Cosmic Variance and the Inhomogeneous UV Luminosity Function of Galaxies During Reionization

Taha Dawoodbhoy, Paul R Shapiro, Pierre Ocvirk, Joseph S.W Lewis, Dominique Aubert, Jenny G Sorce, Kyungjin Ahn, Ilian T Iliev, Hyunbae Park, Romain Teyssier, et al.

Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **524,** 6231–6246 (2023) https://doi.org/10.1093/mnras/stad2331
Advance Access publication 2023 August 02

Cosmic variance and the inhomogeneous UV luminosity function of galaxies during reionization

Taha Dawoodbhoy [1]★ Paul R. Shapiro,[1]★ Pierre Ocvirk [2] Joseph S. W. Lewis [3,4] Dominique Aubert,[2] Jenny G. Sorce [5,6,7] Kyungjin Ahn,[8] Ilian T. Iliev [9] Hyunbae Park,[10,11,12] Romain Teyssier [13] and Gustavo Yepes [14,15]

[1]*Department of Astronomy and Texas Center for Cosmology and Astroparticle Physics, The University of Texas at Austin, Austin, TX 78712-1083, USA*
[2]*Université de Strasbourg, Observatoire astronomique de Strasbourg, UMR 7550, F-67000 Strasbourg, France*
[3]*Zentrum für Astronomie der Universität Heidelberg, Institut für Theoretische Astrophysik, Albert-Ueberle-Straße 2, D-69120 Heidelberg, Germany*
[4]*Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*
[5]*Univ. Lille, CNRS, Centrale Lille, UMR 9189 CRIStAL, F-59000 Lille, France*
[6]*Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany*
[7]*Université Paris-Saclay, CNRS, Institut d'Astrophysique Spatiale, F-91405 Orsay, France*
[8]*Chosun University, 375 Seosuk-dong, Dong-gu, Gwangju 501-759, Korea*
[9]*Astronomy Center, Department of Physics and Astronomy, Pevensey II Building, University of Sussex, Falmer, Brighton BN1 9QH, UK*
[10]*Lawrence Berkeley National Laboratory, CA 94720-8139, USA*
[11]*Berkeley Center for Cosmological Physics, UC Berkeley, CA 94720, USA*
[12]*Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan*
[13]*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08540, USA*
[14]*Departamento de Física Teórica M-8, Universidad Autónoma de Madrid, Cantoblanco, E-28049 Madrid, Spain*
[15]*Centro de Investigación Avanzada en Física Fundamental (CIAFF), Universidad Autónoma de Madrid, E-28049 Madrid, Spain*

Accepted 2023 July 20. Received 2023 June 8; in original form 2023 February 20

ABSTRACT

When the first galaxies formed and starlight escaped into the intergalactic medium to reionize it, galaxy formation and reionization were both highly inhomogeneous in time and space, and fully coupled by mutual feedback. To show how this imprinted the UV luminosity function (UVLF) of reionization-era galaxies, we use our large-scale, radiation-hydrodynamics simulation CoDa II to derive the time- and space-varying halo mass function and UVLF, from $z \simeq 6$–15. That UVLF correlates strongly with local reionization redshift: earlier-reionizing regions have UVLFs that are higher, more extended to brighter magnitudes, and flatter at the faint end than later-reionizing regions observed at the same z. In general, as a region reionizes, the faint-end slope of its local UVLF flattens, and, by $z = 6$ (when reionization ended), the global UVLF, too, exhibits a flattened faint-end slope, 'rolling-over' at $M_{\rm UV} \gtrsim -17$. CoDa II's UVLF is broadly consistent with cluster-lensed galaxy observations of the Hubble Frontier Fields at $z = 6$–8, including the faint end, except for the faintest data point at $z = 6$, based on one galaxy at $M_{\rm UV} = -12.5$. According to CoDa II, the probability of observing the latter is ~ 5 per cent. However, the effective volume searched at this magnitude is very small, and is thus subject to significant cosmic variance. We find that previous methods adopted to calculate the uncertainty due to cosmic variance underestimated it on such small scales by a factor of 2–4, primarily by underestimating the variance in halo abundance when the sample volume is small.

Key words: galaxies: high-redshift – galaxies: luminosity function, mass function – dark ages, reionization, first stars – cosmology: theory.

1 INTRODUCTION

Reionization – the process by which starlight from early galaxies leaks into the surrounding intergalactic medium (IGM), gradually changing its ionization state from almost completely neutral before the first stars formed, at $z \gtrsim 20$, to almost completely ionized at $z \lesssim 5.5$ – was highly inhomogeneous in space and time (see, e.g. Yoshida, Hosokawa & Omukai 2012; Dayal & Ferrara 2018; Bosman et al. 2022). The inhomogeneity was seeded by the gravitational growth of density perturbations in the early Universe, which, when non-linear, formed small overdense regions packed with clusters of dark matter haloes in some places, and vast underdense voids in others. The largest amplitude perturbations formed these non-linear structures first, and so the most overdense regions were the earliest sites of star formation. As such, these regions were the first to reionize, and were the origins of the first H II bubbles that grew radially outward from them, 'exporting' their excess ionizing radiation to nearby lower density regions and reionizing them, as well, in the process (Dawoodbhoy et al. 2018). Over time, regions of

★ E-mail: tahad@astro.as.utexas.edu (TD); shapiro@astro.as.utexas.edu (PS)

progressively smaller amplitude overdensity also reached their non-linear phase, forming more galaxies and stars, making existing H II bubbles larger and forming new ones where none had been before, eventually overlapping and filling all of space to complete the Epoch of Reionization ('EOR'). Thus, the local reionization history of any given region is strongly correlated with its local overdensity, and it is crucially important to take this correlation into account when making predictions for observables that depend on both the density and the ionization state of the observed region. As we stress throughout this paper, the high-redshift UV luminosity function (UVLF; the number density of galaxies per unit UV luminosity or absolute magnitude, denoted Φ) – especially that for faint galaxies (absolute UV magnitudes $M_{\rm UV} \gtrsim -16$) observed in small volumes through high-magnification gravitational lenses – is one such observable (see e.g. Kulkarni & Choudhury 2011, for a semi-analytic study).

Recent analyses of high-z galaxies found via the lensing clusters in the Hubble Frontier Fields (HFF) (Lotz et al. 2017) have led to some debate about the shape of the faint-end of the UVLF, especially at $z \sim 6$. The UVLF is often fit and parametrized with the Schechter function (Schechter 1976), which asymptotes to a power law at the faint end ($\Phi \propto (L/L_*)^{\alpha}$, $\alpha < 0$, for $L \ll L_*$) and an exponential cut-off at the bright end ($\Phi \propto \exp(-L/L_*)$, for $L \gg L_*$). Some HFF studies (Bouwens et al. 2017; Atek et al. 2018) have argued that the faint-end of the $z \sim 6$ UVLF deviates from the power-law behaviour of the Schechter function by gradually flattening in slope faint-ward of $M_{\rm UV} \sim -16$, while others (Livermore, Finkelstein & Lotz 2017; Ishigaki et al. 2018) have argued that the power-law behaviour is maintained well below this luminosity, with no evidence of a change in slope. However, the observations and analyses of these extremely faint galaxies at high redshift are necessarily volume-limited; the galaxies can only be observed if they are located in particular regions such that their magnification by the foreground cluster is sufficient to raise their apparent brightness above the survey's flux limit. These small-volume, high-magnification observations are subject to significant uncertainties, due to both potential errors in the lensing model and cosmic variance. In this work, we seek to understand the latter.

There are several factors that cause the UVLF – or the SFR, on which the UVLF strongly depends – to vary from region to region (see e.g. Dawoodbhoy et al. 2018):

(i) overdense/early-reionizing regions have higher halo number densities than underdense/late-reionizing regions, especially at the high-mass end, and higher mass haloes have higher SFRs than lower mass haloes;
(ii) haloes of a given mass in overdense/early-reionizing regions have higher SFRs on average than haloes of the same mass in underdense/late-reionizing regions;
(iii) low-mass haloes ($M \lesssim 10^{9.5}$ M$_\odot$) have lower SFRs in regions that have already been reionized at a given redshift than regions that have yet to be reionized.

Accurate modelling of these factors and their contribution to the cosmic variance of the UVLF requires the use of large-scale, high-resolution, fully coupled radiation-hydrodynamics simulations, because they are highly contingent on the complex mutual feedback between galaxy formation and reionization. The last factor, in particular, is due to the feedback of ionizing radiation photo-heating the IGM in the vicinity of low-mass haloes, thereby inhibiting their ability to accrete gas and form stars (cf. Shapiro, Giroux & Babul 1994). Since low-mass haloes preferentially occupy the faint end of the UVLF, this reionization-induced suppression of star formation will have a significant impact on the faint end, where observational uncertainty due to cosmic variance is largest. A complete analysis of cosmic variance at the faint end, therefore, requires a simulation that can resolve such low-mass haloes in a box large enough to sample a wide range of local reionization histories, and treat the interplay between their star formation and the backreaction of ionizing radiation self-consistently. Specifically, as we will show, the faint-end HFF observations probe haloes with masses $M \gtrsim 10^{8.5}$ M$_\odot$, so we require a simulation with a dark matter particle mass of $m_{\rm DM} \lesssim 10^6$ M$_\odot$, to resolve the formation of such haloes with at least a few hundred particles each. Furthermore, since the volume searched by an HFF survey is $\sim$3000 cMpc^3, we require a simulation with a box size of $L_{\rm box} \gtrsim 100$ cMpc, so that it contains a statistically meaningful sample of at least a few hundred survey volumes, with a self-consistent distribution of overdensities and reionization histories. While several radiation-hydrodynamics reionization simulations have been produced in recent years (e.g. Gnedin & Fan 2006; Finlator, Davé & Özel 2011; Gnedin 2014; Iliev et al. 2014; So et al. 2014; Ocvirk et al. 2016; Xu et al. 2016; Pallottini et al. 2017; Pawlik et al. 2017; Aubert et al. 2018; Rosdahl et al. 2018; Ocvirk et al. 2020; Kannan et al. 2022; Lewis et al. 2022), only a few meet these size and resolution requirements, simultaneously – namely, the Cosmic Dawn ('CoDa') (Ocvirk et al. 2016; Aubert et al. 2018; Ocvirk et al. 2020; Lewis et al. 2022) and THESAN (Kannan et al. 2022) simulations (see Table 1), though, until now, these have not been analysed for this purpose.

Table 1. A comparison of the dark matter particle masses ($m_{\rm DM}$) and box sizes ($L_{\rm box}$) of high-resolution, large-scale, radiation-hydrodynamics EOR simulations.

Simulation	$m_{\rm DM}$ (M$_\odot$)	$L_{\rm box}$ (cMpc)
CoDa I	3.49×10^5	91.4
CoDa I-AMR	2.79×10^6	91.4
CoDa II	4.07×10^5	94.4
CoDa III	5.09×10^4	94.4
THESAN-1	3.12×10^6	95.5

In what follows, therefore, we present the first study of the inhomogeneous UVLF during the EOR based upon a self-consistent radiation-hydrodynamics simulation of fully coupled galaxy formation and reionization with the required large volume and high mass-resolution described above. We use the second-generation CoDa II simulation (Ocvirk et al. 2020)[1] to determine how the UVLF at $z \geq 6$ varies spatially in correlation with regional variations in the halo mass function (HMF) and the local timing of reionization, to establish the cosmic variance of the UVLF on scales large and small in a statistically meaningful way. We will compare our predicted UVLF's with observations and assess the implications of our results for surveys based upon HFF lensing data.

Our paper is organized as follows. In Section 2, we briefly describe the CoDa II simulation and its relevant post-processing for this work. In Section 3.1, we illustrate the temporal evolution and spatial inhomogeneity of the UVLF in CoDa II across a wide range of local reionization histories and overdensities. In Section 3.2, we compare our results to the HFF observations, and demonstrate a substantial discrepancy between the estimate of uncertainty due

[1]Some results for the high-z UVLF from the CoDa II simulation were presented by us before, in Ocvirk et al. (2020), but here we will analyse its *inhomogeneity* for the first time, while also presenting fitting formulae for the globally averaged UVLF for direct comparison with the observed UVLF, with special attention to evidence for flattening at the faint end.

Table 2. Cosmological parameters from Planck Collaboration XVI (2014), which are used in CoDa II.

Parameter	Value
h	0.677
Ω_m	0.307
Ω_Λ	0.693
Ω_b	0.048
σ_8	0.829
n	0.963

to cosmic variance derived from our simulation and that used in previous studies, when applied to the small-volume lensing results discussed here. In Section 3.3, we fit our globally averaged CoDa II UVLF at $z = 6, 7, 8, 10, 15$ to Schechter functions with and without modifications to the faint-end behaviour, to determine whether our simulation predicts a flattening of the faint-end slope. We conclude and summarize our results in Section 4.

2 CODAII SIMULATION

CoDa II, described in detail in Ocvirk et al. (2020), is the second-generation radiation-hydrodynamics simulation of fully coupled galaxy formation and reionization in a ΛCDM universe by *The Cosmic Dawn ('CoDa') Project*, based upon the massively parallelized, hybrid CPU-GPU code RAMSES-CUDATON. CoDa II has periodic boundary conditions in a cubic volume 94.4 cMpc on a side, with 4096^3 *N*-body particles for the dark matter and 4096^3 grid cells for the baryonic gas and radiation field, resolving the formation of the full range of atomic-cooling halo ('ACH') masses, $M \gtrsim 10^8$ M$_\odot$, and simulating through the end of reionization to $z = 5.8$. The simulation adopts cosmological parameters from Planck Collaboration XVI (2014), which are provided in Table 2.

Hydrodynamics and *N*-body dynamics are handled by the RAMSES code (Teyssier 2002), which uses a second-order Godunov scheme Riemann solver for the gas and a particle-mesh integrator for the dark matter. Radiative transfer ('RT') and thermochemistry are handled by the ATON code (Aubert & Teyssier 2008), which relies on a moment-based description of the radiative transfer equations and uses the M1 closure relation (González, Audit & Huynh 2007). It tracks the out-of-equilibrium ionizations and cooling processes involving atomic hydrogen. Radiative quantities (energy density, flux, and pressure) are described on a fixed, comoving, Eulerian grid – the same grid as is used for its particle-mesh *N*-body gravity solver – and evolved according to an explicit scheme under the constraint of a CFL condition.

The latter condition is especially challenging for the cosmic reionization problem, since weak, R-type ionization fronts, driven by UV starlight emitted inside galaxies, break out of the galaxies and accelerate in the low-density IGM up to velocities that are many thousands of km s^{-1}, even approaching an appreciable fraction of the speed of light. As a result, the time-step upper-limit set by the CFL in the presence of such high-speed I-fronts is orders of magnitude smaller than that set by the CFL condition for hydrodynamics alone (without RT), even if the latter hydro assumes optically thin photoionization that raises the sound speed of ionized gas by heating it to 10^4 K. The small time-steps required by the CFL when hydro and RT are fully coupled has the unfortunate consequence that the number of time-steps required to integrate over a given interval of cosmic time by finite-differencing the hydro, gravity, and RT equations together (with the same time-step) is orders of magnitude larger than for a cosmological simulation of hydro and gravity without RT. It is currently computationally infeasible to do this on the scale required for as large a simulation as CoDa II.

The RAMSES-CUDATON code was specifically developed to overcome this obstacle. It is unique in solving this problem, by being coded to run on a massively parallel, hybrid CPU-GPU supercomputer like Titan at Oak Ridge OLCF, in which each of its thousands of nodes have, not only dozens of CPUs, but also GPUs. In the same wall-clock time it takes to advance the hydro and gravity equations on the CPUs for one hydro-gravity time-step, RAMSES-CUDATON uses the GPUs to advance through ∼100 sub-steps of the RT and ionization rate equations, as well. This enables the net computational time of the problem with RT to approach the computational time of the problem with no RT, by speeding it up by two orders of magnitude.

Other simulations of reionization and galaxy formation with fully coupled hydro and RT that solve the RT equations by a moment method, as RAMSES-CUDATON does, have attempted to side-step this severe requirement of extremely small RT-step-size dictated by the CFL condition by replacing the true speed of light by an artificially reduced value – the so-called reduced-speed-of-light approximation 'RSLA' – to 'trick' the CFL into allowing larger time-steps for the RT. This is not necessary for RAMSES-CUDATON. As a result, CoDa II was able to adopt the full speed of light, thereby avoiding the well-known artefacts introduced in the other reionization simulations by their adoption of the RSLA (see Deparis et al. 2019; Ocvirk et al. 2019). Nevertheless, to simulate through the end of the EOR, down to redshift 5.8, CoDa II had to run for about 6 days on 16 384 nodes of the Titan supercomputer, using 4 cores and 1 GPU per node, for a total of 65 536 cores, with each node hosting 4 MPI processes that each managed a subvolume of $64 \times 128 \times 128$ cells.

Since the mass scale of individual stars is completely unresolved by all reionization simulations, CoDa II included, star formation is modelled by a subgrid algorithm. In CoDa II, star particles are created in each hydro cell in which the baryon overdensity exceeds 50, with the rate of change of the stellar mass density given by

$$\dot{\rho}_\star = \epsilon_\star \frac{\rho_{\rm gas}}{t_{\rm ff}}, \tag{1}$$

where $\rho_{\rm gas}$ is the baryon density, $t_{\rm ff}$ is the free-fall time, and $\epsilon_\star$ is a calibration parameter referred to as the star formation efficiency, which is set to 0.02.

We add to our subgrid star formation algorithm a parametrized ionizing photon efficiency (IPE; the number of ionizing photons released per unit stellar baryon per unit time) into the host grid cell of each star particle. We define this IPE as $\xi_{\rm IPE} \equiv f_{{\rm esc},\star}\xi_{\rm ph,IMF}$, where $f_{{\rm esc},\star}$ is the stellar-birthplace escape fraction and $\xi_{\rm ph,IMF}$ is the number of ionizing photons emitted per Myr per stellar baryon. Each stellar particle is considered to radiate for one massive star lifetime $t_\star = 10$ Myr, after which the massive stars die (triggering a supernova explosion) and the particle becomes dark in the H-ionizing UV. We adopted an emissivity $\xi_{\rm ph,IMF} = 1140$ ionizing photons/Myr per stellar baryon. This is consistent with emission by our assumed $Z = 0.001$ BPASS binary stellar population model (Eldridge et al. 2017), with a Kroupa initial mass function (Kroupa 2001), assuming no dust extinction. We used a mono-frequency treatment of the radiation with an effective frequency of 20.28 eV. Finally, we calibrated $f_{{\rm esc},\star}$ by adjusting the value in a set of smaller box simulations, so as to obtain a reionization redshift close to $z = 6$, which led us to adopt a value of $f_{{\rm esc},\star} = 0.42$.

In order compute the UVLF, we must post-process the CoDa results to find its galaxies by a dark matter halo finding algorithm, assign star particles to each host halo according to their spatial

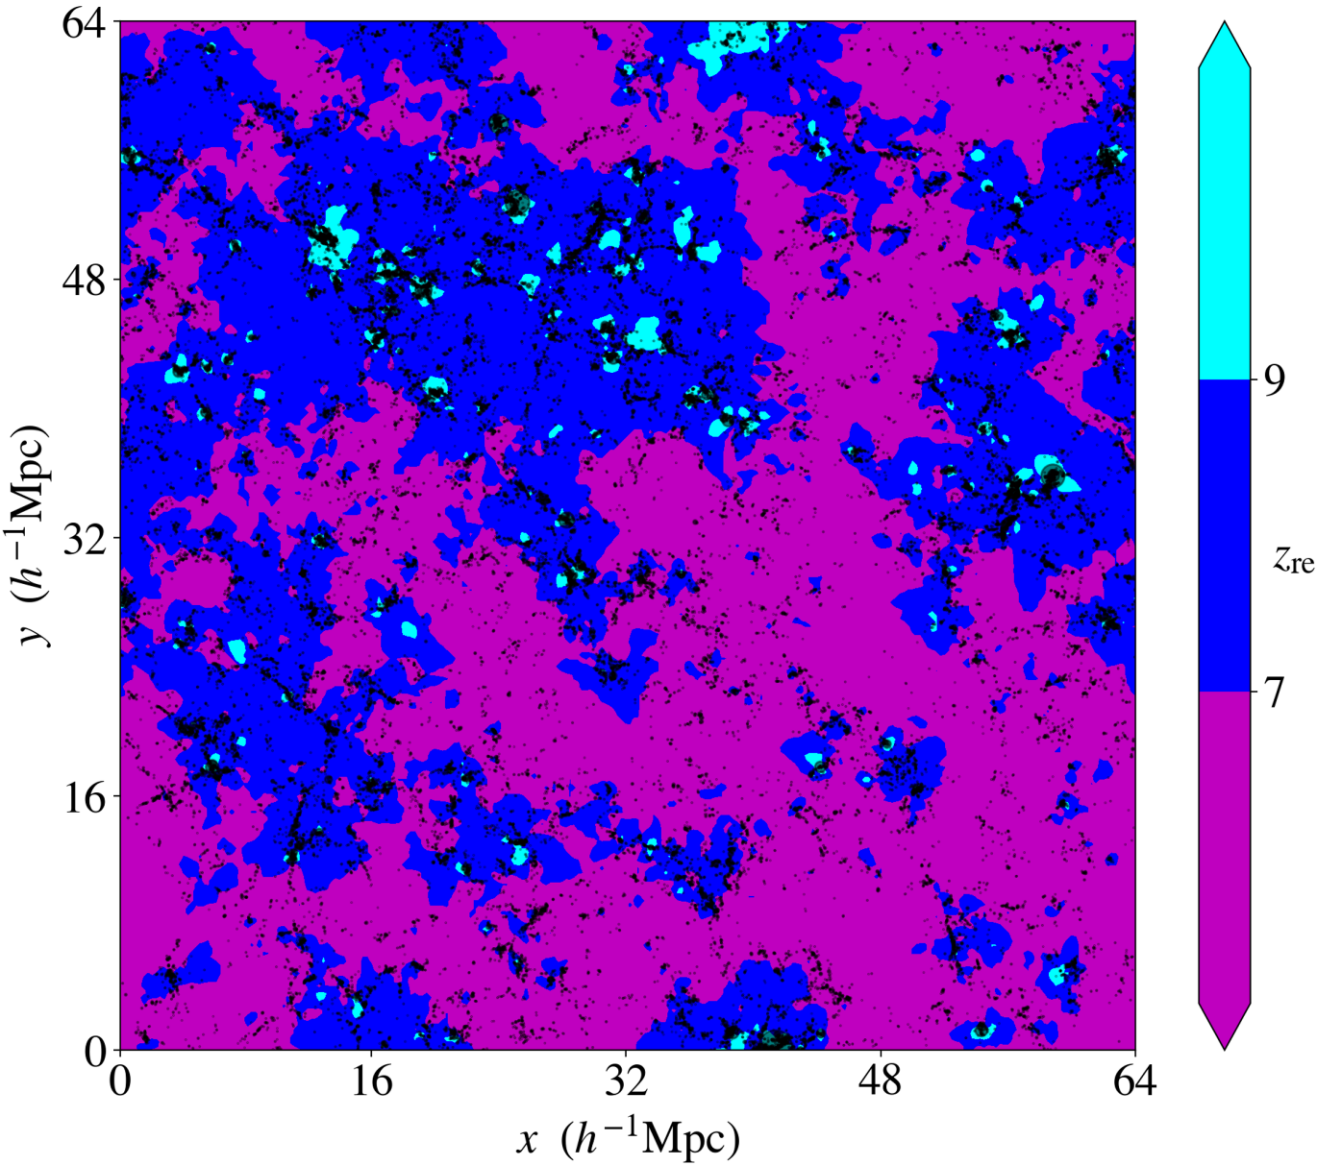

Figure 1. Contour map of a slice through the CoDa II reionization redshift field, 0.25 h^{-1} cMpc thick, divided into regions that reionize relatively early ($z_{re} > 9$; cyan), late ($z_{re} < 7$; magenta), and at intermediate redshifts ($7 < z_{re} < 9$; blue). The black circles show the locations of haloes in the same slice at $z = 10$, with the size of each circle proportional to the halo's mass.

overlap with the halo volumes, and sum the emission of all the star particles associated with a given halo below the Lyman limit of H atoms to compute that galaxy's UV continuum luminosity. Dark matter haloes are identified using a Friends-of-Friends algorithm with a standard linking length parameter of 0.2. The mass of each halo, M, is defined as the total mass of all linked dark matter particles, and the virial radius is estimated as

$$R_{200} = \left(\frac{3M}{4\pi \times 200 \bar{\rho}_{\rm DM}} \right)^{1/3}, \tag{2}$$

where $\bar{\rho}_{\rm DM}$ is the cosmic mean dark matter density. Star particles are then assigned to haloes if they fall within the halo's virial radius, and the masses and ages of each halo's star particles are used to compute the halo's UV luminosity and magnitude ($M_{\rm UV}$) at 1600 Å, according to the $Z = 0.001$ BPASS binary stellar population model described above, again assuming no dust extinction.

Table 3. Volume-weighted global ionized fraction in CoDa II at given z.

	$z = 6$	$z = 7$	$z = 8$	$z = 9$	$z = 10$	$z = 15$
$\langle X_{\rm HII} \rangle_{\rm V}$	1–1.2e-5	5.0e-1	1.7e-1	5.3e-2	1.6e-2	2.4e-4

To track and analyse the progress, patterns, and patchiness of reionization we construct the reionization redshift field of the CoDa II simulation, illustrated in Fig. 1. We start by coarsening the simulated grid to 256^3 cells, and computing the volume-weighted average ionized fraction in each of these cells at each snapshot. (See Table 3 for the global ionized fraction at select redshifts.) The purpose of this coarsening is to smooth over the interiors of haloes, which can be shielded from ionizing radiation due to their high densities, and instead probe the ionization state of the IGM. Then, we identify the redshift at which each coarse-grained cell first reaches an ionized fraction of 90 per cent, which we define as the cell's reionization

Table 4. The fraction of the total volume occupied by early-, intermediate-, and late-reionizing regions, along with the fraction of all haloes contained in these regions.

z_{re} bin	Volume fraction	$z = 6$ Halo fraction	$z = 7$ Halo fraction	$z = 8$ Halo fraction	$z = 10$ Halo fraction
>9	0.022	0.074	0.081	0.089	0.125
7–9	0.395	0.506	0.496	0.517	0.556
<7	0.583	0.420	0.422	0.393	0.319

redshift, z_{re}. Correspondingly, we identify each halo's z_{re} as that of the coarse-grained cell its centre-of-mass belongs to. For the purposes of this work, we consider three ranges of reionization redshift – $z_{re} > 9$, $7 < z_{re} < 9$, and $z_{re} < 7$ – which we refer to as early-, intermediate-, and late-reionizing regions, respectively. Fig. 1 shows a contour map of a slice through the reionization redshift field divided into these three ranges (cyan, blue, and magenta, respectively), along with the positions of $z = 10$ haloes (black dots) in the same slice. There is a clear correlation between halo number density and z_{re}, with the earlier reionizing regions containing a higher density of haloes than the later reionizing regions. For example, while the early-reionizing regions are the rarest and most compact, occupying only around 2 per cent of the total volume, they contain around 13 per cent of all haloes at $z = 10$ (see Table 4). On the other hand, the vast late-reionizing regions, which occupy 58 per cent of the volume, contain around 32 per cent of the haloes at $z = 10$. We explore this correlation further in the following section.

3 CODAII UV LUMINOSITY FUNCTION

3.1 The imprint of patchy reionization

In general, the UVLF can be decomposed into two factors: the HMF, and the star formation rate (SFR) of haloes as a function of mass. As we described in Dawoodbhoy et al. (2018), both of these factors are strongly correlated with the reionization history of the region in which they are observed, and so too must the UVLF be.

The earliest regions to reionize will be those that are the most dense, since these regions will be the first to form a large number of star-forming galaxies – the primary sources of reionization. The latest regions to reionize will be the voids, which typically do not form enough stars to reionize themselves, and so require 'importing' ionizing radiation from external, earlier-reionizing regions nearby, in order for them to become reionized. Therefore, there is a positive correlation between the reionization redshift of a region and its HMF (the number density of haloes per unit mass): higher-z_{re} regions have higher HMFs that extend out to higher mass. We show the combined HMFs in regions binned by their z_{re}, for four different redshifts, in Fig. 2. As can be seen, the early-reionizing regions ($z_{re} > 9$) have the highest and most extended (i.e. the turnover to a steeper decline occurs at a higher mass) HMF at all redshifts, followed by the intermediate- ($7 < z_{re} < 9$) and late-reionizing ($z_{re} < 7$) regions. The globally averaged HMF tracks closest to the intermediate-reionizing regions at the low-mass end.

Naturally, a higher HMF will result in a higher UVLF, overall, so we should expect the correlation between HMF and z_{re} to translate to a correlation between UVLF and z_{re}. However, the actual UV luminosity of each halo is determined by its SFR, which has a more complicated relationship with z_{re}. First, for relatively high-mass

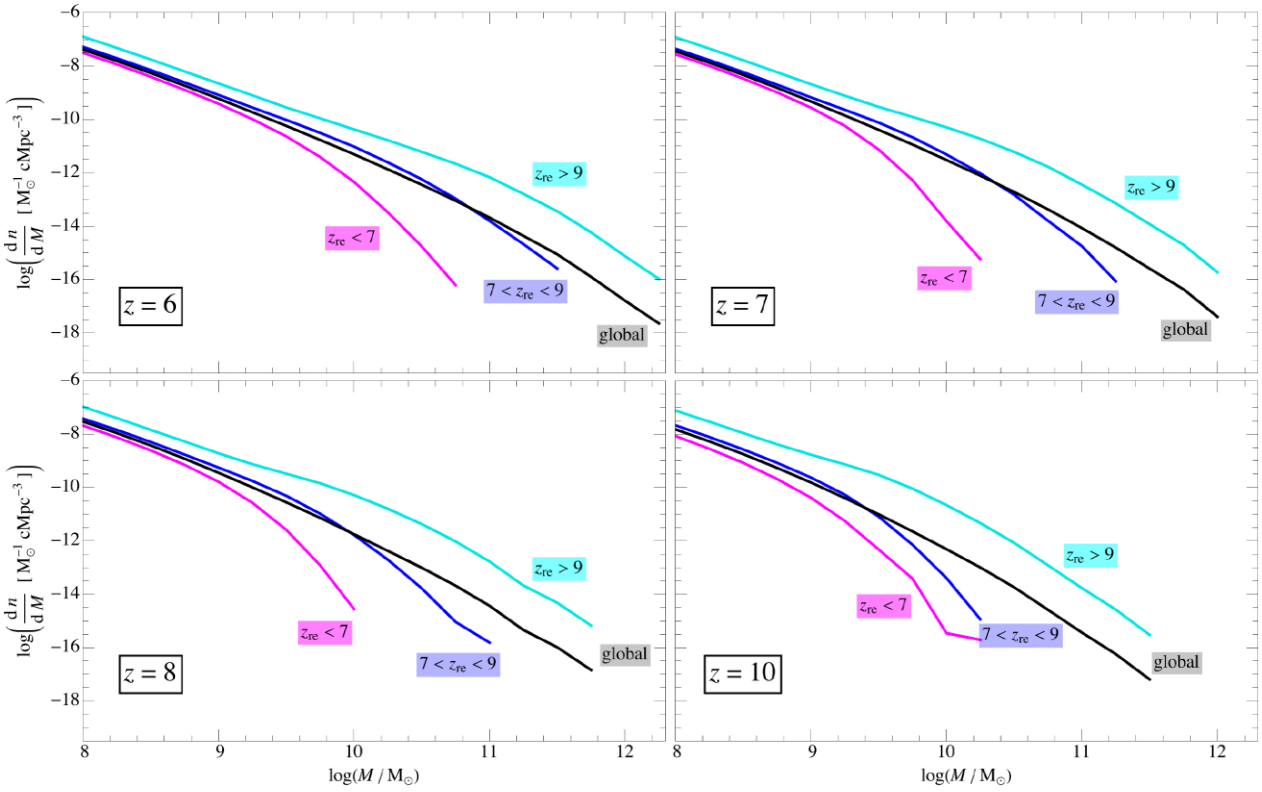

Figure 2. HMFs in CoDa II at $z = 6$, 7, 8, 10. The black solid lines show the HMF in the full simulated volume, while cyan, blue, and magenta lines show the HMF in early-, intermediate-, and late-reionizing regions, respectively, as labelled. Here and elsewhere, log implies $\log_{10}$.

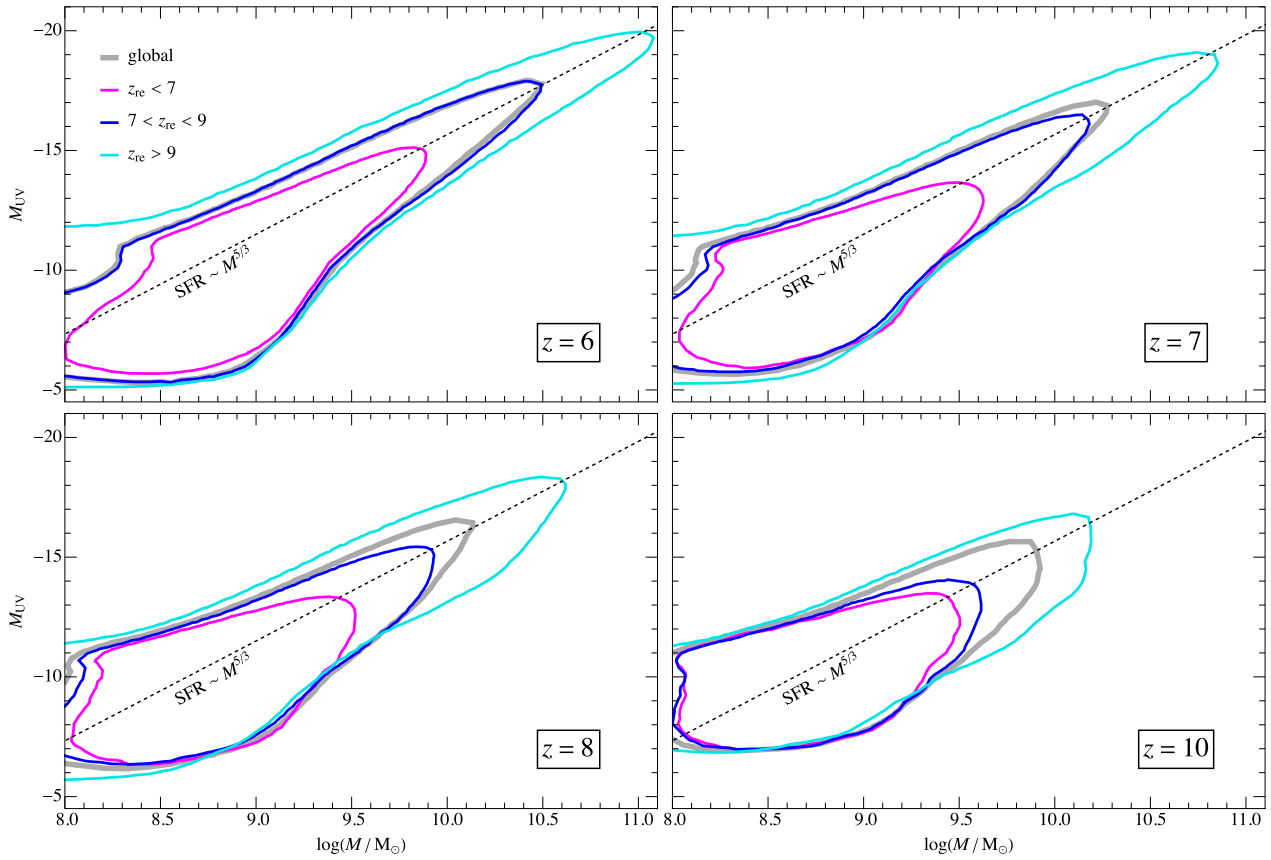

Figure 3. 95 per cent contours for UV magnitude versus halo mass of luminous galaxies in CoDa II, binned by z_{re}, for the same four redshifts as in Fig. 2 and using the same colour lines for different z_{re} bins. The thick grey line is the global median. The black dashed line is the rough expectation given a (pre-suppression) SFR $\propto M^{5/3}$.

haloes ($M \gtrsim 10^9$ M$_\odot$), the SFR scales as $\sim M^{5/3}$ (see e.g. Ocvirk et al. 2016, 2020), which means HMFs that are more extended to high mass (i.e. those of earlier reionizing regions) will correspond to UVLFs that are more extended to the bright end (i.e. the turnover to a steeper decline occurs at a brighter magnitude). We illustrate the effect of this SFR scaling in Fig. 3, which shows 95 per cent contours for the UV magnitude versus halo mass of luminous galaxies in CoDa II, binned by z_{re}, at $z = 6$–10. The rough expectation from the SFR $\sim M^{5/3}$ scaling (i.e. assuming UV luminosity is proportional to SFR) is well-obeyed for $M \gtrsim 10^{9.5}$ M$_\odot$. Furthermore, in addition to the earlier reionizing regions having more haloes at all masses and a HMF that extends out to higher mass, there is also a higher fraction of haloes of a *given* mass at a given redshift with brighter UV magnitudes in earlier reionizing regions than in later reionizing regions (e.g. notice that the bright edge of the contours are 'stacked' by z_{re}), which further contributes to the difference in their UVLFs.

On the other hand, lower mass haloes ($M \lesssim 10^{9.5}$ M$_\odot$)[2] deviate from the $M^{5/3}$ scaling at low redshift, due to the suppression of star formation in low-mass haloes, caused by reionization feedback (e.g. notice that the faint edge of the contours drop more sharply than the $M^{5/3}$ scaling for $M \lesssim 10^{9.5}$ M$_\odot$ at late redshift). After a region becomes reionized, low-mass haloes are unable to accrete the photo-ionized gas in the IGM, due to its increased temperature, and so they will no longer have the fuel required to form stars (see e.g. Dawoodbhoy et al. 2018; Ocvirk et al. 2020). Since these haloes populate the faint-end of the UVLF prior to their local reionization, we should expect to see a reduction at this faint-end over time as reionization occurs and the suppressed haloes move to fainter magnitudes (or disappear entirely), which is usually characterized in terms of a 'turnover' in the faint-end slope. For example, in their semi-analytical study of inhomogeneous reionization feedback, Kulkarni & Choudhury (2011) found such a turnover for $M_{UV} \gtrsim -17$ at $z = 8$, preferentially in the UVLFs of overdense regions (which reionize relatively early). To illustrate this effect in our simulation, we show the UVLF of intermediate-reionizing regions over time in Fig. 4. Notice the change in slope and curvature over time at magnitudes $-16 \lesssim M_{UV} \lesssim -11$. Prior to these regions' local reionization (i.e. $z = 10$), the faint-end slope is fairly steep, roughly following $\Phi \propto M_{UV}^{0.4}$ for $-13 \lesssim M_{UV} \lesssim -11$. During local reionization, however, the faint-end slope gradually flattens out, and by the time reionization has ended for these regions (i.e. $z = 6$), the faint-end power-law index is close to 0 in this magnitude range.

Consequently, the UVLF one observes depends on *where* one looks – an early-reionizing patch of the Universe will have a relatively high and bright-end-extended UVLF, whereas a late-reionizing patch will have a relatively low and bright-end-compressed UVLF– and also *when* one looks – a region that is observed prior to its local reionization will have a relatively steep faint-end slope, whereas a

[2]Note that the flattening seen at $M \lesssim 10^{8.5}$ M$_\odot$ in Fig. 3 is likely a resolution-limit effect.

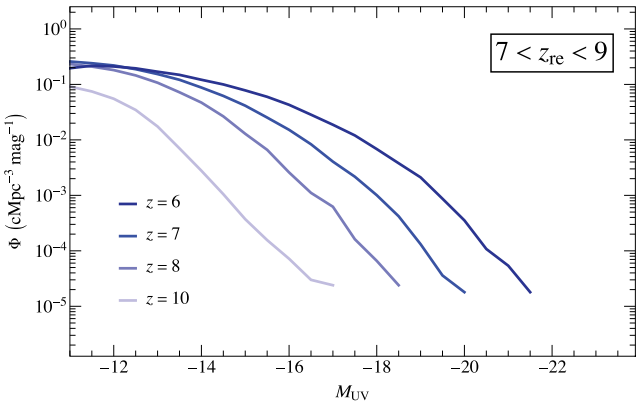

Figure 4. The evolution of the UVLF in intermediate-reionizing regions ($7 < z_{re} < 9$), from pre-reionization ($z = 10$) to post-reionization ($z = 6$). Notice that the latter features a much flatter faint-end than the former.

region observed after its local reionization will have a relatively flat faint-end slope. We show these trends in Fig. 5, which plots the CoDa II UVLFs for early-, intermediate-, and late-reionizing regions at four redshifts, along with the global average.

An important implication of these results is that small-volume observations of the UVLF will necessarily be biased in one way or another, due to the strong correlations with local density and reionization redshift. For example, observations that search in uniformly random volumes are likely to be probing voids, which are underdense regions, since they occupy the most volume. As a result, such observations are likely to return UVLFs that are lower and less extended at the bright end than the cosmic mean. Furthermore, since these regions reionize relatively late, the inferred UVLFs are likely to have steeper-than-average faint-end slopes. On the other hand, observations that preferentially search near the brightest sources are likely to be probing highly overdense regions that reionize relatively early, since these regions have UVLFs that are the most extended to the bright-end. Thus, such observations are likely to return UVLFs that are higher and more extended at the bright end than the cosmic mean, with a flatter-than-average faint-end slope.

3.2 Implications for faint-end HFF observations

The analysis of the previous section illustrates the dramatic variability of the UVLF among regions with different reionization histories. However, the volume of the z_{re}-binned cells in which the UVLF is computed is rather small, only $(250/h \text{ ckpc})^3$. For many observational purposes, it is more useful to assess the variance in the UVLF on larger scales, e.g. characteristic of the size of a galaxy survey. To that end, we divided our CoDa II box into 256 non-overlapping subvolumes, each spanning around 3300 cMpc^3, which is of the order of the survey volumes searched by each of the HFF lensing-cluster fields, which have been used previously to measure the faint-end of the high-z UVLF. Each subvolume contains $32 \times 32 \times 64 = 65\,536$ of the 256^3 coarse-grained cells used in the reionization redshift field of the previous section, and so will encompass a range of reionization histories. To characterize their typical reionization history, we compute the mean reionization redshift of each subvolume in two ways: (1) a halo-weighted average ($\langle z_{re} \rangle_h$), obtained by averaging over the reionization redshifts of all

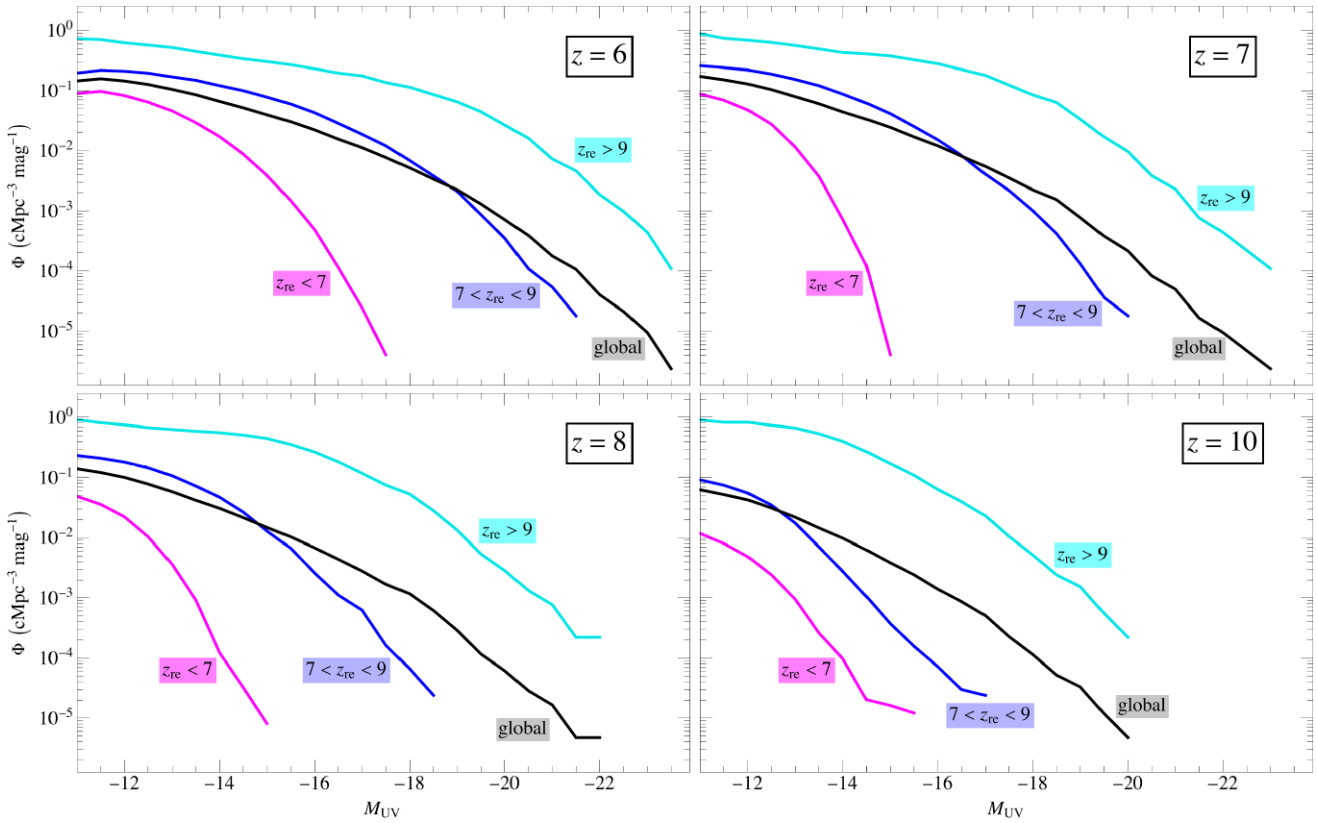

Figure 5. UVLFs in CoDa II at $z = 6$, 7, 8, 10. The black lines represent the full simulated volume, while cyan, blue, and magenta lines show early-, intermediate-, and late-reionizing regions, respectively.

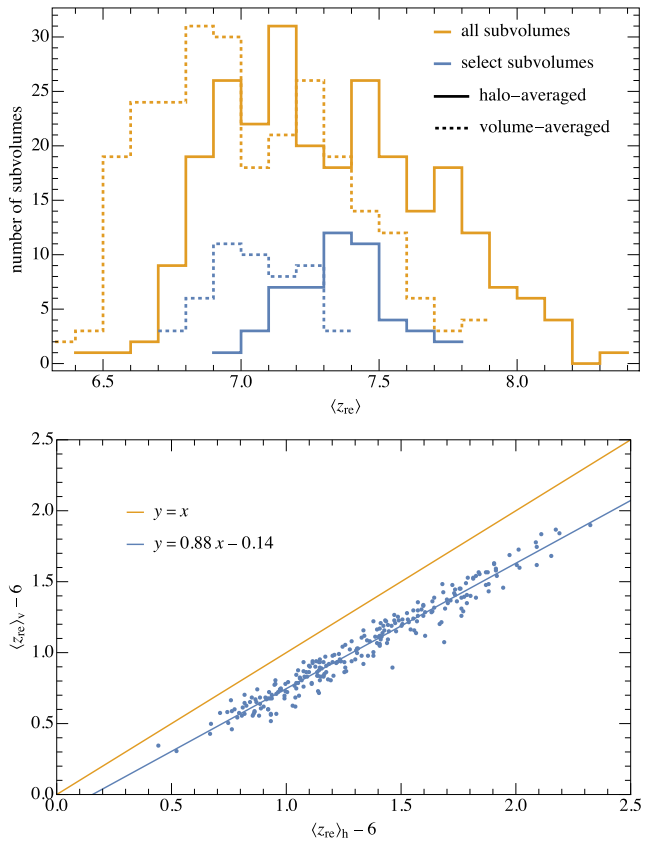

Figure 6. Top: Histograms of the mean reionization redshifts of CoDa II subvolumes (3300 cMpc^3 in size, each). The solid lines represent the halo-weighted average of z_{re} across each subvolume (i.e. the average z_{re} of all haloes in each subvolume), while dashed lines represent the volume-weighted average. The yellow lines show the distribution for all subvolumes, while the blue lines show that of a subset of 50 subvolumes that most closely match the bright-end data from Livermore et al. (2017) (see Fig. 8 and the accompanying text for a description of how these 50 subvolumes are selected). Bottom: Volume-weighted versus halo-weighted averages of z_{re} across each subvolume (blue points) fit to a linear relation (blue line). All points fall below the equality line (yellow), meaning $\langle z_{re} \rangle_v$ is always less (i.e. later) than $\langle z_{re} \rangle_h$.

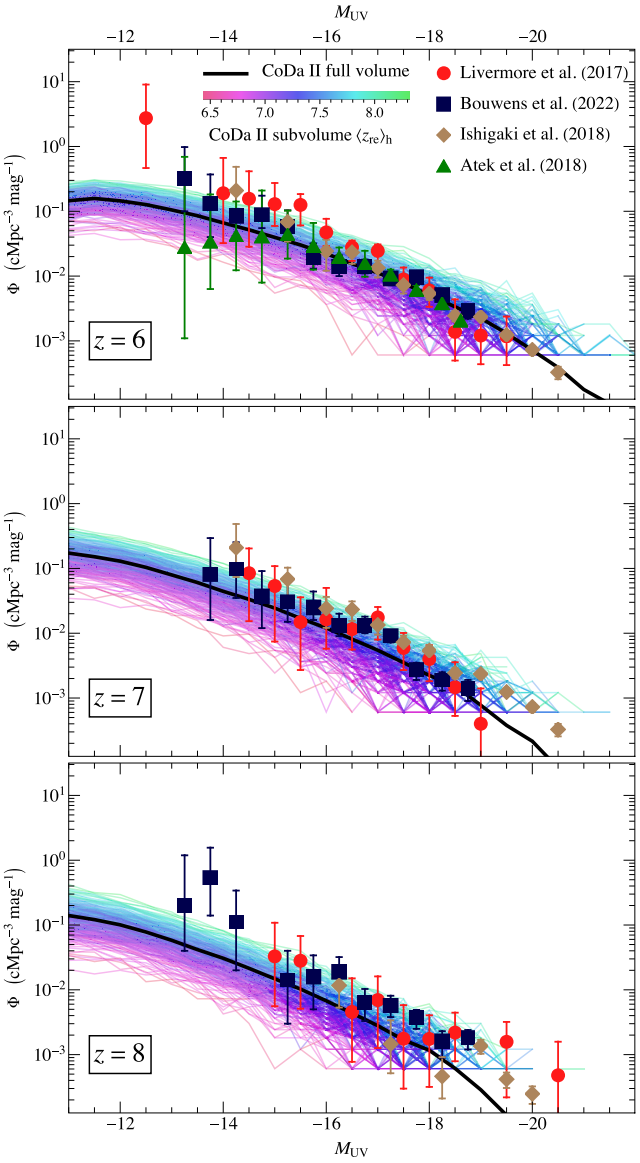

Figure 7. CoDa II UVLFs at $z = 6$, 7, 8 compared to observational HFF data from Livermore et al. (2017); Atek et al. (2018); Ishigaki et al. (2018); Bouwens et al. (2022). In addition to the full volume UVLF (thick black line), we also show the UVLFs of 256 CoDa II subvolumes – each spanning around 3300 cMpc^3, which is of the order of the survey volumes searched by each of the HFF lensing-cluster fields – as thin lines coloured according to the mean reionization redshift of their constituent haloes, as indicated in the legend in the top panel. The spread in the UVLFs of this collection of subvolumes is a measure of the cosmic variance on this scale.

haloes in the subvolume, and (2) a volume-weighted average ($\langle z_{re} \rangle_v$), obtained by averaging over the reionization redshifts of all coarse-grained cells in the subvolume. The distributions of these two means across our subvolumes is shown in the top panel of Fig. 6.

As the histograms show, the distribution of volume-averaged reionization redshifts is skewed towards later redshifts than those of the halo-weighted averages. This is made clear by the plot of $\langle z_{re} \rangle_v$ versus $\langle z_{re} \rangle_h$ for each subvolume, in the bottom panel of Fig. 6. This trend is to be expected, since the effects of reionization tend to propagate 'inside-out', from the neighbourhoods of clustered galaxies to the surrounding, larger volumes of the IGM. Nevertheless, Fig. 6 makes it clear that, even after averaging over the full range of local reionization redshifts within a given survey volume, those survey volumes are small enough that there is still a large variation in this average reionization redshift from one survey volume to another. This means we should expect there to be a corresponding scatter amongst the UVLF's derived for different survey volumes of this size.

We plot the $z = 6$, 7, 8 UVLFs of our subvolumes in Fig. 7. The curves for each subvolume are coloured according to their $\langle z_{re} \rangle_h$, with cyan corresponding to earlier-reionizing regions, magenta corresponding to later-reionizing regions, and blue in between. Naturally, the earlier-reionizing regions have higher UVLFs, and so on.

For comparison, we also plot the HFF observational results from Livermore et al. (2017) (red circles), Bouwens et al. (2022) (blue squares), Ishigaki et al. (2018) (brown diamonds), and Atek et al. (2018) (green triangles).[3] Their inferred UVLFs are broadly

[3]Note that Bouwens et al. (2022), Ishigaki et al. (2018), and Atek et al. (2018) each analyse the full set of six HFF clusters, while Livermore et al. (2017) analyse a subset of two of them.

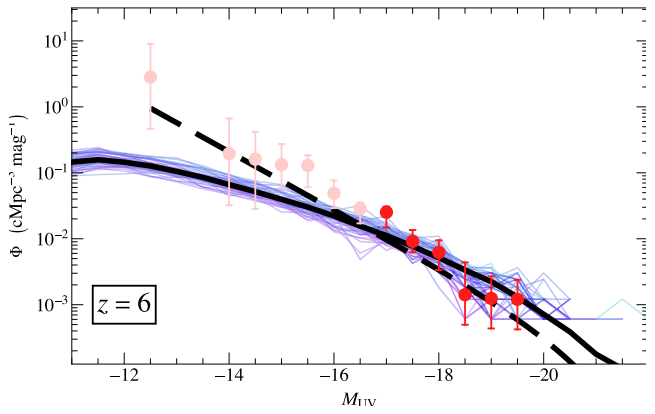

Figure 8. A comparison of the $z = 6$ UVLF data from Livermore et al. (2017) to the UVLFs in a selection of 50 CoDa II subvolumes – a subset of the 256 subvolumes shown in Fig. 7, chosen to most closely match the bright-end of the observational data (saturated red points). Despite their matching at the bright-end, the subvolumes are still discrepant with the data point at $M_{\rm UV} = -12.5$. For reference, the full volume CoDa II UVLF (thick solid line) and best-fitting Schechter function from Livermore et al. (2017) (thick dashed line) are also shown.

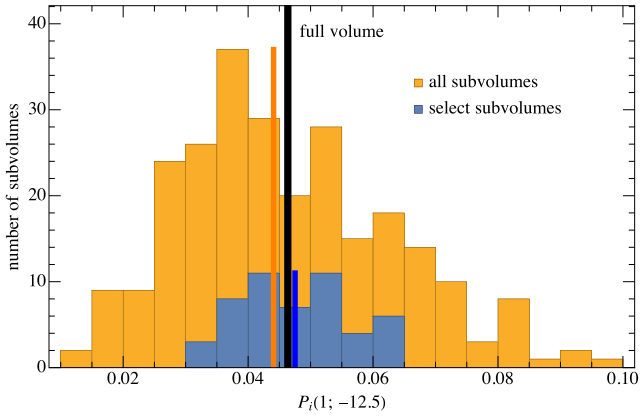

Figure 9. The number of subvolumes in which one $M_{\rm UV} = -12.5$ galaxy may be found in a random 0.73 cMpc^3 region with probability $P_i(1; -12.5)$. The orange bars represent all 256 subvolumes in CoDa II, while the blue bars represent only those select subvolumes that most closely match the bright-end HFF data from Livermore et al. (2017). The orange and blue lines mark the median probability for their colour-coordinated samples, while the black line marks the probability for the globally averaged CoDa II UVLF.

consistent with our CoDa II subvolumes, but there is one noteworthy discrepant data point from Livermore et al. (2017) at the faint-end at $z = 6$.[4] By this late redshift, most of the volume in our CoDa II simulation has been reionized, so most low-mass haloes in the subvolumes have been suppressed, and hence the faint-end slopes of the subvolumes' UVLFs are rather flat. However, due to the fact that they identified a single $M_{\rm UV} = -12.5$ galaxy at $z = 6$, Livermore et al. (2017) inferred a faint-end slope that remains steep down to this low magnitude. Galaxies that faint must be located very close to caustics in the lensing field, in order to be sufficiently magnified so as to be visible, so the effective volume searched for such a galaxy is much smaller than the full volume probed by the entire survey field. According to the lensing models used by Livermore et al. (2017), the effective volume searched for a $M_{\rm UV} = -12.5$ galaxy is around 0.73 cMpc^3. Therefore, identifying even a single galaxy in a randomly sampled volume that small implies a high UVLF at that galaxy's magnitude – high enough to be inconsistent with the average UVLF at that magnitude in all of our CoDa II subvolumes. We note that this data point is similarly discrepant with the UVLF predicted by THESAN (Kannan et al. 2022), another large-scale, high-resolution radiation-hydrodynamics simulation that is otherwise broadly consistent with high-z UVLF observations (as CoDa II is), though they do not explore the inhomogeneity of their UVLF.

To assess the degree of discrepancy between our CoDa II results and the observations of Livermore et al. (2017), we can estimate the probability of observing a single $M_{\rm UV} = -12.5$ galaxy in a randomly sampled 0.73 cMpc^3 region within each CoDa II subvolume. For this purpose, we identified a subset of 50 of our CoDa II subvolumes, chosen because their UVLF at brighter magnitudes most closely matches the UVLF of Livermore et al. (2017) at those magnitudes, for which the effective volume surveyed matches the subvolume size, as shown in Fig. 8. In particular, we selected the subvolumes with the 50 lowest χ^2 values when compared to the number of observed galaxies at each magnitude in the range $-19.5 \leq M_{\rm UV} \leq -17$. The data points at these magnitudes are indicated by the saturated red circles in the figure. Then, given the number density of $M_{\rm UV} = -12.5$ galaxies in the ith subvolume,

$$n_i(-12.5) \simeq \Phi_i(-12.5)\,\Delta M_{\rm UV}, \tag{3}$$

where $\Delta M_{\rm UV} = 0.5$ (as was used by Livermore et al. 2017), the probability of finding N such galaxies in a random search of a $V_{\rm eff}(-12.5) = 0.73$ cMpc^3 volume is given by the Poisson distribution[5]

$$P_i(N; -12.5) = \frac{\left(n_i(-12.5)\,V_{\rm eff}(-12.5)\right)^N e^{-n_i(-12.5)\,V_{\rm eff}(-12.5)}}{N!} \tag{4}$$

Thus, the probability of finding one such galaxy in the limit $n_i(-12.5)\,V_{\rm eff}(-12.5) \ll 1$ is

$$P_i(1; -12.5) \approx n_i(-12.5)\,V_{\rm eff}(-12.5). \tag{5}$$

We find that across all of the selected subvolumes, this probability falls in the range [3.2 per cent, 6.5 per cent], with a median probability of around 4.8 per cent. On the other hand, if we compute the probability across all subvolumes, not just the selected ones, we find a range of [1.3 per cent, 9.5 per cent], with a median around 4.4 per cent. Using the globally averaged CoDa II UVLF, instead, we find a probability of 4.6 per cent. We show a histogram of probability estimates for detecting a $M_{\rm UV} = -12.5$ galaxy in our subvolumes in Fig. 9.

3.2.1 *Cosmic variance in faint-end lensing surveys*

It is worth noting that there is significant uncertainty in the analysis of these extremely high-magnification faint galaxies, due to uncertainties in both the lensing models and cosmic variance, since the

[4]While the second-faintest point from the Bouwens et al. (2022) results at $z = 8$ is also somewhat discrepant with our results, since both data points on either side of it are not, we consider this to be anomalous and do not discuss it further here.

[5]Note that there are at least 143 such galaxies in each of the select $\sim$ 3300 cMpc^3 subvolumes, and around 209 on average. When considering all subvolumes, not just those selected to most closely match the bright-end of Livermore et al. (2017), there are at least 58 such galaxies in each, and around 215 on average.

effective volumes searched are so small. With regards to the former, the analysis of Bouwens et al. (2017), for example, re-interprets the $M_{\rm UV} = -12.5$ galaxy as a brighter galaxy ($M_{\rm UV} = -14.25$) in a larger effective search volume, with a greater uncertainty attributed to the lensing models. With regards to the latter, Livermore et al. (2017) account for uncertainty due to cosmic variance when fitting their data to Schechter functions by following the work of Robertson et al. (2014), which expresses the uncertainty in terms of the galaxy clustering bias, e.g.

$$\sigma_{\rm CV,L}(V) = \frac{\sqrt{\langle (n_{\rm L}(\boldsymbol{x}, V) - \bar{n}_{\rm L})^2 \rangle_{\boldsymbol{x}}}}{\bar{n}_{\rm L}} = b_{\rm L}(V)\,\sigma_{\rm DM}(V), \tag{6}$$

where $\sigma_{\rm CV,L}(V)$ is the fractional uncertainty due to cosmic variance for galaxies of luminosity (or magnitude) L observed in a volume V, $n_{\rm L}(\boldsymbol{x}, V)$ is the local number density of galaxies of the same luminosity in a volume V located at $\boldsymbol{x}$, $\bar{n}_{\rm L}$ is the global mean number density of galaxies of the same luminosity, $\langle\ \ \rangle_{\boldsymbol{x}}$ denotes a global average over all locations $\boldsymbol{x}$, $b_{\rm L}(V)$ is the bias of galaxies of the same luminosity on the scale of volume V, and $\sigma_{\rm DM}(V)$ is the linearly extrapolated RMS of dark matter density fluctuations on the same scale. Robertson et al. (2014) use, as their estimate of the bias, the analysis of dark matter simulations in Tinker et al. (2010). However, since the focus of Tinker et al. (2010) was on *large-scale* bias – i.e. in the limit where the bias is scale-independent, $b_{\rm L}(V) \to b_{\rm L}$ – their analysis was restricted to only the 5–10 largest wavelength modes in each simulation, with simulation box lengths in the range 80–1280 h^{-1} cMpc. Therefore, we believe their results will underestimate the bias on the *very small scales* probed by the extremely high-magnification regions of the HFF data, for which the bias is scale-dependent. As a consequence, the use of this bias in Robertson et al. (2014) will underestimate the uncertainty due to cosmic variance on such small scales, as well. Since detection of the faintest galaxies requires the greatest magnification, the use of observations like the Livermore et al. (2017) $M_{\rm UV} = -12.5$ galaxy to infer the UVLF is especially affected by this underestimation.

To illustrate this, we further sub-divide the CoDa II simulation at $z = 6$ into sets of 'sub-subvolumes' of different scales, all the way down to $V = 0.8$ cMpc^3, which is close[6] to the effective volume searched by Livermore et al. (2017) at $M_{\rm UV} = -12.5$. We then computed the bias of galaxies in the magnitude bin $M_{\rm UV} = -12.5 \pm 0.25$ in each of these sets of sub-subvolumes by computing the variance in their number densities to obtain $\sigma_{\rm CV,L}$, and show the results as a function of volume in Fig. 10 (black points). For comparison, we show the bias as estimated from the (large-scale) analysis of Tinker et al. (2010) for galaxies at this magnitude as a horizontal black line. We obtained this estimate by applying the Tinker et al. (2010) bias fitting function to the halo mass obtained by abundance matching the Tinker et al. (2008) HMF with the Livermore et al. (2017) UVLF at $M_{\rm UV} = -12.5$. As expected, our bias estimates from CoDa II approach the Tinker et al. (2010) estimate on large scales, but deviate substantially from the latter on smaller scales. In particular, near the effective volume searched by Livermore et al. (2017) at $M_{\rm UV} = -12.5$, our bias ($b_{-12.5}$) is larger than the (large-scale) Tinker et al. (2010) estimate by a factor of 4.2. Correspondingly, our estimate of the contribution to the uncertainty due to cosmic variance from an observation at this volume and

[6]We obtained the sub-subvolumes by evenly sub-dividing the 3300 cMpc^3 subvolumes into a sequence of smaller volumes with integer numbers of grid cells per dimension, and arrived at 3300/4096 = 0.8 cMpc^3 as the closest subdivision to $V_{\rm eff} = 0.73$ cMpc^3.

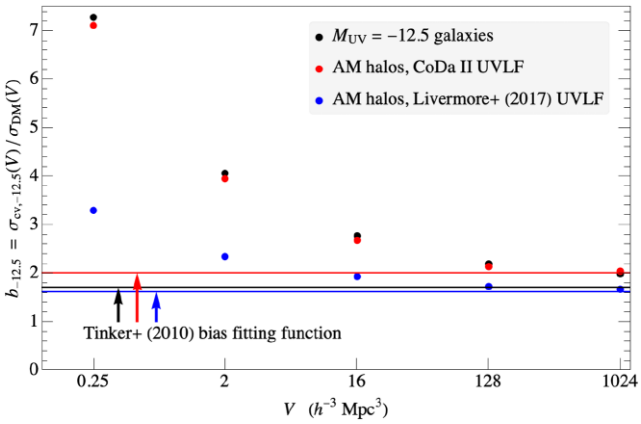

Figure 10. The bias of $M_{\rm UV} = -12.5$ galaxies in CoDa II at $z = 6$, computed by counting galaxies/haloes in volumes of size V. The smallest volume shown, $V = 0.25\ h^{-3}$ cMpc$^3 = 0.8$ cMpc^3, is similar to the effective volume searched by Livermore et al. (2017) at this magnitude, while the largest volume shown, $V = 1024\ h^{-3}$ cMpc$^3 = 3300$ cMpc^3, corresponds to our subvolumes. The black points are obtained by counting galaxies in the magnitude bin $M_{\rm UV} = -12.5 \pm 0.25$ directly, whereas the red and blue points are obtained by counting haloes with masses obtained by abundance matching (AM) to the magnitude bin, using the CoDa II HMF and the UVLFs from CoDa II and Livermore et al. (2017), respectively. For comparison, we show the bias as computed using the fitting function from Tinker et al. (2010) as horizontal lines. The black line is the bias obtained by abundance matching using the HMF from Tinker et al. (2008) and the UVLF from Livermore et al. (2017). The red line is obtained by abundance matching using the CoDa II HMF and UVLF, which is analogous to the red points. The blue line is obtained by abundance matching using the CoDa II HMF and Livermore et al. (2017) UVLF, which is analogous to the blue points. Our results diverge from the Tinker et al. (2010) fitting function estimates at small volumes.

magnitude ($\sigma_{\rm CV,-12.5}$) is a factor of 4.2 larger than that which Robertson et al. (2014) (and, following the former, Livermore et al. 2017) would have found.

This result compares the variance of galaxies in a given *luminosity* range in the CoDa II simulation to that of haloes of a given *mass* according to the scale-independent bias estimate from Tinker et al. (2010). In order to perform a more 'apples-to-apples' comparison (i.e. 'mass-to-mass', rather than 'luminosity-to-mass'), we also computed the bias in two different ways, for both our simulation results and the Tinker et al. (2010) estimate. For our simulation results, rather than calculating the variance in the number density of $M_{\rm UV} = -12.5 \pm 0.25$ galaxies directly, we instead calculate the variance in haloes with masses in a range obtained by abundance matching to this magnitude bin, i.e.

$$\sigma_{\rm CV,M}(V) = \frac{\sqrt{\langle (n_{\rm M}(\boldsymbol{x}, V) - \bar{n}_{\rm M})^2 \rangle_{\boldsymbol{x}}}}{\bar{n}_{\rm M}} = b_{\rm M}(V)\,\sigma_{\rm DM}(V), \tag{7}$$

where the subscript M now refers to the same quantities as before – uncertainty due to cosmic variance, local and global number densities, and bias – but for haloes with mass M, rather than galaxies with a given luminosity. We perform the abundance matching using our numerical CoDa II HMF and either the CoDa II UVLF (red points) or the Livermore et al. (2017) UVLF (blue points). Then, we apply these abundance-matched halo masses to the Tinker et al. (2010) fitting function to obtain new estimates, which are shown in Fig. 10 as red and blue horizontal lines, to be compared with the red and blue points, respectively. We illustrate the relationship between these different abundance matching methods in Fig. 11.

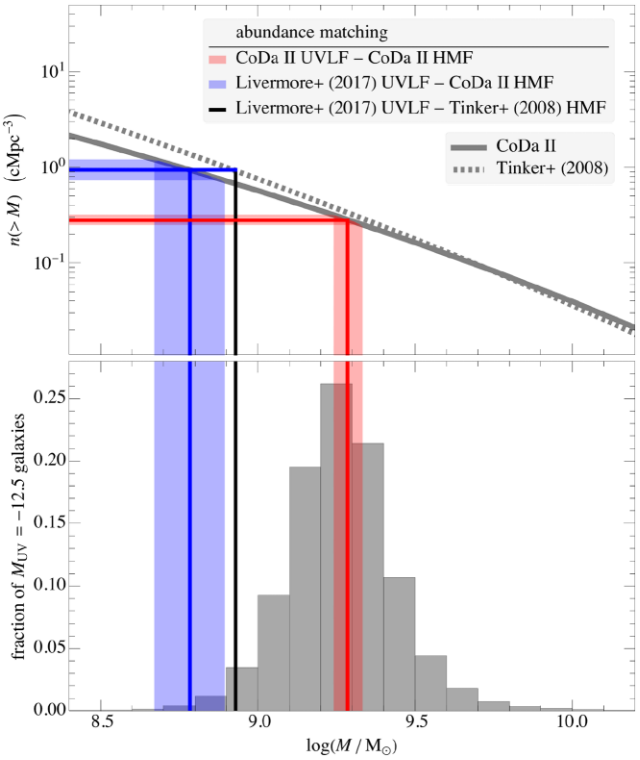

Figure 11. Top: Cumulative number densities of haloes with mass $>M$ for CoDa II (solid grey) compared to that obtained from the HMF of Tinker et al. (2008) (dashed grey), both of which are used for abundance matching. The horizontal red and blue bands represent the cumulative number densities of galaxies with magnitudes brighter than the bin -12.5 ± 0.25, according to the UVLFs from CoDa II and Livermore et al. (2017), respectively. The vertical red and blue bands represent the mass ranges obtained by matching the aforementioned galaxy number densities with the CoDa II cumulative halo number density curve. These mass ranges are used to compute the red and blue points in Fig. 10. The red and blue lines represent abundance matching for the centre of the magnitude bin ($M_{\rm UV} < -12.5$), and the resulting masses are used to obtain the red and blue horizontal lines in Fig. 10. The black vertical line is obtained by instead matching the Livermore et al. (2017) abundance of $M_{\rm UV} < -12.5$ galaxies to the Tinker et al. (2008) cumulative halo number density curve, and the resulting mass is used to obtain the black horizontal line in Fig. 10. Bottom: The distribution of halo masses for $M_{\rm UV} = -12.5 \pm 0.25$ galaxies in CoDa II, which is encoded in the black points in Fig. 10, compared to the abundance-matched masses discussed above.

As before, the points converge to the Tinker et al. (2010) estimate on large scales, but diverge on small scales. Since the red points show the cosmic variance in haloes that are self-consistently abundance-matched to the galaxies represented by the black points, they naturally exhibit roughly the same behaviour as the black points, and are a factor of 3.5 times higher than the corresponding Tinker et al. (2010) estimate at $V = 0.25\,h^{-3}$ cMpc^3. Although the blue points (haloes chosen by abundance matching the CoDa II HMF to the Livermore et al. 2017 UVLF, rather than the CoDa II UVLF) exhibit less of a discrepancy, our CoDa II result is still a factor of 2 larger than the comparable fitting function estimate on this small scale. Thus, we find that the procedure adopted by Livermore et al. (2017) to model their uncertainty due to cosmic variance underestimates the contribution from their observation of a $M_{\rm UV} = -12.5$ galaxy in an effective volume of 0.73 cMpc^3 by at least a factor of 2.

For future small-effective-volume searches that require cosmic variance estimates, we suggest computing the variance directly from simulations, on the scale of the effective search volume, as we have done here, rather than using fitting functions that are only applicable in the large-scale limit. Even using dark-matter-only simulations for this purpose will provide a much more accurate estimate of the uncertainty due to cosmic variance than the latter approach. As we discussed in the introduction and Section 3.1, cosmic variance in the UVLF is a result of both variance in the HMF and variance in the SFRs of haloes of a given mass. While dark-matter-only simulations cannot account for the latter, the proximity of the red (variance in haloes of a given mass) and black (variance in galaxies of a given luminosity) points in Fig. 10 indicates that the variance in SFRs is a sub-dominant contributor to the deviation of our result from the large-scale fitting function, since this variance is accounted for in the black points but not the red ones. Thus, most of the deviation is captured just by accounting for variance in dark matter halo abundances on small scales, which can be approximated from dark-matter-only simulations.

3.3 Fitting functions for the CoDa II UVLF

To parametrize the shape of our CoDa II UVLFs, it is useful to fit them to Schechter functions, or modifications thereof. The Schechter function is defined as

$$\Phi_{\rm L} = \left(\frac{\Phi_*}{L_*}\right)\left(\frac{L}{L_*}\right)^{\alpha} \exp\left(-\frac{L}{L_*}\right), \tag{8}$$

where Φ_*, L_*, and α are free parameters, the last of which is the logarithmic slope of the faint-end. By convention, the Schechter function is usually reparametrized in terms of magnitude as

$$\Phi_{\rm M} = 0.4\ln(10)\Phi_* 10^{0.4(M_*-M_{\rm UV})(\alpha+1)} \exp\left[-10^{0.4(M_*-M_{\rm UV})}\right]. \tag{9}$$

We fit this function to our $z = 6$ CoDa II UVLF in the top panel of Fig. 12 (thin black line). As can be seen there, the shape of the Schechter function fails to capture the curvature of the CoDa II UVLF, due to the fact that the Schechter function maintains a constant power-law slope at the faint end, while the CoDa II UVLF flattens out gradually towards fainter magnitudes, due to the suppression of star formation in low-mass haloes caused by reionization feedback. We can construct a better fit by allowing the power-law slope to change above a certain magnitude, using what we will refer to as the Schechter + turnover function from Jaacks, Thompson & Nagamine (2013):

$$\Phi_{\rm T} = \Phi_{\rm M}\left(1 + 10^{0.4(M_T - M_{\rm UV})\beta}\right)^{-1}, \tag{10}$$

where M_T is the so-called turnover magnitude, above which the power-law slope changes (in luminosity space) from α to $\alpha - \beta$. The fit of this function to the CoDa II UVLF is shown as the red dashed line in Fig. 12. The fitted model exhibits a turnover at

$$M_T = -16.9 \pm 0.4 \tag{11}$$

from a power-law slope of

$$\alpha = -2.23 \pm 0.05 \tag{12}$$

to

$$\alpha - \beta = -1.36 \pm 0.06 \tag{13}$$

Given this turnover magnitude, we can see more clearly the deviation of our CoDa II UVLF from the Schechter function, by fitting only the bright-end of the UVLF ($M_{\rm UV} < -17$) to a Schechter function, and extrapolating that bright-end fit faint-ward. This is shown as the black hashed line in the figure. We can see that the bright-end fit

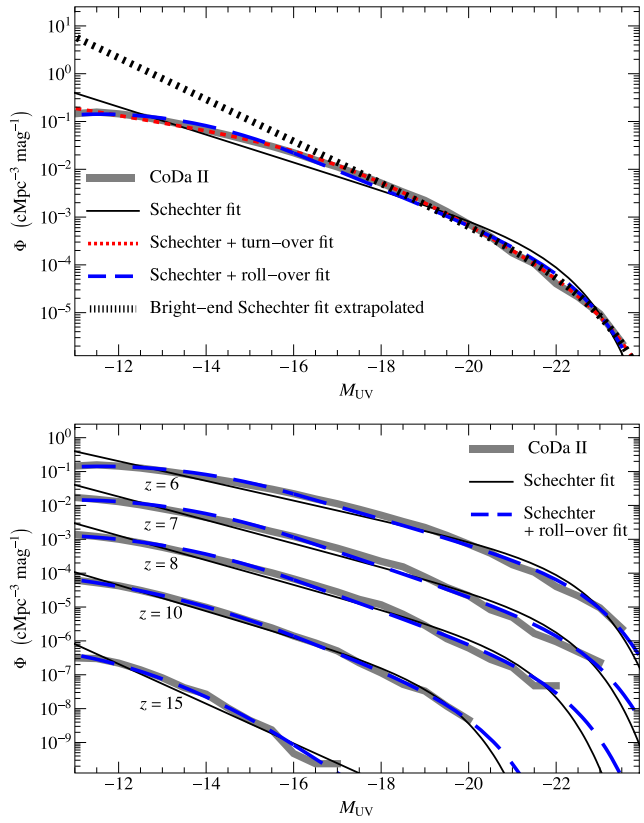

Figure 12. Top: A comparison of various fitting functions to the CoDa II full volume UVLF, as labelled. The black hashed line is a Schechter function fit only to magnitudes bright-wards of −17, but extrapolated faint-wards from there, to illustrate the degree to which the CoDa II UVLF deviates from the faint-end power-law behaviour of the Schechter function. Bottom: Pure Schechter and Schechter + roll-over fits at different redshifts. Curves for $z =$ 7, 8, 10, 15 are shifted downward by 1, 2, 3, and 4 dex, respectively.

clearly starts to deviate from the faint-end of the UVLF for $M_{\rm UV} \gtrsim -16$.

Another modification of the Schechter function was proposed by Bouwens et al. (2017), wherein the faint end gradually 'rolls' over at magnitudes fainter than −16, taking on a smoothly varying parabolic shape, rather than another power law. We will refer to this as the Schechter + roll-over function, which is defined as

$$\Phi_{\rm R} = \begin{cases} \Phi_{\rm M}, & M_{\rm UV} < -16 \\ \Phi_{\rm M} 10^{-0.4(-16-M_{\rm UV})^2 \delta}, & M_{\rm UV} > -16 \end{cases} \quad (14)$$

The advantage of this fitting function is that it captures the continued change in slope that occurs in the UVLF towards fainter magnitudes. We show this fit as the blue dashed line in Fig. 12. In addition to $z =$ 6, we also show the pure Schechter and Schechter + roll-over fits to the CoDa II UVLF at $z = 7$, 8, 10, 15 in the bottom panel of Fig. 12. With the exception of $z = 15$ (where our data are most limited), the pure Schechter fit gets worse with decreasing z, due to the increasing suppression of low-mass haloes as reionization progresses.

Amongst the previously discussed HFF analyses, Livermore et al. (2017) and Ishigaki et al. (2018) preferred pure Schechter function fits, as their data showed no sign of a turnover at the faint-end, while Bouwens et al. (2022) and Atek et al. (2018) preferred Schechter + roll-over fits. We show each of their best-fitting functions, with corresponding error, compared to our CoDa II UVLF in Fig. 13, and provide the best-fitting parameter estimates in Table 5. As can be seen, CoDa II has a clear preference for the roll-over models, with parameters that are most similar to those of Bouwens et al. (2022).

4 CONCLUSIONS

In this work, we analysed the CoDa II simulation to study the spatial and temporal variations in the high-redshift UVLF during the EOR. We find that the UVLF is strongly correlated with the local reionization history of the region in which it is measured. Earlier-reionizing regions, which have higher overdensities and HMFs, have correspondingly higher UVLFs with brighter exponential cut-off magnitudes (M_*). Therefore, the results of small-volume observations of the high-z UVLF, e.g. those made through high-magnification gravitational lenses, will depend on where one looks (e.g. at an early-, intermediate-, or late-reionizing patch of the Universe). In addition, such observations will also depend on *when* one looks – i.e. before or after the observed region has been reionized. Due to the fact that the photoheating of gas in the IGM during reionization suppresses the formation of stars in low-mass haloes, the faint-end of the UVLF in a given region evolves over the course of its local reionization, becoming increasingly flattened over time. The UVLF of a region observed after its local reionization will exhibit a turnover or roll-over at the faint-end, at magnitudes $\gtrsim -17$. By $z = 6$, when most of the Universe is reionized, the global UVLF exhibits this faint-end turnover, as well, and the gradual flattening of the global faint-end slope over time can be seen starting from $z =$ 10.

We find that our CoDa II UVLFs are in good agreement with data from HFF lensing surveys, with the exception of a single data point from Livermore et al. (2017) of a $M_{\rm UV} = -12.5$ galaxy observed at $z = 6$ in an effective search volume of $V_{\rm eff} = 0.73$ cMpc^3. This observation implies a UVLF that is ∼10 times higher at this magnitude than what we predict from CoDa II, as well as a faint-end slope that remains steep down to $z = 6$. This motivated us to ask about the variation of the UVLF with position in CoDa II, to determine the likelihood of the faint-end detection reported by Livermore et al. (2017) given the limitations of the search technique involving gravitational lensing amplification by a foreground cluster. Given the abundance of $M_{\rm UV} = -12.5$ galaxies in our CoDa II simulation, the probability of encountering one in a randomly placed 0.73 cMpc^3 search volume at that redshift is found to be relatively small, at ∼ 4.5 per cent. As such, our results are more consistent with the analysis of Bouwens et al. (2017), wherein this observation is re-interpreted as a brighter galaxy in a larger effective search volume, with a greater uncertainty attributed to the lensing models. Indeed, of the four observational papers whose UVLFs we compare to our simulated one, our results are most similar to those of Bouwens et al. (2022).

Furthermore, given the stark differences in galaxy abundances on small scales resulting from spatial and temporal variations in density and reionization history, we believe that the uncertainty due to cosmic variance attributed to these observations – when using them to infer the global UVLF – has been underestimated. The uncertainty due to cosmic variance for lensing surveys like the HFFs is typically estimated according to Robertson et al. (2014), in which it is proportional to the galaxy clustering bias derived from the large-scale N-body simulations and halo bias analysis of Tinker et al. (2010). However, for these small-scale lensing observations, we find that the bias of $M_{\rm UV} = -12.5$ galaxies in 0.8 cMpc^3 volumes is a factor of 2–4 higher than this large-scale estimate, primarily due

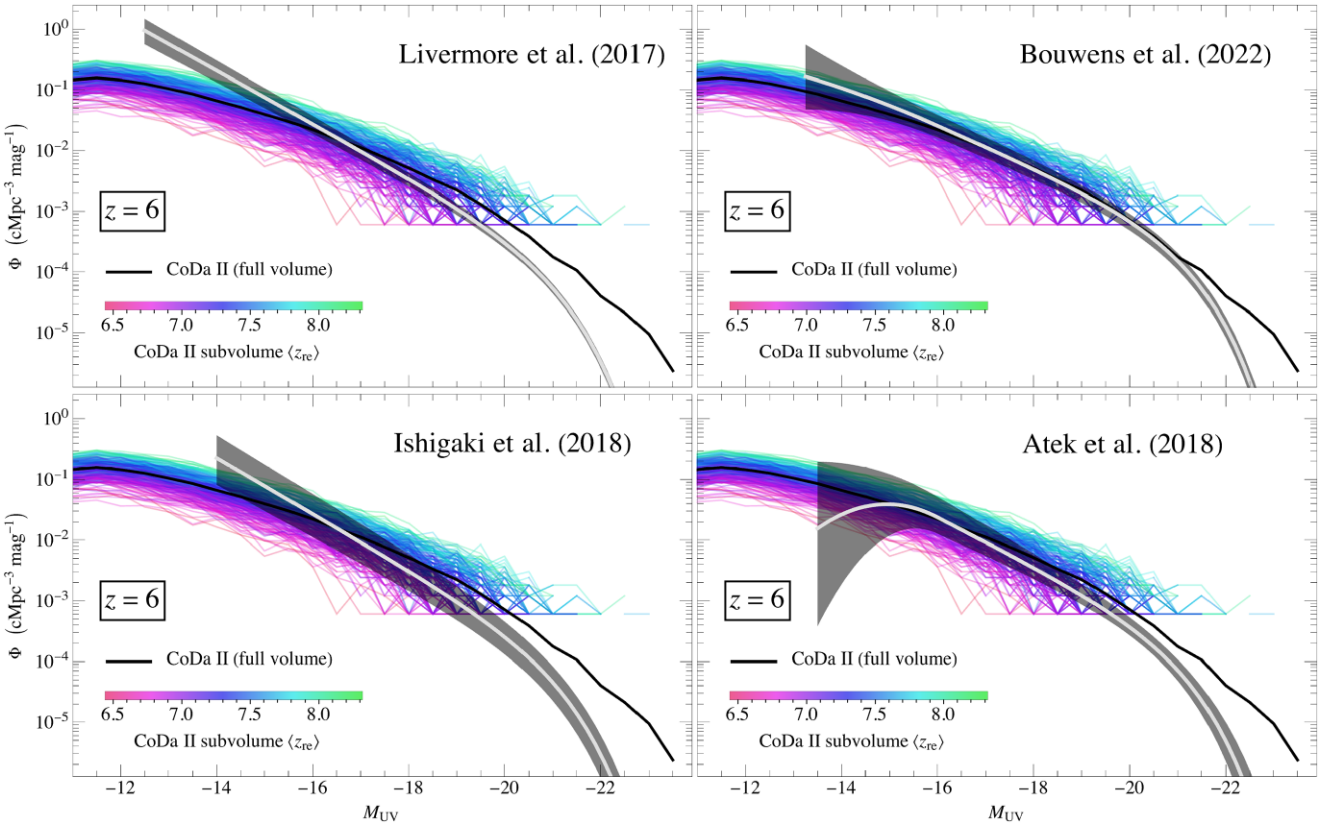

Figure 13. Comparison of our $z = 6$ CoDa II UVLF with the fitting functions (light grey lines with errors shown in dark shading) from Livermore et al. (2017); Atek et al. (2018); Ishigaki et al. (2018); Bouwens et al. (2022). Livermore et al. (2017); Ishigaki et al. (2018) prefer pure Schechter function fits, while Atek et al. (2018); Bouwens et al. (2022) prefer Schechter functions modified to roll-over at the faint end. Our CoDa II results favour the latter.

Table 5. Best-fitting (modified) Schechter function parameters for CoDa II's globally-averaged UVLFs versus those inferred from HFF data by four studies. The HFF results are for $z = 6$.

	$\log \Phi_*$	M_*	α	δ
Livermore et al. (2017)	$-3.647^{+0.039}_{-0.037}$	$-20.825^{+0.055}_{-0.043}$	$-2.10^{+0.05}_{-0.04}$	–
Bouwens et al. (2022)	-3.24 ± 0.08	-20.87 ± 0.07	-1.87 ± 0.04	0.05 ± 0.10
Ishigaki et al. (2018)	$-3.78^{+0.15}_{-0.15}$	$-20.89^{+0.17}_{-0.13}$	$-2.15^{+0.08}_{-0.06}$	–
Atek et al. (2018)	$-3.54^{+0.06}_{-0.07}$	$-20.84^{+0.27}_{-0.30}$	$-2.01^{+0.12}_{-0.14}$	$0.48^{+0.49}_{-0.25}$
CoDa II ($z = 6$)	-3.95 ± 0.086	-22.3 ± 0.099	-1.92 ± 0.026	0.105 ± 0.011
CoDa II ($z = 7$)	-4.60 ± 0.16	-22.3 ± 0.21	-2.08 ± 0.039	0.106 ± 0.016
CoDa II ($z = 8$)	-4.74 ± 0.22	-21.6 ± 0.30	-2.14 ± 0.046	0.100 ± 0.016
CoDa II ($z = 10$)	-4.67 ± 0.11	-19.8 ± 0.13	-2.23 ± 0.026	0.080 ± 0.007
CoDa II ($z = 15$)	-6.50 ± 0.08	-17.8 ± 0.07	-3.06 ± 0.147	0.146 ± 0.034

to the increased variance in dark matter halo abundances on small scales versus large scales. As next-generation space- and ground-based telescopes start to probe the high-z Universe down to fainter magnitudes than ever before, we expect a more thorough accounting of cosmic variance – one that accommodates the full scope of spatial and temporal inhomogeneity during the EOR – to be essential for reconciling competing observational inferences and theoretical models.

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation (NSF) Graduate Research Fellowship Program under Grant No. DGE-1610403. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF. This research used resources of the Oak Ridge Leadership Computing Facility (OLCF) at the Oak Ridge National Lab, which is supported by the Office of Science of the US Dept. of Energy (DOE) under Contract No. DE-AC05-00OR22725. CoDa II was performed on Titan at OLCF under DOE INCITE 2016 award to Project AST031. Data analysis was performed using the computing resources from NSF XSEDE grant TG-AST090005 and the Texas Advanced Computing Center at the University of Texas at Austin. PRS acknowledges support from NASA under Grant No. 80NSSC22K175. JL acknowledges support from the German Research Foundation via the Heidelberg Cluster of Excellence STRUCTURES in the framework of Germany's Excellence Strategy (grant EXC-2181/1-390900948). JS acknowledges support

from the ANR LOCALIZATION project, grant ANR-21-CE31-0019 of the French Agence Nationale de la Recherche. KA is supported by Korea NRF-2021R1A2C1095136, NRF-2016R1A5A1013277, and RS-2022-00197685. KA also appreciates the Asia Pacific Center for Theoretical Physics for its hospitality during completion of this work. ITI was supported by the Science and Technology Facilities Council [grant numbers ST/I000976/1 and ST/T000473/1] and the Southeast Physics Network. HP was supported by the World Premier International Research Center Initiative, MEXT, Japan and JSPS KAKENHI grant No. 19K23455. GY acknowledges financial support from the Ministry of Science and Innovation/FEDER under project grant PGC2018-094975-C21.

DATA AVAILABILITY

The data underlying this article are available in the article and in its online supplementary material.

REFERENCES

Atek H., Richard J., Kneib J.-P., Schaerer D., 2018, MNRAS, 479, 5184
Aubert D., Teyssier R., 2008, MNRAS, 387, 295
Aubert D. et al., 2018, ApJ, 856, L22
Bosman S. E. I. et al., 2022, MNRAS, 514, 55
Bouwens R. J., Oesch P. A., Illingworth G. D., Ellis R. S., Stefanon M., 2017, ApJ, 843, 129
Bouwens R. J., Illingworth G., Ellis R. S., Oesch P., Stefanon M., 2022, ApJ, 940, 55
Dawoodbhoy T. et al., 2018, MNRAS, 480, 1740
Dayal P., Ferrara A., 2018, Phys. Rep., 780–782, 1
Deparis N., Aubert D., Ocvirk P., Chardin J., Lewis J., 2019, A&A, 622, A142
Eldridge J. J., Stanway E. R., Xiao L., McClelland L. A. S., Taylor G., Ng M., Greis S. M. L., Bray J. C., 2017, Publ. Astron. Soc. Austr., 34, e058
Finlator K., Davé R., Özel F., 2011, ApJ, 743, 169
Gnedin N. Y., 2014, ApJ, 793, 29
Gnedin N. Y., Fan X., 2006, ApJ, 648, 1
González M., Audit E., Huynh P., 2007, A&A, 464, 429
Iliev I. T., Mellema G., Ahn K., Shapiro P. R., Mao Y., Pen U.-L., 2014, MNRAS, 439, 725
Ishigaki M., Kawamata R., Ouchi M., Oguri M., Shimasaku K., Ono Y., 2018, ApJ, 854, 73
Jaacks J., Thompson R., Nagamine K., 2013, ApJ, 766, 94
Kannan R., Garaldi E., Smith A., Pakmor R., Springel V., Vogelsberger M., Hernquist L., 2022, MNRAS, 511, 4005
Kroupa P., 2001, MNRAS, 322, 231
Kulkarni G., Choudhury T. R., 2011, MNRAS, 412, 2781
Lewis J. S. W. et al., 2022, MNRAS, 516, 3389
Livermore R. C., Finkelstein S. L., Lotz J. M., 2017, ApJ, 835, 113
Lotz J. M. et al., 2017, ApJ, 837, 97
Ocvirk P. et al., 2016, MNRAS, 463, 1462
Ocvirk P., Aubert D., Chardin J., Deparis N., Lewis J., 2019, A&A, 626, A77
Ocvirk P. et al., 2020, MNRAS, 496, 4087
Pallottini A., Ferrara A., Gallerani S., Vallini L., Maiolino R., Salvadori S., 2017, MNRAS, 465, 2540
Pawlik A. H., Rahmati A., Schaye J., Jeon M., Dalla Vecchia C., 2017, MNRAS, 466, 960
Planck Collaboration XVI 2014, A&A, 571, A16
Robertson B. E., Ellis R. S., Dunlop J. S., McLure R. J., Stark D. P., McLeod D., 2014, ApJ, 796, L27
Rosdahl J. et al., 2018, MNRAS, 479, 994
Sawala T., Frenk C. S., Crain R. A., Jenkins A., Schaye J., Theuns T., Zavala J., 2013, MNRAS, 431, 1366
Schechter P., 1976, ApJ, 203, 297
Shapiro P. R., Giroux M. L., Babul A., 1994, ApJ, 427, 25
Sheth R. K., Mo H. J., Tormen G., 2001, MNRAS, 323, 1
So G. C., Norman M. L., Reynolds D. R., Wise J. H., 2014, ApJ, 789, 149
Teyssier R., 2002, A&A, 385, 337
Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, ApJ, 688, 709
Tinker J. L., Robertson B. E., Kravtsov A. V., Klypin A., Warren M. S., Yepes G., Gottlöber S., 2010, ApJ, 724, 878
Watson W. A., Iliev I. T., D'Aloisio A., Knebe A., Shapiro P. R., Yepes G., 2013, MNRAS, 433, 1230
Xu H., Wise J. H., Norman M. L., Ahn K., O'Shea B. W., 2016, ApJ, 833, 84
Yoshida N., Hosokawa T., Omukai K., 2012, Progr Theoret. Exp. Phys., 2012, 01A305

APPENDIX A: COMPARISON OF THE CODA II GLOBAL HALO MASS FUNCTION TO STANDARD FITS FROM DARK-MATTER-ONLY N-BODY RESULTS

A comparison of the global CoDa II HMF at $z = 6$ to the fitting functions of Sheth, Mo & Tormen (2001); Tinker et al. (2008);

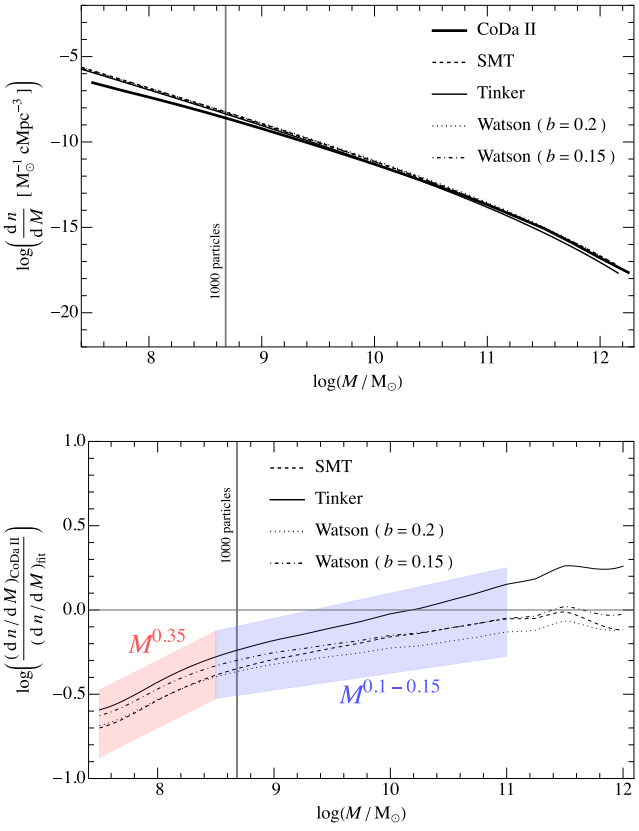

Figure A1. A comparison of the CoDa II HMF at $z = 6$ to the fitting functions from Sheth et al. (2001); Tinker et al. (2008); Watson et al. (2013). We show two versions of the Watson et al. (2013) fit, adopting Friends-of-Friends linking parameters of $b = 0.2$ and $b = 0.15$. The top panel shows the HMFs themselves, while the bottom panel shows the log ratio of the CoDa II HMF to the various fitting functions, as labelled. The shaded region in the bottom panel roughly highlights the trend in the difference between CoDa II and the dark-matter-only simulation fits. At the low-mass end ($M \lesssim 10^{8.5}$ M$_\odot$), we approach the resolution limit of our simulation, and the trend is that the ratios follow a power law with an index of ∼0.35. Above this mass, however, the trend flattens to an index of ∼0.1–0.15. This latter power law continues well above our resolution limit, until around $M \sim 10^{11}$ M$_\odot$. Therefore, we believe the latter difference is a physical consequence of including hydrodynamics in the simulation, rather than a numerical limitation.

Watson et al. (2013), derived from dark-matter-only ('DMO') *N*-body simulations, is shown in Fig. A1. The top panel shows the HMFs themselves, while the bottom panel shows the log of the ratios of the CoDa II HMF to each of these fitting functions, as labeled. The shaded region in the bottom panel roughly highlights the trend in the difference between CoDa II and the DMO simulation fits. There is a tendency for the CoDa II HMF to be a bit lower than the DMO fits over a broad range of masses, by an amount that is comparable to the spread in amplitudes of these different fits. This trend steepens, however, at the low-mass end, below $M \lesssim 10^{8.5}$ $M_\odot$. There, we approach the resolution limit of our simulation, and the trend for the ratios is that of a power law with a relatively steep index of ∼0.35. Above this mass, however, where CoDa II haloes contain more than ∼1000 DM particles, the trend flattens to an index of ∼ 0.1–0.15. This latter power law continues over orders of magnitude in halo mass, all well above our resolution limit, until around $M \sim 10^{11}$ $M_\odot$. Therefore, we believe this lowering of the HMF in CoDa II relative to the DMO fits is a physical consequence of including hydrodynamics in the simulation, rather than a numerical limitation. In short, since our results here for the inhomogeneity of the UVLF, including our use of the CoDa II HMF for abundance-matching down through the faint end of the LF, only depend on the HMF above $M \gtrsim 10^{8.5}$ $M_\odot$, they should be robust with respect to our numerical resolution of the HMF.

As shown here by analysing the CoDa II simulation, the self-consistent treatment of halo formation including baryonic feedback effects from star formation, supernovae, and reionization tends to reduce the HMF relative to that predicted from DMO simulations by a modest amount. This trend was demonstrated previously by Sawala et al. (2013), by directly comparing their GIMIC hydrodynamical simulations with DMO simulations from the same initial conditions, for haloes all the way up to $M \lesssim 10^{12}$ $M_\odot$. As these authors reported, while the two types of simulation agreed well on large scales, objects below this mass scale had systematically lower masses in the GIMIC simulation (i.e. with baryonic hydrodynamics) than in the DMO simulation, resulting in a corresponding shift downwards in the HMF, by a larger amount for smaller mass haloes. This is consistent with the results found here for CoDa II. In fact, the CoDa II simulation strengthens the case for this trend, since it is based on fully coupled radiation-hydrodynamics (i.e. with radiative transfer), while the GIMIC simulation (with coarser particle-mass resolution than ours but comparable length resolution) adopted a uniform, optically thin photoionizing background, instead.

APPENDIX B: RELATED MEASURES OF THE UVLF

We present here some additional quantities related to the UVLF, which may be useful for further theoretical or observational comparisons. In Fig. B1, we plot the number of galaxies in CoDa II that fall within each magnitude and z_{re} bin used throughout this paper. This figure amounts to a renormalization of the various curves shown in Fig. 5, to show the relative contributions of early-, intermediate-, and late-reionizing regions to the global UVLF. For instance, one feature that can be gleaned easily from this figure is that the magnitude at which our early-reionization bin transitions from the dominant contributor to a sub-dominant contributor becomes brighter over time. This is due to the fact that while the early-reionizing regions are relatively overdense, they also occupy a relatively small volume. Therefore, these regions form the dominant share of bright galaxies at early times, but are eventually overtaken as the less-overdense-but-larger-volume regions become increasingly non-linear at later times.

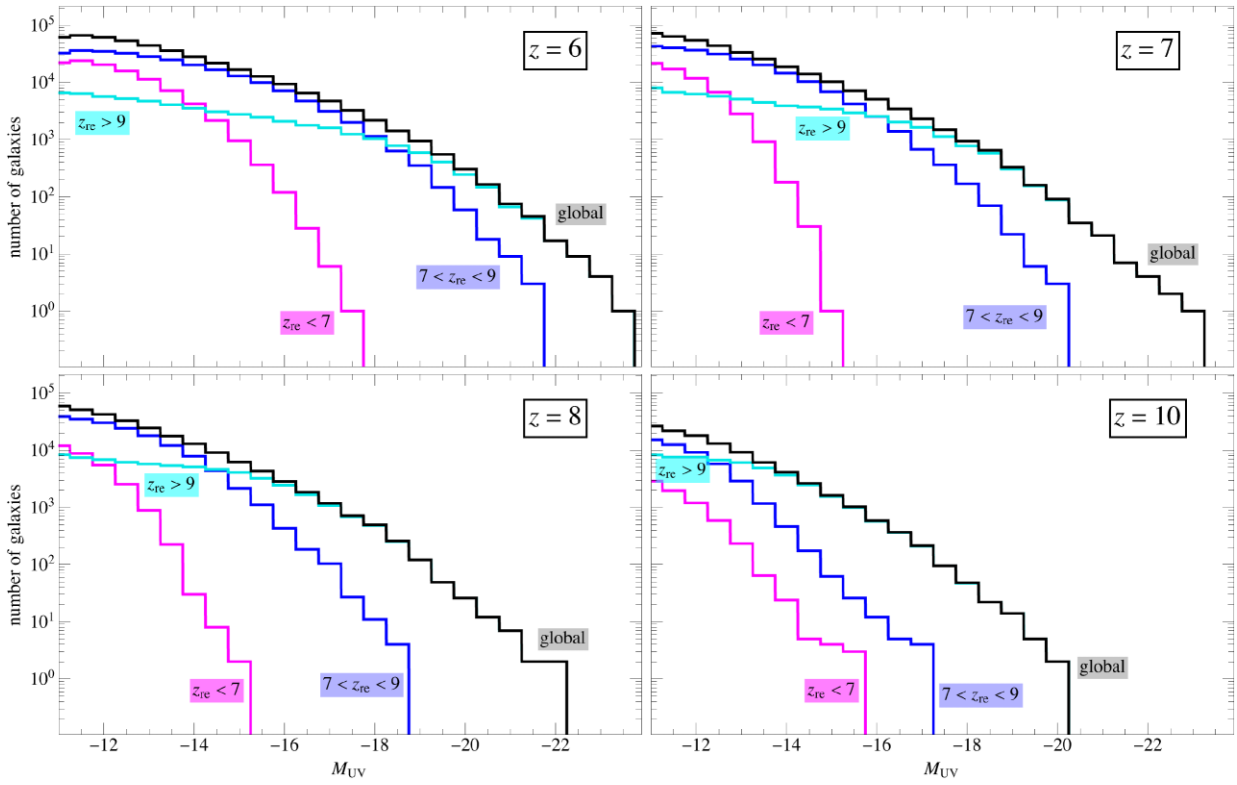

Figure B1. Histogram of CoDa II galaxies as a function of magnitude for the same z_{re} bins and redshifts as previous figures.

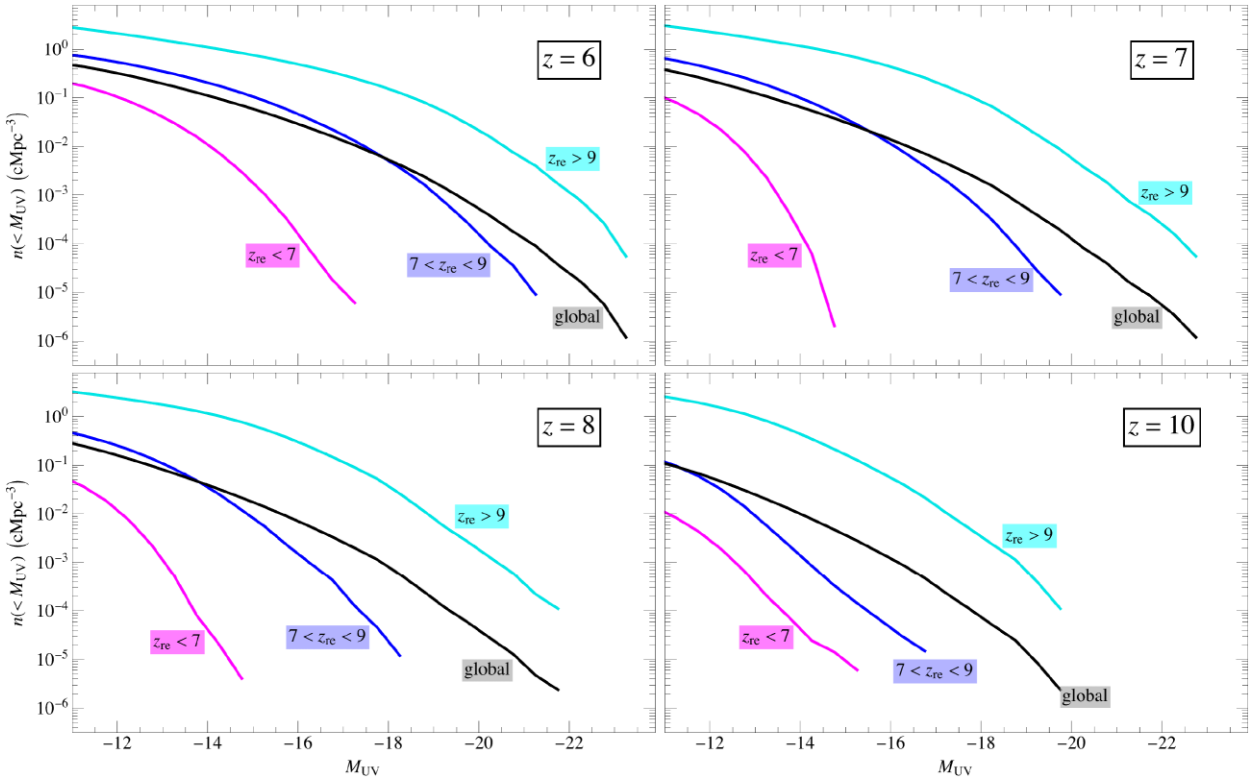

Figure B2. Cumulative number density of CoDa II galaxies brighter than a given magnitude for the same z_{re} bins and redshifts as previous figures.

This plot of the actual numbers of galaxies in the CoDa II simulation volume at each redshift also illustrates why a simulation with as large a volume as CoDa II is required in order to make a statistically meaningful analysis of the UVLF possible. There must be enough galaxies formed in the simulation volume to enable us to bin them in a multi-dimensional space at each redshift, not only by reionization redshift, but also by magnitude over a wide range, from as faint as -11 to as bright as -23. For example, in order to fit our UVLF to Schechter functions as in Table 5, we must have a large enough sample of galaxies in magnitude bins brighter than ~ -21 to model the exponential cut-off without much noise, while also maintaining fine bin-spacing. For our desired bin size of $\Delta M_{UV} = 0.5$, CoDa II contains $\sim$80 galaxies at $M_{UV} = -21$ at $z = 6$. A simulation with half the box size per dimension would, therefore, contain only $\sim$10, and would have almost no galaxies in bins brighter than -22. Thus, apart from their lack of realism in simulating reionization, simulations of smaller volume would undersample the bright end. While the number of galaxies is much higher at the faint end, it is also necessary to have the high mass-resolution of CoDa II in order to represent this number and their luminosities faithfully enough to establish the flattening of the UVLF relative to the bright end, as an effect of the suppression of low-mass haloes due to reionization.

In Fig. B2, we plot the cumulative number density of CoDa II galaxies brighter than a given magnitude, again split into z_{re} bins. The global line in this figure is used for abundance matching to the CoDa II HMF in Section 3.2.1.

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