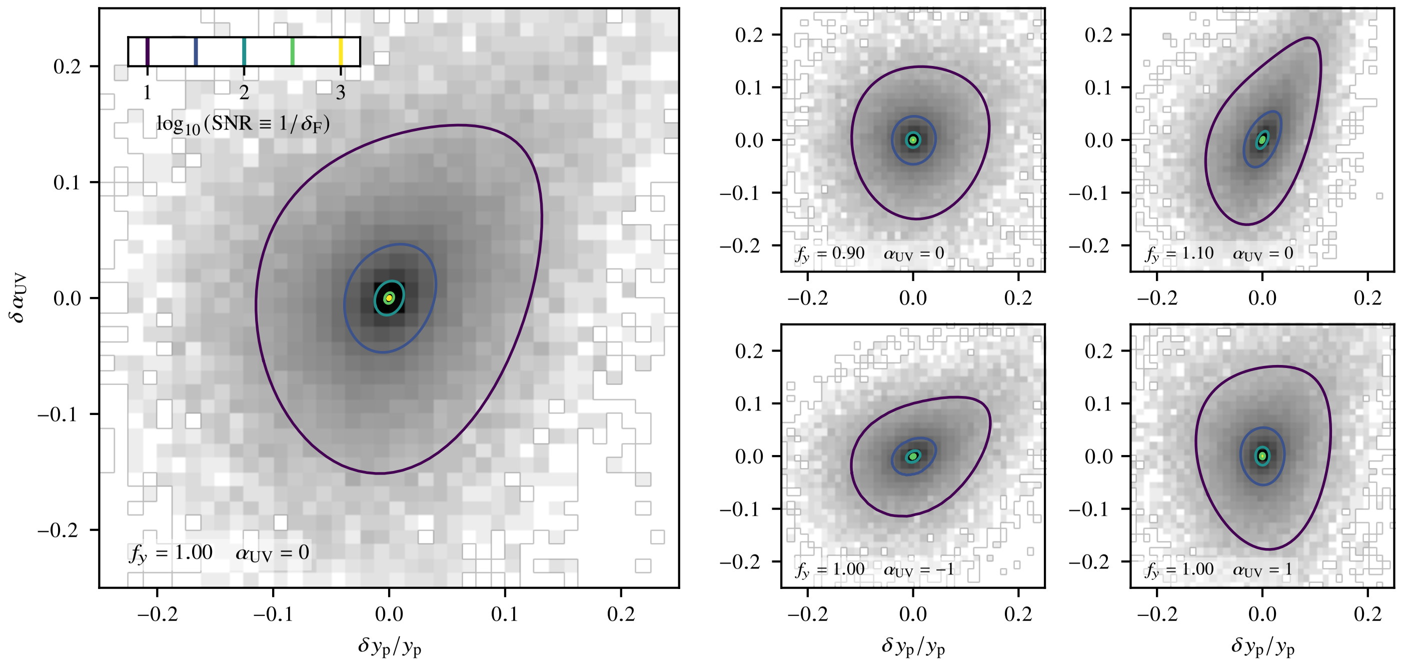

Figure 3.9

Left panel: Histogram of precision with which y_P and α_UV may be recovered from sampled R_F and R_He values. Bins are shaded according to their occupancy using a logarithmic scale. Contours indicate the precision obtained for 1σ flux measurement uncertainties in the range $0.001 \leq \delta_\mathrm{F} \leq 0.1$. The y-axis shows the absolute measurement uncertainty in α_UV, whereas the relative error is plotted for y_P. *Right panels*: As left, but for different underlying y_P and α_UV values, as indicated by the legend in the bottom-left of each panel.

uncertainty into $R_{\rm F}$ and $R_{\rm He}$ as described previously. Taking the underlying values of $y_{\rm P}$ and $\alpha_{\rm UV}$ to be 0.083 and 0 respectively, we sample 1000 uncertain 'measurements' of $R_{\rm F}$ and $R_{\rm He}$ by drawing from a bivariate normal distribution with means given by evaluating Eqs. 3.3.11 and 3.3.12 for the true values. Variances are set to the squares of the chosen flux ratio errors, such that these errors correspond to 1σ uncertainties. We then invert Eqs. 3.3.11 and 3.3.12 to map each of the sampled $R_{\rm F}$–$R_{\rm He}$ values to the $y_{\rm P}$–$\alpha_{\rm UV}$ plane, and repeat this procedure for different values of $\delta_{\rm F}$. In the left-hand panel of Fig. 3.9, we illustrate the resulting collection of $\delta_{y_{\rm P}}$–$\delta_{\alpha_{\rm UV}}$ samples using a two-dimensional histogram. Also shown is a series of contours indicating the 1σ limits on $y_{\rm P}$ and $\alpha_{\rm UV}$ which result from different choices of $\delta_{\rm F}$ in the range $0.001 \leq \delta_{\rm F} \leq 0.1$.

We see that the joint constraints also provide comparable or better precision than the individual ones, particularly for $\alpha_{\rm UV}$. The addition of the $R_{\rm He}$ measurement permits an improvement in the reconstructed precision of this parameter by approximately a factor of 4, with $\delta_{\rm F} = 0.1$ now yielding a constraint on $\alpha_{\rm UV}$ of ± 0.15. Conversely, the precision with which $y_{\rm P}$ may be recovered is $^{+14}_{-12}\%$, effectively unchanged to that obtained from the individual constraints (Fig. 3.5). These results are in agreement with the indication in Fig. 3.8 that $R_{\rm He}$ evolves much more strongly with $\alpha_{\rm UV}$ than with $y_{\rm P}$. Increasing the assumed measurement precision to SNR = 100 improves these constraints significantly, to give $\pm 1.3\%$ and ± 0.015 constraints on the helium abundance and UVB slope respectively.

In the right-hand panels, we repeat the Monte Carlo process outlined above, but using different underlying values of $y_{\rm P}$ and $\alpha_{\rm UV}$. We obtain comparably precise constraints in all cases shown, with SNR = 100 yielding values of $\delta_{y_{\rm P}}/y_{\rm P}$ of $\pm 1.3\%$ or better, and $\delta_{\alpha_{\rm UV}}$ of ± 0.017 or better. Furthermore, a SNR of $10^{1.5} \approx 32$, corresponding to the second-outermost contour, is always sufficient to recover $y_{\rm P}$ to the 5% level at which existing determinations of the abundance differ. The tendency for a positive correlation between $\delta_{y_{\rm P}}$ and $\delta_{\alpha_{\rm UV}}$ is again a consequence of the general shape of Fig. 3.8: a positive value for $\delta_{y_{\rm P}}$ is produced by a negative absolute error on $R_{\rm F}$. When combined with an error on $R_{\rm He}$ which is also ≥ 0, the reconstructed value of $\alpha_{\rm UV}$ exceeds the assumed underlying value, and therefore $\delta_{\alpha_{\rm UV}}$ is positive also. Equivalently, a positive error on $R_{\rm F}$ and a negative error on R_{He} combine to yield inferred values of y_{P} and α_{UV} that lie below the true ones. Curves at constant α_{UV} (defined parametrically by Eqs. 3.3.11 and 3.3.12) flatten toward lower y_{P}, meaning this correlation is not as pronounced and giving the constraint contours shown in Fig. 3.9 their ovoid shape.

3.4 Summary and conclusions

We have examined the properties of hydrogen and helium emission driven by UV background fluorescence in RELHICs, a class of $\sim 10^{9.5}\,\mathrm{M}_{\odot}$ dark matter haloes which fail to form stars, instead retaining a small reservoir of neutral, essentially-pristine gas at redshift $z = 0$. Using results obtained from a specialised radiative transfer code, we have shown that the ratio of hydrogen to helium emission relates directly to the helium abundance of the gas. In particular, we showed that from ratios of integrated quantities, such as the surface brightness and integrated flux, we are able to recover the assumed helium abundance to 1% or better. Hence, these measurements have the potential to provide an independent measurement of the primordial helium abundance.

RELHICs are intrinsically simple systems, making them ideally suited for assessing the presence of systematic errors in the canonical method for measuring y_{P} using metal-poor H II regions. For example, the expected almost-pristine nature of the gas in RELHICs avoids the need to extrapolate observed helium abundance measures down to zero metallicity. Their well-specified structure, in which the majority of emission is produced by gas which settles at its photoionization equilibrium temperature, and at a density dictated by the requirement of hydrostatic equilibrium, reduces the impact of systematics which can arise from uncertainties in the temperature and density structure of H II regions, as well as in the degree to which they are chemically homogeneous (Izotov et al. 2007). Additionally, the much lower typical density of the gas within RELHICs means that all emissivities may be calculated fully in the low-density limit, which can result in up to an order-of-magnitude reduction in their associated uncertainties (Porter et al. 2009). Finally, this approach provides the added bonus of permitting the spectral slope of the $z = 0$ UVB to be inferred, as discussed in Sections 3.3.2 and 3.3.3. Intrinsically ‘dark’ sources like RELHICs are uniquely positioned to allow such a measurement to be made, as any attempt to infer the slope of the UVB using the nebular emission from luminous sources requires the subtle effect of the UVB to be disentangled from the effect of the locally-produced radiation field.

While the strengths of this approach are promising, significant challenges also exist. We make a number of modelling assumptions, such as assuming RELHICs to be spherically-symmetric and in hydrostatic equilibrium with a gravitational potential due solely to their dark matter content. BL17 reports that the first two assumptions are in agreement with the properties of RELHICs identified in APOSTLE, while neglecting the gas self-gravity is justified since $M_{200} \gg M_{\rm g}$ for these systems. More significantly, while the existence of dark matter haloes in the mass range corresponding to RELHICs is a robust result of CDM structure formation, the prediction that they remain star-free but gas-rich is less certain. The limited spatial and mass resolution of cosmological simulations means that they are unable to follow the formation of individual stars. In addition, following the physical processes governing the formation of a cold, molecular gas phase is computationally intensive. Thus, the APOSTLE simulations instead enforce an effective equation of state for cool gas, and consider this gas to be eligible for star formation when it exceeds a metallicity-dependent density threshold, as proposed by Schaye (2004). For the extremely low-metallicity gas RELHICs contain, this threshold is set to $n_{\rm H,th} = 10\,{\rm cm}^{-3}$, which we do predict to be exceeded in the cores of the most-massive RELHICs. However, as noted by BL17, the Schaye (2004) prescription is strictly valid only for metallicity $Z \geq 10^{-4}\,Z_{\odot}$, and diverges for lower metallicities.

The value of $n_{\rm H,th}$ predicted for RELHICs is a somewhat arbitrary value imposed to avoid this behaviour. Consequently, a rigorous investigation of the conditions under which a molecular phase may form in pristine gas would require a self-consistent treatment of the relevant atomic processes in our radiative transfer code, which we do not attempt to implement here. As detailed in §2.3 of S19, we have instead considered H_2 formation as a post-processing step, finding that our adopted column density threshold of $N_{\rm HI,\,max} = 10^{23}\,{\rm cm}^{-2}$ corresponds to the threshold above which formation of H_2 occurs. Additionally, the upper bound on halo mass of $M_{200} \leq 10^{9.6}\,\mathrm{M}_{\odot}$ that this threshold implies is consistent with the masses found for the largest RELHICs in APOSTLE. While more detailed modelling may result in refinements to our predictions, we expect the existence of a window in halo mass for which predominantly 'dark' haloes may contain optically-thick gas to be robust to these changes. Furthermore, we have shown that provided this assumption holds, the results presented here are insensitive to the precise column density threshold (and hence mass scale) chosen.

We must also address the fact that RELHICs are an entirely theoretical prediction, and discuss the prospects for their detection via observations. RELHICs exhibiting the brightest fluorescent rings are expected to be intrinsically rare, due to the narrow range of halo masses these objects may have. In S19, we used APOSTLE to obtain a predicted count of $3^{+2.6}_{-2.0}$ RELHICs with $\Sigma_{\mathrm{H}\alpha,\mathrm{max}} > 10^{-20}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$ and a projected ring diameter $\geq$1 kpc located within a 3 Mpc volume centred on the Milky Way. This rarity, coupled with the fact that even the brightest emission from RELHICs is still very faint for current technology, means detecting them is challenging at present. As such, a blind H α survey using current instrumentation (e. g. the MUSE instrument at the VLT) is likely unfeasible, requiring several tens of hours of integration time per field. However, there remain reasons for optimism. By nature, RELHICs contain substantial reservoirs of neutral hydrogen, making them bright H I 21cm emitters. They are therefore expected to appear in existing deep H I surveys, such as ALFALFA (Giovanelli et al. 2005) and HALOGAS (Heald et al. 2011). In particular, the catalogue of ultra-compact high velocity clouds identified in ALFALFA (Adams et al. 2013) have properties consistent with the expected H I morphology of the most massive RELHICs, as demonstrated by Benítez-Llambay et al. (2017). Furthermore, planned surveys with the Square Kilometre Array and its precursors (e. g. MEERKAT; de Blok et al. 2018) will permit detection of H I sources with column densities down to $N_{\mathrm{HI}} \sim 10^{16}\,\mathrm{cm}^{-2}$ (Popping et al. 2015; Power et al. 2015). This level of sensitivity is sufficient to yield H I detections of all but the lowest-mass RELHICs in the Local Group.

The H I catalogues produced by these surveys may be used in conjunction with deep broad-band imaging to identify 21cm sources with no associated stellar continuum as promising targets. Ultra-deep pointed observations or stacking analysis of objects selected in this way have the potential to reveal the presence of one or more RELHICs, and could be used to obtain the measurements necessary to apply the techniques we have discussed here. Thus, RELHICs remain a promising target for further study. Their successful detection would not only constitute an additional verification of the prevailing CDM cosmological model, but as we have shown in this work, would also provide new insight into properties as disparate as the composition of the primordial Universe and the low-redshift intergalactic radiation environment of the Local Group.

CHAPTER 4

Exploring neutral gas with EAGLE: the ENGinE simulations

4.1 Introduction

As the most abundant element in the Universe, hydrogen comprehensively traces the presence of cosmic structure. The distribution of this gas can be inferred via absorption features, in particular those due to the $n = 2 \rightarrow 1$ Lyman-α (Ly α) transition, seen in the spectra of background UV-bright sources such as quasars. Weak Ly α absorbers situated within the intergalactic medium (IGM), associated with neutral hydrogen column densities N_{HI} in the range 10^{12}–$10^{16}\,\mathrm{cm}^{-2}$, form the ubiquitous Ly α forest, which has become a powerful cosmological probe through enabling a measurement of the baryonic acoustic oscillation feature (des Bourboux et al. 2020; Alam et al. 2021). The Ly α forest also contains a record of the thermal history of the IGM, from which constraints on structure formation and reionisation may be obtained (as reviewed, e. g., by McQuinn 2016). At column densities of $N_{\mathrm{HI}} = 10^{17.2}\,\mathrm{cm}^{-2}$ and above (the so-called Lyman limit, corresponding to optical depth at the H I ionisation edge $\tau_{912} = 1$), Ly α absorbers begin to self-shield their interiors from the background ionising radiation field (Tytler 1982; Sargent et al. 1989). Beyond a further column density threshold of $10^{20.3}\,\mathrm{cm}^{-2}$, the incoming ionising flux has been sufficiently attenuated for the absorbing gas to become predominantly neutral

(Wolfe et al. 2005) and essentially opaque to Ly α photons. The spectra of these strong absorbers show saturated Ly α absorption at line centre, and characteristic Lorentzian damping wings, leading to their classification as damped-Ly α systems (DLAs). They constitute the high-redshift Universe's dominant reservoir of neutral gas (Tytler 1987), and are therefore implicated as the providers of fuel for star formation (Prochaska & Wolfe 2009; Jorgenson et al. 2009; Rafelski et al. 2011).

DLAs have proven challenging to study, with their intrinsic rarity, the requirement for UV spectroscopy at redshift $z < 1.6$, and the need to distinguish them from the surrounding Ly α forest all presenting challenges. Only with the advent of large catalogues of quasar spectra, such as those provided by the Sloan Digital Sky Survey (SDSS), has it become possible to construct statistically robust DLA catalogues (e. g. Noterdaeme et al. 2012; Garnett et al. 2017; Ho et al. 2020, 2021). Furthermore, spectroscopic detections of DLAs provide little insight into the properties of the galaxies with which they are presumed to associate, which has been the source of considerable historic uncertainty as to their nature. Early work favoured the identification of DLAs as high-redshift analogues to present-day spiral galaxies, with rapidly-rotating thick disks of H I (Wolfe et al. 1986; Prochaska & Wolfe 1997, 1998). This was in keeping with the conception of galaxy formation resulting from coherent collapse (e. g. Eggen et al. 1962), but presented a challenge for the hierarchical structure formation models then coming to prominence. Concurrently, pioneering numerical simulations demonstrated that a distribution of irregular protogalactic clumps, more agreeable to hierarchical models, could also reproduce the structure of observed absorption features (Haehnelt et al. 1998; Gardner et al. 1997, 2001), although explaining the observed abundance of high velocity-width DLAs has proven to be a persistent challenge (Prochaska & Wolfe 2001; Razoumov et al. 2008; Tescari et al. 2009; Hong et al. 2010).

For many years, attempts to resolve this uncertainty by imaging the galaxies hosting DLAs in emission were met with relatively little success (e. g. Warren et al. 2001; Møller et al. 2004; Fumagalli et al. 2015), although searches for galaxy counterparts to high-metallicity DLAs conducted with the infrared spectrograph X-shooter have yielded positive results

(Fynbo et al. 2010, 2011). The advent of integral field unit (IFU) instruments, in particular the Multi-Unit Spectroscopic Explorer (MUSE; Bacon et al. 2010) instrument at the Very Large Telescope, revolutionised searches for DLA hosts, by permitting spectroscopic follow-up of multiple candidate hosts at different impact parameters to the quasar with a single telescope pointing (Péroux et al. 2012; Jorgenson & Wolfe 2014; Fumagalli et al. 2017b; Mackenzie et al. 2019). Meanwhile, advances in sensitivity at radio wavelengths due to the Atacama Large Millimetre Array (ALMA; Wootten & Thompson 2009) have enabled complementary observations of DLA hosts via their CO (Neeleman et al. 2018; Fynbo et al. 2018; Kanekar et al. 2018, 2020) or C II (Neeleman et al. 2017, 2019) emission.

Assembling a coherent picture from these observations remains difficult. The radio detections indicate that some DLAs do associate with high-redshift disk galaxies, but these techniques are presently limited to observing luminous, metal-rich systems. IFU observations present a more varied picture, with many detections occurring at relatively large impact parameters, suggesting either the presence of faint low-mass galaxies that remain undetected, or extended circumgalactic material. Evidence has been found for a strong scaling between DLA metallicity and host luminosity, lending some support to the former hypothesis: the low typical metallicity of DLAs implies that many hosts will have Ly α fluxes well below current detection thresholds (Krogager et al. 2017). Thus, observations appear to indicate a heterogeneous population of DLA host galaxies, which sample a wide range of masses, luminosities, and environments.

Hydrodynamical simulations also provide an invaluable tool for understanding H I absorbers, with simulations of cosmological volumes successfully reproducing the overall abundance of HI absorbers, its redshift evolution, and other properties such as the metallicity distribution (Pontzen et al. 2008; Altay et al. 2011, 2013; Bird et al. 2014; Rahmati et al. 2013a, 2015). Concurrently, analyses focusing on individual haloes in greater detail have provided insight into the nature of absorbing gas (Kereš et al. 2005; Fumagalli et al. 2011; Faucher-Giguère & Kereš 2011; van de Voort et al. 2012; Faucher-Giguère et al. 2015). Several model ingredients with significant effects on the simulated DLA population have been identified, chief among them being the implementation of stellar feedback. A number of studies, beginning with Nagamine et al. (2004), have reported that efficient stellar feedback capable of driving galactic winds is required in order to reproduce the overall distribution of DLAs (Nagamine et al. 2007; Bird et al. 2014; Faucher-Giguère et al. 2015; Suresh et al. 2015). In addition, feedback-driven winds have been suggested as a solution to the longstanding issue of reproducing the kinematics of wide DLAs (Cen 2012; Bird et al. 2015). Feedback has also been invoked in order to alleviate the apparent tensions between observed and modelled values of the DLA bias. Measurements of this quantity yield a value of $b_{\rm DLA} = 1.99 \pm 0.11$ (Pérez-Ràfols et al. 2018a), corresponding to a mean DLA host mass above 10^{11} $\rm M_{\odot}$. Taken at face value, simulations favour lower values (Pontzen et al. 2008; Padmanabhan et al. 2016, 2017), implying that strong feedback is needed in lower-mass haloes to suppress the formation of DLAs in these systems (Barnes & Haehnelt 2014, 2015). However, this interpretation is not unique, with alternate explanations for the discrepancy including the effects of non-linear growth on the DLA power spectrum (Bird et al. 2014) or the existence of a correlation between $b_{\rm DLA}$ and metallicity (Pérez-Ràfols et al. 2018b).

Building on these results, which highlight the importance of feedback in shaping the properties of simulated DLAs, in this chapter we explore how the properties of individual galaxies combine to produce the global DLA statistics, and how these contributions are modified by feedback. To do this, we exploit a novel set of hydrodynamical simulations based on the EAGLE galaxy formation model (Schaye et al. 2015; hereafter S15). Through the use of physically-motivated subgrid models providing efficient feedback from stars and active galactic nuclei (AGN), EAGLE has been shown to generate realistic model galaxies that successfully reproduce a variety of galaxy properties and scaling relations (e. g. S15; Furlong et al. 2015; Crain et al. 2015, 2017). Importantly, these successes are achieved independently of the calibration procedure and so represent genuine *ab-initio* predictions of the simulation, permitting us to meaningfully compare the simulated galaxies and their associated neutral gas. Rahmati et al. (2015) have previously calculated H I statistics for the reference EAGLE simulation, finding agreement with observations for the global column density distribution, as well as the covering fractions of strong H I absorption within haloes at $z \sim 2$–3. The existing EAGLE suite also includes runs in which alterations are made to the subgrid parameters governing the feedback implementation, but the numerical resolution and simulation box size varies between these runs, complicating their direct comparison. Accordingly, we have chosen to run a new series of simulations which share a common 50 Mpc box size and an enhanced resolution, which exceeds that of the reference simulation by factors of two and eight in terms of spatial and particle mass resolution respectively. We have named these new simulations the ENGINE (Exploring Neutral Gas in EAGLE) suite, and introduce them fully in Section 4.2. An important additional ingredient for studying the simulated H I distribution is a prescription for the self-shielding of dense absorbers. We accomplish this with a post-processing step in which the radiative transfer code URCHIN (Altay & Theuns 2013) is used to correct the optically-thin behaviour assumed during the simulation, allowing the gas distribution to be partitioned into ionised, atomic, and molecular phases. This procedure is described in detail in Section 4.3. In Section 4.4, we use the post-processed simulation outputs to compute the global absorber distribution, quantified by the column density distribution function (CDDF). We explore the impact of stellar feedback and the atomic–molecular transition on the CDDF, as well as its redshift evolution. Lastly, in Section 4.5 we investigate the absorption properties of individual haloes, again considering their evolution with time and their feedback dependence, and show how these trends are responsible for those that appear in the global statistics.

4.2 Simulations

To be of use for studying the properties of H I absorbers, a cosmological simulation should ideally satisfy a number of criteria. It must possess sufficient resolution to capture the small scales relevant to the dense gas within absorption systems, while also containing sufficient volume to give good statistics over several orders of magnitude in halo mass. Additionally, it must include “sub-grid” recipes to describe relevant processes occurring below the resolution limit, with the treatment of star formation and its associated feedback being of particular importance. Finally, the predominantly neutral ionisation state of the gas making up absorption systems with $N_{\rm HI} \gtrsim 10^{19}\,{\rm cm}^{-2}$ means that, in contrast to the commonly-employed assumption that gas remains optically thin, the effects of self-shielding must be taken into account. In this section, we introduce the simulations we use to study H I absorbers, and describe how they are engineered to meet, as much as possible, these requirements.

We present ENGinE, a suite of six simulations constituting an extension to the Evolution and Assembly of GaLaxies and their Environments (Eagle) project (S15; Crain et al. 2015). The ENGinE simulations were run with the same extensively modified version of the smoothed particle hydrodynamics (SPH) code gadget-3 (last described by Springel 2005) that was used to run Eagle. In addition to improvements to the hydrodynamics and time-stepping schemes (Durier & Dalla Vecchia 2012), the Eagle flavour of gadget-3 incorporates a comprehensive set of subgrid prescriptions for modelling processes below the resolution limit. These include radiative cooling (Wiersma et al. 2009a), chemical enrichment (Wiersma et al. 2009b), star formation (Schaye & Dalla Vecchia 2008) and stellar feedback (Dalla Vecchia & Schaye 2012), and black hole accretion (Rosas-Guevara et al. 2015) and feedback (Booth & Schaye 2009). ENGinE uses the same calibrated values for the free parameters in these models as Eagle, which were chosen by requiring that the simulations reproduce observations of the $z = 0.1$ galaxy stellar mass function (GSMF), the galaxy size–mass relation, and the black hole mass–stellar mass relation, as detailed in S15. We further assume the same flat ΛCDM cosmology employed by S15, with cosmological parameters taken from the Planck Collaboration et al. (2014) results ($\Omega_{\rm M} = 0.307$, $\Omega_\Lambda = 0.693$, $\Omega_{\rm b} = 0.04825$, $h \equiv H_0/(100{\rm km\,s}^{-1}\,{\rm Mpc}^{-1}) = 0.6777$, $\sigma_8 = 0.8288$).

All six ENGinE simulations feature a boxsize of $L_{\rm box} = 50\,{\rm cMpc}$, and with one exception (Recal-0752; see below) were run with 1504^3 DM particles of mass $M_{\rm DM} = 1.21 \times 10^6\,{\rm M}_\odot$. There are initially an equal number of baryonic particles, with mass $M_{\rm bar} = 2.26 \times 10^5\,{\rm M}_\odot$, some of which are converted into star particles or accreted by black holes during the course of the simulation. This configuration results in a equivalent mass and spatial resolution to the Recal-L0025N0752 simulation from S15, while increasing the volume by a factor

Table 4.1

Properties of the ENGɪɴE simulations. From left to right, columns show the name by which we refer to each simulation; the number of dark matter particles; the DM particle mass and initial baryonic particle mass; the Plummer-equivalent gravitational softening; and the minimum and maximum asymptotes of the stellar feedback efficiency (see Section 4.2.1.2).

Name	$N_{\rm part}$	$M_{\rm DM}$ ($M_\odot$)	$M_{\rm bar}$ ($M_\odot$)	ϵ (ckpc)	$f_{\rm th,\,min}$	$f_{\rm th,\,max}$	Notes
Recal	1504^3	1.21×10^6	2.26×10^5	1.33	0.3	3.0	–
Recal-0752	752^3	9.70×10^6	1.81×10^6	2.66	0.3	3.0	–
StrongFB	1504^3	1.21×10^6	2.26×10^5	1.33	0.6	6.0	–
WeakFB	1504^3	1.21×10^6	2.26×10^5	1.33	0.15	1.5	(1)
NoAGN	1504^3	1.21×10^6	2.26×10^5	1.33	0.3	3.0	(2)
WDM	1504^3	1.21×10^6	2.26×10^5	1.33	0.3	3.0	(3)

(1) This run had to be continued from a snapshot at $z = 5.97$ due to a hardware failure, which corrupted the restart-files.
(2) In this run the formation of BH particles was suppressed such that no AGN feedback could occur.
(3) This run used a WDM cosmology with DM particle mass $m_{\rm WDM} = 7.1$ keV.

of eight. Comparing instead to the existing 50 cMpc simulation, Ref-L0050N0752, the five high-resolution ENGɪɴE volumes increase the spatial and mass resolution by factors of two and eight respectively. This combination of large volume and high resolution makes the ENGɪɴE simulations well-suited for our purpose; however, it also makes them computationally expensive. To mitigate this cost, and in light of the fact that HI absorption systems are most readily observed at redshifts where the rest frame ultraviolet Lyman-α line is redshifted into optical wavelengths, we have run the simulations only to $z = 2$.

In addition to a fiducial simulation (Recal-L0050N1504 in the nomenclature of S15; we will simply refer to this run as Recal[1]), we present a number of physics variations, as summarised in Table 4.1. These comprise the WeakFB (StrongFB) runs in which the average efficiency of star formation is decreased (increased) by a factor of 2; the NoAGN run, in which the minimum halo mass required to spawn a black hole particle was set to an unphysically-large value, such that black hole formation and the associated AGN feedback were effectively disabled; and the WDM run, in which an alternate Λ-warm dark matter (WDM) model was assumed. The specific WDM model we implement is a 7.1 keV

---

[1] The term ‘Recal’ is used to indicate that to compensate for the change in numerical resolution, the subgrid parameters have been adjusted such that agreement with the calibration observables is retained.

resonantly-produced sterile neutrino (see Boyarsky et al. 2012 for a review), which is of particular interest due to a number of reported detections of an anomalous 3.55 keV X-ray line consistent with production via the decay of a particle with this mass (Iakubovskyi 2016). Requiring consistency with these observations, as well as with considerations from structure formation and the constraint that the sterile neutrino should constitute the entirety of the dark matter, yields a range of allowed values for the lepton asymmetry parameter L_6, as defined by Lovell et al. (2016). This parameter sets the 'warmness' of the dark matter, and we choose a value of $L_6 = 11.2$, corresponding to the warmest particle allowed by the observations (Lovell et al. 2016, 2017). Finally, we performed the Recal-0752 simulation to investigate the sensitivity of HI absorber statistics to the subgrid calibration procedure. For this run, the particle resolution is returned to the reference EAGLE level, but the subgrid parameters are kept at the recalibrated values appropriate for the higher resolution. Having introduced the full ENGINE suite, in this chapter we will focus on results from the Recal, StrongFB, and WeakFB simulations, deferring analysis of the remaining volumes to future work.

To generate initial conditions for the simulations, the transfer function $T(k)$ for the cosmology we adopt was generated using the code CAMB (version Jan_12; Lewis et al. 2000), from which a matter power spectrum was obtained by multiplying a primordial power spectrum with spectral index $n_s = 0.9611$ by $T(k)^2$ evaluated at $z = 0$. For the WDM simulation, a modified version of CAMB was used to generate a transfer function appropriate for the chosen WDM model. We then extract a Gaussian random field with this power spectrum from PANPHASIA (Jenkins 2013; Jenkins & Booth 2013), and finally use the '2LPT' method of Jenkins (2010), to generate the initial conditions by displacing particles from a glass-like (White 1994) distribution. For each run, we produce a complete snapshot of the simulation state at 15 intervals between $z = 20$ and $z = 2$. We identify haloes by running the friends-of-friends (FoF) algorithm on the DM particles recorded in the snapshots, setting the linking length to 20% of the mean separation between particles. Gas and star particles are assigned to the FoF halo containing the closest DM particle. We then use the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009) to identify individual substructures within each FoF halo, defining ‘galaxies’ as all structures resulting from this procedure that contain a nonzero baryonic component.

4.2.1 Subgrid models

In this study, we aim to explore how different physical processes affect the distribution of neutral hydrogen in and around galaxies, which gives rise to the structures seen in absorption. In particular, we will investigate the effects of star formation and feedback; accordingly, we next provide a brief description of the subgrid implementation used by EAGLE for these phenomena. More thorough explanations may be found in S15.

4.2.1.1 Star formation

At sufficiently high densities, gas is expected to become multiphase and develop cool, dense clumps within which star formation occurs. Cosmological simulations lack the resolution required to resolve the internal structure of these regions, and must therefore impose a star formation threshold above which gas becomes eligible to form stars following a subgrid model. Higher gas metallicity results in increased cooling as well as an increase in the abundance of dust, which increases the effectiveness of shielding from ionising radiation and the formation of molecular hydrogen. This permits the onset of multiphase behaviour to occur at lower typical densities, which EAGLE models with a metallicity-dependent star formation threshold, as proposed by Schaye (2004):

$$n_{\mathrm{H}}{}^{*} = 10^{-1}\mathrm{cm}^{-3}\left(\frac{Z}{0.002}\right)^{-0.64}, \tag{4.2.1}$$

where Z is the gas metallicity. In addition, the unresolved cold phase necessitates the imposition of an effective equation of state $P_{\mathrm{eos}} \propto \rho^{4/3}$. As a consequence, the temperatures stored in the simulation snapshots for star-forming gas only reflect this equation of state, which we must take into account when performing radiative transfer post-processing (see Section 4.3.1).

Star formation itself is implemented following Schaye & Dalla Vecchia (2008), whereby the star formation rate (SFR) $\dot{\rho}_\star$ in gas of density ρ is parameterised in terms of the gas pressure, p, as

$$\frac{\dot{\rho}_*}{\rho} = \frac{A}{\left(1\,\mathrm{M_\odot pc^{-2}}\right)^{-n}} \left(\frac{\gamma p}{G}\right)^{(n-1)/2} . \tag{4.2.2}$$

Here, $\gamma = 5/3$ is the adiabatic index and A and n are the parameters in the Kennicutt-Schmidt star formation law (Kennicutt 1998),

$$\dot{\Sigma}_* = A \left(\frac{\Sigma_g}{1\,\mathrm{M_\odot\,pc^{-2}}}\right)^n , \tag{4.2.3}$$

where EAGLE uses the values $A = 1.515 \times 10^{-4} M_\odot\,\mathrm{yr^{-1}kpc^{-2}}$ (Kennicutt's value converted to a Chabrier stellar initial mass function) and $n = 1.4$. $\dot{\Sigma}_*$ and Σ_g are the surface density of star formation and of gas, respectively. In this implementation, gas particles with $\log_{10}(T) < \log_{10}(T_\mathrm{eos}) + 0.5$ and $n_\mathrm{H} > n_\mathrm{H}{}^*$, where T_eos is the temperature floor obtained from the equation of state, are identified as being eligible for star formation. These particles are then stochastically converted into stars during a time step Δt with probability $\min(\dot{m}_* \Delta t / m_\mathrm{g}, 1)$, where $\dot{m}_*$ is the SFR calculated for the particle, and m_g is its total gas mass, and assuming that $\dot{\rho}_*/\rho = \dot{m}_*/m_g$.

4.2.1.2 Stellar feedback

After formation via the method outlined above, star particles are modelled as simple stellar populations (SSPs) with a Chabrier (2003) initial mass function spanning the range $0.1 - 100\,\mathrm{M_\odot}$. Stars with $M_* > 6\,\mathrm{M_\odot}$ end their lives in Type II supernovae (SNII), which inject energy and momentum into the surrounding interstellar medium (ISM). For sufficiently large rates of star formation, this feedback is capable of driving large scale outflows (Veilleux et al. 2005; Naab & Ostriker 2017). The limited numerical resolution available to cosmological simulations has proved to make reproducing this behaviour challenging, with issues such as 'overcooling' (e. g. Katz et al. 1996) inhibiting the generation of galactic winds from supernova feedback.

To address these issues, cosmological simulations must introduce a subgrid model for feedback, with EAGLE employing the stochastic thermal feedback introduced in Dalla Vecchia & Schaye (2012). In this model, the temperature jump experienced by gas particles which receive feedback energy ΔT, and the expectation value for the feedback energy injected per unit stellar mass $\epsilon_{\rm SN}$, are free parameters. Once specified (we use $\Delta T = 10^{7.5}$ K and $\epsilon_{\rm SN} = 8.73 \times 10^{15}$ erg/g [2]), the efficiency of the feedback may then be controlled by varying the fraction of $\epsilon_{\rm SN}$ that is injected per feedback event $f_{\rm th}$. This parameter is used to determine the probability that each SPH neighbour of the star particle is heated by ΔT when the particle's associated SSP reaches an age of 3×10^7 yr, which is the maximum lifetime of stars which explode as SNII.

For physically motivated reasons (Wiersma et al. 2009b; Dalla Vecchia & Schaye 2012), the EAGLE model allows $f_{\rm th}$ to vary with the local gas properties. Specifically, the expectation of increased thermal losses in metal-rich gas, and the excessive losses at high gas density due to resolution limitations, lead to the following expression for $f_{\rm th}$:

$$f_{\rm th} = f_{\rm th,min} + \frac{f_{\rm th,max} - f_{\rm th,min}}{1 + \left(\frac{Z}{0.1 Z_\odot}\right)^{n_Z} \left(\frac{n_{\rm H,birth}}{n_{\rm H,0}}\right)^{-n_n}} , \tag{4.2.4}$$

where $f_{\rm th,max}$ and $f_{\rm th,min}$ are the maximum and minimum asymptotic values of $f_{\rm th}$, $Z_\odot = 0.0127$ is the solar metallicity, $n_{\rm H,0} = 0.25\,{\rm cm}^{-3}$, and $n_{\rm H,birth}$ is the gas density inherited by the star particle from the gas particle it formed from. The fiducial values of the $f_{\rm th}$ asymptotes, which were determined during the calibration process, are $f_{\rm th,max} = 3.0$ and $f_{\rm th,min} = 0.3$; the exponents for the density and metallicity dependences are set to a value of $1/\ln 10 \approx 0.434$. Allowing $f_{\rm th}$ to exceed unity may be motivated both as a solution to the resolution issues discussed previously, and by the existence of additional, unmodelled sources of energy such as stellar winds or cosmic rays.

4.2.2 Galaxy properties

In this section, we briefly discuss the galaxy stellar mass function (GSMF) obtained from our Recal simulation. Note that because the $z = 0$ GSMF is one of the observables the

---

[2] This value corresponds to that expected for a Chabrier (2003) IMF if all stars with mass $6 - 100\,{\rm M}_\odot$ explode via core collapse supernovae, each releasing 10^{51} erg.

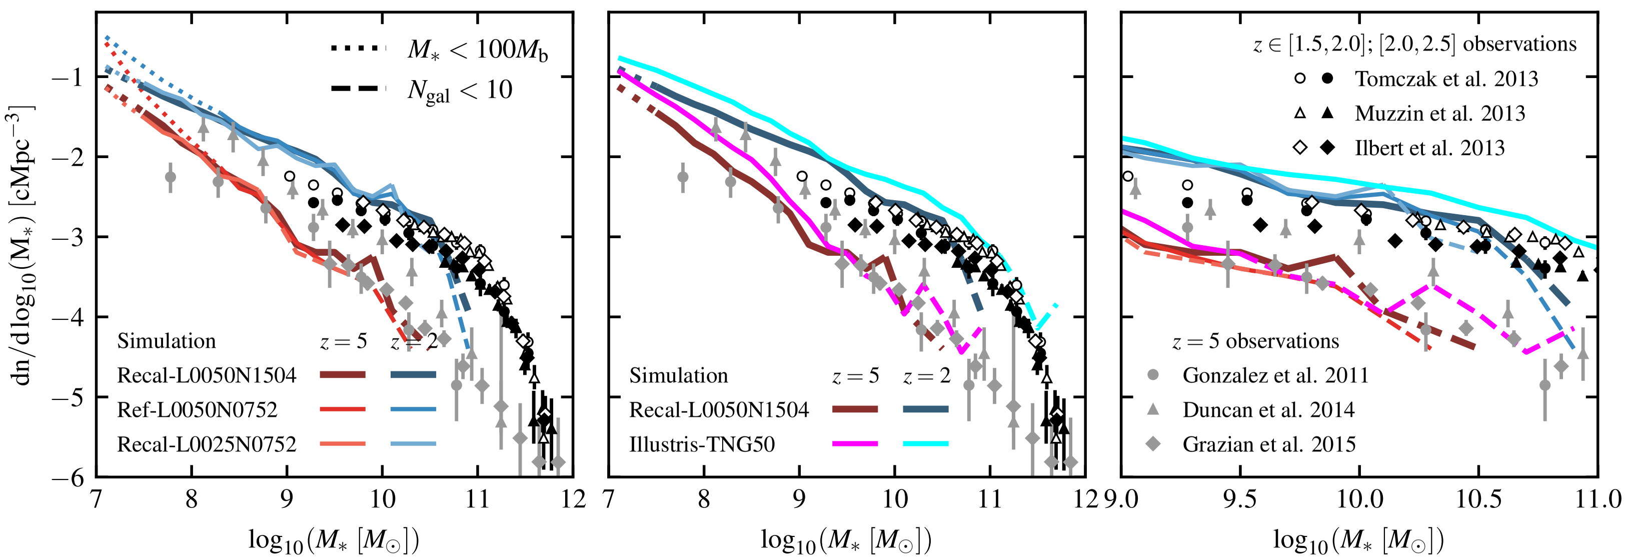

Figure 4.1

The galaxy stellar mass function (GSMF) for the Recal ENGINE simulation at $z = 5$ and 2 (thick dark red and blue lines, respectively). *Left panel*: lighter-shaded lines show the GSMFs obtained from existing EAGLE volumes featuring the same box size at reduced resolution (Ref-L0050N0752), and a smaller box simulated at the same resolution (Recal-L0025N0752). *Centre panel*: in magenta (cyan), the $z = 5$ ($z = 2$) GSMF obtained from the Illustris-TNG50 simulation (Nelson et al. 2019, 2021; Pillepich et al. 2019). *Right panel*: all simulated GSMFs are shown, with the region $9.0 < \log_{10}(M_*) < 11.0$ magnified for clarity. The simulated GSMFs are calculated assuming a 30 pkpc 3D aperture, and are shown with dotted and dashed lines to indicate regimes where the stellar mass falls below the mass corresponding to 100 baryonic particles, and where there are fewer than 10 galaxies per stellar mass bin, respectively. Data points show constraints on the observed GSMF at $z = 1.5 - 2.0$ (empty points) and $z = 2.0 - 2.5$ (black points), in both cases from the ZFOURGE (Tomczak et al. 2014), UltraVISTA (Ilbert et al. 2013) and COSMOS (Muzzin et al. 2013) surveys; and $z \sim 5$ data (grey points) from GOODS-S (González et al. 2011), CANDELS (Duncan et al. 2014) and HUDF (Grazian et al. 2015). All GSMFs assume a Chabrier (2003) IMF.

EAGLE model was calibrated against, the results shown here should not be considered genuine predictions of the simulation. Instead, we present them in order to give a point of comparison with the existing EAGLE models, the stellar mass functions of which have been investigated in detail by Furlong et al. (2015). The left panel of Fig. 4.1 shows the GSMF from the Recal ENGINE volume at $z = 5$ and 2, which we calculate by counting the number of galaxies falling into bins of width 0.2 dex in stellar mass. To determine the stellar mass associated with each galaxy, we use a 30 pkpc 3D aperture as recommended by S15. We also show GSMFs computed from the Ref-L0050N0752 and Recal-L0025N0752 EAGLE volumes, which respectively feature a reduction in resolution and box size relative to ENGINE. At low stellar masses, the limited particle resolution results in poorly-resolved galaxies, which Fig. 4.1 indicates by switching to dotted lines when the stellar mass falls below that equivalent to 100 times the baryonic particle mass. The high-resolution Recal models (i.e. ENGINE and L0025N0752) allow galaxies with lower stellar masses to be resolved, and additionally produce a GSMF with a shallower slope at low stellar masses, which was found to yield better consistency with the observed $z = 0.1$ GSMF in S15. Conversely, high stellar masses are subject to poor statistical sampling due to the finite volume of the simulations, which we show by using dashed lines for stellar mass bins containing fewer than 10 galaxies. Here, the larger volume of the two 50 cMpc simulations proves advantageous, especially at $z = 2$: in contrast to the 25 cMpc volume, they permit the break at $M_* = 10^{10.5}\,\mathrm{M}_\odot$ beyond which the GSMF steepens to be marginally resolved. Furthermore, the initial conditions used for the smaller simulation provide poorer sampling of large-scale density modes, producing a noisier GSMF. These results show that, as expected, the ENGINE GSMF closely matches those from the existing EAGLE volumes, while combining the benefits of both higher resolution and larger computational volume, allowing for the GSMF to be well-resolved over an additional dex in stellar mass.

In the centre panel of Fig. 4.1, we instead compare the stellar mass function from ENGINE with that extracted from the Illustris-TNG50 simulation (Nelson et al. 2019, 2021; Pillepich et al. 2019). At $z = 5$ the two simulations agree well, with TNG50 predicting a slightly higher abundance of galaxies at low stellar mass. This difference is also seen at $z = 2$, where it also applies at higher stellar masses, such that at this redshift the TNG50 GSMF appears to have an overall normalisation that is ~0.2 dex higher than ENGINE. Also plotted on Fig. 4.1 are a selection of observational constraints on the GSMF. At $z = 5$, neither the data nor the simulations clearly indicate the presence of a knee in the GSMF, while the normalisation of the GSMF differs substantially between datasets. Hence, the simulated stellar mass distribution at this redshift agrees well with the data from González et al. (2011) and Grazian et al. (2015), but lies almost one dex below the GSMF found by Duncan et al. (2014). The slopes of the three observed GSMFs remain similar, and they are also consistent with the GSMF from the simulations, except for stellar masses of 10^8 $\mathrm{M}_\odot$ and below where the observed GSMF appears to flatten. Given the inherent difficulty of observing faint objects at high redshift, this effect may be systematic in nature.

At $z \sim 2$ we show data from Tomczak et al. (2014), Ilbert et al. (2013) and Muzzin et al. (2013), using empty and filled points to show observations in the redshift ranges $1.5 < z < 2.0$ and $2.0 < z < 2.5$ respectively. Within each redshift bin, the scatter between datasets is reduced with respect to $z = 5$, although in the higher z bracket there is still some disagreement at low stellar mass. At this redshift, both the data and the 50 cMpc simulations now show a clear Schechter-like break in the GSMF. However, ENGINE under-estimates the stellar mass at which this feature is located by a margin of approximately 0.6 dex, while also predicting a steeper GSMF beyond it and a somewhat greater normalisation than favoured by the observations. The TNG50 GSMF reproduces the observed position of the knee and the high-mass slope more accurately, although it exceeds the observed mass function by a more significant margin.

As discussed by Furlong et al. (2015), the measured stellar masses are susceptible to both random and systematic errors, the presence of which may contribute to the discrepancies between the predicted and observed GSMFs. Random mass errors lead to a net broadening of the break in the GSMF and shift the exponential tail to higher masses, whereas systematic ones maintain the shape of the GSMF while translating it along the stellar mass axis. Meanwhile, despite providing a more complete sample of the stellar mass distribution than the 25 cMpc EAGLE simulation, galaxies in ENGINE with $M_* > 10^{10.5}$ $\mathrm{M}_\odot$ are rare at $z = 2$, and may remain under-sampled. In support of this, the GSMF obtained from the larger Ref-L0100N1504 EAGLE volume (not shown in Fig. 4.1; see Furlong et al. 2015) produces a shallower exponential tail that more closely resembles the one seen in the observations. In light of these potential sources of uncertainty, we confirm that the simulations we present have successfully produced a population of galaxies consistent with the existing EAGLE simulations and in reasonable agreement with observations.

4.3 Radiative transfer

Photoheating and radiative cooling rates in EAGLE are calculated element-by-element following Wiersma et al. (2009a), who use CLOUDY (version 07.02, last described by Ferland et al. 1998) to tabulate the rates as functions of density and temperature. These calculations assume that the gas is optically thin, and that it remains in ionisation equilibrium with a homogeneous background radiation field given by the Haardt & Madau (2001) spectrum. As the majority of gas in the post-reionisation Universe is low-density and highly ionised, this is a good approximation overall. However, in the higher-density gas probed by strong HI absorption, this approximation breaks down and the self-shielding of the gas plays an important role, necessitating a more rigorous treatment. A fully consistent calculation requires radiative transfer to be performed "on-the-fly" as the simulation runs, so that it can be coupled to the hydrodynamics. Such a calculation is computationally expensive, making it impractical for simulations such as ours which evolve a large cosmological volume to relatively low redshifts. Accordingly, we instead compute a radiative transfer solution by post-processing the simulation snapshots with the reverse ray tracing code URCHIN (Altay & Theuns 2013), which has previously been used to study H I in the OWLS simulations (Altay et al. 2011, 2013).

In conventional ray tracing, a large number of rays are cast from a prescribed set of sources in order to build up a background radiation field which reproduces the expected spectrum in the optically thin limit. URCHIN instead starts from the same assumptions used in the simulation (i.e. a spatially uniform background illuminating optically thin gas), and calculates deviations from them. Rays are cast outward from each particle to determine the optical depth at that position, and hence an estimate for the local photoionisation rate accounting for self-shielding. The new photoionisation rate permits the particle's ionisation state to be updated, in the process changing the optical depths along all rays that intersect the particle. The calculation is therefore iterated until optical depths for all particles have converged.

This approach conveys several advantages which make it well-suited for our application: it avoids the need to trace external rays from the simulation box edges inward (which is required by conventional ray tracing to account for the mismatch between the mean free path of ionising radiation and feasible simulation box sizes, but can lead to an unphysical gradient in the photoionisation rate from box edges to centre); it respects the adaptive spatial resolution of the simulation by sampling every SPH particle with an equal number of rays; it permits the transport of arbitrarily detailed spectra with little additional computational cost; and it achieves rapid convergence for particles in both the optically thin and optically thick limits, resulting in a highly-efficient calculation. In addition to the input spectrum, for which we use the same Haardt & Madau (2001) background assumed by the simulations, the principal numerical parameters of URCHIN are the number of rays cast per particle, $N_{\rm ray}$[3], and the distance each ray is traced $l_{\rm ray}$. We use $N_{\rm ray} = 12$ and $l_{\rm ray} = 100\,{\rm pkpc}$, which we have verified leads to converged results for the results presented here.

In Fig. 4.2 we illustrate the H I distribution around a massive ($M_{200}{\sim}10^{12}\,{\rm M}_{\odot}$) halo, obtained from the post-processed Recal simulation at $z = 2$. The image covers a 5 cMpc square region, representing 1% of the total projected area of the simulation box. For the central halo, in addition to a selection of smaller objects, inset panels show their counterparts from the WeakFB and StrongFB simulations at the same redshift. It is immediately apparent that the adopted feedback efficiency has significant effects on the H I distribution, which at low halo masses can dictate whether the halo hosts a DLA or not; in massive haloes it instead influences the H I morphology and covering factor. In the following sections, we will quantify these trends.

---

[3] Ray directions are assigned using the HEALPIX algorithm (Gorski et al. 2005).

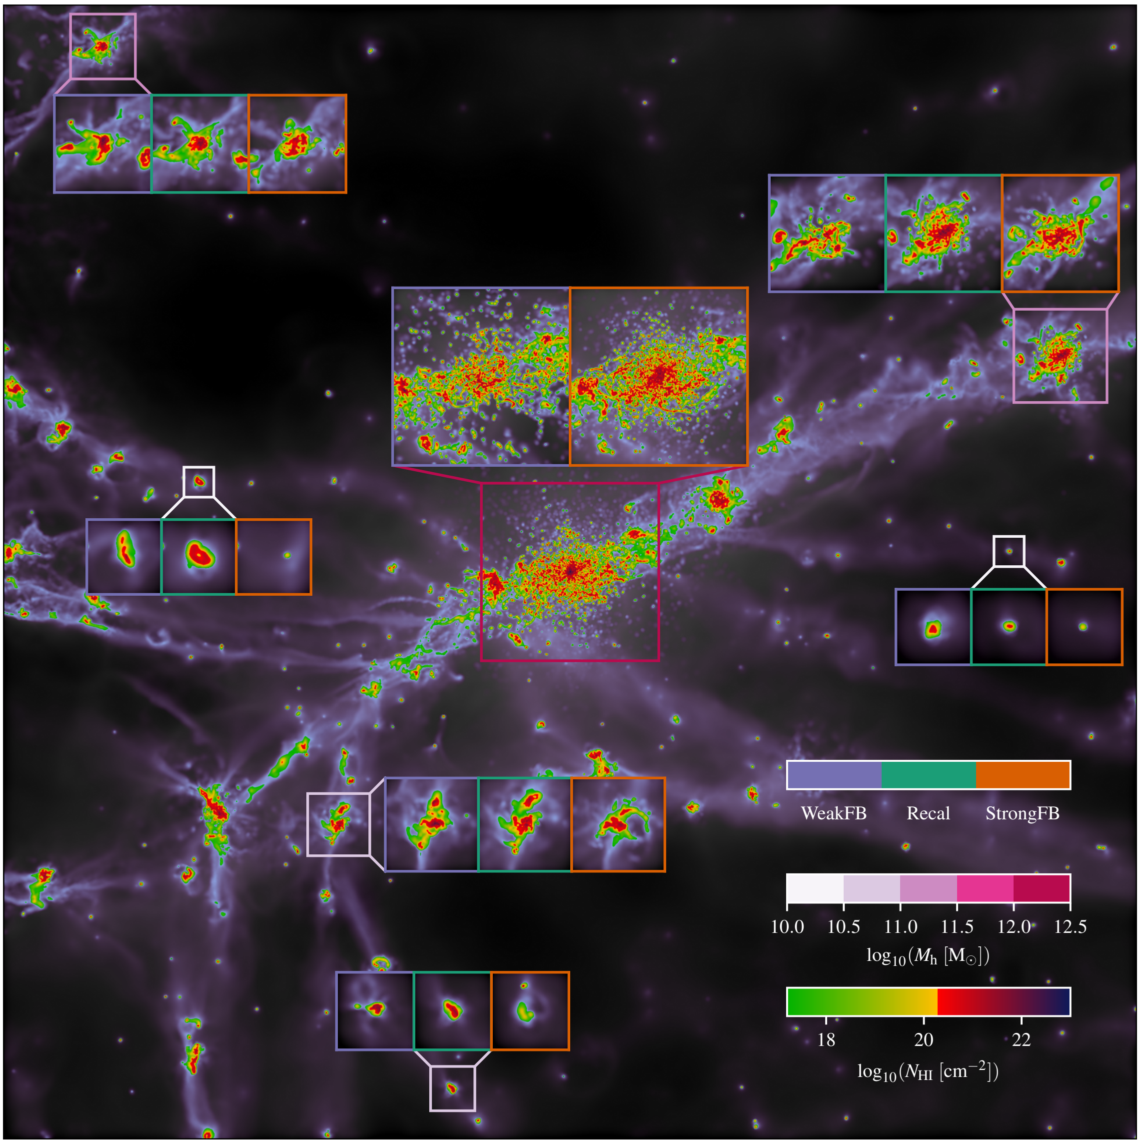

Figure 4.2

SPH-smoothed projection of a 5×5 cMpc region extracted from the Recal simulation at $z = 2$, centred on a halo with virial mass $M_{200} = 2.13 \times 10^{12}\ {\rm M}_\odot$. Total projected hydrogen density is shown in monochrome, while H I above the LLS and DLA column density thresholds is coloured green/orange and red/purple, respectively. Inset panels above the main halo show its counterparts in the WeakFB (blue border) and StrongFB (red border) simulations. Additional insets show zoomed-in views of selected haloes with $10 < \log_{10}(M_{\rm h}/{\rm M}_\odot) < 12$ as they appear across all three volumes. The insets are sized so that they enclose a circle with radius equal to R_{200} for each halo.

4.3.1 Treatment of ISM particles

Here, we highlight some additional considerations regarding the gas particles that make up the simulated interstellar medium. Firstly, the temperatures stored in the simulation snapshot for particles on the star-forming equation of state are not physical, and simply reflect the effective pressure it imposes. To obtain realistic ionisation and recombination rates for these particles, we set a temperature floor $T_{\rm WNM} = 10^4$ K typical of the ISM warm neutral medium. We additionally impose a temperature ceiling of $T_{\rm thick} = 10^4$ K for all particles which become optically thick, because the Wiersma et al. (2009a) photoheating rates are calculated in the optically thin limit, resulting in artificially high temperatures being predicted for these particles. Finally, we include a prescription to model the formation of molecular hydrogen in actively star-forming particles. This takes the form of a pressure-dependent power law giving the molecular mass fraction,

$$f_{\rm H_2} \equiv \frac{2n_{\rm H_2}}{2n_{\rm H_2} + n_{\rm H}} = \left[1 + \left(\frac{P^*}{P_0}\right)\left(\frac{n_{\rm H}}{n_{\rm H}{}^*}\right)^{-\beta}\right]^{-1}, \tag{4.3.1}$$

where P^* is the pressure from the ISM equation of state at the star formation threshold density $n_{\rm H}{}^* = 0.1\,{\rm cm}^{-3}$, and $\beta = \alpha\gamma_{\rm eos}$ with $\gamma_{\rm eos} = 4/3$. The power-law slope $\alpha = 0.92$ and the constant P_0, which satisfies $P_0/k_{\rm B} = 3.5 \times 10^4\,{\rm K\,cm}^{-3}$, are taken from Blitz & Rosolowsky (2006), who fit a power law of this form to observations of local spiral galaxies. This sample has a minimum metallicity one-fifth of the solar value, which exceeds the typical DLA metallicity at $z = 3$. At lower metallicities, H_2 formation becomes less efficient, meaning that the values of $f_{\rm H_2}$ we calculate should be considered an upper limit. We discuss the impact of the molecular correction and the significance of this caveat in Section 4.4.2.

4.4 The HI column density distribution function

In this section we present statistical properties of the H I distribution in ENGINE, via the column density distribution function (CDDF), defined as the number of absorbers N per unit column density and absorption distance:

$$f(N_{\mathrm{HI}}, z) \equiv \frac{\mathrm{d}^2 N}{\mathrm{d}N_{\mathrm{HI}}\mathrm{d}X}, \tag{4.4.1}$$

where the absorption distance X is defined by $\mathrm{d}X/\mathrm{d}z = H_0(1+z)^2/H(z)$, with $H(z)$ the Hubble parameter and H_0 its value at redshift $z = 0$.

We first identify H I absorbers by taking the simulation snapshot, post-processed with URCHIN to obtain neutral fractions $x_{\mathrm{HI}} = n_{\mathrm{HI}}/n_{\mathrm{H}}$ for all gas particles, and projecting all particles onto a 2D grid using a Gaussian approximation to the SPH kernel for each particle. We use a grid with 32768^2 pixels, yielding lines of sight spaced by 0.38 pkpc at $z = 3$.To avoid any artificial smoothing of the contributions from high-density particles over the cell area, we perform the projection operation by computing line integrals along pencil beams located at the centre of each grid cell. Each sightline contributes an absorption distance $\delta X = H_0(1+z)^2 L_{\mathrm{box}}/c$, such that the total absorption distance over all sightlines is $\Delta X = 32768^2 \delta X$. Because we project the entirety of the simulation box at once, it is possible for distinct absorbers located at different positions along the projection axis to overlap. However, the rarity of high-N_{HI} absorbers means the probability of this occurring is low, and we have verified that for the results presented in this section, projection effects are negligible for column densities of $10^{16}\,\mathrm{cm}^{-2}$ and above. We have further verified that our results are converged with respect to the grid resolution.

4.4.1 Comparison with observations

To evaluate Eq. (4.4.1) and calculate the CDDF, we count the number of sightlines which fall into column density bins of width $\mathrm{d}\log_{10} N_{\mathrm{HI}} = 0.1$ dex, and normalise by the bin width and ΔX. The distribution function we obtain is plotted in Fig. 4.3, along with several observational constraints. The shaded pink region indicates the $\langle z \rangle = 2.4$ Ly α forest constraint from Rudie et al. (2013), while the blue region shows a constraint derived by O'Meara et al. (2013) from LLS spectra taken at the same mean redshift. For N_{HI} between 10^{19} and $10^{20}\,\mathrm{cm}^{-2}$, we plot data points from ESO UVES observations in red (Zafar et al.

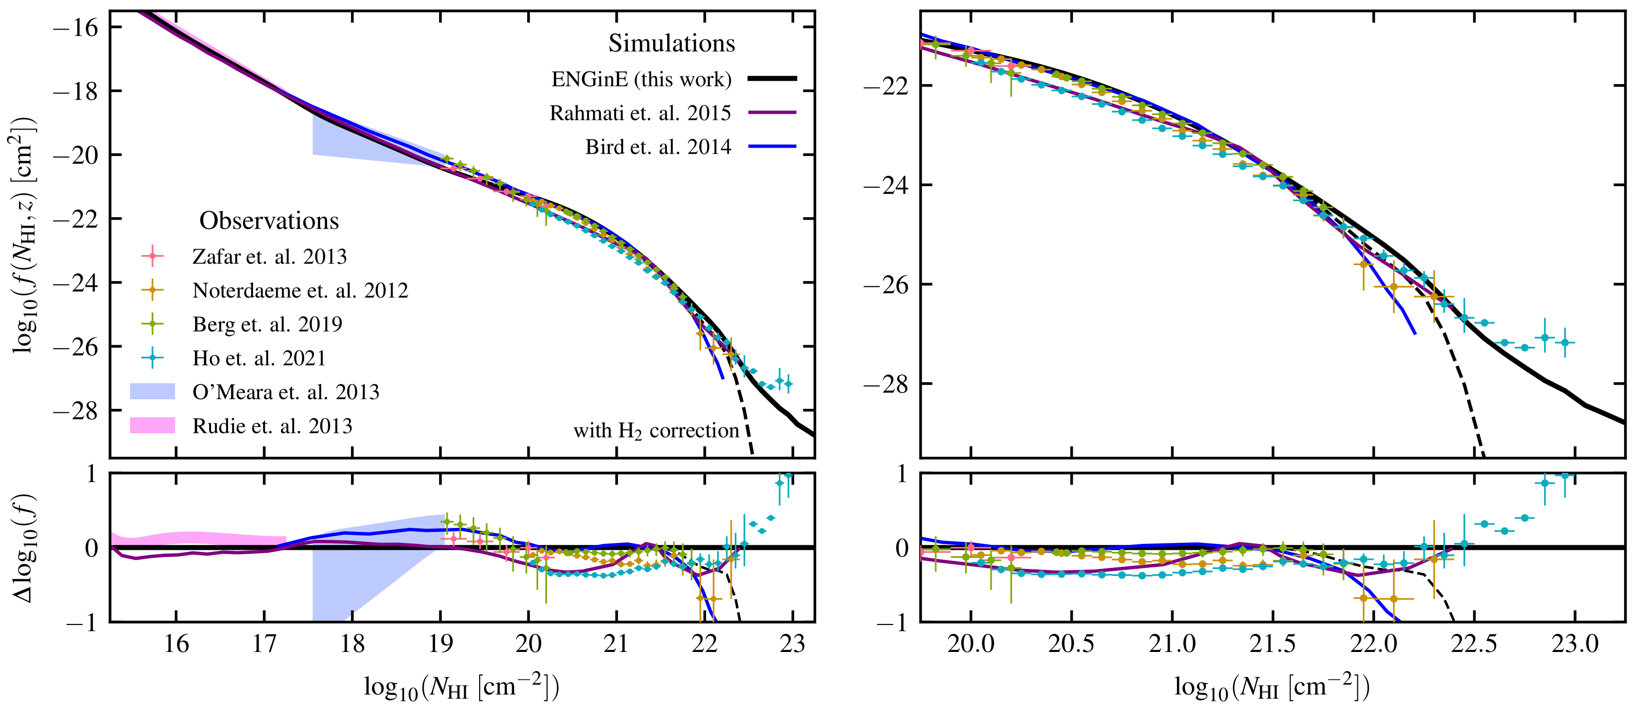

Figure 4.3

Main panels: the H I CDDF $f(N_{\rm HI}, z = 3)$ from the Recal simulation (solid black line). The dashed line shows the same CDDF with the Blitz & Rosolowsky (2006) H_2 correction enabled. Also shown are simulated $z = 3$ CDDFs from Bird et al. 2014 (their '2×UV' model) and Rahmati et al. (2015); low-$N_{\rm HI}$ observational constraints from Rudie et al. (2013) and O'Meara et al. (2013); and spectroscopic sub-DLA/DLA observations from Zafar et al. (2013), Noterdaeme et al. (2012), Berg et al. (2019) and Ho et al. (2021). The left panel shows the entire range of column densities produced in our simulations, demonstrating agreement with the observations to 0.5 dex or better throughout. In the right panel the DLA column density regime is shown magnified for clarity. *Lower panels*: The relative differences Δf between the constraints we report and the $z = 3$ Recal CDDF.

2013; $1.5 < z < 5$), and from the XQ-100 survey in green (Berg et al. 2019; $2.3 < z < 3.2$). The latter observations extend to higher column densities, where we also show results from Noterdaeme et al. (2012; $\langle z \rangle = 2.5$) and Ho et al. (2021; $2.0 < z < 2.5$) with orange and cyan points respectively.

Overall, the values of $f(N_{\mathrm{HI}}, z)$ we obtain compare favourably with these observations, agreeing to 0.5 dex or better throughout. The potential for hydrodynamical simulations to reproduce these observations is well-established, as demonstrated by the column density distributions obtained by Bird et al. (2014) and Rahmati et al. (2015), which we plot in Fig. 4.3 with blue and purple lines, respectively. Nevertheless, it is remarkable that consistency with the observations is maintained over seven decades in column density, and twelve decades in the value of $f(N_{\mathrm{HI}}, z)$.

At column densities below $10^{17}\,\mathrm{cm}^{-2}$, the CDDF derived from ENGINE follows a power law, $f(N_{\mathrm{HI}}, z) \propto N_{\mathrm{HI}}{}^{-\beta}$ with $\beta = 1.61$, which compares well with the value $\beta = 1.65$ reported by Rudie et al. (2013). Good agreement is also achieved with the Rahmati et al. (2015) CDDF, which was obtained by post-processing the EAGLE Ref-L100N1504 simulation with a different radiative transfer scheme first described in Rahmati et al. (2013a). The same scheme is also used by Bird et al. (2014), although in this case it is paired with an unrelated set of simulations which were run with the AREPO moving-mesh code and a different set of subgrid models. In Fig. 4.3 we plot the CDDF corresponding to the '2×UV' model from this suite, which features an increased ultraviolet background (UVB) amplitude relative to their fiducial 'DEF' model. This results in a H I photoionisation rate $\Gamma_{\mathrm{HI}} = 1.04 \times 10^{-12}\,\mathrm{s}^{-1}$, which remains smaller than the value of $\Gamma_{\mathrm{HI}} = 1.50 \times 10^{-12}\,\mathrm{s}^{-1}$ calculated for the Haardt & Madau (2001) spectrum we assume. Lower photoionisation rates are also favoured by more recent models of the UVB (e. g. Haardt & Madau 2012; Puchwein et al. 2019), and the adoption of such a spectrum would result in an increased abundance of H I absorbers with $N_{\mathrm{HI}} < 10^{20}\,\mathrm{cm}^{-2}$. We find that the observational constraints on the CDDF are compatible with a photoionisation rate in the range 0.5–$1.2 \times 10^{-12}\,\mathrm{s}^{-1}$.

At $N_{\mathrm{HI}} > 10^{20}\,\mathrm{cm}^{-2}$ the DLA regime is entered, in which the absorbing gas becomes predominantly neutral. This is expected to result in a 'knee' feature where $f(N_{\mathrm{HI}}, z)$

steepens from a relatively shallow power law slope $\beta \sim 1$, to a steeper one $\beta > 2$ (see Theuns 2021). In the ENGinE CDDF (and also that from Bird et al. 2014), this transition occurs smoothly between $N_{\rm HI} = 10^{20}$ and $10^{22}\,{\rm cm}^{-2}$, whereas the CDDF from Rahmati et al. (2015) steepens abruptly at $N_{\rm HI} = 10^{21.5}\,{\rm cm}^{-2}$. We see similar behaviour in $f(N_{\rm HI}, z)$ obtained from our Recal-0752 run (not shown in Fig. 4.3), which has an equivalent resolution to the reference Eagle model used to produce the Rahmati et al. (2015) CDDF. This suggests that the differences we find in the shape and position of the knee may be interpreted as a resolution effect of the underlying simulation, a possibility that we discuss further in Section 4.4.4. While observations also favour a gradually-steepening column density distribution steepening of the CDDF, they vary in normalisation by up to 0.3 dex, comparable to the maximum value of the difference between the simulated CDDFs. Furthermore, it has been argued that the inclusion of local sources of ionising radiation can suppress the predicted abundance of high–$N_{\rm HI}$ absorbers by a similar margin (Rahmati et al. 2013b).

4.4.2 Molecular hydrogen

As we will show later, the effect of stellar feedback on the distribution of neutral gas is most significant for $N_{\rm HI} \geq 10^{22}\,{\rm cm}^{-2}$. At these column densities, the transition from atomic to molecular gas is expected to result in an exponential cutoff in the H I CDDF (Schaye 2001; Zwaan & Prochaska 2006; Krumholz et al. 2009c).

Using the Blitz & Rosolowsky (2006; hereafter BR06) relation to account for H_2 formation produces the dashed black curve in Fig. 4.3, which displays the expected suppression of the strongest absorbers, and resembles the CDDF from Bird et al. (2014), who include a similar pressure-based model for the formation of H_2. Very few observational sightlines containing extremely strong absorbers have been detected, as demonstrated in Fig. 4.3 by the Noterdaeme et al. (2012) CDDF, which shows tentative signs of a cutoff in the highest-$N_{\rm HI}$ bins. This has been interpreted as evidence for the ubiquity of H_2 formation (e. g. Prochaska & Wolfe 2009); alternatively, it may simply be a consequence of the intrinsic scarceness of these absorbers. Indeed, not all data support the existence of a molecular cutoff: the Ho et al. (2021) CDDF, obtained from the large SDSS/BOSS DR16Q sample, remains non-zero for column densities up to $N_{\text{HI}} = 10^{22.5}\,\text{cm}^{-2}$. If these data are taken at face value, they are matched by the simulated column density distributions from ENGinE and Rahmati et al. (2015), when both are calculated without any H_2 correction. Conversely, including the correction results in a CDDF that is consistent with the observations from Noterdaeme et al. (2012).

To appraise these findings, we first note that the data likely remains incomplete at high column densities, not only because of the intrinsic scarcity of strong absorbers, but also because their increased dust content is likely to significantly attenuate the background quasar continua of the sightlines in which they appear, potentially biasing against their selection in spectroscopic surveys (Krogager et al. 2019). Meanwhile, as mentioned in Section 4.3.1, the BR06 prescription is calibrated using observations of local-Universe galaxies, which are likely to have a higher average metallicity than is typical at $z = 3$. Metal-rich gas cools more readily, and additionally has an increased dust abundance which permits more efficient H_2 formation; as such, the molecular fractions we calculate may be biased high (but see also Fumagalli et al. 2010). This assessment is supported by findings from the Simba simulations (Davé et al. 2019), in which the star formation subgrid recipe explicitly includes a metallicity- and pressure-dependent model for H_2 (Krumholz et al. 2009b; Krumholz 2013). In this model, the transition to molecular gas is not predicted to occur until a median column density above $10^{22.5}\,\text{cm}^{-2}$ at $z = 3$, and consequently the CDDF obtained from Simba shows no evidence for a molecular cutoff (Hassan et al. 2020). It is therefore not unexpected that presently-available measurements of the observed H I column density distribution lack an apparent molecular cutoff.

A complementary source of constraints on the atomic-to-molecular transition is the column density distribution of molecular gas $f(2N_{\text{H}_2}, z)^4$, although measuring this brings its own challenges. The H_2 molecule lacks a permanent dipole moment, and so presents few absorption or emission features. Detection via absorption in the UV Lyman-Werner bands is possible, but the low detection rate ($\lesssim$10%; e. g. Ledoux et al. 2003; Noterdaeme et al.

2008; Balashev et al. 2014) makes achieving a statistically robust measurement expensive. This difficulty may again be attributed to a bias against sightlines containing the strong H I absorbers which are most likely to show molecular absorption features, or alternatively may indicate that the majority of the molecular gas is contained in dense clumps with a small covering factor. Balashev & Noterdaeme (2018) demonstrate a technique that aims to circumvent these issues by stacking many DLA-containing absorption spectra in order to improve the signal-to-noise ratio with which the weak H_2 absorption features are detected. By using the combined spectrum to constrain the incidence rate of H_2 at $z \sim 3$, $f(2N_{\mathrm{H_2}})$ can be inferred from the corresponding H I CDDF. Alternatively, H_2 gas can be detected indirectly in emission, via a tracer such as CO. This method has been used to constrain the H_2 CDDF at $z = 0$ (Zwaan & Prochaska 2006), but statistical samples of sufficiently high-resolution spectra at higher redshift do not yet exist.

We plot these observations in Fig. 4.4, noting that as pointed out by Balashev & Noterdaeme (2018), their close agreement implies a remarkable lack of evolution in $f(2N_{\mathrm{H_2}})$ between $z = 3$ and 0. We do not attempt to compute an equivalent H_2 CDDF from ENGINE, because of the expectation that our molecular fractions represent upper limits, as well as the fact that our simulations do not resolve the small-scale structures (i.e. giant molecular clouds) within which H_2 formation occurs. Instead, we observe that for $N_{\mathrm{H}} > 10^{22.5}\,\mathrm{cm^{-2}}$, the $f(N_{\mathrm{HI}}, z)$ calculated from ENGINE without any H_2 correction (which therefore represents the total abundance of neutral gas), closely reproduces the Zwaan & Prochaska (2006) observations, albeit with an approximately constant offset of 0.3 dex. In the lower panel of Fig. 4.4, we also plot the mean H_2 fraction predicted by the BR06 prescription, averaged over all sightlines within each N_{H} bin. This shows that the column density at which the two CDDFs come into agreement coincides to the point where $\langle f_{\mathrm{H_2}} \rangle$ exceeds 0.5 (i.e. the majority of H I has been converted to H_2). This threshold is in agreement with expectations from more advanced H_2 models (e. g. Krumholz et al. 2009a,b; Krumholz 2013), implying that in the most molecule-rich regions (which are also those with the highest typical

---

[4] The factor of 2 multiplying the column density is introduced so that the abundances given by the H_2 CDDF refer to the surface density of H atoms, for consistency with $f(N_{\mathrm{HI}})$.

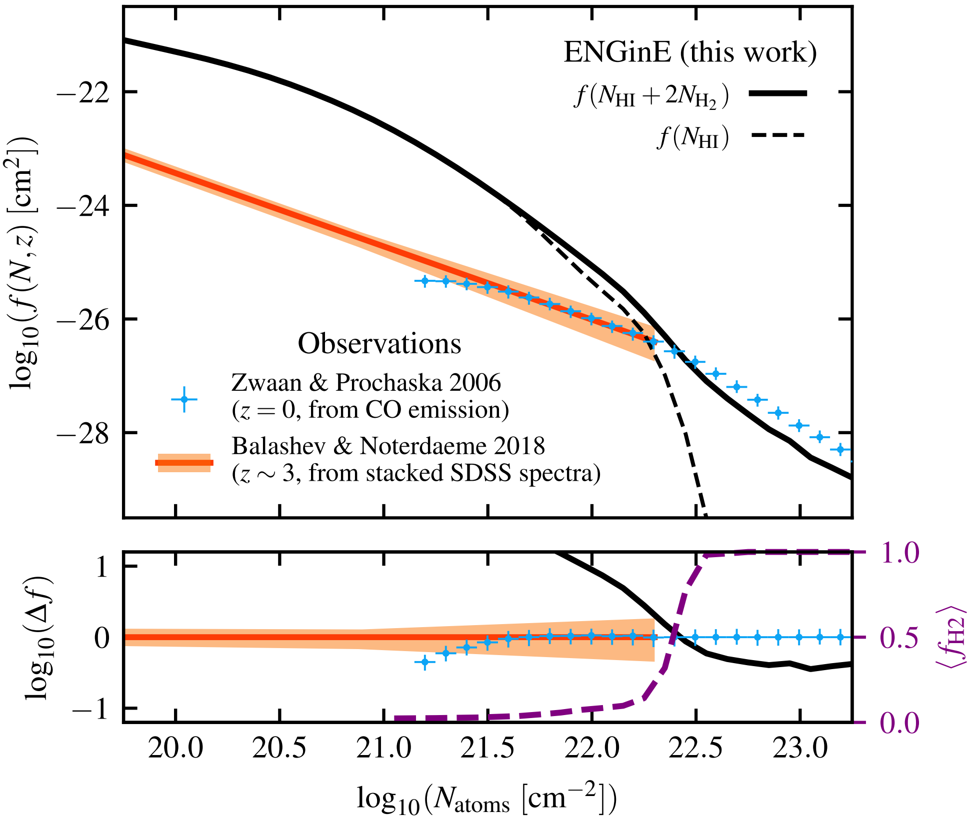

Figure 4.4

Main panel: constraints on the H_2 CDDF $f(2N_{\rm H_2}, z)$. The orange line indicates the Balashev & Noterdaeme (2018) constraint at $z \sim 3$, with the shaded region denoting 1σ uncertainties. Blue points show the $z = 0$ $f(H_2)$ derived from CO measurements by Zwaan & Prochaska (2006), which we have recalculated using the conversion factor $X_{\rm CO}$ from Bolatto et al. (2013). The H I CDDF is reproduced for reference, with a solid (dashed) line showing the abundance without (with) H_2 correction included. Note that the total column density of neutral H atoms is plotted on the x-axis. *Lower panel*: for $N_{\rm atoms} > 10^{21}$ cm^{-2}, we show the relative difference between $f(N_{\rm HI}+N_{\rm H_2})$ calculated from ENGinE and the Zwaan & Prochaska (2006) constraint on $f(H_2)$. The purple dashed line, corresponding to the right-hand axis, instead shows the average H_2 fractions for all sightlines in each column density bin. Above $N_{\rm atoms} \sim 10^{22.3}$ cm^{-2}, $\langle f_{\rm H_2}\rangle$ approaches 1, such that the CDDF for atomic hydrogen from ENGinE agrees well with the data.

metallicities), the BR06 relation yields the correct limiting behaviour. Thus, we conclude that the CDDF from ENGINE appears to faithfully reproduce the overall gas distribution up to the highest column densities probed by observations, although the partitioning of this gas into atomic and molecular components remains uncertain.

4.4.3 Stellar feedback

In Fig. 4.5 we again show the $z = 3$ CDDF in the DLA column density range, but now include results from our StrongFB and WeakFB simulations, in which the efficiency of stellar feedback was respectively doubled and halved relative to the Recal model. Here we see that while the impact of feedback is negligible for $N_{\rm HI} < 10^{21}\,{\rm cm}^{-2}$, it significantly affects the shape of $f(N_{\rm HI}, z)$ in the strong DLA regime. The CDDF obtained from the WeakFB simulation shows an enhancement with respect to Recal of greater than 0.7 dex for $N_{\rm HI} > 10^{22.5}\,{\rm cm}^{-2}$. In the StrongFB simulation, the effect is even more pronounced, with $f(N_{\rm HI}, z)$ being suppressed by >1 dex at the same column density.

These results resemble those of Altay et al. (2013), who investigated the H I CDDF obtained from the OWLS project (Schaye et al. 2010), which include simulations with a wide variety of feedback variations. They found that modifications to the assumed efficiency of stellar feedback (which in the OWLS model is explicitly parameterised in terms of the outflow velocity and mass-loading factor $\eta \equiv \dot{M}_{\rm wind}/\dot{M}_{*}$), as well as other subgrid properties such as the assumed IMF and the normalisation constant A appearing in the Kennicutt-Schmidt law (Eq. 4.2.3), only produced significant differences in $f(N_{\rm HI}, z)$ for $N_{\rm HI} > 10^{21}\,{\rm cm}^{-2}$. As discussed by Altay et al. (2013), this can be interpreted as evidence that strong DLA systems are associated with the innermost regions of galaxies, and in particular the simulated ISM, where subgrid implementations of feedback have the greatest effect (see also van de Voort et al. 2012). We note that this conclusion is not universally supported, with some feedback implementations (e. g. the HVEL model of Bird et al. 2014) instead modifying $f(N_{\rm HI}, z)$ at all column densities. We will investigate the degree to which this conclusion applies to our simulations in Section 4.5.4.

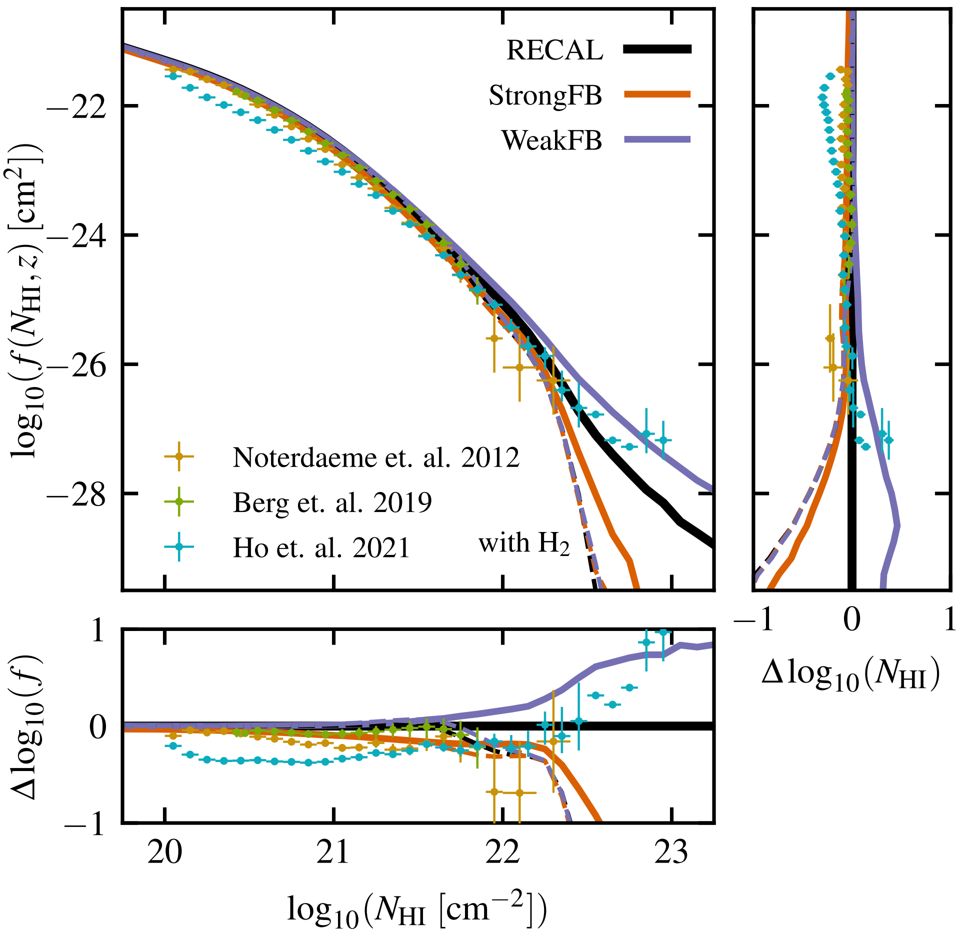

Figure 4.5

The H I CDDF $f(N_{HI}, z = 3)$ from the Recal, StrongFB and WeakFB simulations (black, red, and blue respectively). As in Fig. 4.3, dashed lines show the effect of accounting for the formation of molecular hydrogen, and the fractional change in abundance relative to the Recal CDDF is plotted in the lower panel. The right-hand panel instead shows the fractional change in column density at a given abundance.

Since the same initial conditions were used for all three simulations, they are expected to form the same haloes, as shown explicitly by Fig. 4.2. Therefore, the differences in their column density distributions do not arise from an intrinsic change in the number of absorbers, but instead indicate a change in the typical values of N_{HI} associated with absorbers of a given abundance. We demonstrate this in the right-hand panel of Fig. 4.5, which shows the fractional change in N_{HI} for absorbers in the StrongFB and WeakFB simulations required to obtain an abundance equal to the Recal value. Thus, we see that the relatively large changes in the abundance of absorbers may instead be interpreted as less dramatic shifts in the column density attained by an absorber of given abundance. For example, the values of $f(N_{\mathrm{HI}}, z)$ at $N_{\mathrm{HI}} = 10^{22.5}\,\mathrm{cm}^{-2}$ differ by more than a factor of 10 between the three runs, but this change may be attributed to a more modest ±0.5 dex shift in the column density that corresponds to an absorber with abundance of order $10^{-28}\,\mathrm{cm}^2$.

When H_2 formation is accounted for, we obtain the same exponential cutoff seen in the Recal CDDF, regardless of feedback strength. The resulting suppression of high-N_{HI} absorbers is comparable in magnitude to that found in the StrongFB CDDF without any correction for H_2, but qualitatively different. Modifying the feedback strength changes the degree to which the CDDF steepens above $10^{22}\,\mathrm{cm}^{-2}$, but does not affect its power-law character, in contrast to the sharp exponential cutoff that appears when the molecular gas is removed.

4.4.4 Redshift evolution

Finally, we consider the redshift evolution of the H I distribution in the Recal simulation. In the left panel of Fig. 4.6, we plot the CDDF computed at unit-spaced redshifts between 5 and 2. In agreement with previous work (e. g. Rahmati et al. 2013a), we find remarkably little evolution over much of the DLA column density range. At column densities near the DLA threshold, $f(N_{\mathrm{HI}}, z)$ decreases by 0.4 dex between $z = 5$ and 2, whereas at the highest column densities, the non-linear growth which gives rise to strong H I absorbers results in an increase in $f(N_{\mathrm{HI}}, z)$ of 0.9 dex over the same redshift interval. Fig. 4.6 also

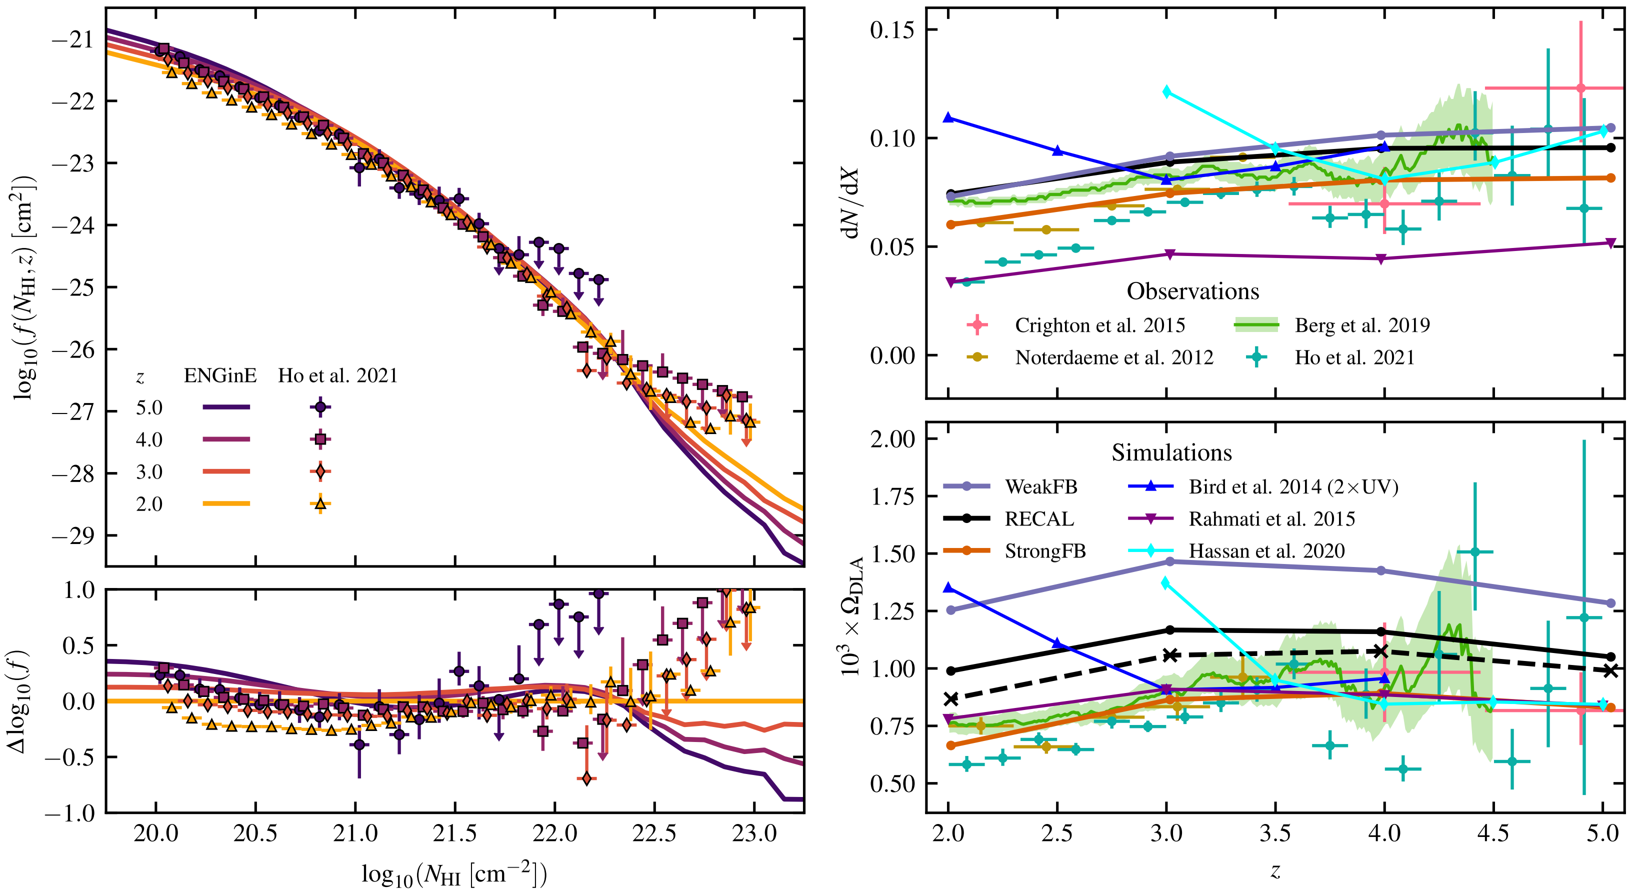

Figure 4.6

Left panel: the H I CDDF $f(N_{\rm HI}, z)$ from the Recal simulation at $z = \{5, 4, 3, 2\}$. As in Figs 4.3 and 4.5, in the lower panel we plot $\log_{10}$ of the deviation from a 'reference' curve, here the $z = 2$ CDDF. Binned observations from Ho et al. (2021) are also plotted, with downward-pointing arrows indicating 1σ upper limits. *Upper- and lower-right panels*: the redshift evolution of the DLA incidence rate dN/dX, and the DLA mass density $\Omega_{\rm HI}^{\rm DLA}$, respectively. Observed values are taken from Noterdaeme et al. 2012 (red), Crighton et al. 2015b (yellow), Berg et al. 2019 (green) and Ho et al. 2021 (blue).

shows the constraints on the CDDF obtained by Ho et al. (2021), when the observed DLAs are divided into redshift bins. At column densities near the DLA threshold, the binned data shows evidence for a similar redshift evolution to the simulations, although there is a small, consistent difference in the normalisation, which was also apparent in Fig. 4.3. Detections of strong DLAs are extremely rare at higher redshift, resulting in the majority of the binned data points representing upper limits; in this regime observations remain unable to strongly constrain the H I CDDF at $z > 2$.

In the right hand panels of Fig. 4.6, we plot the DLA incidence rate and mass density obtained from ENGinE, along with observed constraints on these quantities. The incidence rate $\mathrm{d}N/\mathrm{d}X$ is calculated as the zeroth moment of the CDDF:

$$\frac{\mathrm{d}N}{\mathrm{d}X} \equiv \int_{10^{20.3}}^{\infty} f(N_{\mathrm{HI}}, z)\, \mathrm{d}N_{\mathrm{HI}}\,, \tag{4.4.2}$$

while the H I mass density in DLAs, $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$, is obtained from the first moment:

$$\Omega_{\mathrm{HI}}^{\mathrm{DLA}} \equiv \frac{m_{\mathrm{p}} H_0}{c \rho_{\mathrm{crit}}} \int_{10^{20.3}}^{\infty} N_{\mathrm{HI}} f(N_{\mathrm{HI}}, z)\, \mathrm{d}N_{\mathrm{HI}}\,, \tag{4.4.3}$$

where m_{p} is the proton mass and $\rho_{\mathrm{crit}} \equiv 3{H_0}^2/8\pi G$ is the FLRW critical density. Because the CDDF falls steeply with increasing N_{HI}, $\mathrm{d}N/\mathrm{d}X$ is primarily set by absorbers with column densities close to the DLA threshold, whereas $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$ is more sensitive to higher-N_{HI} systems.

All observations suggest a degree of evolution in $\mathrm{d}N/\mathrm{d}X$, which is found to increase by up to a factor of 2 between $z = 2$ and 5, although there is substantial uncertainty at $z \geq 4$. At $z = 2$, the scatter between observations is increased, with Ho et al. (2021) finding a stronger decline in the incidence rate. In ENGinE, $\mathrm{d}N/\mathrm{d}X$ shows only weak evolution, rising from a value of 0.074 at $z = 2$ to 0.096 at $z = 5$. We find that the incidence rate is minimally sensitive to the assumed feedback efficiency, in line with Fig. 4.5, which indicates that the impact of feedback on the weak DLA systems which principally contribute to $\mathrm{d}N/\mathrm{d}X$ is minor. Above $z = 3$, Bird et al. (2014) also find an incidence rate consistent with observations, but after this point their simulations predict a sharp increase in $\mathrm{d}N/\mathrm{d}X$ not seen in the real data. We also show the line densities found by Hassan et al. (2020) for the “Technicolor Dawn” simulation, which features an on-the-fly self-consistent treatment of radiative transfer. These results show a similar trend to Bird et al. (2014), but with the ‘pivot’ redshift at which they start to diverge from the observed trend instead being at $z = 4$; this may reflect the fact that this simulation is primarily calibrated to match observational constraints from higher redshifts. Meanwhile, the values of $\mathrm{d}N/\mathrm{d}X$ taken from Rahmati et al. (2015) show the correct evolutionary trend, but lie substantially below the data. Finally, Hassan et al. (2020) also compute $\mathrm{d}N/\mathrm{d}X$ for the SIMBA simulation, obtaining very similar results to Rahmati et al. (2015); we choose to omit these from Fig. 4.6 for clarity.

Observational determinations of $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$ also show some evidence for evolution, although there are again large uncertainties at higher redshift. The DLA mass densities predicted by Bird et al. (2014) and Hassan et al. (2020) behave similarly to the corresponding incidence rates, being in good agreement with observations at higher redshifts, but diverging afterwards. Our results and those from Rahmati et al. (2015) are instead consistent with an essentially non-evolving $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$, providing better agreement with observations. This can be understood as resulting from the use of the same feedback prescriptions in the underlying simulations, leading to similar conditions in the dense gas that the $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$ measurement is most sensitive to. In fact, the self-shielding correction we use, and the differences in resolution and computational volume between ENGINE and the reference EAGLE volume used by Rahmati et al. (2015), are the only factors which differentiate between their calculations and ours. Despite this, the DLA mass densities we obtain are roughly 30% higher, which follows directly from the slightly higher values of $f(N_{\mathrm{HI}}, z)$ that we find at $N_{\mathrm{HI}} \sim 10^{20.5}\,\mathrm{cm}^{-2}$. The incidence rates we calculate differ by an even larger margin, with the ENGINE value of $\mathrm{d}N/\mathrm{d}X = 0.089$ at $z = 3$ almost double the corresponding result from Rahmati et al. (2015).

Two factors appear to combine to create these discrepancies. Firstly, they indicate systematic differences between URCHIN and the method used in Rahmati et al. (2015), which involves the use of a fitting formula for the neutral fraction x_{HI} first introduced in Rahmati et al. (2013a). Applying this method to the $z = 3$ ENGINE snapshot yields $\mathrm{d}N/\mathrm{d}X = 0.051$ and $\Omega_{\mathrm{HI}}^{\mathrm{DLA}} = 0.73 \times 10^{-3}$, both of which are significantly smaller than our fiducial results.

We have not made the opposite comparison (i. e. running URCHIN on the 100 Mpc EAGLE volume) due to computational expense. Second, they provide evidence that convergence in the H I distribution is incomplete at the reference EAGLE resolution: independent of the self-shielding correction used, we find that values of $\mathrm{d}N/\mathrm{d}X$ and $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$ differ by up to 20% between the Recal and lower-resolution Recal-0752 ENGINE volumes. This non-convergence is not unexpected; indeed, Rahmati et al. (2015) reach this conclusion for the lower-N_{HI} LLSs they investigate. Likewise, Crain et al. (2017) find that the $z = 0$ H I mass function is systematically underestimated in the reference simulations. The picture is further complicated by the fact that $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$ is strongly sensitive to the subgrid physics, with varying the feedback efficiency and removing the molecular gas component (the effect of which is shown in Fig. 4.6 by the dashed black curve) both having significant effects. Further investigation is clearly required to establish the significance of each of these potential factors. In the meantime, the degree to which we reproduce the observed evolutionary trends remains encouraging, and serves to further highlight the benefits of the EAGLE subgrid models in comparison to more ‘ad-hoc’ implementations used previously e.g. in OWLS or the simulations of Bird et al. (2014).

4.5 Connecting absorbers with haloes

In this section we investigate the nature of H I absorption as a function of the properties of the haloes within which absorbers are found. We begin by briefly reviewing the analytical model of Theuns (2021; hereafter TT21), which provides a useful ruler with which to interpret the results from the simulations that follow.

4.5.1 The dynamics of DLA gas in haloes

The essence of the TT21 model is that the majority of gas giving rise to a DLA is primordial gas that is accreting from the IGM onto a galaxy through its host dark matter halo, rather than a static reservoir of gas in the galaxy’s ISM. Thus, most sightlines containing DLAs intersect the circumgalactic medium (CGM) of the galaxy, rather than its denser ISM. This assumption is motivated both by the frequently large projected separations between absorbers and their identified hosts (Fumagalli et al. 2017b; Mackenzie et al. 2019), and the extended, filamentary distributions seen in simulations (our Fig. 4.2; also e. g. Cen 2012; van de Voort et al. 2012; Rahmati et al. 2015). The CGM gas spends only on the order of a dynamical time $T_{\rm dyn} = H(z)^{-1}/\sqrt{200}$ in the halo, before accreting onto the galaxy. Subsequently, a small fraction of the gas is converted into stars, with the majority instead ejected by stellar feedback. We will test this aspect of the model directly in Section 4.5.5.

The next important aspect of the model is guided by the similarity solution of Bertschinger (1985), who showed that in an Einstein-de Sitter universe[5], spherically symmetric accretion gives rise to a power-law density profile $\rho(r) \propto r^{-\alpha}$ with exponent $\alpha = 2.2$. By definition, the mean overdensity within R_{200} satisfies $\rho/\langle\rho\rangle \approx 200$, where $\langle\rho\rangle$ is the cosmic mean density; the mass enclosed within this radius is M_{200}. Assuming also that the halo contains its cosmological share of baryons, such that the ratio of baryonic to total matter is $\omega_{\rm b} \equiv \Omega_{\rm b}/\Omega_{\rm M}$, and that the gas is primordial in composition, the model stipulates a gas density profile

$$n_{\rm H}(r) = n_{\rm H,0} \left(\frac{R_{200}}{r}\right)^{\alpha} , \tag{4.5.1}$$

where

$$n_{\rm H,0} = f_{\rm g}\omega_{\rm b} \frac{200}{3} \frac{X_{\rm H}\rho_{\rm crit}}{m_{\rm H}} . \tag{4.5.2}$$

Here, $X_{\rm H} = 0.752$ is the primordial hydrogen abundance, and we have included an additional factor $f_{\rm g} \leq 1$, to account for the fact that simulated dark matter haloes typically have baryon fractions that are somewhat below the cosmological value. We will set $f_{\rm g} = 0.5$.

To partition the gas distribution given by Eq. (4.5.1) into ionised and neutral components, TT21 performed simplified radiative transfer, assuming the gas is in photoionisation equilibrium with the UVB at constant temperature $T = 10^4$ K; approximating the pho-

---

[5] Such a cosmological model provides an excellent approximation of an ΛCDM universe at $z \gtrsim 1$.

toionisation cross-section $\sigma(\nu)$ with its value at the threshold frequency $\nu_{\rm th} = 13.6\,{\rm eV}/h$; ignoring collisional ionisation and the effects of helium and metals; and neglecting the angular dependence of the optical depth arising from spherical symmetry by assuming a slab geometry. With these simplifications, the neutral fraction $x_{\rm HI}$ may be computed as a function of radius as:

$$x_{\rm HI}(r) \equiv \frac{n_{\rm HI}(r)}{n_{\rm H}(r)} = \frac{\alpha_{\rm B}}{\Gamma(r)} n_{\rm H}(r)\,, \tag{4.5.3}$$

where $\alpha_{\rm B}$ is the case-B recombination rate, and

$$\Gamma(r) = \Gamma_0 \exp(-\tau(r)) \tag{4.5.4}$$

is the H I photoionisation rate, given by the incident value Γ_0 attenuated by the optical depth

$$\tau(r) = \sigma(\nu_{\rm th}) \int_r^\infty x_{\rm HI}(r') n_{\rm H}(r') {\rm d}r'\,. \tag{4.5.5}$$

4.5.2 Radial and projected density profiles from ENGinE

To extract the density profile for a halo of mass $10^M\,{\rm M}_\odot$ from the simulation, we identify all particles located within FoF haloes for which $M \leq \log_{10} M_{200} < M + 0.1$. We compute the distance from each particle to the centre of potential of the halo it resides in, divide the particles into 50 logarithmically-spaced radial bins, and calculate the median density within each bin. The top rows of panels in Figs 4.7 and 4.8 show density profiles obtained in this way for haloes with $\log_{10}(M_{200}/{\rm M}_\odot) \in \{9.5, 10.5, 11.5, 12.5\}$, along with corresponding evaluations of the TT21 density profile (Eq. 4.5.1), at $z = 5$ and 2 respectively. We find good agreement overall, with both the logarithmic slope of the density profiles and the location of the ionisation front typically agreeing well with the model prediction. The width of the shaded region in each panel, which indicates the 1σ range around the median $n_{\rm HI}$ profile from ENGinE, shows that there is considerable halo-to-halo scatter.

The model predicts somewhat lower densities than we find in the simulation, although they remain within the 1σ range in almost all cases. This discrepancy is partly is due to our choice of $f_{\rm g} = 0.5$, towards the bottom of the range of reasonable values, but also results

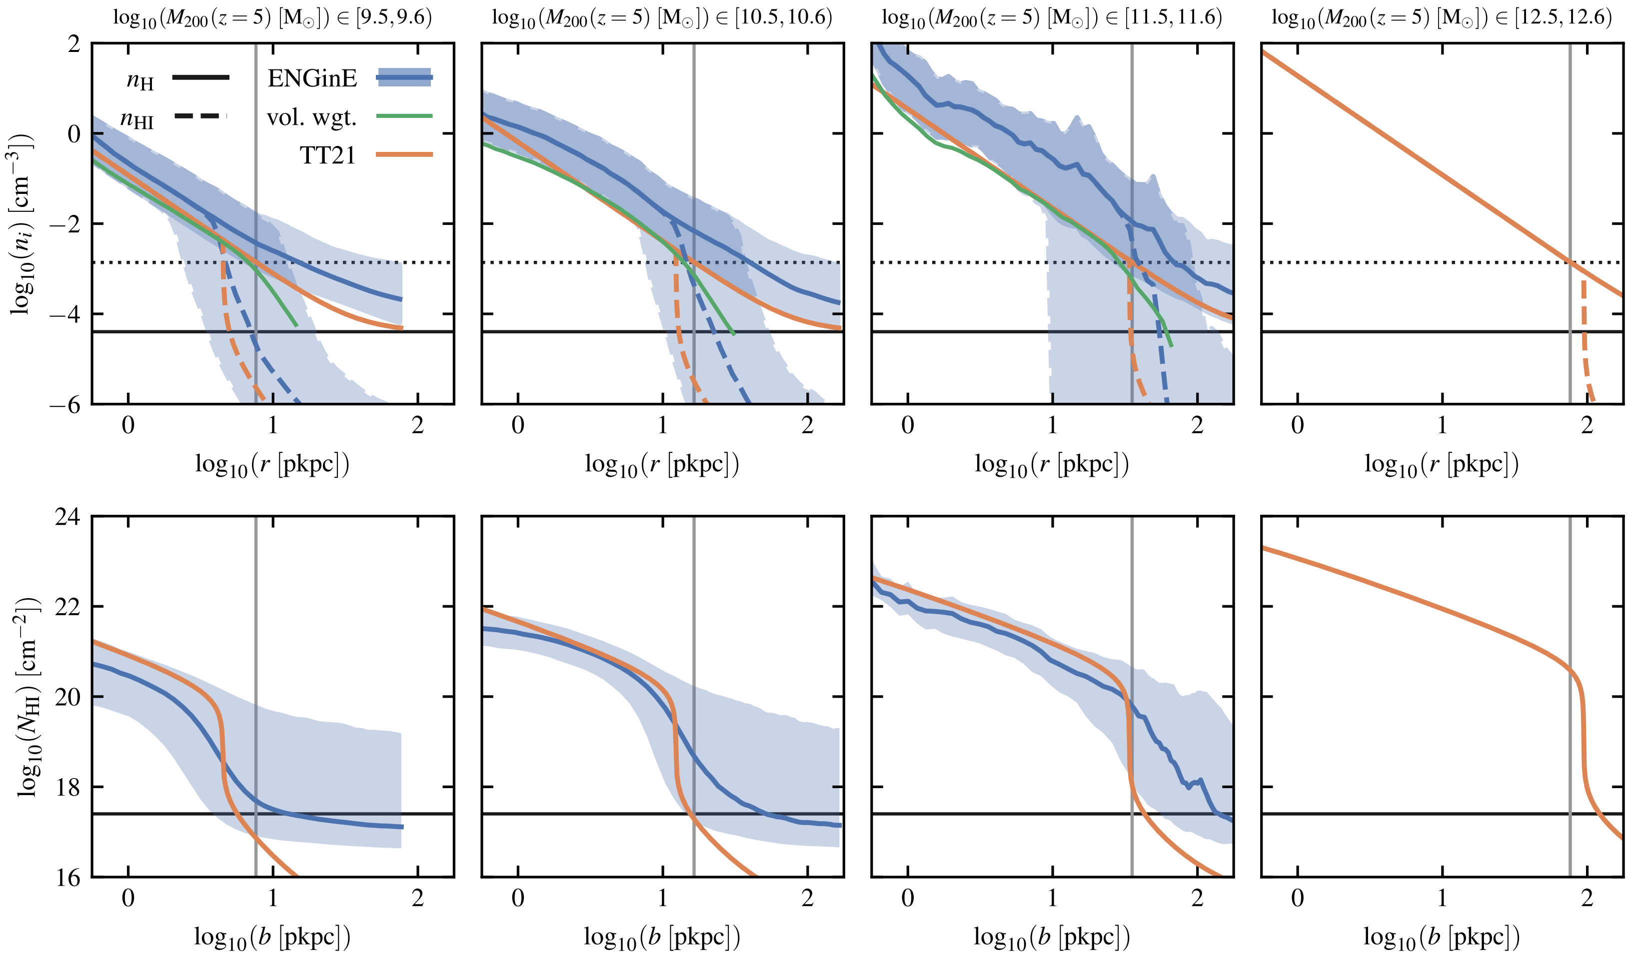

Figure 4.7

Density profiles from ENGɪɴE (blue) and the TT21 model (orange) at $z = 5$. Shaded regions indicate the 16th-84th percentile range for the ENGɪɴE haloes. *Top row*: solid and dashed lines indicate the total and neutral hydrogen density profiles respectively. Green curves show volume-weighted averages for n_{H} (see text). Also indicated are R_{200} (grey vertical line), the mean hydrogen density $\bar{n}_{\mathrm{H}}$ (solid black line), and the value of $n_{\mathrm{H},0}$ (Eq. 4.5.2) (dotted black line). *Bottom row*: neutral hydrogen column density profiles at the same redshifts. The grey vertical line is now the projected position of R_{200}, and the solid horizontal line is the mean column density along all sightlines at this redshift.

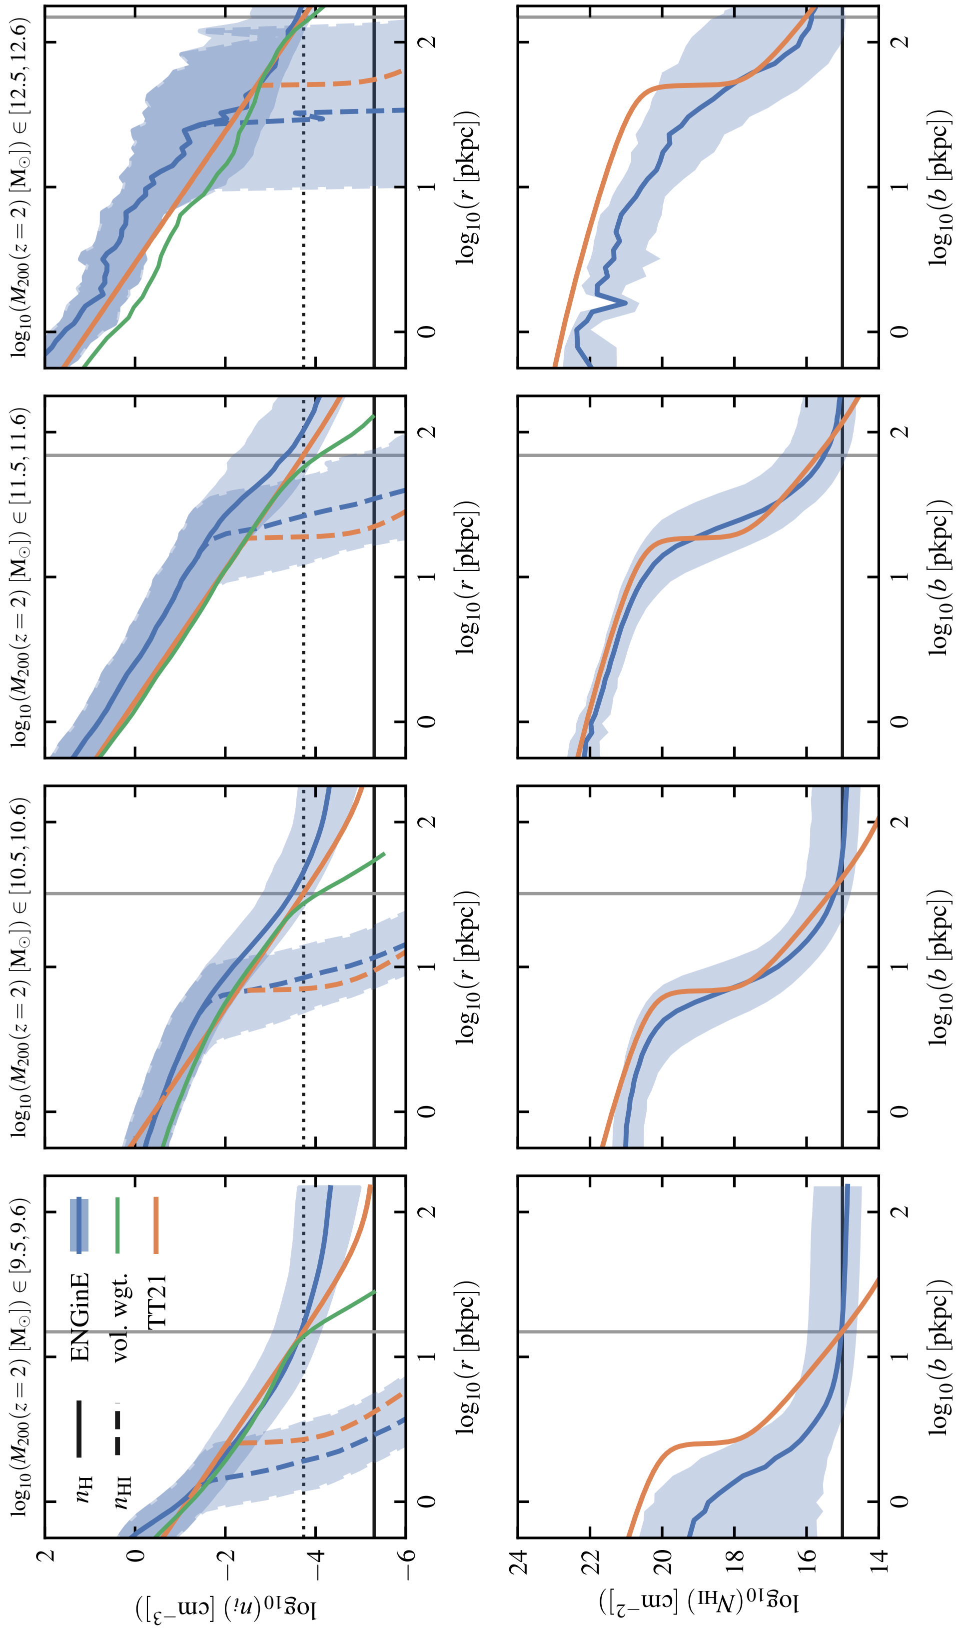

Figure 4.8
As Fig. 4.7, but at $z = 2$.

from deviations from spherical symmetry in the ENGɪɴE haloes. These are reduced but not eliminated by our procedure of combining together all particles from haloes within a given mass bin. As discussed by Altay & Theuns (2013), any clumpiness that remains will bias the mass-weighted densities that we show high with respect to the volume-weighted equivalent (i. e. the sum of particle masses within concentric radial shells, divided by the shell volumes). The latter quantity is shown with green lines in Figs 4.7 and 4.8, where we see that it tracks the lower edge of the mass-weighted n_{H} distribution, and also agrees well with the TT21 profile.

Beyond the ionisation front, the model also predicts a more abrupt decrease in the neutral gas density than seen in the simulated haloes. This primarily occurs because the slab geometry it assumes neglects the angular dependence of the optical depth arising from the three-dimensional structure of the simulated haloes, the net effect of which is to smooth out the transition from highly-ionised ($\tau < 1$) to neutral ($\tau \gg 1$) gas. The slab approximation is also responsible for the offset in the position of the ionisation front in the lowest-mass bin at $z = 2$: these haloes are close to the critical mass below which the gas remains photoionised by the UVB throughout (e. g. Okamoto et al. 2008), and so changes to the behaviour of the optical depth significantly affect their ionisation balance.

The column density N_{HI} as a function of impact parameter b may be obtained from the model density profiles by computing the Abel integral

$$N_{\mathrm{HI}}(b) = 2\int_{b}^{\infty} n_{\mathrm{HI}}(r)\frac{r}{\sqrt{r^2 - b^2}}\,\mathrm{d}r\,, \tag{4.5.6}$$

where in practice a finite value $R_{\mathrm{max}} = 10R_{200}$ is used for the upper limit of integration. For the simulated haloes, we instead extract column density profiles starting from the N_{HI} map which we previously used to determine $f(N_{\mathrm{HI}}, z)$. There is an ambiguity in how to assign the column density along each sightline in this map to individual haloes, because multiple haloes may overlap in projection. Counting every sightline with a projected separation less than R_{200} for every halo would therefore result in the 'double-counting' of absorption associated with overlapping haloes, as well as the association of anomalously high column densities with low-mass haloes found in front or behind larger ones. To resolve this, we assign sightlines to the most massive haloes they intersect, on the basis that the majority of the absorption along the sightline will generally be due to the highest-mass halo. Analogously to the procedure used to compute the density profiles, the set of sightlines associated with all haloes in a 0.1 dex range in $\log_{10} M_{200}$ are then binned by projected radius, and the $N_{\rm HI}$ profile is calculated as the median value for each radial bin. The two sets of column density profiles are compared in the bottom rows of panels in Figs 4.7 and 4.8, again for $z = 5$ and 2 respectively.

The overall normalisation of the profiles is consistent to the 1σ level or better, with exceptions being found at $z = 2$ that are due to disagreements in the position of the ionisation front. These arise because haloes become increasingly photoionised (for $M_{200} = 10^{9.5}\,{\rm M}_\odot$) and from asymmetry (at $M_{200} = 10^{12.5}\,{\rm M}_\odot$, where only three haloes have been stacked together). The slopes of the profiles are in very good agreement while the gas remains largely neutral; once the ionisation front is reached, the model column densities drop by a factor $\gtrsim$100 due to the precipitous fall in the neutral gas density predicted to occur there. In contrast, the profiles extracted from the simulation fall more gradually, in concordance with the gentler slope found for $n_{\rm HI}(r)$. This is particularly true at $z = 5$, where substantial filamentary absorption (e. g. connecting adjacent haloes) results in large scatter in $N_{\rm HI}$ and contributes to the shallower run of the median value with b. Finally, at large impact parameters the simulated column densities approach a (somewhat arbitrary) asymptotic value, which is associated with the ‘background’ column density accumulated along a sightline resulting from the residual neutral fraction in the IGM.

4.5.3 The differential absorber cross-section

The absorption associated with an individual halo may be quantified in terms of the cross-section σ, which measures the projected area within which the column density exceeds a certain threshold. The cross-section may be connected to the CDDF by considering its derivative, $\mathrm{d}\sigma/\mathrm{d}N_{\rm HI}$, which represents the probability of an absorber with given column density arising within a halo of given mass. To compute this, we start from $\mathrm{d}N$, the number of absorbers due to haloes of mass M_h with column density N_HI found along a sightline of proper length $\mathrm{d}\ell$. Such absorbers have proper cross-section $\sigma(M_\mathrm{h}, N_\mathrm{HI})$ and comoving number density $n(M_\mathrm{h}, z)$, so that:

$$\mathrm{d}N = n(M_\mathrm{h}, z)(1+z)^3\sigma(M_\mathrm{h}, N_\mathrm{HI})\,\mathrm{d}\ell\,, \tag{4.5.7}$$

where the factor $(1+z)^3$ is included to convert the number density to proper units. Given that $\mathrm{d}\ell = c\,\mathrm{d}z/[(1+z)H(z)]$, and $\mathrm{d}X$ is defined as in Section 4.4, we have:

$$\mathrm{d}N = n(M_\mathrm{h}, z)\sigma(M_\mathrm{h}, N_\mathrm{HI})\frac{c}{H_0}\mathrm{d}X \tag{4.5.8}$$

and so, the number of absorbers with given column density arising in haloes of a given mass is:

$$\begin{aligned} g(N_\mathrm{HI}, M_\mathrm{h}) &\equiv \frac{\mathrm{d}^3 N}{\mathrm{d}N_\mathrm{HI}\,\mathrm{d}\log_{10} M_\mathrm{h}\,\mathrm{d}X} && (4.5.9)\\ &= \frac{c}{H_0}\frac{\mathrm{d}n}{\mathrm{d}\log_{10} M_\mathrm{h}}\frac{\mathrm{d}\sigma(M_\mathrm{h}, N_\mathrm{HI})}{\mathrm{d}N_\mathrm{HI}}\,. && (4.5.10)\end{aligned}$$

Finally, the quantity $g(N_\mathrm{HI}, M_\mathrm{h})$ may be identified as the differential contribution to the CDDF per dex in halo mass:

$$g(N_\mathrm{HI}, M_\mathrm{h}) \equiv \frac{\mathrm{d}f(N_\mathrm{HI}, z)}{\mathrm{d}\log_{10} M_\mathrm{h}}\,. \tag{4.5.11}$$

We compute Eq. (4.5.9) from the simulations directly, by binning the projected N_HI map by N_HI and $\log_{10} M_\mathrm{h}$ (as discussed previously, each sightline is associated with the mass of the most massive halo it intersects) and normalising with $\mathrm{d}X$. The TT21 model instead obtains $g(N_\mathrm{HI}, M_\mathrm{h})$ from Eq. (4.5.10), in which the differential cross-section appears as the third term. This may be calculated from the radial N_HI profiles produced by the model; meanwhile, the second term of Eq. (4.5.10) is the halo mass function $\phi(M_\mathrm{h})$, which we evaluate using the Tinker et al. (2008) parameterisation, as implemented in the COLOSSUS package (Diemer 2018). We determine values of $\mathrm{d}\sigma/\mathrm{d}N_\mathrm{HI}$ and $g(N_\mathrm{HI}, M_\mathrm{h})$ for different redshifts, halo masses, and N_HI bins, and show the results in Fig. 4.9. Here, each column of panels corresponds to a single N_HI bin, and solid (dashed) lines are plotted for the simulation (model), at $z = \{5, 4, 3, 2\}$.

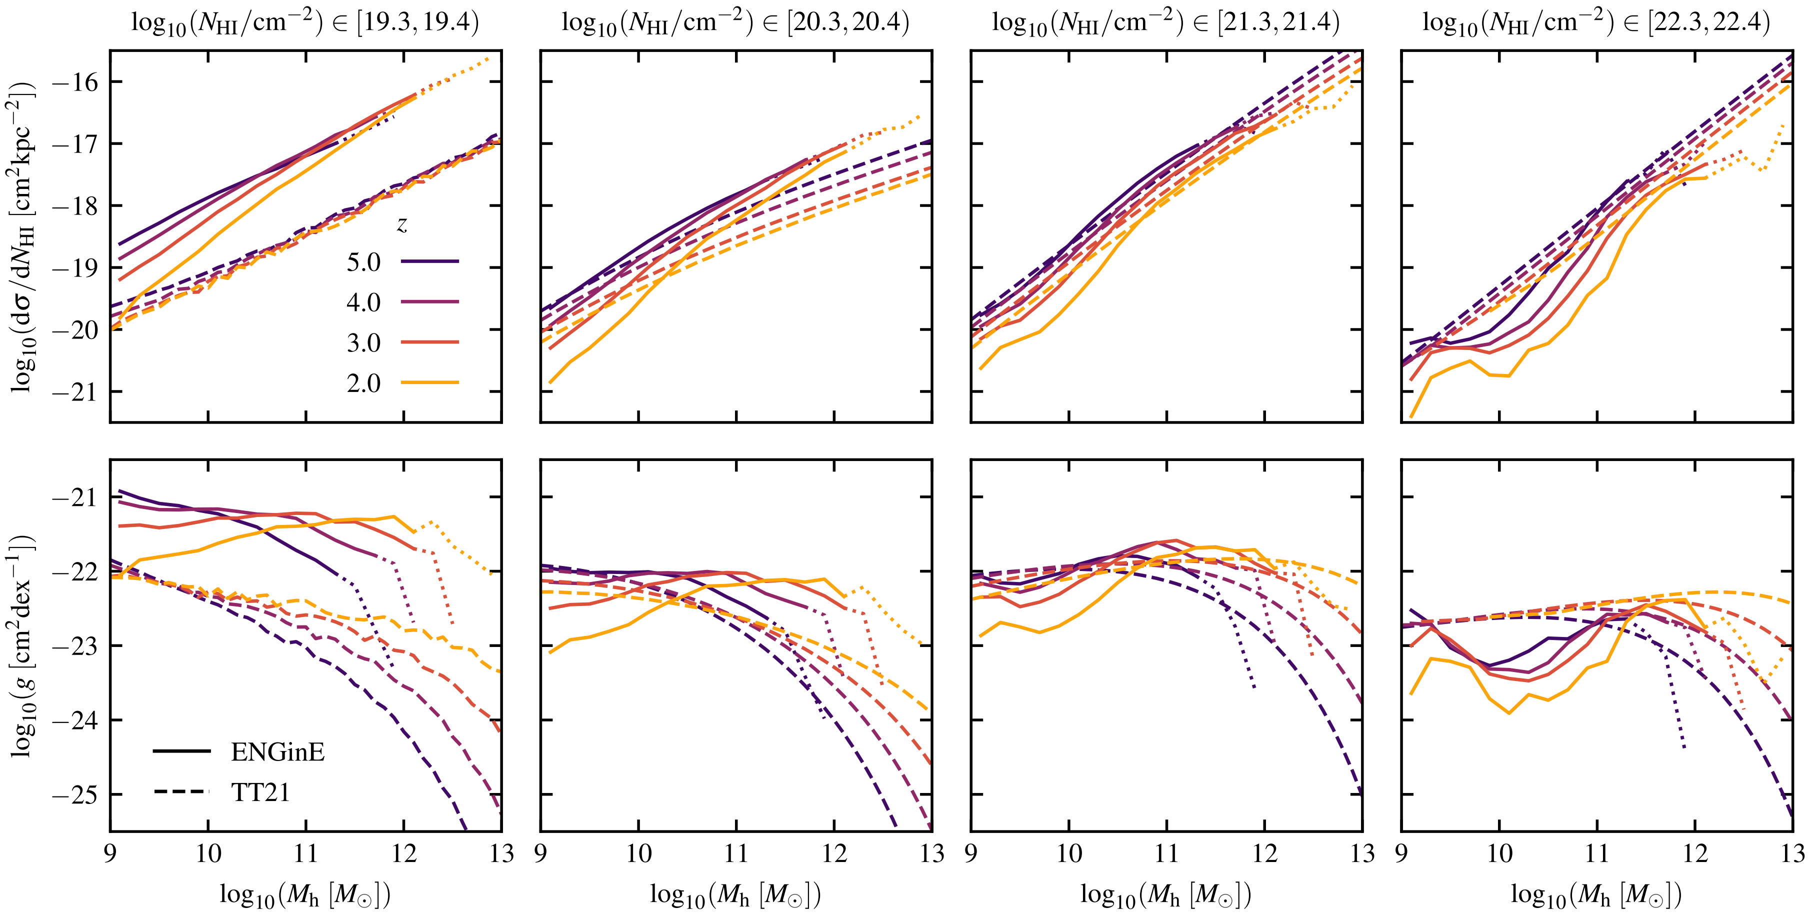

Figure 4.9

Top row: the differential cross-section $\mathrm{d}\sigma/\mathrm{d}N_{\mathrm{HI}}$ obtained from ENGINE (solid lines) and predicted by the TT21 model (dashed lines). *Bottom row*: the differential contribution to the CDDF $g(N_{\mathrm{HI}}, M_{\mathrm{h}})$. Each column of panels corresponds to a single bin in N_{HI}; within panels line colours indicate z. In the third and fourth columns, all curves have been shifted upward by 2 and 4 dex, respectively, to allow the axis range to be compressed. Curves from the simulation are dotted for mass bins containing fewer than 10 haloes.

First, we consider the behaviour of $\mathrm{d}\sigma/\mathrm{d}N_{\mathrm{HI}}$, shown in the top row of panels in Fig. 4.9. The steeper slope of the TT21 N_{HI}–b relation at impact parameters beyond the ionisation front leads to predicted values of $\mathrm{d}\sigma/\mathrm{d}N_{\mathrm{HI}}$ at $N_{\mathrm{HI}} = 10^{19.3}\,\mathrm{cm}^{-2}$ which are smaller than those derived from the simulations. Otherwise, we obtain good qualitative agreement: in both the simulation and the model, the cross-section shows a gradual increase with z at all column densities, indicating that haloes have generally higher H I content at earlier times. The shape of $\mathrm{d}\sigma/\mathrm{d}N_{\mathrm{HI}}$ is approximately a power-law function of M_{h} for every N_{HI} bin, with the slope of this relationship consistent between ENGInE and the model, particularly at higher N_{HI} where the model predicts $\mathrm{d}\sigma/\mathrm{d}N_{\mathrm{HI}} \propto {M_{\mathrm{h}}}^{1.22}$, in close agreement with the simulation. At lower N_{HI}, the cross-section becomes a steeper function of mass at later times, because the increased UVB amplitude results in an increasing degree of photoionisation, as was noted in the discussion of Fig. 4.8.

Moving on to the bottom row of Fig. 4.9, the difference in normalisation at $N_{\mathrm{HI}} \sim 10^{19}\,\mathrm{cm}^{-2}$ is again apparent, indicating that the CDDF predicted from the model (see Fig. 2 of TT21) would lie below the ENGInE result at this column density. We find that, as predicted by the model, $g(N_{\mathrm{HI}}, M_{\mathrm{h}})$ is only weakly dependent on halo mass in the DLA regime, implying that haloes of radically different masses nevertheless make similar contributions to $f(N_{\mathrm{HI}}, z)$. This is a consequence of a coincidental cancellation between the halo mass function, which scales as roughly $\phi(M_{\mathrm{h}}) \propto {M_{\mathrm{h}}}^{-1.1}$, and the differential cross-section, which shows the opposite scaling: in other words, more massive haloes are less common, but make greater individual contributions to the total cross-section for absorption.

In the model, the halo mass dependence becomes stronger at lower N_{HI}, whereas the simulations show a similarly flat $g(N_{\mathrm{HI}}, M_{\mathrm{h}})$ for all four N_{HI} bins. The main cause of this is a larger-than-predicted contribution from haloes with $M_{200} \gtrsim 10^{10.5}\,\mathrm{M}_{\odot}$ at column densities near the DLA threshold. The approximate nature of the TT21 model means that we do not *a priori* expect it to give the 'correct' answer, but in this case we note that improving the agreement here would also lead to improved consistency with observations of $f(N_{\mathrm{HI}}, z)$ and $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$. This may hint at a need for more efficient feedback in massive haloes; alternatively, it may indicate that we are misattributing some absorption arising in

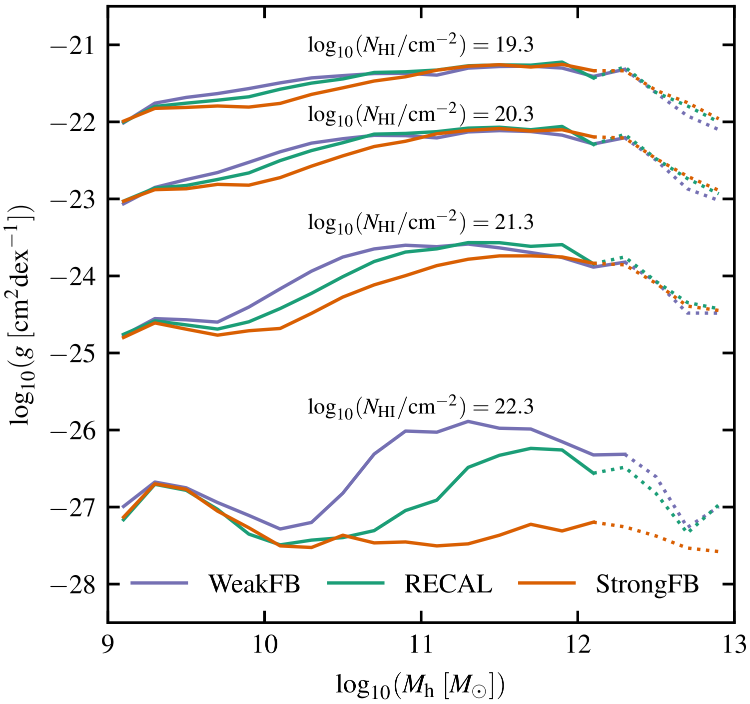

Figure 4.10

The differential contribution to the CDDF $g(N_{\rm HI}, M_{\rm h})$, for the RECAL, StrongFB, and WeakFB simulations at $z = 2$. Curves are shown for 0.1 dex $N_{\rm HI}$ bins and are dotted for mass bins containing fewer than 10 haloes.

smaller haloes to larger ones along the same line of sight.

The importance of haloes of different masses does evolve with redshift: at $z = 5$, low-mass haloes make the largest contribution to the CDDF at column densities near the DLA threshold; by $z = 2$, this is instead provided by haloes with $\log_{10}(M_{\rm h}) \geq 11$. This is again the result of competition between ϕ, which increases at fixed halo mass over cosmic time as structure formation continues, and the cross-section per halo, which decreases because the typical density of the gas within haloes falls with z. Finally, we also find the reason for the evolution in $f(N_{\rm HI}, z)$ at the high-density end of the $N_{\rm HI}$ range: the abundance of massive haloes, which are always the dominant source of very strong absorption, is heavily suppressed by the exponential tail of the halo mass function at $z = 5$, but increases toward lower redshift.

4.5.4 The characteristic column density for feedback

In Fig. 4.10, we instead show values of $g(M_\mathrm{h}, N_\mathrm{HI})$ from the Recal, StrongFB and WeakFB simulations at $z = 2$. As seen in Fig. 4.5, the impact of feedback only becomes significant at $N_\mathrm{HI} > 10^{21}\,\mathrm{cm}^{-2}$. However, for gas in intermediate-mass haloes we find evidence for a small but systematic change in $g(N_\mathrm{HI}, M_\mathrm{h})$ at all column densities, likely because the overall baryon fraction inside haloes scales with the feedback efficiency. At the extremes of the halo mass range, the effect of changing the feedback efficiency is minimal: below $10^{9.5}\,\mathrm{M}_\odot$, the majority of haloes avoid stellar feedback entirely by remaining 'dark' (i. e. star-free; these objects are in fact high-redshift analogues of the RELHICs discussed in Chapters 2 and 3). At high mass, $M_\mathrm{h} > 10^{12}\,\mathrm{M}_\odot$, the H I distribution is more extended and less affected by feedback events, while other processes (e. g. shock heating or AGN feedback) will play an increasingly important role in setting the gas conditions. This again results in minimal change in $g(N_\mathrm{HI}, M_\mathrm{h})$ with feedback, except for the highest N_HI bin which likely corresponds to sightlines that directly probe the ISM of massive galaxies.

We can make a general test of the idea that feedback begins to influence the cross-section when sightlines intersect a given galaxy's ISM by using the TT21 model to estimate the column density above which this occurs. If the slope of the gas density profile is taken to be $\alpha = 2$ (rather than the fiducial $\alpha = 2.2$), then the column density at given impact parameter b is

$$\begin{aligned} N_\mathrm{HI}(b) &= 2n_\mathrm{H,0}\,{R_{200}}^2 \int_b^{r_I} \frac{\mathrm{d}r}{r(r^2 - b^2)^{1/2}} \\ &= \frac{2n_\mathrm{H,0}\,{R_{200}}^2}{b} \arccos\left(\frac{b}{r_I}\right) , \end{aligned} \tag{4.5.12}$$

where we have further approximated by ignoring the column density contributed by gas beyond the ionisation front r_I (which is negligible for $b \ll r_I$); r_I is itself given by (Eq. 8 of TT21)

$$r_I = \left(\frac{R_{200}}{r_c}\right)^{1/(2\alpha-1)} R_{200}\,, \tag{4.5.13}$$

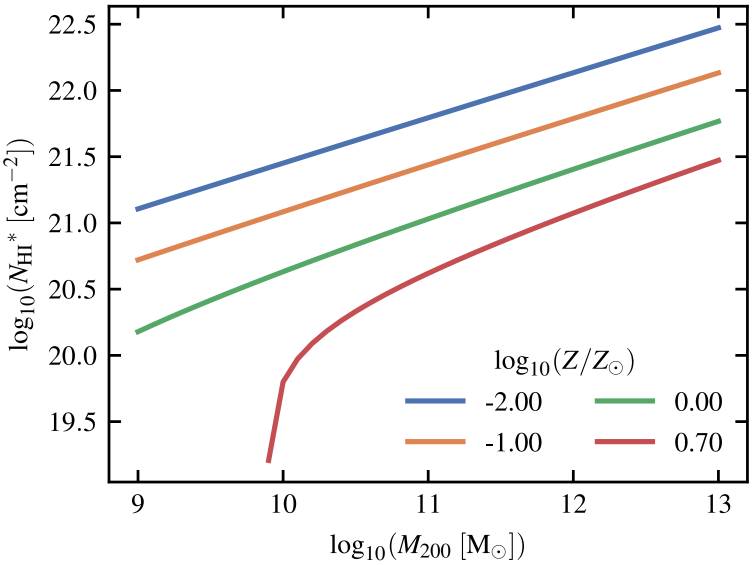

Figure 4.11

The threshold column density above which sightlines are predicted to intersect the ISM at $z = 2$, plotted as a function of host halo mass. This is also a function of metallicity, because of the dependence of the star formation threshold (Eq. 4.2.1), here shown by curves for $Z \in \{0.01, 0.1, 1, 5\}\, Z_\odot$.

where

$$r_c = \frac{(2\alpha - 1)\Gamma_0}{\sigma_{\nu_{\rm th}} \alpha_{\rm B}\, n_{\rm H,0}^2} \,.$$

The impact parameter b^* corresponding to a sightline along which the maximum density just reaches the star formation threshold ${n_{\rm H}}^*$ satisfies

$$ {n_{\rm H}}^* = n_{\rm H,0} \left(\frac{R_{200}}{b^*} \right)^\alpha \,. \tag{4.5.14}$$

Combining Eqs. (4.5.13) and (4.5.14) yields for the minimum column density associated with star-forming gas

$$ {N_{\rm HI}}^* = 2{n_{\rm H}}^* \left(\frac{n_{\rm H,0}}{{n_{\rm H}}^*} \right)^{1/2} R_{200} \arccos \left\{ \left(\frac{n_{\rm H,0}}{{n_{\rm H}}^*} \right)^{1/2} \left(\frac{r_c}{R_{200}} \right)^{1/(2\alpha-1)} \right\} \,. \tag{4.5.15}$$

values for which are plotted in Fig. 4.11 as a function of host halo mass. Because the EAGLE model takes ${n_{\rm H}}^*$ to be a function of metallicity (Eq. 4.2.1), curves are shown for four different choices of the gas metallicity between 0.01 and 5 times the solar value. First considering a sightline at the DLA threshold, $N_{\rm HI} = 10^{20.3}{\rm cm}^{-2}$, we see that it will only intersect star-forming gas in low-mass halos, and even then, only if the metallicity of the gas is high. Thus, the majority of such sightlines do not intersect star-forming gas, and are unlikely to be strongly affected by stellar feedback. At a factor-10 larger column density, sightlines may instead intersect star-forming gas over a range of halo masses, corresponding to lower metallicities in lower-mass haloes, and vice versa in larger ones. Thus, we identify $N_{\rm HI} \sim 10^{21}\,{\rm cm}^{-2}$ as the threshold where a typical DLA sightline may first encounter star-forming gas, and hence at which feedback begins to significantly affect the CDDF.

4.5.5 The lifecycle of DLA gas

In the previous sections, we have shown that the TT21 model can successfully predict many of the properties seen in the ENGinE DLAs. However, we have yet to directly test the most important assumption of the model, namely that DLAs are quasi-static structures, with which gas from the IGM is only transiently associated as it accretes onto haloes. In this section, we will establish whether this interpretation of the "life-cycle" of DLA gas is accurate, by tracing gas which forms part of a DLA at a given redshift forwards through cosmic time.

To accomplish this, we identify all particles at $z = 5$ that have neutral fractions $x_{\rm HI} > 0.5$, and classify them as "DLA gas". These particles are additionally tagged with the mass of the FoF halo in which they are found at $z = 5$ (or that they do not belong to any halo). For each of $z = \{4, 3, 2\}$, we then extract the same set of particles, and determine the fractions of them which remain in DLAs, have become (re)-ionised, and have been transmuted into star particles. We also record whether the particles are found inside a FoF halo at each later time, and the value of the `OnEquationOfState` flag, which indicates whether the particle is currently on the star-forming equation of state, is not currently but has been at some earlier time, or has never been star-forming.

The result of this categorisation is shown in Fig. 4.12. In this figure, each panel corresponds to one of the halo mass bins identified at $z = 5$, and the segmented bars indicate the proportion of particles falling into each category at later times. Focusing first on overall

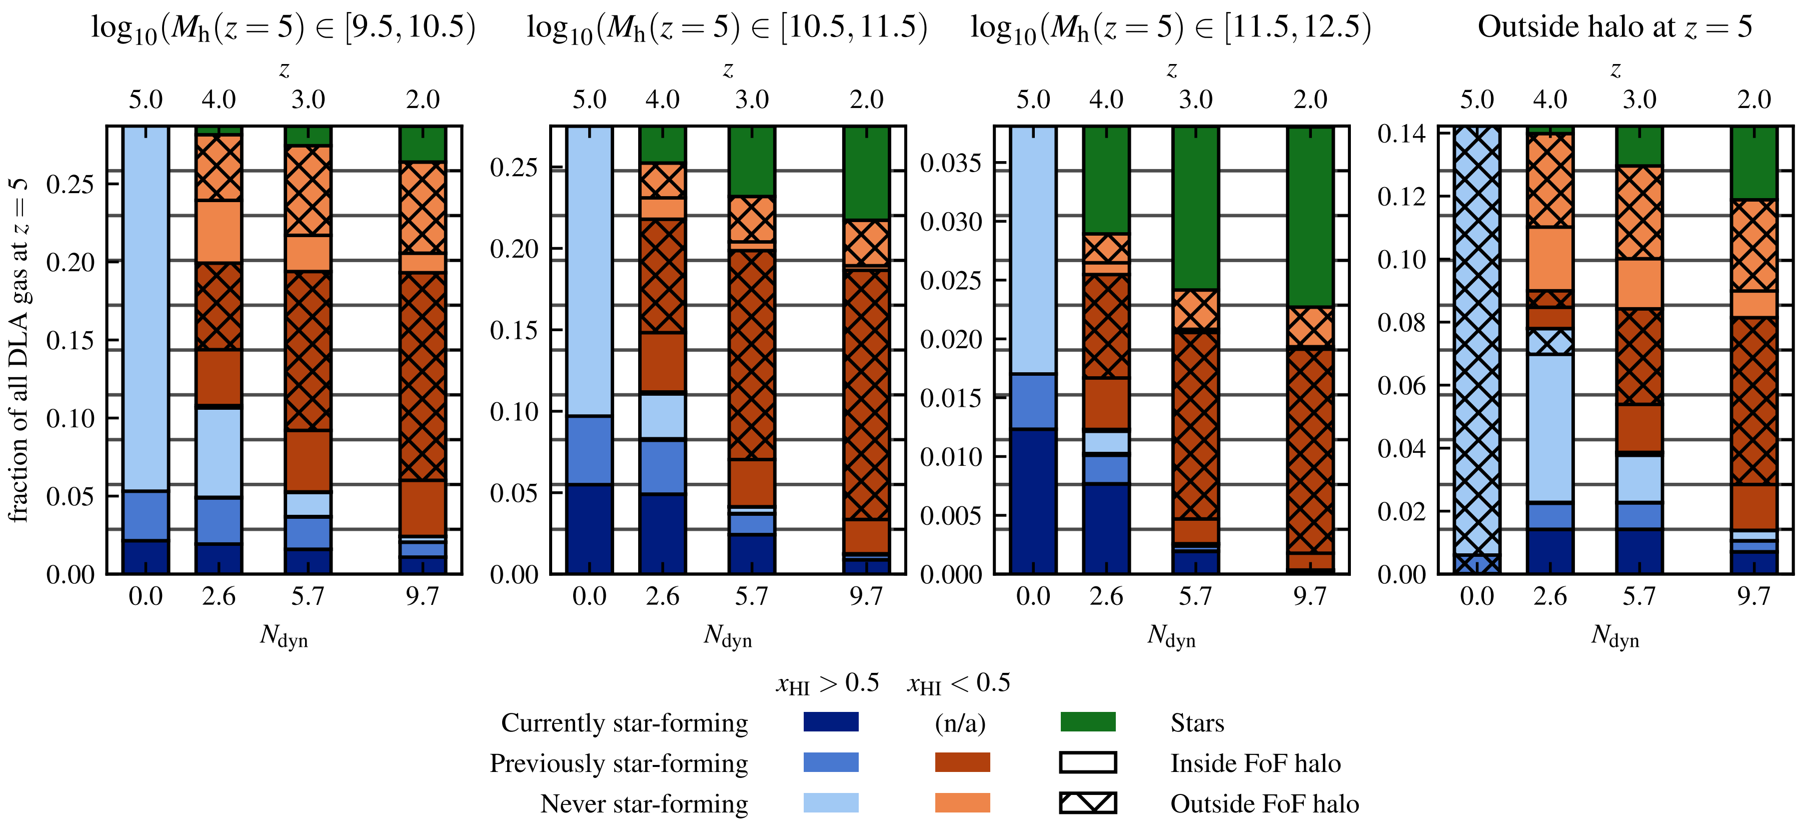

Figure 4.12

The lifecycle of DLA gas. Particles belonging to a DLA at $z = 5$ (defined as those with -neutral fraction $x_{HI} > 0.5$) are divided into one of three halo mass bins, plus one for particles which do not belong to any FoF halo. At $z = 4$, 3, and 2, the same particles are identified and partitioned into those which still form part of a DLA, have become ionised gas, and have been converted into stars. In the first two of these cases, the particles are further categorised by whether they are found within any FoF halo at the current time, and whether they are or have ever been actively star-forming. Each bar segment indicates the number of particles falling into the corresponding category, as a fraction of the total number of neutral gas particles at $z = 5$, such that the y-axis range in each panel indicates the fraction of all $z = 5$ DLA particles falling into that mass bin. The x-coordinates of the bars correspond to N_{dyn}, the cumulative number of dynamical times elapsed since $z = 5$, which we calculate by summing the values of $T_{dyn} \equiv H(z)^{-1}/\sqrt{200}$ averaged between successive pairs of redshifts. To guide the eye, gridlines are drawn in each panel at 10% intervals.

behaviour, we find that at all halo masses, the majority of DLA gas at $z = 5$ is indeed not associated with star formation. By $z = 2$ no more than 10% of this gas remains neutral, with the majority becoming ionised as soon as $z = 4$, and no longer associated with any FoF halo. However, more than three quarters of this expelled gas was cool and dense enough to join the star-forming equation of state at some point between $z = 4$ and 5; moreover, at $z = 4$ there is more currently-ionised but previously star-forming gas than the total $z = 5$ ISM content. These are exactly the trends we would expect if DLAs are primarily composed of actively-accreting gas, which subsequently reaches the innermost region of the halo it falls into, briefly joins the ISM, and is then expelled by feedback (with the exception of the small fraction that successfully undergoes star formation). Meanwhile, gas which is still found inside a DLA at $z = 2$ is predominantly star-forming, suggesting that this is gas which avoided being supplied with energy from feedback, instead settling into the ISM deep in the host halo's potential well. Some of this gas may also represent material that has undergone at least one complete "fountain" cycle, by being ejected by feedback before $z = 2$ and subsequently re-accreting. The typical timescale for such a cycle is dominated by the accretion phase, which has a duration on the order of the halo dynamical time. As indicated by the axis labels in Fig. 4.12, which show the cumulative number of dynamical times elapsed since $z = 5$, it is possible for a large number of feedback cycles to occur, providing a mechanism for ongoing enrichment of the halo gas.

There is a marked increase in star formation efficiency with halo mass, which results in the proportion of $z = 5$ DLA gas that has transformed into stars by $z = 2$ rising from $\sim$10% for haloes of mass $10^{9.5} - 10^{10.5}\,\mathrm{M}_{\odot}$, to 40% for haloes two orders of magnitude more massive. The proportion of ejected gas which has previously been star-forming also increases with mass, suggesting that in deeper halo potentials, gas on average cools for longer before receiving an injection of energy sufficient to eject it. This trend also indicates a mass-loading factor ($\eta \equiv \dot{M}_{\mathrm{wind}}/\dot{M}_{*}$, the ratio of the total mass of gas expelled by feedback to the mass of stars formed) that declines with halo mass, in agreement with the findings of Mitchell et al. (2020), who calculated this quantity for the reference EAGLE simulation explicitly. The final panel of Fig. 4.12 shows the evolution of DLA gas which does not belong to a FoF halo at $z = 5$, finding that ~40% has been accreted by $z = 2$. The relatively high efficiency with which this gas is converted into stars implies that this material is dominated by gas in the vicinity of massive haloes, which is not unexpected, given that they will have the most extended H I distributions.

4.6 Summary

We have presented the ENGinE simulations, which extend the EAGLE project to include new volumes that simultaneously offering large computational volume and relatively high numerical resolution. In addition to a fiducial model, we have performed additional runs with variations to modelling assumptions regarding the nature of dark matter and the subgrid feedback processes. All six volumes will be made available for public use shortly.

The analysis presented here has focused on the properties of damped-Ly α absorbers at $z = 2 - 5$, which we have investigated by post-processing the simulation outputs with an accurate radiative transfer code. Our findings are summarised as follows:

(i) We calculated the H I column density distribution function (CDDF) $f(N_{\rm HI}, z = 3)$ from our fiducial Recal simulation, and demonstrated its agreement with observations over seven orders of magnitude in column density.

(ii) By also considering observational constraints on the column density distribution of molecular gas, we concluded that the total neutral gas distribution remains consistent with data up to the highest column densities encountered, although uncertainties in the modelling are such that definitive predictions of the split between atomic and molecular phases remain difficult.

(iii) We confirmed that the fiducial EAGLE stellar feedback model, which is calibrated to reproduce the observed galaxy stellar mass function, is also optimal for the CDDF. Factor-2 variations in the feedback efficiency lead to large ($\gtrsim$1 dex) shifts in $f(N_{\rm HI}, z)$ at high column densities.

(iv) We further calculated the DLA abundance $\mathrm{d}N/\mathrm{d}X$ and mass density $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$, finding that their evolution between $z = 5$ and 2 is broadly consistent with observations, although especially for $\Omega_{\mathrm{HI}}^{\mathrm{DLA}}$ our results are in some tension with the data at $z = 2$, primarily as a result of excess H I in massive haloes.

(v) We verified a recent analytic model for DLAs, finding that it successfully reproduces many of the properties of our simulated absorbers. In particular, we used it to show that the characteristic column density at which feedback begins to significantly affect the absorption statistics is that associated with lines of sight that begin to probe galaxies' interstellar media.

(vi) Through the overall success of the model, and further explicit investigation of the life-cycle of neutral gas, we demonstrated evidence for an extended gaseous halo surrounding galaxies being the primary source of DLA absorption. These haloes are quasi-static structures, formed of gas that is actively accreting onto the galaxy. Predominantly this gas is pristine material originating from the IGM, but also contains a significant component of gas that has previously been ejected by feedback as part of the baryon cycle.

CHAPTER 5

Summary and outlook

In this thesis, we have followed two distinct lines of inquiry, concerning phenomena on substantially different physical scales and occurring at different phases of the Universe's history. The common thread linking these topics is their focus on the properties of the diffuse gas found within cosmic structures, as revealed observationally through emission and absorption lines, and modelled using sophisticated numerical simulations.

5.1 RELHICs

The first topic we discussed involved RELHICs, gas-rich but star-free dark matter haloes first identified in high-resolution hydrodynamical simulations of the local Universe. In Chapter 2, we presented an ionisation balance code designed to model RELHICs, in which their gas content is set by the imposition of hydrodynamic equilibrium with the gravitational potential of the host dark matter halo, and ionisation equilibrium with the extragalactic ultraviolet background (UVB). After validating this code against the existing CLOUDY software, we demonstrated that it reproduces the density and temperature structures of RELHICs seen in the simulations. The gas within RELHICs absorbs energy from high-energy photons in the UVB, which is subsequently radiated away in recombination lines such as H α. We found that the surface brightness of this fluorescent emission peaks sharply at the projected radius of the ionisation front, giving RELHICs a characteristic ring-shaped morphology when seen in projection. Finally, we showed that observable (in principle) characteristics of these rings such as their projected size and peak surface brightness are sensitive to modelling assumptions including the UVB spectrum and the density profile of the underlying dark matter halo. We presented possible ways this dependence could be used to extract information from observations of a RELHIC with fluorescent emission in $\mathrm{H}\,\alpha$. In Chapter 3, we showed that the primordial composition and generally well-constrained physics of RELHICs presents another intriguing possibility: they are ideal environments in which to measure the primordial helium abundance y_{P}. We found that even in the absence of a spatially-resolved detection of the ring structure, measurements of the total flux in hydrogen and helium lines can provide an accurate determination of y_{P}, as well as a novel constraint on the (poorly-known) spectral slope of the UVB.

All these possibilities are enticing, but it is–arguably–more than a little premature to discuss them. At present, RELHICs are purely hypothetical objects. The simulations which predict them are not reality, and their approximate modelling of physical processes renders the existence and detailed properties of RELHICs somewhat uncertain. Meanwhile, we showed that RELHICs exhibiting fluorescent emission are relatively rare objects, with no more than half a dozen expected in a 3 Mpc volume surrounding the Local Group. Furthermore, their small angular size and extremely faint emission combine to present formidable challenges for detection.

So, what avenues for progress exist? From the theoretical perspective, the key uncertainty relates to the ability of dark matter haloes to avoid ubiquitous star formation before reionisation, while also retaining their baryonic content until the present day. This was investigated by Benitez-Llambay & Frenk (2020), who examined RELHICs in a more cosmological context. They demonstrated that under the favoured 'late' scenario for reionisation, in which it begins no earlier than $z = 9$ (Aghanim et al. 2020, 2021), galaxies may form in less massive haloes, such that the majority of haloes at the present-day mass scale $M_{200} \sim 10^{9.5}\,\mathrm{M}_{\odot}$ that we considered will have luminous counterparts. Nevertheless, a population of 'dark' haloes (i. e. RELHICs) is still expected to exist, given the observational constraint that reionisation completes no later than $z \sim 6$ (e. g. Becker et al. 2015; Bouwens et al. 2015). Further work is required in order to determine the extent to which fluorescent H α ring emission is expected for this or any other reionisation model, since the baryon fraction inside RELHICs, which forms an input parameter to our calculations, is sensitive to the reionisation timeline. The relative importance of atomic and molecular hydrogen cooling processes forms an additional source of uncertainty, with the latter potentially allowing haloes with much lower present-day masses to host luminous galaxies. Modelling of the effects these cooling processes is currently seeing rapid progress, due to its relevance for determining the conditions of Population III star formation. There is some consensus that atomic cooling dominates (e. g. Hu et al. 2016; Xu et al. 2016), but this is not yet universal: indeed, in some simulations which explicitly couple together the formation of H_2 and stars, the introduction of an implicit metallicity dependence on star formation results in fewer luminous structures than in models that only consider a density dependence for star formation (Kuhlen et al. 2012; Munshi et al. 2019).

On the observational front, we have identified the most plausible route to detecting RELHICs as being the deep H I surveys being undertaken by SKA precursors (e. g. MeerKAT; Jarvis et al. 2018) using the 21 cm hyperfine transition. These surveys have sufficient sensitivity to detect the H I content of any RELHICs in the nearby ($d \lesssim 20$ Mpc) Universe. Candidate RELHICs may be identified from the source catalogues they will produce via comparison with the H I properties predicted for RELHICs by Benítez-Llambay et al. (2017), and followed up with broad-band optical imaging to rule out the presence of a stellar component. Some preliminary work along these lines has been conducted already. Any detection of the fluorescent H α emission would then require extremely deep spectrally-resolved observations, requiring many tens of hours of integration time on an instrument such as MUSE. This would undeniably be a challenging task, but we do not consider it outside the realms of possibility (see e. g. Fumagalli et al. 2017a). An interesting alternative is presented by the Dragonfly telescope (Abraham & van Dokkum 2014; Lokhorst et al. 2020), which is designed to conduct narrow-band imaging at extremely low surface brightnesses. In fact, this instrument has been specifically shown to be capable of reaching the surface brightness threshold of $10^{-20}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$ required to detect the H α emission from RELHICs (Lokhorst et al. 2019).

5.2 DLAs in cosmological simulations

In Chapter 4, we changed direction, to instead focus on the properties of damped Ly α (DLA) absorbers at $z \sim 3$ using a set of cosmological simulations, which we named ENGinE (Exploring Neutral Gas in Eagle). These extend the existing Eagle simulation suite by evolving a large 50 cMpc volume to $z = 2$ at double the spatial resolution of existing runs, providing a uniquely wide dynamic range. In addition to a fiducial (aka "Recal") simulation, they include runs in which the efficiency of stellar feedback is altered, in which feedback from AGN is disabled, and in which a sterile neutrino warm dark matter cosmology is assumed.

After demonstrating that the galaxy stellar mass function extracted from our fiducial simulation is in agreement with expectations, we described our use of the reverse ray-tracing code Urchin to model the effects of self-shielding in dense regions, and hence partition gas into ionised, atomic, and molecular components. From the simulation snapshots post-processed with Urchin, we computed a two-dimensional column density map, and hence calculated the H i column density distribution function (CDDF). We demonstrated that the CDDF we obtain compares favourably with constraints on it derived from observations of both atomic and molecular gas, and showed how it is affected by the strength of stellar feedback, as parameterised by the efficiency parameter of the feedback model. We also calculated the DLA line and mass densities, and compared these metrics with observations from $2 < z < 5$, finding that the observed evolutionary trends are well-reproduced by the simulations. However, the mass densities we obtained are approximately 30% higher than those obtained from the observations, implying that the simulations produce slightly too many absorbers with $N_{\rm HI} \sim 10^{20.5}\,{\rm cm}^{-2}$.

We proceeded to investigate the relationships between absorbers and the haloes that host them, with the aid of an analytical model for H i in haloes. Despite the simplicity of this model, we showed that it produces volumetric and projected density profiles in good agreement with the equivalents extracted from the simulation, giving us confidence in its use as a yardstick to compare against. We determined the contribution of haloes with different masses to the CDDF, gaining insight into the remarkably minimal evolution the CDDF undergoes between $z = 5$ and 2, and potentially identifying an excess of H I in massive haloes as the underlying cause of the discrepant values of the mass density. We also identified the column density at which the CDDF becomes sensitive to the feedback model, $N_{\mathrm{HI}} \sim 10^{21}\,\mathrm{cm}^{-2}$, as the threshold above which the majority of sightlines pass through a galaxy's ISM. Finally, by tracking the simulated gas distributions through cosmic time we showed explicitly that they support the viewpoint that the gas making up DLAs is primarily composed of material accreting onto galaxies from the IGM, and consequently that DLAs are quasi-static structures whose content evolves with time.

Of the results presented here, the convincing support we find for this model is probably the most significant, given the historical debate regarding the closeness of the association between DLAs and galaxies. We find that the majority of DLAs arise in a gaseous halo that extends beyond the stellar component of galaxies, which provides natural explanations for observational trends such as the large projected separations between DLAs and their putative hosts, and the anti-correlation of both column density and metallicity with impact parameter. From a practical perspective, it also provides some confidence in the reliability of DLA properties computed from simulations, since as we have shown, only the highest–column density sightlines are affected by uncertainties arising from the inevitably approximate treatment of the ISM in current-generation simulations.

However, there are many interesting phenomena we have yet to investigate. Reproducing the observed DLA bias (i. e. the degree to which DLAs cluster more strongly than the underlying mass distribution) has proven somewhat challenging for cosmological simulations (e. g. Bird et al. 2014; Rahmati & Schaye 2014; Padmanabhan et al. 2016, 2017). These simulations favour lower values than observed, which corresponds to an underestimate of the typical DLA host mass, or equivalently, to a DLA cross-section weakly dependent on halo mass, $\sigma(M_{\mathrm{h}}) \propto {M_{\mathrm{h}}}^{\beta}$ with $\beta \lesssim 1$. We instead find $\beta \sim 1.15$, which agrees well with the value required to reproduce the observed bias (Pérez-Ràfols et al. 2018b).

Likewise, simulations have historically struggled to match the observed distribution of DLA line widths, as measured by the v_{90} metric, which quantifies the velocity width enclosing 90% of the optical depth for an absorption line. Typically, simulations have predicted an excess of narrow absorbers, while failing to reproduce the high-v_{90} tail of the observed distribution. Proposed resolutions to the former issue generally invoke efficient feedback to suppress the formation of DLAs in lower-mass haloes. We see some evidence for this, in that $\mathrm{d}\sigma/\mathrm{d}N_{\mathrm{HI}}$ evaluated at the DLA threshold steepens at low masses; however, it remains to be seen to what extent this would produce a velocity width distribution compatible with observations.

Of the ENGinE simulations we presented, we focused on the fiducial Recal volume, as well as the two runs with modifications to the efficiency of stellar feedback. The additional runs, NoAGN in which AGN feedback is disabled, and WDM in which a sterile neutrino dark matter cosmology is assumed, have yet to be analysed. The WDM model is chosen to be (marginally) consistent with the Ly α forest observations which constrain the column density distribution at low N_{HI}. However, the halo mass function generated by this model remains significantly different to the CDM case, and the effects this has on both the DLA bias and velocity width distributions will be interesting to study.

Meanwhile, preliminary investigations suggest that the NoAGN run produces a CDDF almost indistinguishable from the fiducial one. This is surprising, as Altay et al. (2013) found that incorporating AGN feedback in the OWLS model significantly suppressed $f(N_{\mathrm{HI}}, z)$ in the strong-DLA regime. Furthermore, Rahmati et al. (2015), who use a reference EAGLE simulation with the same subgrid model for AGN feedback as ours, find that it significantly modifies the covering fraction of LLSs in haloes with $M_{200} > 10^{12}\,\mathrm{M}_{\odot}$. The reason for the different behaviour we see may simply reflect the fact that in a 50 Mpc volume at $z = 2$, the massive haloes which host AGN remain rare, and those that do form may not be undergoing sufficient accretion to drive strong feedback. Alternatively, the absence of feedback from AGN may have been replaced by a larger contribution from stellar feedback, although the different mass regimes in which these two mechanisms are most effective means it would still be unexpected that this substitution leaves $f(N_{\mathrm{HI}}, z)$

unmodified.

Finally, there are several opportunities for connecting the simulation results to observations that we have not followed. The analyses we have presented have primarily been phrased in terms of halo masses, but these cannot generally be determined from observations. Therefore, it would be valuable to investigate the properties of DLA absorption also as a function of stellar mass. One direction we plan to explore is to extend the analytic model of TT21 into this regime by coupling it to the "$I\kappa\epsilon\alpha$" model (Sharma & Theuns 2020), which describes galaxy properties in the context of self-regulation between feedback and cosmological accretion.

We have also neglected metals more or less entirely. In reality, observations of metal lines provide an important complementary source of information about absorption line systems, and allow distinct phases within circumgalactic gas to be distinguished. Moreover, the abundance and distribution of metals should be sensitive to the nature of outflows. We have shown that gas which is at one time star-forming (and hence will be enriched with metals) can reappear in a DLA at some later time, having been ejected by feedback. It will be interesting to investigate how the resulting metallicity distribution in DLAs is affected by this cycle.

5.3 Final remarks

It is remarkable to note the pace of advancements in our understanding of the Universe. Throughout the body of work comprising the modern science of cosmology, there are no elements much more than one hundred years old. Within the author's lifetime, new instruments and observational techniques, along with the transition into an era of "big data", have generated a quantity and quality of astronomical data that the field's pioneers could scarcely have dreamed of. Meanwhile, the development of new computational methods and the steady progression of Moore's law have allowed the nascent field of cosmological simulations to become a powerful sandbox, within which competing models may be contrasted and evaluated against data.

A central pillar of the field is the ΛCDM cosmological model, many aspects of which have been exhaustively validated against observations. The RELHICs discussed in this thesis can provide a further such test, on a scale yet to be probed by observations. Although their detection is undoubtedly challenging, barring a cosmological paradigm shift we are optimistic that it is eventually guaranteed. It is less clear whether real RELHICs will closely resemble the idealised, star-free objects seen in simulations, but even if not, the insights into structure formation and galaxy evolution at the lowest masses they would deliver will be no less valuable.

Meanwhile, future prospects are bright for the study of gas around galaxies. The new generation of multi-object spectrographs, e. g. DESI (DESI Collaboration et al. 2016); WEAVE (Pieri et al. 2016), will further enlarge the catalogue of quasar absorption line systems, providing more robust statistics especially for rare objects. Building on the transformative impact of MUSE, shorter-wavelength IFUs such as KCWI (Morrissey et al. 2018) and the conceptual Blue-MUSE (Richard et al. 2019) will extend similar capabilities to lower redshifts. The long-awaited launch of JWST will instead drive a new wave of developments in high-redshift science. It will provide an unprecedented view on galaxy formation and evolution during the era of reionisation, which will eventually be complemented by radio observations of neutral gas conducted by the SKA. And lastly, the unparalleled light gathering power offered by 30 m-class telescopes will push detection limits to new low levels, perhaps facilitating the detection of the elusive population of faint DLA hosts.

In the face of all these developments, the principal challenge for simulations is keeping pace. The push to ever-larger simulated volumes is supported by algorithmic improvements (e. g. SWIFT; Schaller et al. 2018), while ongoing work into the development of more sophisticated modelling will extend robust simulation predictions into the new parameter space targeted by potential future observations.

APPENDIX A

Additional material from Chapter 2

In this Appendix we present a derivation of an analytic result for the H α surface brightness in the optically-thick limit (see Fig. 2.1), following the method discussed by Gould & Weinberg (1996). We also report the results of further comparisons between the ionization balance code introduced in Section 2.2 and the existing code CLOUDY.

A.1 Analytic estimate of H α surface brightness in optically-thick limit

For a constant-density, plane-parallel slab of pure hydrogen composition, photoionization equilibrium gives

$$\alpha(1 - x_{\rm HI})^2 {n_{\rm H}}^2 = \Gamma x_{\rm HI} n_{\rm H} \, , \tag{A.1.1}$$

where $x_{\rm HI} = n_{\rm HI}/n_{\rm H}$, and α is the recombination coefficient for H^+. At a depth ℓ inside the cloud, the photoionization rate is given by its value at the illuminated face of the cloud Γ_0, attenuated by the total optical depth $\tau(\ell)$, i.e.

$$\Gamma(\ell) = \Gamma_0 \exp\left[-\tau(\ell)\right] \, . \tag{A.1.2}$$

The Ly α surface brightness in units of photons cm^{-2} s^{-1} is given by

$$\Sigma_{{\rm Ly}\alpha} = f_{{\rm Ly}\alpha} \int_0^{\ell} \alpha(1 - x_{\rm HI})^2 {n_{\rm H}}^2 \, {\rm d}\ell' \tag{A.1.3}$$

where $f_{\mathrm{Ly}\alpha} \sim 0.6$ is the fraction of recombinations that lead to emission of a Ly α photon (Osterbrock & Ferland 2006), and α is the recombination coefficient. Substituting Eqs. A.1.1 and A.1.2 into the above then gives

$$\Sigma_{\mathrm{Ly}\alpha} = f_{\mathrm{Ly}\alpha} \int_0^{\ell} x_{\mathrm{HI}} n_{\mathrm{H}} \Gamma_0 \exp\left(-\tau\right) \, \mathrm{d}\ell' \,. \tag{A.1.4}$$

We next define a photon number-weighted average of the photoionization cross section $\langle\sigma\rangle$ as

$$\langle\sigma\rangle \equiv \frac{\int_{\nu_{\mathrm{th}}}^{\infty} \sigma(\nu)\phi(\nu)\,\mathrm{d}\nu}{\int_{\nu_{\mathrm{th}}}^{\infty} \phi(\nu)\,\mathrm{d}\nu} \,. \tag{A.1.5}$$

where $\phi(\nu) \equiv J(\nu)/h\nu$ gives the number of photons per second with frequency ν. With this definition, the total optical depth becomes

$$\tau \equiv \int_0^{\ell} x_{\mathrm{HI}} \langle\sigma\rangle n_{\mathrm{H}} \, \mathrm{d}\ell' \quad \text{and so} \quad \mathrm{d}\tau = x_{\mathrm{HI}} \langle\sigma\rangle n_{\mathrm{H}} \, \mathrm{d}\ell' \,. \tag{A.1.6}$$

Inserting these results into Eq. A.1.4 gives

$$\Sigma_{\mathrm{Ly}\alpha} = f_{\mathrm{Ly}\alpha} \frac{\Gamma_0}{\langle\sigma\rangle} \int_0^{\tau'=\tau} \exp(-\tau') \, \mathrm{d}\tau' \,. \tag{A.1.7}$$

In the optically thick limit, $\tau \longrightarrow \infty$ and so $\Sigma_{\mathrm{Ly}\alpha} = f_{\mathrm{Ly}\alpha} \, \Gamma_0/\langle\sigma\rangle$. Rewriting this result using Eq. A.1.5 and the definition of Γ_0 previously given in Eq. 2.3.1 yields

$$\Sigma_{\mathrm{Ly}\alpha} = 4\pi f_{\mathrm{Ly}\alpha} \int_{\nu_{\mathrm{th}}}^{\infty} \phi(\nu) \, \mathrm{d}\nu \,, \tag{A.1.8}$$

which is the Gould & Weinberg (1996) result: in the optically thick limit, all photons are absorbed and therefore the Ly α surface brightness tends to a constant fraction of the illuminating flux set by $f_{\mathrm{Ly}\alpha}$. To obtain an equivalent result for H α, we divide by the case B flux ratio $F_{\mathrm{Ly}\alpha}/F_{\mathrm{H}\alpha} = 8.5$ for $T = 10^4$ K (Osterbrock & Ferland 2006). The resulting prediction, $\Sigma_{\mathrm{H}\alpha,\mathrm{GW}} = 2.15 \times 10^{-20} \mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$, provides a useful check of our code.

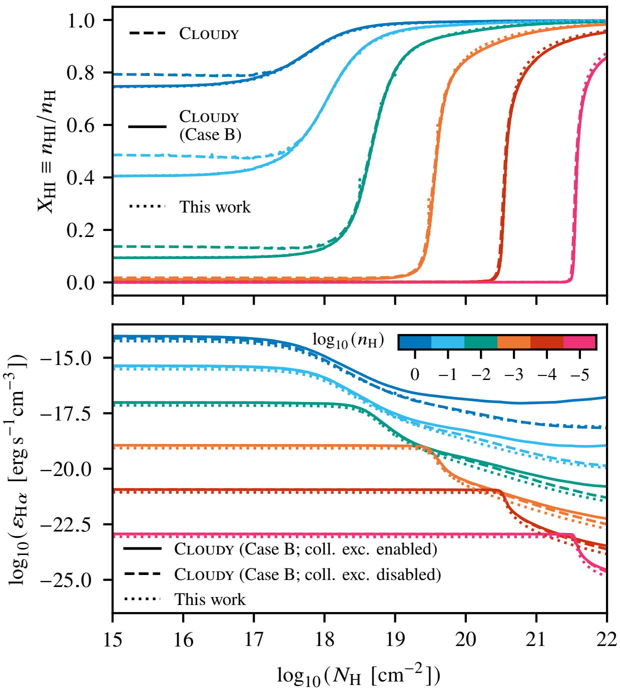

Figure A.1

Upper panel: H I neutral fraction as a function of total H column density for isothermal, H-only plane parallel models. Results from our code are shown by dotted lines; over much of the plotted $N_{\rm H}$ range, these are indistinguishable from the solid lines, which show CLOUDY results with the Case B assumption switched on. Relaxing this assumption results in the curves shown by dashed lines. *Lower panel:* Hα volume emissivity for the same models. Our models are again shown with dotted lines, while solid and dashed lines correspond to Case B CLOUDY models with collisional excitation processes enabled and disabled respectively.

A.2 Additional tests of ionization balance

Firstly, we consider the effects of our use of the case B recombination assumption, by comparing the hydrogen neutral fractions $x_{\rm HI} \equiv n_{\rm HI}/n_{\rm H}$ that we obtain with corresponding predictions from CLOUDY, both using its default self-consistent treatment of recombination radiation, and overriding this behaviour to also assume the case B limit. The upper panel panel of Fig. A.1 indicates that this approximation results in a slight underestimate of the neutral fraction at column densities $N_{\rm H} < 10^{18}\,{\rm cm}^{-2}$.

The volume emissivities plotted in the bottom panel of Fig. A.1 appear to show that this optically thin gas contributes the majority of the Hα emission; however, this is an effect of the artificial constant-density assumption used in the slab models. For the

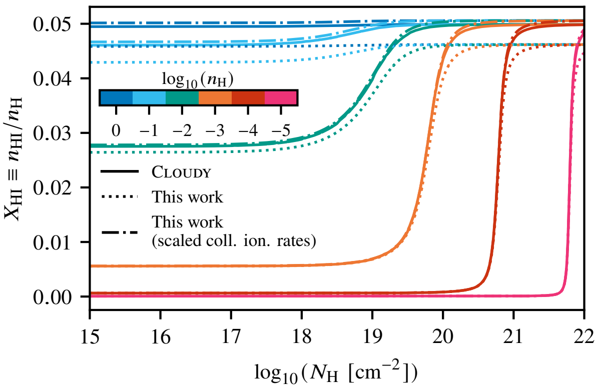

Figure A.2

H I neutral fraction as a function of total H column density for isothermal, H-only plane parallel models with $T_{\rm gas} = 20,000$ K. Solid and dotted curves correspond to output generated by CLOUDY and our code respectively, while the dot-dashed curves show our predictions if the collisional ionization rates we use are rescaled to match those in CLOUDY.

physically-motivated models we consider, the gas density profile is dictated by hydrostatic equilibrium such that n_H rapidly decreases with increasing radius and the H α emissivity peaks near the ionization front. This is located at $N_{\rm H} \gtrsim 10^{18}$ cm^{-2}, where the case B limit is applicable; accordingly, we do not expect the absence of self-consistent diffuse radiation in our code to significantly affect our results. Fig. A.1 also indicates emissivities below the CLOUDY values at high column and volume densities. This is likely caused by the exclusion of collisional excitation processes in our code, as disabling these processes in CLOUDY produces the dashed curves in Fig. A.1, with which our results are a better match. Since the emissivity of the predominantly neutral gas found at high $N_{\rm H}$ is at least two orders of magnitude lower than that of gas near the ionization front, we expect the net impact of this discrepancy on the total surface brightness to be minimal.

Finally, in Fig. A.2 we show results for a set of plane parallel models run at a higher gas temperature of 20,000 K. At this temperature, collisional ionization dominates throughout the slab, resulting in neutral fractions $\lesssim 0.05$. There is a systematic offset of ~10% in the neutral fractions predicted by our code and by CLOUDY, which is due to a difference of similar magnitude in the collisional ionization rate coefficients assumed by the two codes. Rescaling the CHIANTI rates we employ to match those reported by CLOUDY produces the dot-dashed curves in Fig. A.2, reducing the offset in the neutral fractions to below 1%.

Hence, excluding differences in the input data used, we find excellent agreement between our code and CLOUDY in the collisionally-ionized limit.

In summary, while the simplicity of the code we present here as compared to CLOUDY results in some differences in the predicted values of x_{HI} and $\varepsilon_{\mathrm{H}\alpha}$, these occur for regimes that are either absent in our physically-motivated models, or contribute little to the observed surface brightness. Therefore, our code's predictions appear to be robust in its intended domain of applicability.

Bibliography

Abraham R. G., van Dokkum P. G., 2014, PASP, 126, 55

Abruzzo M. W., Bryan G. L., Fielding D. B., 2021, preprint (arXiv:2101.10344)

Adams J. J., Uson J. M., Hill G. J., MacQueen P. J., 2011, ApJ, 728, 107

Adams E. A. K., Giovanelli R., Haynes M. P., 2013, ApJ, 768, 77

Aghanim N., et al., 2020, A&A, 641, A6

Aghanim N., et al., 2021, A&A, 652, C4

Alam S., et al., 2017, MNRAS, 470, 2617

Alam S., et al., 2021, Phys. Rev. D, 103, 083533

Allen S. W., Evrard A. E., Mantz A. B., 2011, ARA&A, 49, 409

Alpher R. A., Herman R., 1948, Nature, 162, 774

Alpher R. A., Bethe H., Gamow G., 1948, Phys. Rev., 73, 803

Altay G., Theuns T., 2013, MNRAS, 434, 748

Altay G., Theuns T., Schaye J., Crighton N. H. M., Vecchia C. D., 2011, ApJL, 737, L37

Altay G., Theuns T., Schaye J., Booth C. M., Dalla Vecchia C., 2013, MNRAS, 436, 2689

Arons J., Wingert D. W., 1972, ApJ, 177, 1

Aver E., Olive K. A., Skillman E. D., 2015, J. Cosmol. Astropart. Phys., 2015, 011

Babcock H. W., 1939, Lick Obs. Bull., 498, 41

Babul A., Rees M. J., 1992, MNRAS, 255, 346

Bacon D. J., Refregier A. R., Ellis R. S., 2000, MNRAS, 318, 625

Bacon R., et al., 2010, in Ground-Based and Airborne Instrumentation for Astronomy III. International Society for Optics and Photonics, p. 773508, doi:10.1117/12.856027

Bahcall N. A., 2015, PNAS, 112, 3173

Bahcall J. N., Salpeter E. E., 1965, ApJ, 142, 1677

Bahcall N. A., Ostriker J. P., Perlmutter S., Steinhardt P. J., 1999, Science, 284, 1481

Balashev S. A., Noterdaeme P., 2018, MNRASL, 478, L7

Balashev S. A., Klimenko V. V., Ivanchik A. V., Varshalovich D. A., Petitjean P., Noterdaeme P., 2014, MNRAS, 440, 225

Bardeen J. M., Bond J. R., Kaiser N., Szalay A. S., 1986, ApJ, 304, 15

Barkana R., Loeb A., 2001, Phys. Rep., 349, 125

Barkana R., Loeb A., 2007, Rep. Prog. Phys., 70, 627

Barnes L. A., Haehnelt M. G., 2014, MNRAS, 440, 2313

Barnes L. A., Haehnelt M. G., 2015, MNRAS, 454, 218

Becker G. D., Bolton J. S., Madau P., Pettini M., Ryan-Weber E. V., Venemans B. P., 2015, MNRAS, 447, 3402

Behnel S., Bradshaw R., Citro C., Dalcin L., Seljebotn D. S., Smith K., 2011, Comp. Sci. Eng., 13, 31

Benitez-Llambay A., 2015, Py-Sphviewer: Py-SPHViewer v1.0.0, Zenodo, doi:10.5281/zenodo.21703

Benitez-Llambay A., Frenk C., 2020, MNRAS, 498, 4887

Benítez-Llambay A., Navarro J. F., Abadi M. G., Gottlöber S., Yepes G., Yehuda Hoffman Steinmetz M., 2013, ApJL, 763, L41

Benítez-Llambay A., et al., 2017, MNRAS, 465, 3913

Benítez-Llambay A., Frenk C. S., Ludlow A. D., Navarro J. F., 2019, MNRAS, 488, 2387

Berg T. A. M., et al., 2019, MNRAS, 488, 4356

Bergeron J., 1986, A&A, 155, L8

Bergeron J., Boissé P., 1991, A&A, 243, 344

Bertschinger E., 1985, ApJS, 58, 39

Bird S., Vogelsberger M., Haehnelt M., Sijacki D., Genel S., Torrey P., Springel V., Hernquist L., 2014, MNRAS, 445, 2313

Bird S., Haehnelt M., Neeleman M., Genel S., Vogelsberger M., Hernquist L., 2015, MNRAS, 447, 1834

Bland-Hawthorn J., Freeman K. C., Quinn P. J., 1997, ApJ, 490, 143

Blitz L., Rosolowsky E., 2006, ApJ, 650, 933

Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J., 1984, Nature, 311, 517

Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J., 1985, Nature, 313, 72

Bolatto A. D., Wolfire M., Leroy A. K., 2013, ARA&A, 51, 207

Bond J. R., Szalay A. S., 1983, ApJ, 274, 443

Bond J. R., Efstathiou G., Silk J., 1980, Phys. Rev. Lett., 45, 1980

Bond J. R., Kofman L., Pogosyan D., 1996, Nature, 380, 603

Bondi H., Gold T., 1948, MNRAS, 108, 252

Booth C. M., Schaye J., 2009, MNRAS, 398, 53

Borgani S., et al., 2001, ApJ, 561, 13

Bose S., Hellwing W. A., Frenk C. S., Jenkins A., Lovell M. R., Helly J. C., Li B., 2016, MNRAS, 455, 318

Bose S., et al., 2017, MNRAS, 464, 4520

Bouwens R. J., et al., 2015, ApJ, 803, 34

Boyarsky A., Iakubovskyi D., Ruchayskiy O., 2012, Phys. Dark Univ., 1, 136

Bregman J. N., 2007, ARA&A, 45, 221

Bromm V., Larson R. B., 2004, ARA&A, 42, 79

Bryan G. L., et al., 2014, ApJS, 211, 19

Bull P., et al., 2016, Phys. Dark Univ., 12, 56

Bullock J. S., Kravtsov A. V., Weinberg D. H., 2000, ApJ, 539, 517

Burbidge E. M., Burbidge G. R., 1960, ApJ, 132, 30

Burkert A., 1995, ApJ, 447, L25

Cantalupo S., Porciani C., Lilly S. J., Miniati F., 2005, ApJ, 628, 61

Cantalupo S., Arrigoni-Battaia F., Prochaska J. X., Hennawi J. F., Madau P., 2014, Nature, 506, 63

Cen R., 1992, ApJS, 78, 341

Cen R., 2012, ApJ, 748, 121

Cen R., Ostriker J., 1992, ApJ, 393, 22

Cen R., Ostriker J. P., 1999, ApJ, 514, 1

Cen R., Miralda-Escudé J., Ostriker J. P., Rauch M., 1994, ApJ, 437, L9

Chabrier G., 2003, PASP, 115, 763

Charlton J. C., Churchill C. W., 2001, preprint (arXiv:astro-ph/0006002)

Chen H.-W., Helsby J. E., Gauthier J.-R., Shectman S. A., Thompson I. B., Tinker J. L., 2010, ApJ, 714, 1521

Churchill C. W., Charlton J. C., 1999, AJ, 118, 59

Ćirković M. M., Bland-Hawthorn J., Samurović S., 1999, MNRAS, 306, L15

Clowe D., Bradač M., Gonzalez A. H., Markevitch M., Randall S. W., Jones C., Zaritsky D., 2006, ApJ, 648, L109

Cole S., 1991, ApJ, 367, 45

Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, MNRAS, 319, 168

Collette A., 2013, Python and HDF5: Unlocking Scientific Data. O'Reilly Media, Inc., Sebastopol, CA

Cooke R. J., Fumagalli M., 2018, Nat. Astron., 2, 957

Cooke R., Pettini M., 2016, MNRAS, 455, 1512

Crain R. A., et al., 2015, MNRAS, 450, 1937

Crain R. A., et al., 2017, MNRAS, 464, 4204

Crighton N. H. M., Hennawi J. F., Simcoe R. A., Cooksey K. L., Murphy M. T., Fumagalli M., Prochaska J. X., Shanks T., 2015a, MNRAS, 446, 18

Crighton N. H. M., et al., 2015b, MNRAS, 452, 217

DESI Collaboration et al., 2016, preprint (arXiv:1611.00036)

Dalla Vecchia C., Schaye J., 2012, MNRAS, 426, 140

Davé 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

Davis M., Huchra J., Latham D. W., Tonry J., 1982, ApJ, 253, 423

Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371

Dayal P., Ferrara A., 2018, Phys. Rep., 780–782, 1

Del Popolo A., Le Delliou M., 2017, Galaxies, 5, 17

Dere K. P., 2007, A&A, 466, 771

Dere K. P., Landi E., Mason H. E., Fossi B. C. M., Young P. R., 1997, A&AS, 125, 149

Dicke R. H., Peebles P. J. E., Roll P. G., Wilkinson D. T., 1965, ApJ, 142, 414

Diemer B., 2018, ApJS, 239, 35

Diemer B., Kravtsov A. V., 2014, ApJ, 789, 1

Dolag K., Borgani S., Murante G., Springel V., 2009, MNRAS, 399, 497

Donahue M., Aldering G., Stocke J. T., 1995, APJL, 450, L45

Dove J. B., Shull J. M., 1994, ApJ, 423, 196

Duncan K., et al., 2014, MNRAS, 444, 2960

Durier F., Dalla Vecchia C., 2012, MNRAS, 419, 465

Dyson F. W., Eddington A. S., Davidson C., 1920, Philos. Trans. R. Soc, 220, 291

Efstathiou G., 1992, MNRAS, 256, 43P

Efstathiou G., Davis M., White S. D. M., Frenk C. S., 1985, ApJS, 57, 241

Efstathiou G., Kaiser N., Saunders W., Lawrence A., Rowan-Robinson M., Ellis R. S., Frenk C. S., 1990, MNRAS, 247, 10P

Efstathiou G., Bond J. R., White S. D. M., 1992, MNRAS, 258, 1P

Eggen O. J., Lynden-Bell D., Sandage A. R., 1962, ApJ, 136, 748

Einasto J., Kaasik A., Saar E., 1974a, Nature, 250, 309

Einasto J., Kaasik A., Saar E., 1974b, Nature, 250, 790

Einstein A., 1914, Ber. König. Preuß. Akad. Wiss., pp 1030–1085

Einstein A., 1915a, Ber. König. Preuß. Akad. Wiss., pp 778–786

Einstein A., 1915b, Ber. König. Preuß. Akad. Wiss., pp 831–839

Einstein A., 1917, Ber. König. Preuß. Akad. Wiss., pp 142–152

Einstein A., 1922, Z. Phys., 11, 326

Einstein A., 1931, Ber. König. Preuß. Akad. Wiss., 11, 235

Eisenstein D. J., et al., 2005, ApJ, 633, 560

Enßlin T. A., Pfrommer C., Springel V., Jubelgas M., 2007, A&A, 473, 41

Erkal D., Belokurov V., 2015, MNRAS, 454, 3542

Fardal M. A., Giroux M. L., Shull J. M., 1998, AJ, 115, 2206

Faucher-Giguère C.-A., 2020, MNRAS, 493, 1614

Faucher-Giguère C.-A., Kereš D., 2011, MNRASL, 412, L118

Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, ApJ, 703, 1416

Faucher-Giguère C.-A., Hopkins P. F., Kereš D., Muratov A. L., Quataert E., Murray N., 2015, MNRAS, 449, 987

Feldmann R., Spolyar D., 2015, MNRAS, 446, 1000

Feng J. L., 2010, ARA&A, 48, 495

Ferland G. J., Korista K. T., Verner D. A., Ferguson J. W., Kingdon J. B., Verner E. M., 1998, PASP, 110, 761

Ferland G. J., et al., 2017, Rev. Mex. Astron. Astrofis., 53, 385

Fernández V., Terlevich E., Díaz A. I., Terlevich R., Rosales-Ortega F. F., 2018, MNRAS, 478, 5301

Fernández V., Terlevich E., Díaz A. I., Terlevich R., 2019, MNRAS, 487, 3221

Fielding D., Quataert E., McCourt M., Thompson T. A., 2017, MNRAS, 466, 3810

Finkelstein S. L., et al., 2015, ApJ, 810, 71

Flores R. A., Primack J. R., 1994, ApJ, 427, L1

Friedmann A., 1922, Z. Phys., 10, 377

Fumagalli M., Krumholz M. R., Hunt L. K., 2010, ApJ, 722, 919

Fumagalli M., Prochaska J. X., Kasen D., Dekel A., Ceverino D., Primack J. R., 2011, MNRAS, 418, 1796

Fumagalli M., O'Meara J. M., Prochaska J. X., Rafelski M., Kanekar N., 2015, MNRAS, 446, 3178

Fumagalli M., O'Meara J. M., Prochaska J. X., 2016, MNRAS, 455, 4100

Fumagalli M., Haardt F., Theuns T., Morris S. L., Cantalupo S., Madau P., Fossati M., 2017a, MNRAS, 467, 4802

Fumagalli M., et al., 2017b, MNRAS, 471, 3686

Furlanetto S. R., Oh S. P., 2009, ApJ, 701, 94

Furlong M., et al., 2015, MNRAS, 450, 4486

Fynbo J. P. U., et al., 2010, MNRAS, 408, 2128

Fynbo J. P. U., et al., 2011, MNRAS, 413, 2481

Fynbo J. P. U., et al., 2018, MNRAS, 479, 2126

Gardner J. P., Katz N., Weinberg D. H., Hernquist L., 1997, ApJ, 486, 42

Gardner J. P., Katz N., Hernquist L., Weinberg D. H., 2001, ApJ, 559, 131

Garnett R., Ho S., Bird S., Schneider J., 2017, MNRAS, 472, 1850

Geller M. J., Huchra J. P., 1989, Science

Georgi H., Glashow S. L., 1974, Phys. Rev. Lett., 32, 438

Giovanelli R., et al., 2005, AJ, 130, 2598

Giroux M. L., Shapiro P. R., 1996, ApJS, 102, 191

González V., Labbé I., Bouwens R. J., Illingworth G., Franx M., Kriek M., 2011, ApJL, 735, L34

Gorski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759

Gould A., Weinberg D. H., 1996, ApJ, 468, 462

Grand R. J. J., et al., 2017, MNRAS, 467, 179

Grazian A., et al., 2015, A&A, 575, A96

Gronke M., Oh S. P., 2018, MNRASL, 480, L111

Gunn J. E., Gott III J. R., 1972, ApJ, 176, 1

Gunn J. E., Peterson B. A., 1965, ApJ, 142, 1633

Guo Q., et al., 2011, MNRAS, 413, 101

Gurvich A., Burkhart B., Bird S., 2017, ApJ, 835, 175

Guth A. H., 1981, Phys. Rev. D, 23, 347

Haardt F., Madau P., 1996, ApJ, 461, 20

Haardt F., Madau P., 2001, in Neuman D., Trân Thanh Vân J., eds, Clusters of Galaxies and the High Redshift Universe Observed in X-rays. Savoie, France (arXiv:astro-ph/0106018)

Haardt F., Madau P., 2012, ApJ, 746, 125

Haehnelt M. G., Steinmetz M., Rauch M., 1998, ApJ, 495, 647

Harris C. R., et al., 2020, Nature, 585, 357

Hassan S., Finlator K., Davé R., Churchill C. W., Prochaska J. X., 2020, MNRAS, 492, 2835

Hayes M., Melinder J., Östlin G., Scarlata C., Lehnert M. D., Mannerström-Jansson G., 2016, ApJ, 828, 49

Heald G., et al., 2011, A&A, 526, A118

Hellwing W. A., Frenk C. S., Cautun M., Bose S., Helly J., Jenkins A., Sawala T., Cytowski M., 2016, MNRAS, 457, 3492

Henriques B. M. B., Yates R. M., Fu J., Guo Q., Kauffmann G., Srisawat C., Thomas P. A., White S. D. M., 2020, MNRAS, 491, 5795

Hernquist L., Katz N., Weinberg D. H., Miralda-Escudé J., 1996, ApJL, 457, L51

Hezaveh Y. D., et al., 2016, ApJ, 823, 37

Ho M.-F., Bird S., Garnett R., 2020, MNRAS, 496, 5436

Ho M.-F., Bird S., Garnett R., 2021, MNRAS, 507, 704

Hong S., Katz N., Davé R., Fardal M., Kereš D., Oppenheimer B. D., 2010, preprint (arXiv:1008.4242)

Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, MNRAS, 445, 581

Hoyle F., 1948, MNRAS, 108, 372

Hoyle F., Tayler R. J., 1964, Nature, 203, 1108

Hu C.-Y., Naab T., Walch S., Glover S. C. O., Clark P. C., 2016, MNRAS, 458, 3528

Hubble E. P., 1925, Obser., 48, 139

Hubble E., 1929, PNAS, 15, 168

Hui L., Gnedin N. Y., 1997, MNRAS, 292, 27

Hummels C. B., et al., 2019, ApJ, 882, 156

Hunter J. D., 2007, Comp. Sci. Eng., 9, 90

Iakubovskyi D., 2016, Adv. Astron. Spa. Phys., 6, 3

Ilbert O., et al., 2013, A&A, 556, A55

Izotov Y. I., Thuan T. X., Stasińska G., 2007, ApJ, 662, 15

Izotov Y. I., Thuan T. X., Guseva N. G., 2014, MNRAS, 445, 778

Jarvis M., et al., 2018, in Proceedings of MeerKAT Science: On the Pathway to the SKA — PoS(MeerKAT2016). SISSA Medialab, Cape Town, p. 006, doi:10.22323/1.277.0006

Jenkins A., 2010, MNRAS, 403, 1859

Jenkins A., 2013, MNRAS, 434, 2094

Jenkins A., Booth S., 2013, preprint (arXiv:1306.5771)

Jorgenson R. A., Wolfe A. M., 2014, ApJ, 785, 16

Jorgenson R. A., Wolfe A. M., Prochaska J. X., Carswell R. F., 2009, ApJ, 704, 247

Jubelgas M., Springel V., Enßlin T., Pfrommer C., 2008, A&A, 481, 33

Kahn F. D., Woltjer L., 1959, ApJ, 130, 705

Kanekar N., et al., 2018, ApJL, 856, L23

Kanekar N., Prochaska J. X., Neeleman M., Christensen L., Møller P., Zwaan M. A., Fynbo J. P. U., Dessauges-Zavadsky M., 2020, ApJL, 901, L5

Kapteyn J. C., 1922, Contr. Mount Wilson Obser., 230, 1

Katz N., 1992, ApJ, 391, 502

Katz N., Gunn J. E., 1991, ApJ, 377, 365

Katz N., Weinberg D. H., Hernquist L., 1996, ApJS, 105, 19

Kay S. T., Thomas P. A., Theuns T., 2003, MNRAS, 343, 608

Kennicutt J. R. C., 1998, ApJ, 498, 541

Kereš D., Katz N., Weinberg D. H., Davé R., 2005, MNRAS, 363, 2

Khaire V., Srianand R., 2015, MNRASL, 451, L30

Kilbinger M., 2015, Rep. Prog. Phys., 78, 086901

Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, ApJ, 522, 82

Kollmeier J. A., et al., 2014, ApJL, 789, L32

Komatsu E., et al., 2011, ApJS, 192, 18

Kravtsov A. V., Klypin A. A., Khokhlov A. M., 1997, ApJS, 111, 73

Krogager J.-K., Møller P., Fynbo J. P. U., Noterdaeme P., 2017, MNRAS, 469, 2959

Krogager J.-K., Fynbo J. P. U., Møller P., Noterdaeme P., Heintz K. E., Pettini M., 2019, MNRAS, 486, 4377

Krumholz M. R., 2013, MNRAS, 436, 2747

Krumholz M. R., McKee C. F., Tumlinson J., 2009a, ApJ, 693, 216

Krumholz M. R., McKee C. F., Tumlinson J., 2009b, ApJ, 699, 850

Krumholz M. R., Ellison S. L., Prochaska J. X., Tumlinson J., 2009c, ApJ, 701, L12

Kuhlen M., Krumholz M. R., Madau P., Smith B. D., Wise J., 2012, ApJ, 749, 36

Lacey C., Silk J., 1991, ApJ, 381, 14

Lacey C. G., et al., 2016, MNRAS, 462, 3854

Lam S. K., et al., 2021, Numba/Numba: Version 0.54.0, Zenodo, doi:10.5281/zenodo.5524874

Ledoux C., Petitjean P., Srianand R., 2003, MNRAS, 346, 209

Lemaître G., 1927, Ann. Soc. Sci. Bruxelles, 47, 49. [Translated: Lemaître G., 1931, MNRAS, 91, 483]

Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473

Linde A. D., 1982, Phys. Lett. B, 108, 389

Lokhorst D., Abraham R., van Dokkum P., Wijers N., Schaye J., 2019, ApJ, 877, 4

Lokhorst D. M., Abraham R. G., van Dokkum P., Chen S., 2020, in Ground-Based and Airborne Telescopes VIII. SPIE, pp 492–505, doi:10.1117/12.2561162

Lopez S., et al., 2018, Nature, 554, 493

Lopez S., et al., 2020, MNRAS, 491, 4442

Lovell M. R., et al., 2016, MNRAS, 461, 60

Lovell M. R., et al., 2017, MNRAS, 468, 4285

Lubimov V. A., Novikov E. G., Nozik V. Z., Tretyakov E. F., Kosik V. S., 1980, Phys. Lett. B, 94, 266

Ludlow A. D., Bose S., Angulo R. E., Wang L., Hellwing W. A., Navarro J. F., Cole S., Frenk C. S., 2016, MNRAS, 460, 1214

Lynds R., 1971, ApJ, 164, L73

Mackenzie R., et al., 2019, MNRAS, 487, 5070

Madau P., Haardt F., 2015, ApJL, 813, L8

Maddox S. J., Efstathiou G., Sutherland W. J., Loveday J., 1990, MNRAS, 242, 43

Madsen G. J., Reynolds R. J., Haffner L. M., Tufte S. L., Maloney P. R., 2001, ApJ, 560, L135

Maloney P., 1993, ApJ, 414, 41

Martin D. C., Matuszewski M., Morrissey P., Neill J. D., Moore A., Cantalupo S., Prochaska J. X., Chang D., 2015, Nature, 524, 192

Mayall N. U., 1951, in Publications of Michigan Observatory, Vol. 10, The Structure of the Galaxy. University of Michigan Press, Ann Arbor, MI, p. 19

McGreer I. D., Mesinger A., D'Odorico V., 2015, MNRAS, 447, 499

McQuinn M., 2016, ARA&A, 54, 313

McQuinn M., Upton Sanderbeck P. R., 2016, MNRAS, 456, 47

Meiksin A. A., 2009, Rev. Mod. Phys., 81, 1405

Meiksin A., Madau P., 1993, ApJ, 412, 34

Milne E. A., 1935, Relativity, Gravitation and World Structure. Oxford University Press, Oxford

Miralda-Escude J., Ostriker J. P., 1990, ApJ, 350, 1

Miralda-Escudé J., Rees M. J., 1994, MNRAS, 266, 343

Mitchell P. D., Schaye J., Bower R. G., Crain R. A., 2020, MNRAS, 494, 3971

Møller P., Fynbo J. P. U., Fall S. M., 2004, A&A, 422, L33

Monaghan J. J., 1992, ARA&A, 30, 543

Moore B., 1994, Nature, 370, 629

Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, ApJL, 524, L19

Morrissey P., et al., 2018, ApJ, 864, 93

Munshi F., Brooks A. M., Christensen C., Applebaum E., Holley-Bockelmann K., Quinn T. R., Wadsley J., 2019, ApJ, 874, 40

Muzzin A., et al., 2013, ApJ, 777, 18

Naab T., Ostriker J. P., 2017, ARA&A, 55, 59

Nagamine K., Springel V., Hernquist L., 2004, MNRAS, 348, 421

Nagamine K., Wolfe A. M., Hernquist L., Springel V., 2007, ApJ, 660, 945

Navarro J. F., Steinmetz M., 1997, ApJ, 478, 13

Navarro J. F., White S. D. M., 1993, MNRAS, 265, 271

Navarro J. F., White S. D. M., 1994, MNRAS, 267, 401

Navarro J. F., Frenk C. S., White S. D. M., 1995, MNRAS, 275, 56

Navarro J. F., Eke V. R., Frenk C. S., 1996a, MNRAS, 283, L72

Navarro J. F., Frenk C. S., White S. D. M., 1996b, ApJ, 462, 563

Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493

Neeleman M., Kanekar N., Prochaska J. X., Rafelski M., Carilli C. L., Wolfe A. M., 2017, Science, 355, 1285

Neeleman M., Kanekar N., Prochaska J. X., Christensen L., Dessauges-Zavadsky M., Fynbo J. P. U., Møller P., Zwaan M. A., 2018, ApJL, 856, L12

Neeleman M., Kanekar N., Prochaska J. X., Rafelski M. A., Carilli C. L., 2019, ApJL, 870, L19

Nelson D., et al., 2019, MNRAS, 490, 3234

Nelson D., et al., 2021, preprint (arXiv:1812.05609)

Nicastro F., et al., 2018, Nature, 558, 406

Noterdaeme P., Ledoux C., Petitjean P., Srianand R., 2008, A&A, 481, 327

Noterdaeme P., et al., 2012, A&A, 547, L1

O'Meara J. M., Prochaska J. X., Worseck G., Chen H.-W., Madau P., 2013, ApJ, 765, 137

O'Raifeartaigh C., O'Keeffe M., Nahm W., Mitton S., 2017, EPJ H, 42, 431

Okamoto T., Eke V. R., Frenk C. S., Jenkins A., 2005, MNRAS, 363, 1299

Okamoto T., Gao L., Theuns T., 2008, MNRAS, 390, 920

Oort J. H., 1932, Bull. Astronom. Inst. Netherlands, 6, 249

Oppenheimer B. D., Davé R., 2008, MNRAS, 387, 577

Osterbrock D. E., Ferland G. J., 2006, Astrophysics of Gaseous Nebulae and Active Galactic Nuclei, second edn. University Science Books, Sausalito, CA

Ostriker J. P., Peebles P. J. E., 1973, ApJ, 186, 467

Ostriker J. P., Peebles P. J. E., Yahil A., 1974, ApJL, 193, L1

Padmanabhan H., Choudhury T. R., Refregier A., 2016, MNRAS, 458, 781

Padmanabhan H., Refregier A., Amara A., 2017, MNRAS, 469, 2323

Pearce F. R., et al., 1999, ApJ, 521, L99

Peebles P. J. E., 1982a, ApJ, 258, 415

Peebles P. J. E., 1982b, ApJ, 263, L1

Peebles P. J. E., 1993, Principles of Physical Cosmology. Princeton University Press

Peeples M. S., et al., 2019, ApJ, 873, 129

Peimbert A., Peimbert M., Luridiana V., 2016, Rev. Mex. Astron. Astrofis., 52, 419

Pen U.-L., 1998, ApJS, 115, 19

Penzias A. A., Wilson R. W., 1965, ApJ, 142, 419

Perez F., Granger B. E., 2007, Comp. Sci. Eng., 9, 21

Pérez-Ràfols I., et al., 2018a, MNRAS, 473, 3019

Pérez-Ràfols I., Miralda-Escudé J., Arinyo-i-Prats A., Font-Ribera A., Mas-Ribas L., 2018b, MNRAS, 480, 4702

Perivolaropoulos L., Skara F., 2021, preprint (arXiv:2105.05208)

Perlmutter S., et al., 1999, ApJ, 517, 565

Péroux C., Bouché N., Kulkarni V. P., York D. G., Vladilo G., 2012, MNRAS, 419, 3060

Pieri M. M., et al., 2016, in Proceedings of the Annual Meeting of the French Society of Astronomy and Astrophysics. Lyon, France, pp 259–266

Pillepich A., et al., 2018, MNRAS, 473, 4077

Pillepich A., et al., 2019, MNRAS, 490, 3196

Pitrou C., Coc A., Uzan J.-P., Vangioni E., 2018, Phys. Rep., 754, 1

Planck Collaboration et al., 2014, A&A, 571, A16

Pontzen A., Governato F., 2012, MNRAS, 421, 3464

Pontzen A., et al., 2008, MNRAS, 390, 1349

Popping A., Meyer M., Staveley-Smith L., Obreschkow D., Jozsa G., Pisano D. J., 2015, in Proceedings of Advancing Astrophysics with the Square Kilometre Array —

PoS(AASKA14). SISSA Medialab, Giardini Naxos, p. 132, doi:10.22323/1.215.0132

Porter R. L., Ferland G. J., MacAdam K. B., Storey P. J., 2009, MNRASL, 393, L36

Porter R. L., Ferland G. J., Storey P. J., Detisch M. J., 2012, MNRASL, 425, L28

Porter R. L., Ferland G. J., Storey P. J., Detisch M. J., 2013, MNRASL, 433, L89

Power C., et al., 2015, in Proceedings of Advancing Astrophysics with the Square Kilometre Array — PoS(AASKA14). SISSA Medialab, Giardini Naxos, p. 133, doi:10.22323/1.215.0133

Press W. H., Schechter P., 1974, ApJ, 187, 425

Primack J. R., 2009, Proc. American Phys. Soc., 1192, 101

Prochaska J. X., Wolfe A. M., 1997, ApJ, 487, 73

Prochaska J. X., Wolfe A. M., 1998, ApJ, 507, 113

Prochaska J. X., Wolfe A. M., 2001, ApJ, 560, L33

Prochaska J. X., Wolfe A. M., 2009, ApJ, 696, 1543

Prochaska J. X., et al., 2017, ApJ, 837, 169

Puchwein E., Haardt F., Haehnelt M. G., Madau P., 2019, MNRAS, 485, 47

Quinn T., Katz N., Efstathiou G., 1996, MNRAS, 278, L49

Rafelski M., Wolfe A. M., Chen H.-W., 2011, ApJ, 736, 48

Rahmati A., Schaye J., 2014, MNRAS, 438, 529

Rahmati A., Pawlik A. H., Raičevič M., Schaye J., 2013a, MNRAS, 430, 2427

Rahmati A., Schaye J., Pawlik A. H., Raičevič M., 2013b, MNRAS, 431, 2261

Rahmati A., Schaye J., Bower R. G., Crain R. A., Furlong M., Schaller M., Theuns T., 2015, MNRAS, 452, 2034

Rauch M., 1998, ARA&A, 36, 267

Razoumov A. O., Norman M. L., Prochaska J. X., Sommer-Larsen J., Wolfe A. M., Yang Y.-J., 2008, ApJ, 683, 149

Reed D. S., Bower R., Frenk C. S., Jenkins A., Theuns T., 2007, MNRAS, 374, 2

Rees M. J., 1986, MNRAS, 218, 25P

Rees M. J., Ostriker J. P., 1977, MNRAS, 179, 541

Richard J., et al., 2019, preprint (arXiv:1906.01657)

Riess A. G., 2020, Nat. Rev. Phys., 2, 10

Riess A. G., et al., 1998, AJ, 116, 1009

Roberts M. S., Rots A. H., 1973, A&A, 26, 483

Robertson H. P., 1935, ApJ, 82, 284

Robertson H. P., 1936a, ApJ, 83, 187

Robertson H. P., 1936b, ApJ, 83, 257

Rogstad D. H., Shostak G. S., 1972, ApJ, 176, 315

Rollinde E., Theuns T., Schaye J., Pâris I., Petitjean P., 2013, MNRAS, 428, 540

Rosas-Guevara Y. M., et al., 2015, MNRAS, 454, 1038

Rubin V. C., Ford Jr. W. K., 1970, ApJ, 159, 379

Rubin V. C., Ford Jr. W. K., Thonnard N., 1980, ApJ, 238, 471

Rudie G. C., Steidel C. C., Shapley A. E., Pettini M., 2013, ApJ, 769, 146

Rybicki G. B., Lightman A. P., 1991, Radiative Processes in Astrophysics. John Wiley & Sons

Ryle M., Scheuer P. a. G., 1955, Proc. R. Soc., 230, 448

Sachs R. K., Wolfe A. M., 1967, ApJ, 147, 73

Sargent W. L. W., Young P. J., Boksenberg A., Tytler D., 1980, ApJS, 42, 41

Sargent W. L. W., Steidel C. C., Boksenberg A., 1989, ApJS, 69, 703

Sawala T., et al., 2016, MNRAS, 457, 1931

Schaller M., Gonnet P., Draper P. W., Chalk A. B. G., Bower R. G., Willis J., Hausammann L., 2018, Astrophysics Source Code Library, p. ascl:1805.020

Schaye J., 2001, ApJ, 562, L95

Schaye J., 2004, ApJ, 609, 667

Schaye J., Dalla Vecchia C., 2008, MNRAS, 383, 1210

Schaye J., et al., 2010, MNRAS, 402, 1536

Schaye J., et al., 2015, MNRAS, 446, 521

Shapiro P. R., Giroux M. L., 1987, ApJ, 321, L107

Shapley H., Curtis H. D., 1921, Bull. Natl. Res. Council, 2, 171

Sharma M., Theuns T., 2020, MNRAS, 492, 2418

Shull J. M., Moloney J., Danforth C. W., Tilton E. M., 2015, ApJ, 811, 3

Slipher V. M., 1915, Pop. Astron., 23, 21

Slipher V. M., 1917, Proc. American Phil. Soc., 56, 403

Smoot G. F., et al., 1992, ApJ, 396, L1

Sommer-Larsen J., Gelato S., Vedel H., 1999, ApJ, 519, 501

Spergel D. N., Steinhardt P. J., 2000, PRL, 84, 3760

Spergel D. N., et al., 2003, ApJS, 148, 175

Spitzer L., 1978, Physical Processes in the Interstellar Medium. Wiley, New York, doi:10.1002/9783527617722

Springel V., 2005, MNRAS, 364, 1105

Springel V., 2010a, ARA&A, 48, 391

Springel V., 2010b, MNRAS, 401, 791

Springel V., Hernquist L., 2003, MNRAS, 339, 289

Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726

Springel V., et al., 2005, Nature, 435, 629

Steidel C. C., Sargent W. L. W., 1992, ApJS, 80, 1

Sternberg A., McKee C. F., Wolfire M. G., 2002, ApJS, 143, 419

Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, MNRAS, 373, 1074

Suresh J., Bird S., Vogelsberger M., Genel S., Torrey P., Sijacki D., Springel V., Hernquist L., 2015, MNRAS, 448, 895

Suzuki N., et al., 2012, ApJ, 746, 85

Sykes C., Fumagalli M., Cooke R., Theuns T., Benítez-Llambay A., 2019, MNRAS, 487, 609

Tescari E., Viel M., Tornatore L., Borgani S., 2009, MNRAS, 397, 411

Teyssier R., 2002, A&A, 385, 337

Thacker R. J., Couchman H. M. P., 2001, ApJ, 555, L17

Theuns T., 2021, MNRAS, 500, 2741

Theuns T., Leonard A., Efstathiou G., Pearce F. R., Thomas P. A., 1998, MNRAS, 301, 478

Thompson R., Davé R., Nagamine K., 2015, MNRAS, 452, 3030

Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, ApJ, 688, 709

Tomczak A. R., et al., 2014, ApJ, 783, 85

Tripp T. M., Sembach K. R., Bowen D. V., Savage B. D., Jenkins E. B., Lehner N., Richter P., 2008, ApJS, 177, 39

Tulin S., Yu H.-B., 2018, Phys. Rep., 730, 1

Tumlinson J., et al., 2011, Science, 334, 948

Tumlinson J., Peeples M. S., Werk J. K., 2017, ARA&A, 55, 389

Tytler D., 1982, Nature, 298, 427

Tytler D., 1987, ApJ, 321, 49

Tytler D., et al., 2004, ApJ, 617, 1

Valerdi M., Peimbert A., Peimbert M., Sixtos A., 2019, ApJ, 876, 98

Van Rossum G., Drake Jr F. L., 1995, Python Tutorial. Department of Computer Science [CS] Vol. 620, Centrum voor Wiskunde en Informatica Amsterdam

Van Waerbeke L., et al., 2000, A&A, 358, 30

Vegetti S., Koopmans L. V. E., Bolton A., Treu T., Gavazzi R., 2010, MNRAS, 408, 1969

Veilleux S., Cecil G., Bland-Hawthorn J., 2005, ARA&A, 43, 769

Viel M., Haehnelt M. G., 2006, MNRAS, 365, 231

Viel M., Haehnelt M. G., Bolton J. S., Kim T.-S., Puchwein E., Nasir F., Wakker B. P., 2017, MNRASL, 467, L86

Virtanen P., et al., 2020, Nat. Meth., 17, 261

Vogel S. N., Weymann R., Rauch M., Hamilton T., 1995, ApJ, 441, 162

Wadepuhl M., Springel V., 2011, MNRAS, 410, 1975

Wagoner R. V., 1967, ApJ, 149, 465

Walker A. G., 1937, Proc. London Math. Soc., 42, 90

Warren S. J., Møller P., Fall S. M., Jakobsen P., 2001, MNRAS, 326, 759

Weinberger R., et al., 2017, MNRAS, 465, 3291

Werk J. K., et al., 2016, ApJ, 833, 54

Weymann R. J., Vogel S. N., Veilleux S., Epps H. W., 2001, ApJ, 561, 559

White S. D. M., 1994, preprint (arXiv:astro-ph/9410043)

White S. D. M., Frenk C. S., 1991, ApJ, 379, 52

White S. D. M., Rees M. J., 1978, MNRAS, 183, 341

White S. D. M., Frenk C. S., Davis M., 1983, ApJL, 274, L1

White S. D. M., Navarro J. F., Evrard A. E., Frenk C. S., 1993, Nature, 366, 429

Wiersma R. P. C., Schaye J., Smith B. D., 2009a, MNRAS, 393, 99

Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009b, MNRAS, 399, 574

Wisotzki L., et al., 2018, Nature, 562, 229

Wolfe A. M., Turnshek D. A., Smith H. E., Cohen R. D., 1986, ApJS, 61, 249

Wolfe A. M., Gawiser E., Prochaska J. X., 2005, ARA&A, 43, 861

Wootten A., Thompson A. R., 2009, Proc. IEEE, 97, 1463

Wright E. L., et al., 1992, ApJL, 396, L13

Xu H., Wise J. H., Norman M. L., Ahn K., O'Shea B. W., 2016, ApJ, 833, 84

Zafar T., Péroux C., Popping A., Milliard B., Deharveng J.-M., Frank S., 2013, A&A, 556, A141

Zanna G. D., Dere K. P., Young P. R., Landi E., Mason H. E., 2015, A&A, 582, A56

Zwaan M. A., Prochaska J. X., 2006, ApJ, 643, 675

Zwicky F., 1929, PNAS, 15, 773

Zwicky F., 1933, Helvetica Phys. Acta, 6, 110. [Translated: Heinz A., 2017, arXiv:1711.01693 [astro-ph]]

Zwicky F., 1937, ApJ, 86, 217

de Blok W. J. G., et al., 2018, in Proceedings of MeerKAT Science: On the Pathway to the SKA — PoS(MeerKAT2016). SISSA Medialab, Cape Town, p. 007, doi:10.22323/1.277.0007

des Bourboux H. d. M., et al., 2020, ApJ, 901, 153

van de Voort F., Schaye J., Altay G., Theuns T., 2012, MNRAS, 421, 2809

van de Voort F., Springel V., Mandelker N., van den Bosch F. C., Pakmor R., 2019, MNRASL, 482, L85