Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **492,** 2151–2160 (2020)
Advance Access publication 2019 December 20

doi:10.1093/mnras/stz3573

# Determining the primordial helium abundance and UV background using fluorescent emission in star-free dark matter haloes

Calvin Sykes [1,2]★ Michele Fumagalli [1,2,3] Ryan Cooke [2] and Tom Theuns [1]

[1] *Institute for Computational Cosmology, Durham University, Durham DH1 3LE, UK*
[2] *Centre for Extragalactic Astronomy, Durham University, Durham DH1 3LE, UK*
[3] *Dipartimento di Fisica G. Occhialini, Università degli Studi di Milano Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy*

Accepted 2019 December 13. Received 2019 December 13; in original form 2019 November 15

## ABSTRACT

Observational measures of the primordial helium mass fraction, $Y_{\rm P}$, are of interest for cosmology and fundamental particle physics. Current measures obtained from H II regions agree with the Standard Model prediction to approximately 1 per cent precision, although these determinations may be affected by systematic uncertainties. This possibility can only be tested by independently measuring the helium abundance in new ways. Here, we propose a novel method to obtain a measurement of $Y_{\rm P}$ using hydrogen and helium recombination line emission from REionization-Limited HI Clouds (RELHICs): pristine, gas-rich but star-free low-mass dark matter haloes whose existence is predicted by hydrodynamical simulations. Although expected to be uncommon and intrinsically faint in emission, the primordial composition and simple physical properties of these objects make them an ideal laboratory to determine $Y_{\rm P}$. We present radiative transfer simulations to demonstrate the effectiveness of this approach, finding that a comparison of the emission in H and He lines, either via their volumetric emissivities, or integrated properties such as the surface brightness and total flux, may be used to infer $Y_{\rm P}$. Furthermore, we show that RELHICs can be used to provide an entirely novel constraint on the spectral slope of the ultraviolet background, and discuss the possibility of measuring this slope and the primordial helium abundance simultaneously.

**Key words:** radiative transfer – galaxies: dwarf – primordial nucleosynthesis.

## 1 INTRODUCTION

Almost all helium atoms in the Universe were synthesized in the first few minutes after the Big Bang, during the period of Big Bang nucleosynthesis (BBN; Alpher, Bethe & Gamow 1948; Hoyle & Tayler 1964). The primordial helium mass fraction $Y_{\rm P}$, or equivalently the abundance by number $y_{\rm P}$,[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 $Y_{\rm P} = 0.24672 \pm 0.00017$ ($y_{\rm P} = 0.08188 \pm 0.00008$) is reliable (Pitrou et al. 2018).

Comparing observational measures of $Y_{\rm P}$ with the BBN prediction allows the presence of any new physics beyond the Standard Model to be investigated. To date, the leading 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, Thuan & Guseva 2014; Aver, Olive & Skillman 2015; Peimbert, Peimbert & Luridiana 2016; Fernández et al. 2018, 2019; Valerdi et al. 2019). These studies select H II regions in metal-poor galaxies ($Z/Z_\odot \lesssim 0.1$), to minimize 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, Thuan & Stasińska 2007; Porter et al. 2009), the characterization of which becomes increasingly important as statistical errors on the

★ E-mail: calvin.v.sykes@durham.ac.uk

[1] These two quantities are related by $Y_{\rm P} = 4y_{\rm P}/(1 + 4y_{\rm P})$. We note that $Y_{\rm P}$ is defined as $Y_{\rm P} \equiv 4\,n(^4{\rm He})/n_{\rm b}$, where $n_{\rm b}$ is the baryon density. $Y_{\rm P}$ 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 $y_{\rm P} \equiv n_{\rm He}/n_{\rm H}$ in this paper.

measurements improve. Consequently, it is beneficial to consider independent techniques for determining $Y_P$. 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 $Y_P = 0.250^{+0.033}_{-0.025}$ (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 $Y_P = 0.248 \pm 0.001$.

In this paper, we discuss a novel method for determining $Y_P$. 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 $10^{10}$ 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 $\sim 10^4$ K, inhibiting star formation in haloes with potential wells too shallow to confine the heated gas (Miralda-Escudé & Rees 1994; Okamoto, Gao & Theuns 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 a previous paper (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$\alpha$. 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 the intrinsic 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 $y_P$, albeit one with substantial observational challenges given current instrumentation, as we will show. In this paper, 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.

This paper is organized as follows. In Section 2, we provide a brief description of our numerical method and describe how we have extended the calculations in S19 to additionally predict surface brightnesses for He I and He II emission lines. We then present our results in Section 3, considering constraints on the primordial helium abundance $y_P$ and on the UVB slope in turn (Sections 3.1 and 3.2), and then combined constraints on both parameters (Section 3.3). We conclude by discussing our results and their implications in Section 4. Throughout, we assume a set of cosmological parameters ($H_0 = 67.3$ 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 XVI et al. 2014).

## 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.[2] Similar in functionality to photoionization codes such as CLOUDY (Ferland et al. 2017), our code includes two important additional features. First, it applies the condition of hydrostatic equilibrium to determine the gas density profile, using the background gravitational of a Navarro–Frenk–White (NFW; Navarro, Frenk & White 1996) 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. Consequently, we perform our calculations in (projected) spherical coordinates.

Our calculations take the following form, which is similar to that described by Sternberg, McKee & Wolfire (2002). A dark matter potential is defined by choosing a virial mass $\rm M_{200}$ 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 $\rm R_{200}$, 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 $\rm M_{200}$. A total baryonic gas mass $M_g$ is then assigned to the halo using the analytic model employed by BL17; the gas is split into $N_r = 1000$ radial cells and initialized 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 $n_{\rm He} = y_P n_{\rm H}$.

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, $\bar{n}_{\rm H} \simeq 10^{-6.7}$cm$^{-3}$ at $z = 0$. From the resulting gas density profile, we next determine the intensity of the radiation

[2]The code is made available at https://github.com/calvin-sykes/spherical_cloudy

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.[3] By equating the relevant processes for each atomic species, we enforce ionization equilibrium and thus determine the fractional ionization $X_{\mathrm{A}^{\mathrm{i}+}} \equiv n_{\mathrm{A}^{\mathrm{i}+}}/n_{\mathrm{A}}$ 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 ($n_{\mathrm{H}} > 10^{-4.8}\,\mathrm{cm}^{-3}$), where this threshold is set by the condition that the time-scale 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 time-scale 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 per cent between successive iterations. With a converged ionization structure found, we then compute the volume emissivity $\varepsilon_\nu$ for an emission line with frequency $\nu$ as

$$\varepsilon_\nu(r) = h\nu\, n_{\mathrm{ion}}(r)\, n_{\mathrm{e}}(r)\, \alpha_{\mathrm{eff}}(T(r)), \qquad [\varepsilon] = \mathrm{erg\,s^{-1}\,cm^{-3}} \tag{1}$$

where $h$ is Planck's constant, $n_{\mathrm{ion}}$ is the density of H II, He II, or He III as appropriate, $n_{\mathrm{e}}$ is the electron density, and $\alpha_{\mathrm{eff}}$ is a temperature-dependent coefficient expressing the rate per unit ion and electron densities at which the relevant transitions occur. For H$\alpha$ and the analogous He II 4686 Å line, we obtain $\alpha_{\mathrm{eff}}$ 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 $n_{\mathrm{e}}$ and $T$. Finally, we calculate the projected surface brightness $\Sigma_\nu$ as the integral of $\varepsilon_\nu(r)$ along lines of sight corresponding to an impact parameter $b$:

$$\Sigma_\nu(b) = \frac{1}{2\pi}\int_b^{\mathrm{R_{200}}} \frac{r}{\sqrt{r^2-b^2}}\varepsilon_\nu(r)\,\mathrm{d}r$$

$$[\Sigma] = \mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}} \tag{2}$$

## 3 RESULTS

We first discuss the qualitative properties of the emission lines using a fiducial model, which takes the primordial helium abundance to be $y_{\mathrm{P}} = 0.083$, corresponding to a primordial mass fraction $Y_{\mathrm{P}} = 0.249$.

[3]We further justify this assumption in appendix B of S19.

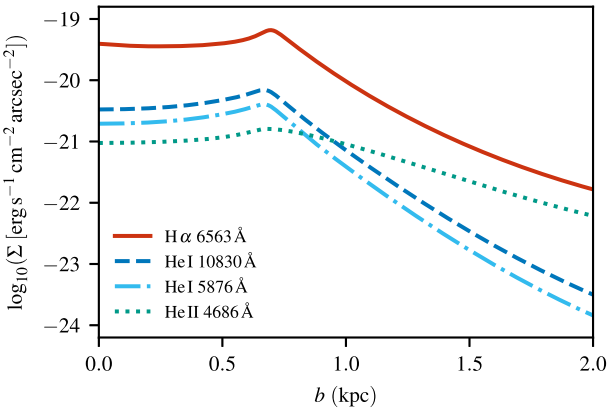


**Figure 1.** Surface brightness $\Sigma$ as a function of impact parameter $b$, for the 4686 Å He II line and two He I lines, in addition to the hydrogen H$\alpha$ line.

In Fig. 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, section 2.3). In addition to the hydrogen H$\alpha$ line, we plot the surface brightness of the two brightest He I lines (at 10830 and 5876 Å), and the He II 4686Å line. For this halo, the H$\alpha$ surface brightness reaches a peak intensity of $\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}} = 6.58 \times 10^{-20}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$. The helium line surface brightnesses are significantly dimmer, reaching maximum values of $4.06 \times 10^{-21}$ and $1.61 \times 10^{-21}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}}$ for the He I and He II lines, respectively.

$\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}}$ 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$\alpha$ 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.[4] We find that the peak helium surface brightnesses occur at a similar radial position to $\Sigma_{\mathrm{H}\alpha,\,\mathrm{max}}$, despite the respective ionization fronts being located at different radii. This occurs because the requirement of hydrostatic equilibrium produces gas densities that decrease rapidly with radius, such that the helium emissivities are affected more strongly by the falling electron density than by the helium ion densities.

From equation (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 $y_{\mathrm{P}}$ from this ratio involves some additional considerations. First, the densities in equation (1) refer to single ionization stages, whereas $y_{\mathrm{P}}$ is set by the overall atomic abundances. Secondly, the rate coefficients $\alpha_{\mathrm{eff}}$ are temperature dependent. Since the temperature and ionization structure of the gas (which influence $\alpha_{\mathrm{eff}}$ and the fractional ionizations, respectively) depend on both $y_{\mathrm{P}}$ and each other, we proceed by performing the iterative procedure outlined in Section 2 for a number of different $y_{\mathrm{P}}$ values. By comparing the results obtained in each case, we may determine the effects of changing $y_{\mathrm{P}}$ in relative terms.

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

### 3.1 Determining $y_P$

To quantify how varying $y_P$ changes the predicted emission (and hence the sensitivity for determining $y_P$), we first consider the ratio of H $\alpha$ to helium emissivity $R_\varepsilon$, defined as

$$R_\varepsilon \equiv \frac{\varepsilon(\mathrm{H}\,\alpha)}{\varepsilon(\mathrm{He\,II}\,4686\,\text{Å}) + \varepsilon(\mathrm{He\,I}\,10830\,\text{Å}) + \varepsilon(\mathrm{He\,I}\,5876\,\text{Å})}, \quad (3)$$

such that the net He I emissivity is the sum of the brightest two He I lines. These comprise the near-infrared 10830 Å line and the 5876 Å line, which is intrinsically fainter but lies in the optical part of the spectrum along with the H $\alpha$ 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 $y_P$ is altered by a factor $f_y$:

$$f_y \equiv \frac{y_P}{y_{P,\,\mathrm{fid}}}, \quad (4)$$

which we allow to take the values $f_y = (1.01, 0.99, 1.10, 0.90)$, corresponding to $\pm 1$ per cent and $\pm 10$ per cent changes in $y_P$.

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 that 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 $N_{\mathrm{H\,I,\,max}}$. This condition is appropriate because $N_{\mathrm{H\,I,\,max}}$ provides a proxy for $\tau_\infty$, 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 $N_{\mathrm{H\,I,\,max}}$ is not known a priori, so for each model variation we use an iterative procedure to determine the halo mass that results in the desired value of $N_{\mathrm{H\,I,max}}$. While we must choose a value of $N_{\mathrm{H\,I,max}}$ 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 ($N_{\mathrm{H\,I,max}} > 10^{19}$ cm$^{-2}$). We adopt $N_{\mathrm{H\,I,max}} = 10^{23}$ 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 $N_{\mathrm{H\,I}}$ with radius means that this column density is achieved only for sightlines passing through the very centre of the RELHIC.[5]

We plot the resulting emissivity ratios, as a function of the hydrogen number density $n_H$, in the upper panel of Fig. 2. As expected, $R_\varepsilon$ decreases (i.e. $\varepsilon_{He}$ is larger relative to $\varepsilon_{H\alpha}$) for models with a larger value of $y_P$, and conversely for smaller $y_P$. In the middle panel of Fig. 2, we additionally scale the obtained $R_\varepsilon$ curves by $f_y$ to demonstrate that these profiles remain almost self-similar, particularly near the ionization front where the emissivities are maximized. We quantify this property in the lower panel, where the

[5]See Section 4 of this paper and section 2.3 of S19 for further discussion.

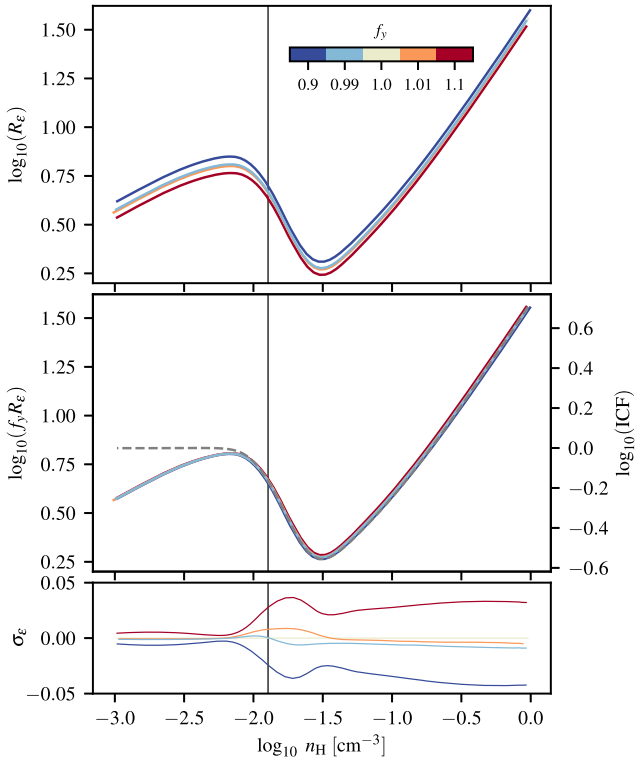


**Figure 2.** *Upper panel*: hydrogen to helium emissivity ratio $R_\varepsilon$ (equation 3) for models with different $f_y$ (equation 4), as a function of $n_H$. *Middle panel*: emissivity ratios, scaled additionally by $f_y$, for the same models. The ionization correction factor for $f_y = 1$, defined in the text, is shown with a dashed grey line using the right-hand scale. *Lower panel*: deviation from self-similarity $\sigma_\varepsilon$ (equation 5), 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 $\sigma_\varepsilon$, defined as

$$\sigma_\varepsilon \equiv \frac{f_y R_\varepsilon - R_{\varepsilon,\,\mathrm{fid}}}{R_{\varepsilon,\,\mathrm{fid}}}, \quad (5)$$

is plotted for each value of $f_y$. We find larger values of $\sigma_\varepsilon$ 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 $<1$ per cent for 1 per cent changes in $y_P$ and $\sim 2$ per cent for 10 per cent changes. Hence, they may be safely discounted, and we are justified in interpreting the variations in $R_\varepsilon$ as being solely caused by changes in $y_P$. This direct relationship between the observable values of $R_\varepsilon$ 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 as follows:

$$\mathrm{ICF} \equiv \frac{n_{\mathrm{H\,II}}}{n_{\mathrm{H}}} \frac{n_{\mathrm{He}}}{n_{\mathrm{He\,II}} + n_{\mathrm{He\,III}}} = \frac{1 - X_{\mathrm{H\,I}}}{1 - X_{\mathrm{He\,I}}}, \quad (6)$$

noting that this expression has a similar functional form to the emissivity ratio (by using equation 1 to expand each of the terms in equation 3). We illustrate this correspondence in the middle panel of Fig. 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 $R_\varepsilon$ 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 $R_\varepsilon$. Nevertheless, at densities of $n_{\mathrm{H}} \sim 10^{-2}\,\mathrm{cm}^{-3}$ associated with the peak emissivity, assuming a direct proportionality between $R_\varepsilon$ and the ICF allows a measurement of the former to be translated to an ionization correction, and hence a value of $y_{\mathrm{P}}$, to a precision of within 5 per cent.

However, the volume emissivity is not itself an observable quantity, but rather 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 $R_\varepsilon$, as

$$R_\Sigma \equiv \frac{\Sigma(\mathrm{H}\,\alpha)}{\Sigma(\mathrm{He\,II}\,4686\,\text{Å}) + \Sigma(\mathrm{He\,I}\,10830\,\text{Å}) + \Sigma(\mathrm{He\,I}\,5876\,\text{Å})}. \quad (7)$$

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

$$\sigma_\Sigma \equiv \frac{f_y R_\Sigma - R_{\Sigma,\,\mathrm{fid}}}{R_{\Sigma,\,\mathrm{fid}}}. \quad (8)$$

We plot $R_\Sigma$, $f_y R_\Sigma$, and $\sigma_\Sigma$ for the same models shown previously in Fig. 3, where the $x$-axis now shows projected neutral hydrogen column densities $N_{\mathrm{H\,I}}$ (rather than $n_{\mathrm{H}}$ as used in Fig. 2). We find that despite the projection effects, self-similarity is closely preserved when moving to surface brightnesses, with values of $\sigma_\Sigma$ at the position of the peak surface brightness in fact being smaller than the equivalent quantity for $R_\varepsilon$. 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 $y_{\mathrm{P}}$ from $R_\Sigma$ requires accurate measurements of the peak surface brightness $\Sigma_{\mathrm{max}}$ in H $\alpha$ and the three helium lines we consider. In absolute terms, this emission is still extremely faint, particularly for the He II 4686 Å 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 $y_{\mathrm{P}}$ 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 of $R_\Sigma$ are sensitive to our modelling of the complete temperature and density structure of the gas, whereas $R_\varepsilon$ depends only on local values of $T$ and $n_{\mathrm{e}}$. 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 $y_{\mathrm{P}}$ 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 appendix A of S19). We integrate the surface brightness over impact parameter to calculate total line luminosities, finding typical values of $L_{\mathrm{i}} \approx (6 \times 10^{35},\ 9 \times 10^{34},\ 3 \times 10^{34})\,\mathrm{erg\,s^{-1}}$ for the H $\alpha$, total He I, and He II 4686 Å 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 $\alpha$ to helium flux ratio as

$$R_{\mathrm{F}} \equiv \frac{F(\mathrm{H}\,\alpha)}{F(\mathrm{He\,II}\,4686\,\text{Å}) + F(\mathrm{He\,I}\,10830\,\text{Å}) + F(\mathrm{He\,I}\,5876\,\text{Å})}, \quad (9)$$

but in practice, compute the luminosity ratio $R_{\mathrm{L}}$ instead. In the main panel of Fig. 4, we show the values of $R_{\mathrm{F}}$ we obtain for different values of $y_{\mathrm{P}}$, finding the expected trend of decreasing $R_{\mathrm{F}}$ with increasing $y_{\mathrm{P}}$. More quantitatively, we expect that $R_{\mathrm{F}} \propto 1/y_{\mathrm{P}}$, with the normalization of this relation being set by the relative intrinsic emissivities of the three lines we consider. We use standard

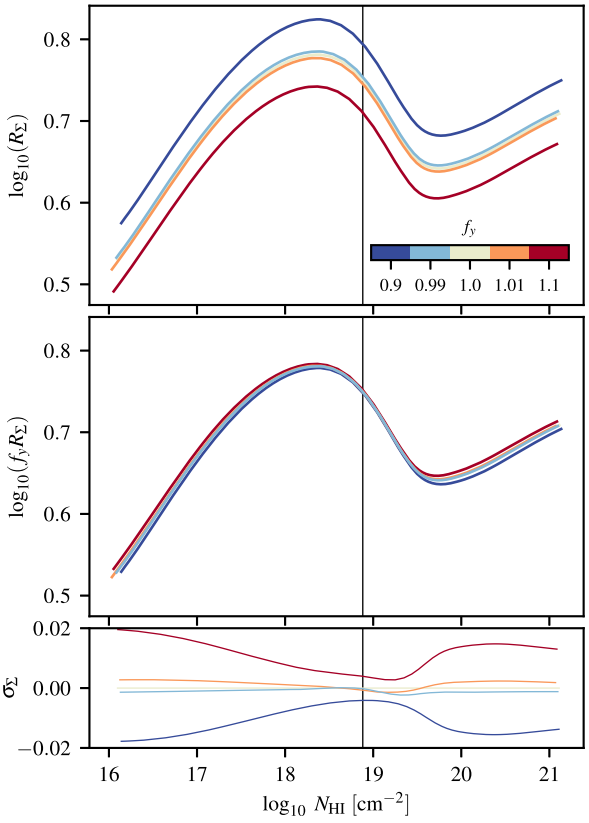


**Figure 3.** *Upper panel:* hydrogen to helium surface brightness ratio $R_\Sigma$ (equation 8) for the same models shown in Fig. 2, as a function of $N_{\mathrm{H\,I}}$. *Middle panel:* surface brightness ratios scaled by $f_y$. *Lower panel*: deviation from self-similarity $\sigma_\Sigma$ (equation 8), i.e. the error incurred in assuming $f_y R_\Sigma$ is independent of $y_{\mathrm{P}}$. The solid vertical lines indicate the H I column density at the location of the ionization front.

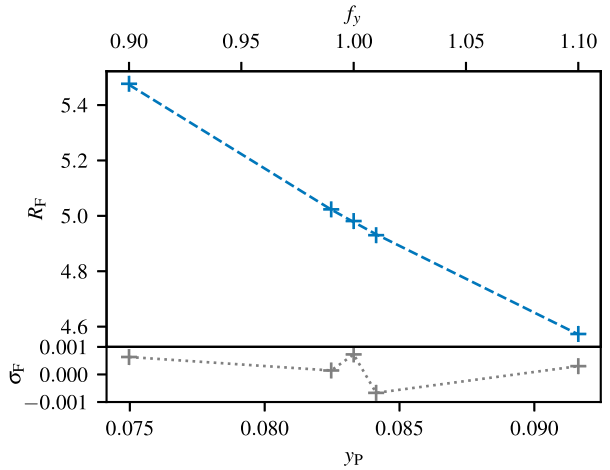


**Figure 4.** *Upper panel*: flux ratio $R_F$ (equation 9) as a function of $y_P$. The dashed line indicates a fit of the form $R_F = A/y_P + b$, where $A = 0.372$ and $b = 0.516$ are arbitrary scaling constants. *Lower panel*: the relative error $\sigma_F$ (equation 10) between the calculated flux ratios and the fitted line.

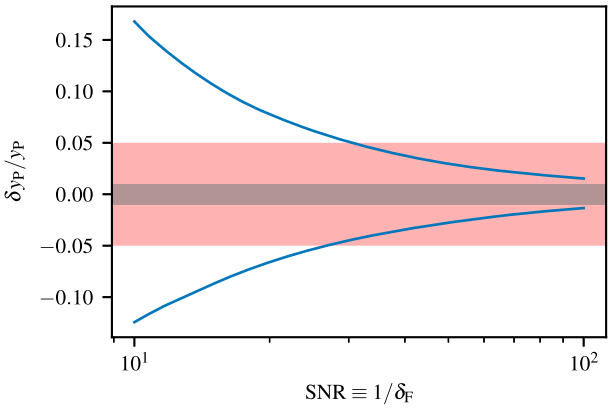


**Figure 5.** The relative precision to which $y_P$ can be inferred (blue curves) as a function of $\delta_F$, the precision to which the individual line fluxes are measured. The true helium abundance is assumed to be $y_P = 0.083$. Red and grey shaded bands show 5 and 1 per cent errors on $y_P$, respectively.

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. 4. To improve the fit, we additionally compute values of $R_F$ for $f_y = \pm 20$ per cent (not shown in Fig. 4). We define $\sigma_F$ as the normalized residual of the data with respect to this fit:

$$\sigma_F \equiv \frac{R_F - R_{F,\,\mathrm{fit}}}{R_{F,\,\mathrm{fit}}}, \tag{10}$$

and plot this as a function of $y_P$ in the lower panel, finding that $\sigma_F \ll 1$ per cent 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 measurements of the fluxes from which the flux ratio is calculated will have an associated uncertainty. We assume that this uncertainty is described by a single relative error value $\delta_F$ 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. 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 $\delta_{y_P}$, the error in $y_P$. In Fig. 5, we plot normalized values of $\delta_{y_P}$ as a function of the flux SNR, defined as SNR $= 1/\delta_F$. We find that for SNRs of 10 and 100, corresponding to $\delta_F = 10$ and 1 per cent, $y_P$ may be inferred to a precision of $^{+13}_{-10}$ and $^{+1.2}_{-1.0}$ per cent, respectively, where we have assumed that the underlying 'true' helium abundance is the fiducial value $y_P = 0.083$; repeating these calculations assuming different values of $y_P$ does not significantly affect the obtained values of $\delta_{y_P}$. The asymmetry in these limits results from the non-linearity of the function $R_F(y_P)$, which means that for a flux error of fixed magnitude $|\delta_F|$, the magnitude of $\delta_{y_P}$ will vary depending on the sign of $\delta_F$. Fig. 5 may also be used to determine the SNR required to achieve a constraint on $y_P$ of a given precision. The 5 per cent constraint indicated by the red shaded region corresponds to the range of reported values of $y_P$ (see references in Section 1); in the absence of a uniform systematic offset in these measurements, this is the minimum level of precision that must be reached for an independent measurement to provide additional information. Conversely, the grey region indicates a 1 per cent 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 $\delta_F \approx 3$ per cent or better. As discussed, we expect this conclusion to be insensitive to the details of the gas distribution given that it remains optically thick.

### 3.2 Determining the UVB spectral slope

In Section 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 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$\alpha$ 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 $y_P$. As discussed in Section 1, existing measurements of $y_P$ approach a precision of 1 per cent, 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 parametrize the UVB slope using the shape parameter $\alpha_{UV}$ introduced by Crighton et al. (2015), which modifies the slope of a given reference spectrum as follows:

$$J_\nu(E) = \begin{cases} N_\Gamma \times J_{\nu,\mathrm{ref}}(E) & E \le E_0 \\ N_\Gamma \times J_{\nu,\mathrm{ref}}(E) \times (E/E_0)^{\alpha_{UV}} & E_0 < E \le E_1 \\ N_\Gamma \times J_{\nu,\mathrm{ref}}(E) \times (E_1/E_0)^{\alpha_{UV}} & E > E_1, \end{cases} \tag{11}$$

where $J_{\nu,\,\mathrm{ref}}(E)$ 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 $\alpha_{UV}$. We set $E_0 = 1$ Ryd and $E_1 = 10$ Ryd, and introduce an additional factor $N_\Gamma \equiv \Gamma_{\mathrm{H\,I,fid}}/\Gamma_{\mathrm{H\,I}}$. This acts to renormalize 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 that corresponds to $\alpha_{UV} = 0$, we compute models for RELHICs illuminated by UVBs with $\alpha_{UV} = (-2, -1, 1)$. These values are chosen to cover all

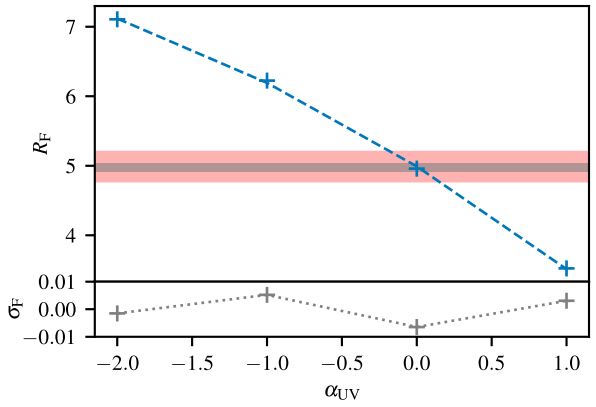


**Figure 6.** *Upper panel*: flux ratio $R_F$ as Fig. 4, but for models with fiducial value $y_P = 0.833$ and different UVB slopes $\alpha_{UV}$. The dashed line indicates a quadratic fit $R_F = A\alpha_{UV}{}^2 + B\alpha_{UV} + C$, where $A = -0.138$, $B = -1.34$, and $C = 4.99$. Red and grey shading indicate the range of $R_F$ values consistent with $\pm5$ and $\pm1$ per cent variations of $y_P$, respectively. *Lower panel*: the relative error $\sigma_F$ between the calculated flux ratios and the fitted line.

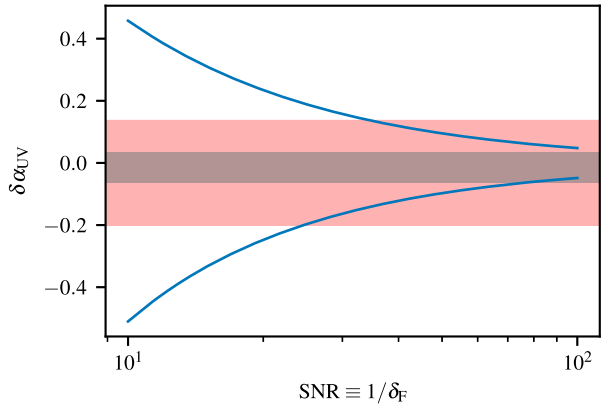


**Figure 7.** The absolute error with which $\alpha_{UV}$ can be inferred as a function of $\delta_F$. The true UVB slope parameter is assumed to be $\alpha_{UV} = 0$. Red shading and grey shading indicate the values of $\delta_{\alpha_{UV}}$ for which the expected change in $R_F$ is degenerate with that caused by changes in $y_P$ of 5 and 1 per cent, respectively.

realistic UVB spectra between the extremes of a soft, starburst-driven spectrum ($\alpha_{UV} = -2$) and a hard, AGN-dominated spectrum ($\alpha_{UV} = 1$). As described at the start of Section 3, for each value of $\alpha_{UV}$ we iteratively perform calculations for haloes of different masses to obtain a model RELHIC with H I column density equal to the adopted threshold $N_{HI,max} = 10^{23}$ cm$^{-2}$.We calculate values for $R_F$ as described previously, and plot these as a function of $\alpha_{UV}$ in Fig. 6. In red (grey) shading, we show the range of values of $R_F$ resulting from $\pm5$ per cent (1 per cent) variations in $y_P$, as plotted in Fig. 5. We see that changing $\alpha_{UV}$ 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 $\alpha_{UV}$, which reproduces the data to a precision of 1 per cent or better, as shown in the bottom panel of Fig. 6. In the same manner as was done for $y_P$, we use this fit to compute the precision with which $\alpha_{UV}$ may be reconstructed from uncertain measurements of $R_F$. This is shown in Fig. 7, where we plot $\delta_{\alpha_{UV}}$, the absolute measurement uncertainty in $\alpha_{UV}$, as a function of $\delta_F$. We find that SNRs of 10 and 100 yield constraints on $\alpha_{UV}$ 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 $\pm5$ per cent changes in $y_P$.

### 3.3 Combined constraints on $y_P$ and $\alpha_{UV}$

We have demonstrated that it is feasible to determine either $y_P$ or $\alpha_{UV}$ using the hydrogen-to-helium flux ratio $R_F$, assuming perfect knowledge of the other property. However, this is not representative of the real-world scenario in which both $y_P$ and $\alpha_{UV}$ are uncertain, as illustrated by the degeneracy visible in Fig. 7 and discussed above. In this section, we investigate the possibility of simultaneously constraining $y_P$ and $\alpha_{UV}$.

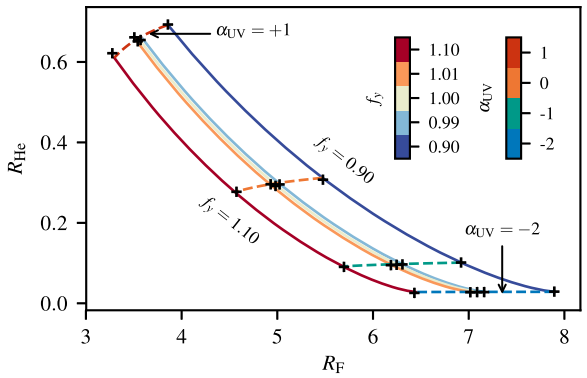


**Figure 8.** $R_F$ (equation 9) versus $R_{He}$ (equation 12) for models with different $y_P$ and $\alpha_{UV}$. Black points indicate values calculated from runs of our ionization balance code; curves show best fits to this data using equations (13) and (14). Curves at constant $y_P$ and different values of $\alpha_{UV}$ are shown with solid lines, coloured according to the left-hand colour bar. Conversely, curves at constant $\alpha_{UV}$ and varying $y_P$ are shown as dashed lines and coloured according to the right-hand colour bar.

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

$$R_{He} \equiv \frac{F(\text{He II } 4686\text{Å})}{F(\text{He I } 10830\text{ Å}) + F(\text{He I } 5876\text{ Å})} \quad (12)$$

to increase with increasing $\alpha_{UV}$. In contrast, changing $y_P$ scales all the helium ionic abundances equally, and so will only have a minor, indirect effect on $R_{He}$ 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 $y_P$ or $\alpha_{UV}$ while keeping the other parameter constant, with additional runs of our photoionization code in which both $y_P$ and $\alpha_{UV}$ are varied. We calculate values of $R_F$ and $R_{He}$ for this grid of models, which we show as the black points in Fig. 8. As in the 1D cases presented previously, we next fit a 2D surface to the calculated flux ratios, in order to allow interpolation of the ratios for arbitrary

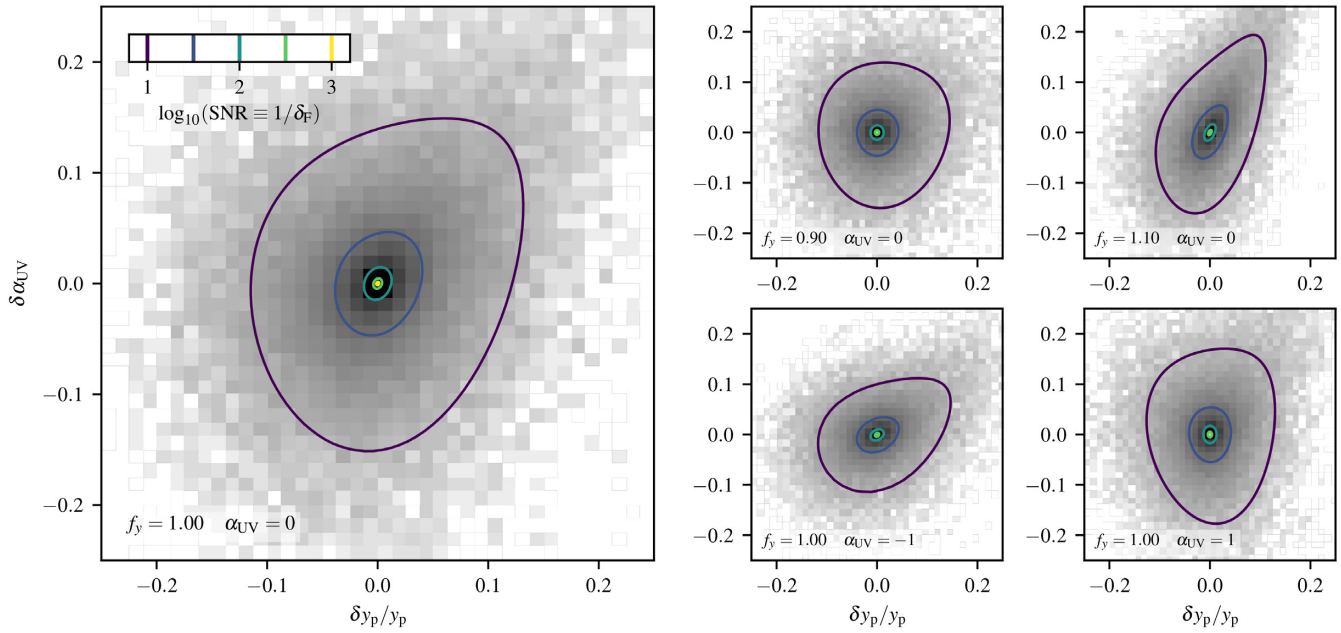


**Figure 9.** *Left-hand panel*: histogram of precision with which $y_P$ and $\alpha_{UV}$ may be recovered from sampled $R_F$ and $R_{He}$ values. Bins are shaded according to their occupancy using a logarithmic scale. Contours indicate the precision obtained for $1\sigma$ flux measurement uncertainties in the range $0.001 \leq \delta_F \leq 0.1$. The $y$-axis shows the absolute measurement uncertainty in $\alpha_{UV}$, whereas the relative error is plotted for $y_P$. *Right-hand panels*: as left, but for different underlying $y_P$ and $\alpha_{UV}$ values, as indicated by the legend in the bottom-left of each panel.

values of $y_P$ and $\alpha_{UV}$. We define these fits as follows:

$$R_F = f_1(y_P)\, g_1(\alpha_{UV}) \quad (13)$$

$$R_{He} = f_2(y_P)\, g_2(\alpha_{UV}), \quad (14)$$

where $g_1$, $g_2$, and $f_2$ are cubic polynomials, and $f_1 \propto 1/y_P$ as in Section 3.1. We emphasize 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 per cent or better across the range of $y_P$ and $\alpha_{UV}$ values we consider. We evaluate these fits and plot the resulting curves in Fig. 8 in order to show the degree to which they reproduce the data.

As a result of the non-linear mapping between the $R_F$–$R_{He}$ and $y_P$–$\alpha_{UV}$ 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 $R_F$ and $R_{He}$ are measured, which we derive by choosing a single flux error $\delta_F$ and propagating this uncertainty into $R_F$ and $R_{He}$ as described previously. Taking the underlying values of $y_P$ and $\alpha_{UV}$ to be 0.083 and 0, respectively, we sample 1000 uncertain 'measurements' of $R_F$ and $R_{He}$ by drawing from a bivariate normal distribution with means given by evaluating equations (13) and (14) for the true values. Variances are set to the squares of the chosen flux ratio errors, such that these errors correspond to $1\sigma$ uncertainties. We then invert equations (13) and (14) to map each of the sampled $R_F$–$R_{He}$ values to the $y_P$–$\alpha_{UV}$ plane, and repeat this procedure for different values of $\delta_F$. In the left-hand panel of Fig. 9, we illustrate the resulting collection of $\delta_{y_P}$–$\delta_{\alpha_{UV}}$ samples using a 2D histogram. Also shown is a series of contours indicating the $1\sigma$ limits on $y_P$ and $\alpha_{UV}$, which result from different choices of $\delta_F$ in the range $0.001 \leq \delta_F \leq 0.1$.

We see that the joint constraints also provide comparable or better precision than the individual ones, particularly for $\alpha_{UV}$. The addition of the $R_{He}$ measurement permits an improvement in the reconstructed precision of this parameter by approximately a factor of 4, with $\delta_F = 0.1$ now yielding the constraint on $\alpha_{UV}$ of $\pm 0.15$. Conversely, the precision with which $y_P$ may be recovered is $^{+14}_{-12}$ per cent, effectively unchanged to that obtained from the individual constraints (Fig. 5). These results are in agreement with the indication in Fig. 8 that $R_{He}$ evolves much more strongly with $\alpha_{UV}$ than with $y_P$. Increasing the assumed measurement precision to SNR = 100 improves these constraints significantly, to give $\pm 1.3$ per cent and $\pm 0.015$ 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 $y_P$ and $\alpha_{UV}$. We obtain comparably precise constraints in all cases shown, with SNR = 100 yielding values of $\delta_{y_P}/y_P$ of $\pm 1.3$ per cent or better, and $\delta_{\alpha_{UV}}$ of $\pm 0.017$ or better. Furthermore, an SNR of $10^{1.5} \approx 32$, corresponding to the second-outermost contour, is always sufficient to recover $y_P$ to the 5 per cent level at which existing determinations of the abundance differ. The tendency for a positive correlation between $\delta_{y_P}$ and $\delta_{\alpha_{UV}}$ is again a consequence of the general shape of Fig. 8: a positive value for $\delta_{y_P}$ is produced by a negative absolute error on $R_F$. When combined with an error on $R_{He}$ which is also $\geq 0$, the reconstructed value of $\alpha_{UV}$ exceeds the assumed underlying value, and therefore $\delta_{\alpha_{UV}}$ is positive also. Equivalently, a positive error on $R_F$ and a negative error on $R_{He}$ combine to yield inferred values of $y_P$ and $\alpha_{UV}$ that lie below the true ones. Curves at constant $\alpha_{UV}$ (defined parametrically by equations 13 and 14) flatten towards lower $y_P$, meaning this correlation is not as pronounced and giving the constraint contours shown in Fig. 9 their ovoid shape.

## 4 SUMMARY AND CONCLUSIONS

We have examined the properties of hydrogen and helium emission driven by UV background fluorescence in RELHICs, a class of $\sim 10^{9.5}\,M_\odot$ dark matter haloes that fail to form stars and instead

retain a small reservoir of neutral, essentially pristine gas at redshift $z = 0$. Using results obtained from a specialized 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 per cent 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 that settles at its photoionization equilibrium temperature, and at a density dictated by the requirement of hydrostatic equilibrium, reduces the impact of systematics that 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 3.2 and 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 report 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 $M_{200} \gg M_g$ 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 that RELHICs contain, this threshold is set to $n_{H,\,th} = 10\,cm^{-3}$, 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 $n_{H,\,th}$ 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 section 2.3 of S19, we have instead considered $H_2$ formation as a post-processing step, finding that our adopted column density threshold of $N_{H\,I,\,max} = 10^{23}\,cm^{-2}$ corresponds to the threshold above which formation of $H_2$ occurs. Additionally, the upper bound on halo mass of $M_{200} \leq 10^{9.6}\,M_\odot$ 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 $\Sigma_{H\alpha,\,max} > 10^{-20}\,erg\,s^{-1}\,cm^{-2}\,arcsec^{-2}$ and a projected ring diameter $\geq$1 kpc located within a 3 Mpc volume centred on the Milky Way. This rarity, coupled with the fact that even the brightest emission from RELHICs is still very faint for current technology, means detecting them is challenging at present. As such, a blind H$\alpha$ survey using current instrumentation (e.g. the Multi-Unit Spectrographic Explorer (MUSE) instrument at the Very Large Telescope) 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 21 cm 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, Giovanelli & Haynes 2013) has 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 $N_{H\,I} \sim 10^{16}\,cm^{-2}$ (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 21 cm sources with no associated stellar continuum as promising targets. Ultradeep 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.

## ACKNOWLEDGEMENTS

CS, MF, and TT acknowledge support by the Science and Technology Facilities Council (STFC; grant numbers ST/R504725/1 and ST/P000541/1). This project has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 757535). During this work, RC was supported by a Royal Society University

Research Fellowship. This work 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, and the Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. We thank the authors of the public PYTHON packages NUMPY, SCIPY, ASTROPY, and MATPLOTLIB, from which this work has benefited. We also thank an anonymous referee for their time and attention.

## REFERENCES

Adams E. A. K., Giovanelli R., Haynes M. P., 2013, ApJ, 768, 77
Alpher R. A., Bethe H., Gamow G., 1948, Phys. Rev., 73, 803
Aver E., Olive K. A., Skillman E. D., 2015, J. Cosmol. Astropart. Phys., 2015, 011
Benítez-Llambay A. et al., 2017, MNRAS, 465, 3913
Cooke R. J., Fumagalli M., 2018, Nat. Astron., 2, 957
Cooke R., Pettini M., 2016, MNRAS, 455, 1512
Crighton N. H. M., Hennawi J. F., Simcoe R. A., Cooksey K. L., Murphy M. T., Fumagalli M., Prochaska J. X., Shanks T., 2015, MNRAS, 446, 18
de Blok W. J. G., et al., 2018, Proc. Sci., An Overview of the MHONGOOSE Survey: Observing Nearby Galaxies with MeerKAT. SISSA, Trieste, PoS#007
Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, ApJ, 703, 1416
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
Giovanelli R. et al., 2005, AJ, 130, 2598
Gould A., Weinberg D. H., 1996, ApJ, 468, 462
Heald G. et al., 2011, A&A, 526, A118
Hoyle F., Tayler R. J., 1964, Nature, 203, 1108
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
Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, ApJ, 522, 82
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
Madau P., Haardt F., 2015, ApJ, 813, L8
Meiksin A. A., 2009, Rev. Mod. Phys., 81, 1405
Miralda-Escudé J., Rees M. J., 1994, MNRAS, 266, 343
Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, ApJ, 524, L19
Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563
Okamoto T., Gao L., Theuns T., 2008, MNRAS, 390, 920
Osterbrock D. E., Ferland G. J., 2006, Astrophysics of Gaseous Nebulae and Active Galactic Nuclei, 2nd edn. University Science Books, Sausalito, CA
Peimbert A., Peimbert M., Luridiana V., 2016, Rev. Mex. Astron. Astrofis., 52, 419
Pitrou C., Coc A., Uzan J.-P., Vangioni E., 2018, Phys. Rep., 754, 1
Planck Collaboration XVI et al., 2014, A&A, 571, A16
Popping A., Meyer M., Staveley-Smith L., Obreschkow D., Jozsa G., Pisano D. J., 2015, Proc. Soc., Observations of the Intergalactic Medium and the Cosmic Web in the SKA Era. SISSA, Trieste, PoS#132
Porter R. L., Ferland G. J., MacAdam K. B., Storey P. J., 2009, MNRAS, 393, L36
Porter R. L., Ferland G. J., Storey P. J., Detisch M. J., 2012, MNRAS, 425, L28
Porter R. L., Ferland G. J., Storey P. J., Detisch M. J., 2013, MNRAS, 433, L89
Power C. et al., 2015, Proc. Soc., Galaxy Formation & Dark Matter Modelling in the Era of the Square Kilometre Array. SISSA, Trieste, PoS#133
Puchwein E., Haardt F., Haehnelt M. G., Madau P., 2019, MNRAS, 485, 47
Sawala T. et al., 2016, MNRAS, 457, 1931
Schaye J., 2004, ApJ, 609, 667
Sternberg A., McKee C. F., Wolfire M. G., 2002, ApJS, 143, 419
Sykes C., Fumagalli M., Cooke R., Theuns T., Benítez-Llambay A., 2019, MNRAS, 487, 609
Valerdi M., Peimbert A., Peimbert M., Sixtos A., 2019, ApJ, 876, 98

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