Exploiting synergies between *JWST* and cosmic 21-cm observations to uncover star formation in the early Universe

Jiten Dhandha [1,2]★ Anastasia Fialkov,[1,2] Thomas Gessey-Jones [2,3,4] Harry T. J. Bevins [2,3]
Sandro Tacchella [2,3] Simon Pochinda [2,3] Eloy de Lera Acedo,[2,3] Saurabh Singh[5] and Rennan Barkana[6]

[1]*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB30HA, UK*
[2]*Kavli Institute for Cosmology, Madingley Road, Cambridge CB30HA, UK*
[3]*Astrophysics Group, Cavendish Laboratory, J. J. Thomson Avenue, Cambridge CB30HE, UK*
[4]*PhysicsX, Victoria House, 1 Leonard Circus, London EC2A 4DQ, UK*
[5]*Raman Research Institute, Bangalore, Karnataka 560080, India*
[6]*School of Physics and Astronomy, Tel-Aviv University, Tel-Aviv 69978, Israel*

ABSTRACT

In the current era of *James Webb Space Telescope* (*JWST*), we continue to uncover a wealth of information about the Universe deep into the Epoch of Reionization. In this work, we use a suite of simulations with 21CMSPACE to explore the astrophysical properties of early galaxies and their imprint on high-redshift observables. Our analysis incorporates a range of multiwavelength data sets including ultraviolet luminosity functions (UVLFs) from *Hubble Space Telescope* and *JWST* spanning $z = 6$–14.5, the 21-cm global signal and power spectrum limits from SARAS 3 (Shaped Antenna measurement of the background RAdio Spectrum) and HERA (Hydrogen Epoch of Reionization Array), respectively, as well as present-day diffuse X-ray and radio backgrounds. We constrain a flexible halo-mass and redshift dependent model of star formation efficiency (SFE), defined as the fraction of gas converted into stars, and find that it is best described by minimal redshift evolution at $z \approx 6$–10, followed by rapid evolution at $z \approx 10$–15. Using Bayesian inference, we derive functional posteriors of the SFE, inferring that haloes of mass $M_h = 10^{10}\,\mathrm{M_\odot}$ have efficiencies of ≈ 1–2 per cent at $z \lesssim 10$, ≈ 8 per cent at $z = 12$, and ≈ 21 per cent at $z = 15$. We also highlight the synergy between UVLFs and global 21-cm signal from SARAS 3 in constraining the minimum virial conditions required for star formation in haloes. In parallel, we find the X-ray and radio efficiencies of early galaxies to be $f_X = 0.8^{+9.7}_{-0.4}$ and $f_r \lesssim 16.9$, respectively, improving upon previous works that exclude UVLF data. Our results underscore the critical role of UVLFs in constraining early galaxy properties, and their synergy with 21-cm and other mutliwavelength observations.

Key words: galaxies: high-redshift – galaxies: star formation – (cosmology:) dark ages, reionization, first stars – (cosmology:) early Universe .

1 INTRODUCTION

The early Universe remains an exciting frontier due to its observational and theoretical prospects. Observations of the cosmic microwave background (CMB) at redshift $z \approx 1100$ ($t_{\mathrm{age}} \sim 0.4\,\mathrm{Myr}$ after the big bang) have provided us with a snapshot of the Universe in its infancy, while large-scale spectroscopic surveys like Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013) and Dark Energy Spectroscopic Instrument survey (DESI; DESI Collaboration 2024) are mapping the nearby Universe in unprecedented detail. The epochs intervening these bookends of our cosmic history, however, remain largely unexplored. Recently, the *James Webb Space Telescope* (*JWST*; Gardner et al. 2006) has opened a new window into the distant Universe, with its high sensitivity and resolution, allowing us to probe some of the earliest galaxies at $z \gtrsim 10$ ($t_{\mathrm{age}} \lesssim 500\,\mathrm{Myr}$) deep into the Epoch of Reionization (EoR; Whitler et al. 2025). Prior

to the *JWST*, the *Hubble Space Telescope* (*HST*) had provided us with a wealth of information on the high-redshift Universe, including the ultraviolet luminosity functions (UVLF) of galaxies and the cosmic star formation history (e.g. Bouwens et al. 2021a).

Now, in the era of the *JWST*, hundreds of candidate galaxies have been observed at very early epochs (e.g. Austin et al. 2023; Castellano et al. 2023; Leung et al. 2023; Harikane et al. 2023a; Finkelstein et al. 2024; Hainline et al. 2024; Pérez-González et al. 2025; Whitler et al. 2025), with the highest spectroscopically confirmed galaxy JADES-GS-z14-0 at $z = 14.32$ (Carniani et al. 2024), alongside a host of distant supermassive black holes, bright AGN and quasars (Larson et al. 2023; Bogdán et al. 2024; Furtak et al. 2024; Maiolino et al. 2024; Yue et al. 2024). Indeed, the abundance of these bright galaxies at such high redshifts given the small survey volumes exceed those expected by simple extrapolations of the UVLF (e.g. fig. 10 of Finkelstein et al. 2024; or fig. 6 of Whitler et al. 2025), or based on theoretical predictions (e.g. fig. 17 of Harikane et al. 2023a; or fig. 9 of Whitler et al. 2025). Several suggestions have been made in the literature to resolve this discrepancy, including a top-heavy

★ E-mail: jvd29@cam.ac.uk

initial mass function (IMF) of early stars (e.g. Inayoshi et al. 2022; Wang et al. 2023; Trinca et al. 2024; Hutter et al. 2025; Mauerhofer et al. 2025), an increased star formation efficiency (SFE; Dekel et al. 2023; Li et al. 2024; Boylan-Kolchin 2025; Mauerhofer et al. 2025; Somerville et al. 2025), the presence of hidden AGN (Hegde, Wyatt & Furlanetto 2024), or star formation variability/stochasticity (e.g. Mason, Trenti & Treu 2023; Mirocha & Furlanetto 2023; Shen et al. 2023; Sun et al. 2023a, b; Gelli, Mason & Hayward 2024; Kravtsov & Belokurov 2024).

Despite its grand successes, the era of the very first generation of stars at $z \sim 20-50$ ($t_{\mathrm{age}} \sim 50-200$ Myr), the so-called metal-free Population III stars (Pop III), will likely be too faint for *JWST* to detect. The cosmological 21-cm signal, on the other hand, is a powerful probe of this epoch. The signal is a unique signature of neutral hydrogen in the intergalactic medium (IGM) that is sensitive to the thermal and ionization history of the Universe (Furlanetto, Oh & Briggs 2006; Pritchard & Loeb 2012) resulting from an interplay between the radiation sourced by stars, stellar remnants, or exotic sources. In particular, the 21-cm signal is sensitive to the properties of the CMB/radio background, early Lyα sources such as Pop III stars at Cosmic Dawn ($z \lesssim 30$, e.g. Yajima & Khochfar 2015; Schauer, Liu & Bromm 2019; Mebane, Mirocha & Furlanetto 2020; Gessey-Jones et al. 2022, 2025), heating of the IGM through X-ray photons and Lyα scattering ($z \lesssim 20$, e.g. Chen & Miralda-Escudé 2004; Pritchard & Furlanetto 2007; Fialkov, Barkana & Visbal 2014; Pacucci et al. 2014; Reis, Fialkov & Barkana 2021), and reionization from stellar/quasar ultraviolet emissions ($z \lesssim 10$, e.g. Madau, Meiksin & Rees 1997; Ciardi & Madau 2003).

Detection of this signal is, however, challenging due to its faintness compared to the bright Galactic foregrounds and instrumental noise. Current experiments aimed at the detection of the 21-cm signal typically focus on two summary statistics: the sky-averaged global signal such as EDGES (Bowman et al. 2018), SARAS (Singh et al. 2022), REACH (de Lera Acedo et al. 2022), MIST (Monsalve et al. 2024), RHINO (Bull et al. 2024), PRIZM (Philip et al. 2019); or the two-point correlation function of spatial variations (power spectrum) such as MWA (Tingay et al. 2013), LOFAR (van Haarlem et al. 2013), HERA (DeBoer et al. 2017), and others.[1] Future experiments will be able to go a step further and make 21-cm tomographic maps (e.g. SKA; Madau et al. 1997), or probe even deeper into the Dark Ages from moon-based missions (see e.g. Fialkov, Gessey-Jones & Dhandha 2024, and references therein).

In 2018, the EDGES collaboration reported a deeper than expected absorption signal at $z \approx 17$ (Bowman et al. 2018). This tentative detection spurred a flurry of exciting theoretical works on early Universe physics. To add to this mystery, measurements of the diffuse cosmic radio background (CRB) from ARCADE2 (Fixsen et al. 2011) and LWA1 (Dowell & Taylor 2018) hint at an excess above the known CMB spectrum and known Galactic/extragalactic sources. Standard astrophysical mechanisms of galactic emissions (Mirocha & Furlanetto 2019), either synchrotron emission from Pop III supernovae (Jana, Nath & Biermann 2019) or radio-loud accretion of supermassive blackholes (e.g. Ewall-Wice et al. 2018; Ewall-Wice, Chang & Lazio 2020) could provide such a diffuse contribution. They could also be sourced by exotic agents like annihilating dark matter (DM; Fraser et al. 2018), super-conducting cosmic

strings (Brandenberger, Cyr & Shi 2019), or primordial blackholes (Mittal & Kulkarni 2022; Acharya, Dhandha & Chluba 2022). The 21-cm signal, by its nature, is sensitive to the radio background; an excess there would naturally boost the amplitude of the signal (Feng & Holder 2018; Fialkov & Barkana 2019), thus explaining both the EDGES detection and radio measurements. Although promising, the EDGES detection has been disputed as potential instrumental systematics (Hills et al. 2018; Singh & Subrahmanyan 2019; Bradley et al. 2019; Sims & Pober 2020) and follow up observations from the SARAS 3 experiment failed to detect the signal (Singh et al. 2022).

Recently, Pochinda et al. (2024) conducted the most complete work to date utilizing multiwavelength synergies alongside the 21-cm signal to constrain properties of the early Universe. Their work uses the SARAS 3 non-detection of the 21-cm global signal (Singh et al. 2022), the latest HERA upper limits on the 21-cm power spectrum (HERA Collaboration 2023), and collated measurements of the CRB (Dowell & Taylor 2018) and cosmic X-ray background (CXB; Hickox & Markevitch 2006; Harrison et al. 2016). In this work, we extend their analysis by exploring a more flexible star formation model, and include the UVLF measurements from large *HST* and *JWST* surveys as an additional observable. In particular, we constrain the SFE of galaxies, an important quantity widely used in analytic galaxy formation models (e.g. Sun & Furlanetto 2016; Tacchella et al. 2018; Sipple & Lidz 2024), as well as the efficiency of X-ray and radio sources. In Section 2, we discuss the theoretical framework, including the enhanced SFE model introduced in this work, followed by an introduction of the observational data used in Section 3. We then describe our methodology for Bayesian inference in Section 4, and present our results in Section 5. We conclude in Section 6.

2 THEORY OF THE EARLY UNIVERSE

In order to simulate the observables from the early Universe, we use the code 21-cm Semi-numerical Predictions Across Cosmic Epochs (21CMSPACE[2]; Fialkov et al. 2012; Visbal et al. 2012). We briefly describe the simulations in Section 2.1, but otherwise refer the readers to the most recent code development papers for details (Gessey-Jones et al. 2023, 2025). In this work, we add two new features to the code: an upgraded Pop II SFE model (described in Section 2.2), and a calculation of UVLFs (described in Section 2.4) for constraining astrophysics using *HST*/*JWST* data. All simulations are performed assuming *Planck 2013* best-fitting Lambda cold dark matter (ΛCDM) cosmology (Planck Collaboration XVI 2014, 'Planck + WP'): $H_0 = 67.04$ km s^{-1} Mpc^{-1}, $\Omega_{\mathrm{b}} = 0.0490$, $\Omega_{\mathrm{c}} = 0.2678$, $n_s = 0.9619$.

2.1 Overview of 21CMSPACE

The simulations begin at $z_{\max} = 50$ with cosmological initial conditions generated using CAMB (Lewis & Challinor 2011) for large-scale Eulerian matter overdensity δ and baryon–DM relative velocity/streaming velocity v_{bc}, and RECFAST (Seager, Sasselov & Scott 2011) for gas kinetic temperature T_{K} and ionized fraction x_{e}. The initial conditions are set over a cosmological box split into 128^3 cubic cells, assuming periodic boundary conditions, and a resolution of 3 comoving Mpc, giving a total volume of $(384 \, \mathrm{cMpc})^3$. The fields δ and v_{bc} are evolved forward in time using linear perturbation theory (see, e.g. Barkana 2016), while T_K and x_{e} are evaluated at each step as we describe below. Halo formation is modelled analytically in each pixel using a hybrid DM

---

[1]For a comprehensive list of 21-cm experiments, see github.com/JitenDhandha/21cmExperiments, a community-driven public resource which the author invites the readers to contribute to. The list contains details of individual telescopes, their status and associated bibliography for each.

[2]For an overview, see cosmicdawnlab.com/21cmSPACE.

halo mass function (HMF) $dn(M_h, z|\delta, v_{bc})/dM_h$ (Barkana & Loeb 2004; Tseliakhovich, Barkana & Hirata 2011; Fialkov et al. 2012) that depends on the large-scale density fluctuations and streaming velocity resolved at 3 cMpc scales. This large-scale overdensity is capped at $\delta_{max} = 1.3$, which is technically quasi-linear but still below the critical overdensity $\delta_{crit} \approx 1.69$ often used in HMF formalisms. This particular value was chosen to preserve physical validity of the hybrid HMF we employ.[3]

We model star formation to begin with the formation of Pop III stars in metal-free mini-haloes following the semi-analytic prescription of Magg et al. (2022). In each halo, Pop III stars are assumed to form in a single burst when the halo reaches the critical mass for star formation M_{crit} (described in more detail in Section 2.2), with a constant, fixed efficiency of $f_{\star,III}$ and with stellar masses distributed according to the IMF (Gessey-Jones et al. 2022, 2025). After the first generation of stars reach the end of their lives, haloes undergo a period of recovery t_{delay} due to Pop III supernovae during which they do not form stars. Eventually, the metal-enriched gas re-collapses, and the haloes start-forming Pop II stars. We assume that Pop II stars form continuously with a star formation rate density (SFRD) modelled using the subgrid prescription:

$$\dot{\rho}_{\star,II}(\boldsymbol{x}, z) = \left(1 + \delta(\boldsymbol{x})\right) \times \frac{1}{t_{\star,II} H(z)^{-1}} \times$$

$$\int_{M_{crit}}^{M_{max}} \left(f_{\star,II}(\boldsymbol{x}) f_{II}(\boldsymbol{x}, t_{delay}) \frac{f_g}{f_b}(\boldsymbol{x}, M_h) \frac{\Omega_b}{\Omega_m} M_h \right) \frac{dn(\boldsymbol{x}, M_h, z)}{dM_h} dM_h,$$

where $\boldsymbol{x}$ is the spatial coordinate in the simulation box (with varying δ, v_{bc} and M_{crit}), and $(1 + \delta)$ accounts for conversion from Lagrangian to Eulerian space (see e.g. Mesinger, Furlanetto & Cen 2011; Tseliakhovich et al. 2011; Muñoz 2023). Here, $t_{\star,II} = 0.2$ is the fraction of Hubble time over which Pop II stars are assumed to form (corresponding to the typical dynamical time for DM haloes; Park et al. 2019; Reis, Barkana & Fialkov 2022b),[4] f_{II} is mass fraction of haloes forming Pop II stars which depends on t_{delay} (Magg et al. 2022), f_g/f_b is the gas-to-baryon mass fraction affected by v_{bc} in small haloes (Gnedin 2000; Naoz, Barkana & Mesinger 2009; Tseliakhovich et al. 2011), and $f_{\star,II}$ is the Pop II SFE (the quantity of interest in this work). Since we intend to focus on Pop II stars, Pop III star formation parameters are fixed to $f_{\star,III} = 0.002$, with a log-flat IMF in the mass range $2\,M_\odot$ to $180\,M_\odot$ (Gessey-Jones et al. 2022), and a $t_{delay} = 30\,Myr$ recovery time (e.g. Chiaki, Susa & Hirano 2018; Magg et al. 2022).

---

[3]In order to derive this bound, we first calculate the conditional HMF $\psi(M_h, z) = dn(M_h, z|\delta, v_{bc} = 0)/dM_h$ for $\delta \in [-1, 2]$, and numerically find the values of δ_{max} at which the $d\psi/dM_h > 0$ for any given M_h (i.e. the HMF increases, instead of decreasing, with increasing halo mass which is unphysical behaviour). A value of $\delta_{max} = 1.3$ ensures physical validity for all redshifts $z \geq 6$. In the simulation, no pixels reach this threshold at $z = 15$, 5 pixels (0.0002 per cent) at $z = 10$, and 3387 pixels (0.16 per cent) at $z = 6$, which is a reasonable approximation expected in a semi-numerical model.

[4]The dynamical (half-crossing) time for DM halo is $t_{DM,dyn} = V_c/R_{vir} = \sqrt{3/(4\pi G \rho_{vir})} = \sqrt{6/(\Delta_c H(z)^2)}$ for $\rho_{vir} = \Delta_c \rho_{crit}$. Different halo radius/mass definitions depending on the overdensity threshold $\Delta_c = 50 - 500$ (e.g. White 2001) give values $t_{DM,dyn} H(z) \sim 0.10 - 0.35$. The Schmidt law (Schmidt 1959), extended to general star formation in proto-galaxies (Silk 1997), yields the star formation (or gas depletion) time-scale $t_{SF,dyn} = f_\star \rho_{gas}/\dot{\rho}_{star}$. The two are closely linked via gas accretion rate onto galaxies in haloes so that $t_{DM,dyn} \sim t_{SF,dyn}$ (see e.g. analytic models of Davé, Finlator & Oppenheimer 2012; Dekel & Mandelker 2014, which are in agreement with simulations).

Star formation in haloes is modelled through the analytic means described above. However, the radiative fields from stellar and galactic emission propagate to cosmological distances larger than the pixel size and thus need to be handled numerically. This is done by calculating the comoving emissivity $\epsilon(\boldsymbol{x}, \nu, z)$ of the various radiative species, and propagating them through window functions taking into account the lightcone effect and redshifting of source spectra. The sources are modelled as follows:

(i) X-rays: Emitted by Pop III and Pop II high-mass X-ray binaries using the SEDs from Fragos et al. (2013b), as we shall describe in Section 2.3.2. We do not differentiate between the two X-ray populations here, but this distinction will be made in a future work (e.g. following the prescription developed in Gessey-Jones et al. 2025).

(ii) Lyman: Lyman-α and Lyman–Werner bands from Pop II stars as described in Barkana & Loeb (2005) calculated using STARBURST99 (Leitherer et al. 1999) and from Pop III stars using the spectra from Gessey-Jones et al. (2022). We account for Ly α heating and multiple scattering effects (Reis et al. 2021).

(iii) Radio: Emitted by Pop II and Pop III galaxies using the power-law spectrum from Gürkan et al. (2018), Mirocha & Furlanetto (2019), which we will describe in Section 2.3.3. Here again, we do not differentiate radio emission between the two populations. We also exclude the effect of CMB heating (Venumadhav et al. 2018) in this work.

(iv) Ionization: Ionization of the neutral IGM is computed using excursion set formalism (Furlanetto, Zaldarriaga & Hernquist 2004) with a spherical top-hat filter, where UV photons ionize the immediate surrounding of galaxies (Madau et al. 1997; Barkana & Loeb 2004). We use the following criterion for reionization:

$$\exists R < R_{mfp}, \quad \text{s.t.} \quad \zeta_{ion} f_{coll}(\boldsymbol{x}, R) > 1 - x_{e,oth}(\boldsymbol{x}, R) \quad (2)$$

where ζ_{ion} is the ionization rate per baryon, $R_{mfp} = 50\,cMpc$ is the maximum mean free path of ionizing photons (based on end of EoR QSO observations; Wyithe & Loeb 2004; Qin et al. 2021; Zhu et al. 2023), $f_{coll}(\boldsymbol{x}, R)$ is the collapse fraction of baryons into galaxies averaged over a volume of radius R centred at $\boldsymbol{x}$, and $x_{e,oth}(\boldsymbol{x}, R)$ is the fractional ionization from long-range agents like X-rays. If the criterion in equation (2) is not satisfied down to cell resolution (i.e. 3 cMpc), the cell is treated as a two-phase medium with an ionized and non-ionized fraction (in agreement with cosmological radiative transfer algorithms; Zahn et al. 2011; Hutter 2018).

Using the radiation fields calculated above, their effect on the gas kinetic temperature T_K is calculated by solving for the thermal differential equation including X-ray heating, Ly α heating/cooling, adiabatic cooling, ionization cooling, and structure formation heating. The simulations end at $z_{min} = 6$, close to the end of the EoR (e.g. McGreer, Mesinger & D'Odorico 2015; Planck Collaboration VI 2020; Jin et al. 2023).

2.2 Enhanced Pop II star formation efficiency

The star formation prescription described in the previous section contains two important parameters: M_{crit} and $f_{\star,II}$ (for Pop II star formation).

Haloes of mass $\sim 10^{5-10}\,M_\odot$ are sensitive to the threshold for star formation M_{crit}, affected by feedback from streaming velocities, Lyman-Werner photons, and UV photons. The equations governing the threshold are as follows:

$$M_{vir} = \frac{R_{vir} V_c^2}{G} \approx 5.8 \times 10^5 M_\odot \left(\frac{V_c}{4.2\,km\,s^{-1}} \right)^3 \left(\frac{1+z}{21} \right)^{-3/2}; \quad (3)$$

$$M_{\rm mol} \approx f_{v_{\rm bc}} f_{\rm LW} \left(5.8 \times 10^5 {\rm M}_\odot\right) \left(\frac{1+z}{21}\right)^{-3/2};\tag{4}$$

$$M_{\rm atm} \approx \left(3.5 \times 10^7 \, {\rm M}_\odot\right) \left(\frac{1+z}{21}\right)^{-3/2}.\tag{5}$$

Here, the first equation gives the virial mass of a halo with circular velocity V_c (the minimum required for efficient cooling of gas in haloes[5]), the second equation is the molecular cooling threshold for haloes (corresponding to $V_c \approx 4.2 \, {\rm km \, s^{-1}}$), and the third is the atomic cooling threshold (corresponding to $V_c \approx 16.5 \, {\rm km \, s^{-1}}$). The factors $f_{v_{\rm bc}}$ and $f_{\rm LW}$ account for the effects of streaming velocity and Lyman–Werner feedback on the molecular cooling threshold, respectively, as defined in Gessey-Jones et al. 2023 (see also Fialkov et al. 2012, 2013; Muñoz et al. 2022).

The variable V_c encapsulates our uncertainty on the feedback processes in small haloes (e.g. from SNe) and is a free parameter in our model, so that $M_{\rm crit}(\boldsymbol{x}, z) = \max\left(\min\left(M_{\rm mol}, M_{\rm atm}\right), M_{\rm vir}\right)$. During the EoR, UV photons irradiate star-forming gas in ionized bubbles, further increasing $M_{\rm crit}$ (to $\sim 10^9 \, {\rm M}_\odot$) via photoheating feedback (described in Sobacchi & Mesinger 2013; Cohen, Fialkov & Barkana 2016, accounting for ionized fractions in partially ionized cells).

In previous works using 21CMSPACE, Pop II star formation was modelled using a fixed, constant efficiency in the atomic cooling regime of DM halo masses, and a log-suppression in the molecular cooling regime as the gas cooling rate declines smoothly with virial temperature (Machacek, Bryan & Abel 2001; Fialkov et al. 2013). This is sufficient for modelling the poorly constrained high-redshift cosmological 21-cm signal, but at lower redshifts this becomes unrealistic. There is a strong SFE–halo mass dependence across redshifts for high-mass haloes, seen in both observations (Guo et al. 2010; Behroozi & Silk 2015; Stefanon et al. 2021) and simulations (Ma et al. 2018; Ceverino, Klessen & Glover 2018; Lovell et al. 2021; Kannan et al. 2022; Pallottini et al. 2022). Indeed, the current understanding of SFE at low redshifts ($z \lesssim 6$) is that stellar feedback and AGN feedback effects lead to a suppression at the low halo mass and high halo mass end respectively, resulting in a double power-law function that peaks at $\sim 10^{11-12} \, {\rm M}_\odot$ (e.g. UNIVERSEMACHINE; Behroozi et al. 2019).

In this work, we build upon the existing model in the code by introducing a new four parameter star formation prescription, inspired by Park et al. (2019), which is *halo-mass and redshift dependent* as follows:

---

[5]Note that this is simply related to a halo's virial temperature often used in the literature via:

$$T_{\rm vir} = \frac{\mu m_{\rm p} V_c^2}{2 k_{\rm B}} \approx 635 \, {\rm K} \left(\frac{\mu}{0.6}\right) \left(\frac{V_c}{4.2 \, {\rm km \, s^{-1}}}\right)^2,\tag{6}$$

where $m_{\rm p}$ is the proton mass, $k_{\rm B}$ is the Boltzmann constant, and μ is the mean molecular weight with $\mu = 1.22$ for neutral primordial gas and $\mu = 0.6$ for ionized hydrogen gas. We adopt the latter for consistency with Barkana & Loeb (2001) and other works, but the choice does not affect our results and simulations in any way as we directly deal with mass and velocity. It only affects our derived constraints on $T_{\rm vir}$ presented in Table 3 and in Section 5.1.

$$f_{\star,\rm II}(\boldsymbol{x}, M_h, z) = \tilde{f}_{\star,\rm II} \begin{cases} 0 & M_h < M_{\rm crit} \\ \dfrac{\log\left(M_h/M_{\rm crit}\right)}{\log\left(M_{\rm atm}/M_{\rm crit}\right)} & M_{\rm crit} \leq M_h < M_{\rm atm} \\ 1 & M_{\rm atm} \leq M_h < M_{\rm high}, \\ \left(\dfrac{M_h}{M_{\rm high}}\right)^{\alpha_\star} & M_{\rm high} \leq M_h \end{cases}\tag{7}$$

where the turning point is defined as

$$M_{\rm high} = M_0 \left(\frac{1+z}{7}\right)^{-\beta_\star}.\tag{8}$$

and $f_{\star,\rm II}$ is naturally capped at 1. The particular form of $M_{\rm high}$ anchors the SFE at $z_{\rm min} = 6$, where our simulations end and the UVLF data is strongest (i.e. essentially calibrating the SFE magnitude and evolution to be with respect to redshift $z = 6$). Fig. 1 shows a schematic of this SFE model. Similar models with a flattening below a given mass threshold have been explored in the literature before (Sun & Furlanetto 2016; Sipple & Lidz 2024), the latter of whom find that a shallow-slope best fits the lensed UVLFs from *HST* at low halo masses, as opposed to a single power law (dash–dotted green in Fig. 1) that assumes a steep decline. The parametrization used in our work, due to its flexibility, relaxes the assumption of declining SFE made by the single power-law model for small mass haloes. We also expect our model to be reasonable compared to the double power-law model (solid pink line in Fig. 1) for two reasons: (i) AGN feedback effects are generally expected for large haloes $\sim 10^{12} \, {\rm M}_\odot$. However, few haloes reach $\sim 10^{11} \, {\rm M}_\odot$ at $z \geq 6$ in our simulation volume of interest $\sim 10^8 \, {\rm cMpc}^3$. For comparison, multifield determinations of UVLFs combining many *HST*/*JWST* surveys typically probe volumes of $\lesssim 10^6 \, {\rm cMpc}^3$ across redshift bins, making large haloes even rarer. Since the 21-cm signal, and in general reionization, is sensitive to the cumulative effect of all haloes down to $10^5 \, {\rm M}_\odot$, our relevant halo mass range is wider than other abundance matching and analytic inferences that usually consider $M_h \gtrsim 10^9 \, {\rm M}_\odot$. (ii) The role of AGN at high redshifts is still quite uncertain. Although pre-*JWST* studies suggest a decline in AGNs at $z \gtrsim 6$ (Kulkarni, Worseck & Hennawi 2019), more recent works indicate a non-negligible faint AGN population (Harikane et al. 2023b; Fujimoto et al. 2024). The contribution of these AGN to star formation quenching remains unclear since the bright end of the luminosity function is poorly sampled at high redshifts.

The redshift evolution factor $\beta_\star$ is a timely addition, motivated by the abundance of bright galaxies in recent high-redshift *JWST* observations. Physically, the term captures an enhancement in the SFE in the early Universe (for $\beta_\star > 0$), meaning DM haloes of the same mass host more massive and brighter galaxies in the early Universe than today.

2.2.1 Phenomenological comment on the SFE

The SFE defined here is the fraction of the accreted gas in a DM halo that is converted into *stars*, sometimes referred to as the integrated baryon conversion efficiency (e.g. Moster, Naab & White 2018; Zhang et al. 2022). To first order, the SFE is equal to the stellar to halo mass fraction divided by the baryon fraction $f_b \approx 0.16$.

We note here that the SFE implementation in our subgrid model is essentially an ensemble average since we do not track individual haloes. For a given halo mass and redshift, the SFE of the galaxy in the host halo is fixed (i.e. there is no galaxy-to-galaxy scatter around the mean). This approach is reasonable as both the 21-cm signal and UVLF are statistical quantities probed at large scales (but
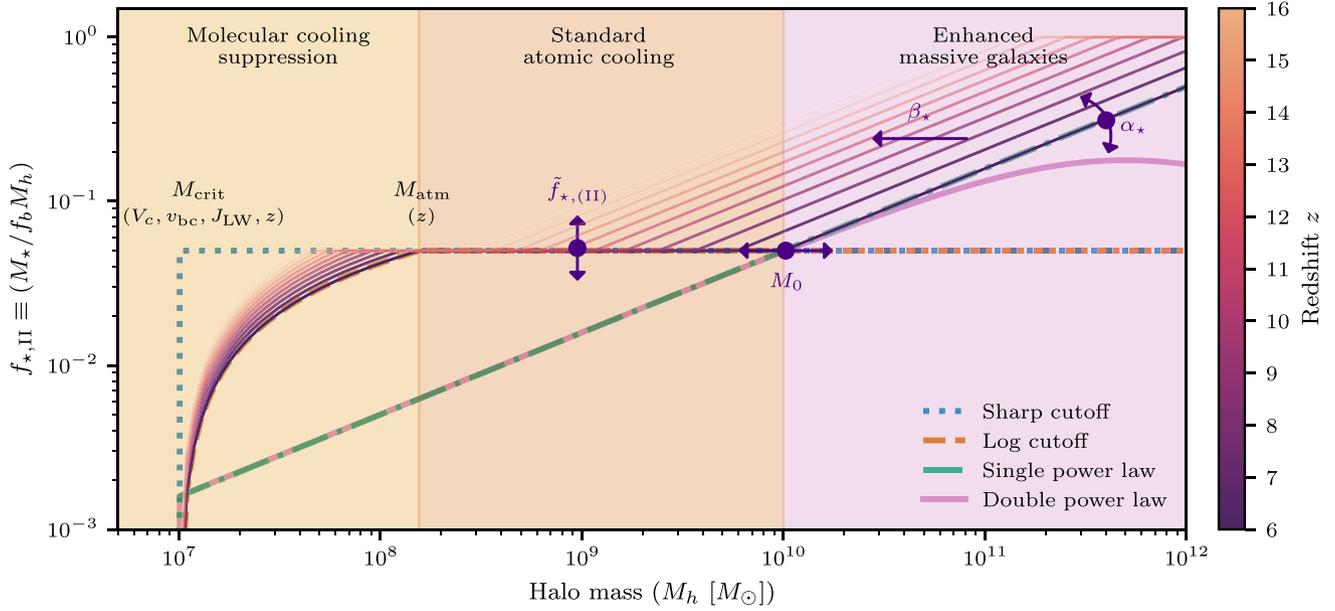
Figure 1. Schematic of the Pop II SFE model $f_{\star,\mathrm{II}}(M_h, z)$ used in this work (purple-brown gradient, at various redshifts), described in equation (7), as a function of halo mass M_h and redshift z. The model is a four-parameter function that is zero below the critical mass M_{crit} (set to $10^7\,\mathrm{M}_\odot$ here for illustration, but can be as low as $10^5\,\mathrm{M}_\odot$ at $z = 30$), logarithmically increasing in the molecular cooling regime (left, light orange shaded region), constant in the atomic cooling regime (center, dark orange shaded region), and a power law in the high-mass regime (right, purple region). The turning point $M_{\mathrm{high}} = M_0\,[(1+z)/7]^{-\beta_\star}$ is redshift dependent and the evolution rate $\beta_\star$ captures the enhanced SFE at fixed halo mass in the early Universe. The SFE is naturally capped at 1. The example illustrated here corresponds to $\tilde{f}_{\star,\mathrm{II}} = 0.05$, $M_0 = 10^{10}\,\mathrm{M}_\odot$, $\alpha_\star = 0.5$, $\beta_\star = 4$. Previous works using 21CMSPACE (e.g. Pochinda et al. 2024) utilize the log cut-off model (dashed orange line, shown at $z = 6$), where the SFE is a fixed, constant value in the atomic cooling regime which is too simplistic for the observed mass dependency of SFE for high-mass haloes. The single power law (dash–dotted green line), on the other hand, provides a better fit to large haloes, but assumes a low SFE in unobserved faint, low-mass haloes which may not be accurate. Finally, the double power law (solid pink line) accounts for feedback from AGN in the highest mass haloes, which are abundant at $z \le 6$, but their impact is less clear at high redshifts.

see also Reis, Barkana & Fialkov 2022a; Nikolić et al. 2024, on the importance of stochasticity). Observations of individual galaxies can significantly differ from the 'global mean', and indeed any observational data set is biased to the brightest and most massive galaxies which will be inherently stochastic (e.g. Mirocha & Furlanetto 2023). Nonetheless, rare galaxies manifest in our simulations via our pixel-level formalism. The abundance of haloes and their gas-to-baryon mass fraction (f_g / f_b) is conditional on the halo mass, and pixel-level overdensity and streaming velocity. The minimum mass for star-forming haloes M_{crit} is further affected by local feedback mechanisms. Thus, the stellar mass and UVLF still varies between pixels across the simulation. Rare overdensities at high-z will produce rare high-mass galaxies, which can have an elevated SFE if they fall in the power-law regime of equation (7). The main challenges of SFE enhancement (at any redshift) is feedback, and so there are two interpretations of this redshift evolution owing to differing star formation histories:

(i) The first interpretation is one of bursty star formation, where episodes of high-SFE can be induced by large-scale inflows (Mc-Clymont et al. 2025), positive feedback from AGN (Silk et al. 2024), or delayed stellar feedback (Furlanetto & Mirocha 2022, which lead to a flattened SFE slope in the low halo mass regime; see their fig. 5). n extreme version of the last case are the 'feedback-free bursts' (FFB) recently discussed in Dekel et al. (2023) and Li et al. (2024) where burstiness occurs on very short time-scales of $\sim 10\,\mathrm{Myr}$ (i.e. multiple cycles/generations within a single half-crossing time of a virialized DM halo). This happens through rapid infall of gas via cold streams in low metallicity, high density environments, quicker

than SNe/wind time-scales, thereby essentially creating feedback-free conditions. The halo mass threshold above which these bursty galaxies form decreases with increasing redshift in this scenario, leading to an early time enhancement (see fig. 2 of Li et al. 2024).

(ii) The second interpretation is the standard uniform star formation, through sustained gas supply and accretion to the galaxy over the halo half-crossing time. If the conditions in the early Universe are viable for compact galaxies hosting a large number of high density stellar clusters which have short dynamical time-scales, they may be resilient to feedback processes essentially creating a 'feedback failure' environment where stars can form with high efficiencies. These conditions could be due to increased binding energy (or deeper potential wells) of small haloes at high-z (Dayal et al. 2014; Furlanetto et al. 2017); dark-matter accelerated gas collapse (Boylan-Kolchin 2025); or a SFE-scaling that increases with gas surface density analogous to giant molecular clouds [density modulated SFE or density-modulated SFE (DMSFE); Somerville et al. 2025].

2.3 Simulating observables of interest

2.3.1 The 21-cm signal

The 21-cm spectral line arises from the forbidden transition of a neutral hydrogen atom from the hyperfine ground state (electron and proton spins anti-aligned) to an excited state (spins aligned). The vast abundance of neutral hydrogen at cosmic scales, however, ensures that enough atoms undergo this process creating observable signals. Whether the signal is seen in emission or absorption depends on the relative occupancy of the hyperfine levels (Purcell & Field 1956;

Scott & Rees 1990):

$$\frac{n_1}{n_0} = 3 \exp\left(\frac{-h\nu_{21}}{k_B T_S}\right), \tag{9}$$

where n_0 and n_1 are number densities of hydrogen atoms in the lower and higher energy states, the factor of 3 accounts for the statistical degeneracy of the higher energy triplet state, $\nu_{21} = 1420$ MHz is the rest-frame frequency of the 21-cm line, and T_S is the spin temperature of the hydrogen atoms.

As the Universe expands, this 21-cm signal is observed today as a distortion to the CMB black body spectrum (or radio background, with temperature T_r) across the radio frequency range $\nu_{21,\mathrm{obs}} = \nu_{21}/(1 + z)$, corresponding to the different epochs at redshifts z. The strength of the signal is usually quantified in terms of the differential brightness temperature

$$T_{21}(z) = \left(1 - e^{-\tau_{21}}\right)\frac{T_S - T_r}{1 + z}, \tag{10}$$

where τ_{21} is the 21-cm radiation optical depth given by

$$\tau_{21}(z) \approx \frac{3}{32\pi}\frac{hc^3 A_{10}}{k_B \nu_{21}^2}\frac{x_{\mathrm{HI}}(z)n_H(z)}{(1 + z)\mathrm{d}v_\parallel/\mathrm{d}r_\parallel}\frac{1}{T_S(z)}, \tag{11}$$

with x_{HI} being the fractional abundance of neutral hydrogen, n_H the hydrogen number density, A_{10} the spontaneous emission rate of the 21-cm transition, and $\mathrm{d}v_\parallel/\mathrm{d}r_\parallel$ the proper velocity gradient along the observer line of sight. If $T_S > T_r$, the signal is seen in net emission against the radio background, otherwise it is seen in net absorption.

In order to calculate this brightness temperature, one needs the spatially and temporally varying fractional abundance x_{HI} and the spin temperature T_S. The former depends on the distribution of ionizing stellar/quasar sources, while the latter is an interplay between several competing influences that can cause a spin-flip transition in hydrogen. Coupling of T_S to the background radiation temperature T_r is mediated via scattering of background photons, while coupling to the kinetic temperature of matter T_K is through collisions between atoms/electrons in dense IGM, and the Wouthusen-Field (WF) effect when star formation begins (Wouthuysen 1952; Field 1958). Thus, we get the equation (Furlanetto et al. 2006):

$$T_S(z)^{-1} = \frac{T_r^{-1} + x_c T_K^{-1} + x_\alpha T_C^{-1}}{1 + x_c + x_\alpha} \tag{12}$$

where x_c and x_α are the coupling strengths of the latter two processes respectively, and T_C ($\approx T_K$) is the colour temperature of the Lyα radiation field that intermediates the WF effect.

Current experiments are either aimed at detecting the sky-averaged global signal (e.g. Bowman et al. 2018; Philip et al. 2019; Singh et al. 2022; de Lera Acedo et al. 2022):

$$\langle T_{21}(z)\rangle = \int T_{21}(\hat{\boldsymbol{r}}, z)\mathrm{d}\Omega, \tag{13}$$

where $\hat{\boldsymbol{r}}$ is a line-of-sight on the sky and Ω is the solid angle; or the power spectrum $P_{21}(k, z)$ (e.g. Ewall-Wice et al. 2016; Mertens et al. 2020; HERA Collaboration 2022b):

$$\langle \tilde{T}_{21}(\boldsymbol{k}, z)\tilde{T}_{21}^*(\boldsymbol{k}', z)\rangle = (2\pi)^3\delta^D(\boldsymbol{k} - \boldsymbol{k}')P_{21}(k, z) \tag{14}$$

where $\tilde{T}_{21}(\boldsymbol{k}, z)$ is the Fourier-transform of $T_{21}(\boldsymbol{x}, z)$, k is the comoving wavevector and δ^D is the 3D Dirac delta function. In this work, we use the power spectrum in its conventional 'dimensionless' form:

$$\Delta_{21}^2(k, z) = (k^3/2\pi^2)P_{21}(k, z). \tag{15}$$

For detailed pedagogical reviews on 21-cm cosmology, we refer the readers to Barkana & Loeb (2001), Furlanetto et al. (2006),

Pritchard & Loeb (2012), Barkana (2016), Mesinger (2019). The calculation of the 21-cm global signal and power spectrum is done in post-processing in 21CMSPACE.

2.3.2 Diffuse X-ray background

X-ray emissions are another key ingredient of the early Universe, contributing to both the heating of the IGM and the ionization of the neutral hydrogen. Despite its importance, much remains uncertain regarding its properties: soft X-rays ($E \sim 20$ eV) in the form of thermal emission from supernovae remnants (Oh 2001; Venkatesan, Giroux & Shull 2001; Ricotti & Ostriker 2004) and mini-quasars (Kuhlen, Madau & Montgomery 2006; Ciardi, Salvaterra & Di Matteo 2010), hard X-rays ($E \sim 3$ keV) from black hole X-ray binaries (XRBs; Mirabel et al. 2011; Fragos et al. 2013a), and AGN have all been suggested as dominant sources of X-ray emission. The properties of the X-ray background such as its strength, inhomogeneous heating effect (Pritchard & Furlanetto 2007; Fialkov & Barkana 2014), and shape of its spectral energy distribution (SED; Fialkov et al. 2014; Fialkov & Barkana 2014; Pacucci et al. 2014) affect the evolution of the IGM temperature, and thus the 21-cm signal coupled to it. In addition to this, the contribution of ancient X-ray sources to the unresolved X-ray background seen today could act as a useful observable of the early Universe (Fialkov et al. 2017), setting upper limits on their emission efficiency.

XRBs may be a dominant source of X-ray emissions at $z \gtrsim 6$, in particular HMXBs (HMXBs; Fragos et al. 2013b; Lehmer et al. 2016). A simple model for their X-ray luminosity, following local starburst-like galaxies, assumes proportionality to the halo star formation rate (SFR; e.g. Grimm, Gilfanov & Sunyaev 2003; Lehmer et al. 2010; Mineo, Gilfanov & Sunyaev 2012; Mineo et al. 2014):

$$\frac{L_X}{\mathrm{SFR}} = f_X \times \left(3 \times 10^{40}\ \mathrm{erg\ s^{-1}\ M_\odot^{-1}\ yr}\right) \tag{16}$$

where f_X is the X-ray emission efficiency of sources at high-redshifts, normalized to the theoretical prediction for low metallicity HMXBs (Fragos et al. 2013a, b; Lehmer et al. 2021) as explained in Fialkov et al. (2017). The comoving X-ray emissivity throughout the simulation can then be calculated as $\epsilon_X(\boldsymbol{x}, E, z) = L_X/\mathrm{SFR} \times \dot{\rho}_{\star,\mathrm{II+III}}(\boldsymbol{x}, z) \times \hat{\epsilon}_X(E)$ where $\dot{\rho}_{\star,\mathrm{II+III}}$ is the total Pop II and Pop III SFRD, and $\hat{\epsilon}_X(E)$ is the normalized X-ray SED in eV^{-1} (the mean model from Fragos et al. 2013b, at $z = 15.34$ in the energy range $0.2 - 95$ keV, including interstellar absorption). The present-day angle-averaged specific intensity from these sources at $z > z_0$ is thus (e.g. Pritchard & Furlanetto 2007):

$$J_X(E, z = 0) = \frac{1}{4\pi}\int_{z_0}^\infty \epsilon_X(E', z')e^{-\tau_X(E', z')}\left|\frac{c\mathrm{d}t}{\mathrm{d}z'}\right|\mathrm{d}z', \tag{17}$$

usually quoted in units of erg s^{-1} eV^{-1} cm^{-2} sr^{-1}. In this equation, $E' = E(1 + z')$ is the energy at the emission redshift z', $\epsilon_X(E', z')$ is averaged across the simulation box,[6] $\tau_X(E', z')$ is the optical depth of X-rays calculated using the weighted cross section of hydrogen and helium species (Verner et al. 1996), and $H(z')$ is the Hubble parameter at emission. The observed integrated flux in some energy band is then given by:

$$S_X(E_{\mathrm{band}}) = \int_{E_{\mathrm{band}}} J_X(E, z = 0)\mathrm{d}E \tag{18}$$

---

[6]There is no need for shell radiative transfer here since we're interested in the sky-averaged diffuse background.

in units of $\mathrm{erg\,s^{-1}\,cm^{-2}\,sr^{-1}}$. In the case of our simulations, we use $z_0 = z_{\mathrm{min}} = 6$, which means that only high-redshift sources are included in the redshift integral. This complements the idea that the X-ray background measurements serve as an absolute upper limit on the contribution from high-z sources. Note that we assume $f_X = f_{X,\mathrm{II}} = f_{X,\mathrm{III}}$ in the constraints here; the largest contribution to the integral comes from lower redshifts, where Pop II stars dominate the star formation rate.

2.3.3 Diffuse radio background

In a similar manner to X-rays, one can estimate the radio luminosity per unit frequency from galaxies as (Gürkan et al. 2018; Mirocha & Furlanetto 2019; Reis et al. 2020):

$$\frac{L_r(\nu)}{\mathrm{SFR}} = f_r \times \left(10^{22}\,\mathrm{W\,s^{-1}\,Hz^{-1}\,M_\odot^{-1}\,yr}\right) \left(\frac{\nu}{150\,\mathrm{MHz}}\right)^{-\alpha_r} \qquad (19)$$

where f_r is the emission efficiency with respect to the present-day radio galaxies, and the spectral index $\alpha_r = 0.7$ matches the observed power spectrum of synchrotron emission from low redshift radio galaxies (Hardcastle et al. 2016; Gürkan et al. 2018). As discussed in the introduction, the excess radio background seen today (Fixsen et al. 2011; Dowell & Taylor 2018), if non-Galatic, could be due to a variety of exotic sources. The above formalism attributes this, at least in part since the background is an upper limit, to astrophysical origins in galaxies.

To estimate the contribution of high-redshift galaxies at $z > z_0$ to the present-day sky-averaged excess radio temperature, we use the equation (e.g. Ewall-Wice et al. 2020; Reis et al. 2020):

$$T_r(\nu, z = 0) = T_{\mathrm{CMB}} + \frac{c^2}{2k_B\nu^2}\frac{1}{4\pi}\int_{z_0}^{\infty}\epsilon_r(\nu', z')\left|\frac{c\,\mathrm{d}t}{\mathrm{d}z'}\right|\mathrm{d}z', \qquad (20)$$

where $\nu' = \nu(1 + z')$ is the frequency at emission redshift z', $\epsilon_r(\nu, z')$ is the comoving radio emissivity averaged across the simulation volume, $z_0 = z_{\mathrm{min}} = 6$ as before, and we assume that $f_r = f_{r,\mathrm{II}} = f_{r,\mathrm{III}}$. Unlike X-rays, the radio emission is not absorbed as it propagate through space and time. We include the effect of line-of-sight fluctuations introduced recently in 21CMSPACE in Sikder et al. (2024a). We also note that we have not included the effect of soft photon heating from the excess background, which can have leading order effects on the matter temperature and thus the 21-cm signal (Cyr, Acharya & Chluba 2024), and leave this for future work.

2.4 Ultraviolet luminosity function

2.4.1 Analytic model

We now discuss the modelling of UVLFs used in this work, which is a new output in 21CMSPACE. UVLFs are a key statistic of galaxy surveys, which quantify the abundance of galaxies as a function of the rest-frame luminosity measured in the band $\lambda = 1500 \pm 50\,\text{Å}$. This is a useful tool to study the evolution/properties of galaxy populations. The main features of UVLFs can be reproduced within a framework of the HMFs (see e.g. Vale & Ostriker 2004; Park et al. 2019):

$$\Phi(\mathbf{x}, M_{\mathrm{UV}}, z) = \left(1 + \delta(\mathbf{x})\right)\frac{\mathrm{d}n(\mathbf{x}, M_h, z)}{\mathrm{d}M_h}\left|\frac{\mathrm{d}M_h}{\mathrm{d}M_{\mathrm{UV}}}\right|, \qquad (21)$$

where Φ is the number density of galaxies at UV magnitude M_{UV} (defined later in equation 23), $\mathrm{d}n/\mathrm{d}M_{\mathrm{UV}}$ is the hybrid HMF discussed in Section 2.1, and the $(1 + \delta)$ accounts for the conversion of the HMF from Lagrangian to Eulerian space as in equation (1) (e.g. Tseliakhovich et al. 2011). In order to map $M_h \leftrightarrow M_{\mathrm{UV}}$, we assume

that the SFR is proportional to the rest-frame UV luminosity of galaxies L_{UV},

$$\dot{M}_{\star,\mathrm{II}}(\mathbf{x}, M_h, z) = \frac{M_{\star,\mathrm{II}}(\mathbf{x}, M_h, z)}{t_{\star,\mathrm{II}}H(z)^{-1}} = \kappa_{\mathrm{UV}} \times L_{\mathrm{UV}}(\mathbf{x}, M_h, z) \qquad (22)$$

where $M_{\star,\mathrm{II}}$ is the term in parentheses in the integrand of equation (1), $t_{\star,\mathrm{II}} = 0.2$ is the characteristic time for formation of Pop II stars in units of the Hubble time, as discussed before in Section 2.1, and $\kappa_{\mathrm{UV}} = 1.15 \times 10^{-28}\,\mathrm{M_\odot\,yr^{-1}\,erg^{-1}\,s\,Hz}$ is a constant conversion factor. We adopt this value, following Madau & Dickinson (2014) who calculate it assuming a continuous star formation history, a Salpeter IMF in the range $0.1 - 100\,\mathrm{M_\odot}$, and an evolving metallicity of $Z_\star = 10^{-0.15z}Z_\odot$. Note that these assumptions regarding the stellar population are not necessarily justified (e.g. Tacchella et al. 2018; Harikane et al. 2023a) and indeed Madau & Dickinson (2014) caution about the metallicity dependence being calibrated to low redshifts only. At high redshifts, the metallicity could be drastically lower, with a different IMF leading to a much larger κ_{UV}. We use the value above for consistency with the rest of literature, but briefly explore its impact on the UVLF, together with those of δ_{max} and $t_{\star,\mathrm{II}}$, in Appendix A. We find that the difference between assuming a fixed Chabrier or Salpeter IMF is negligible compared to modelling uncertainties, although noting that more extreme IMFs like top-heavy Pop III stars that are not explored here could generate larger differences (e.g. Inayoshi et al. 2022).

The UV luminosity thus calculated is then related to the absolute magnitude (in the monochromatic AB magnitude system) through the standard relation (Oke & Gunn 1983):

$$\log_{10}\left(\frac{L_{\mathrm{UV}}}{\mathrm{erg\,s^{-1}\,Hz^{-1}}}\right) = 0.4 \times (51.63 - M_{\mathrm{UV}}). \qquad (23)$$

Although equation (21) is an analytic prescription, we stress that it is applied at a pixel level (3 cMpc scale) before averaging over the simulation volume as shown in Fig. 2. This simulation-based approach has two benefits over simple analytic models: (i) we can capture the 'extreme statistics' of rare overdensities in finite simulation volume, and (ii) each pixel has a unique M_{crit} from photoheating feedback during the EoR, giving rise to a natural suppression of faint galaxies. The model can also be easily extended to simulate larger volumes and study large-scale overdensity biases, or include spatially varying metallicities which would alleviate the need for a constant κ_{UV}. We leave such investigations for future work.

2.4.2 Dust correction

Dust plays a crucial role in the evolution of a galaxy, as a key component in chemical reactions across the ISM, and also as an absorber of UV starlight. Formed from stellar outflows such as supernova ejecta or AGB winds (Gehrz 1989), the dust alters the observed galactic spectrum by emitting the absorbed energy back in the infrared wavelengths (e.g. Calzetti 2001; Weingartner & Draine 2001). This attenuation in the UV regime is especially important at the low redshift end of our simulations, $z \sim 6$ since *HST* provides the tightest UVLF constraints there.

In order to correct for dust in our UVLF model, we add an attenuation term to the magnitudes:

$$M_{\mathrm{UV}}^{\mathrm{dust}} = M_{\mathrm{UV}}^{\mathrm{dust\text{-}free}} + A_{\mathrm{UV}}(M_{\mathrm{UV}}^{\mathrm{dust}}). \qquad (24)$$

A common simple method for estimating the correction term A_{UV} is to use the UV spectral slope β as a proxy for the dust distribution across galaxies, i.e. the flatter the slope, the redder the galaxy and the more dust there is. Based on the properties of over 50 starburst
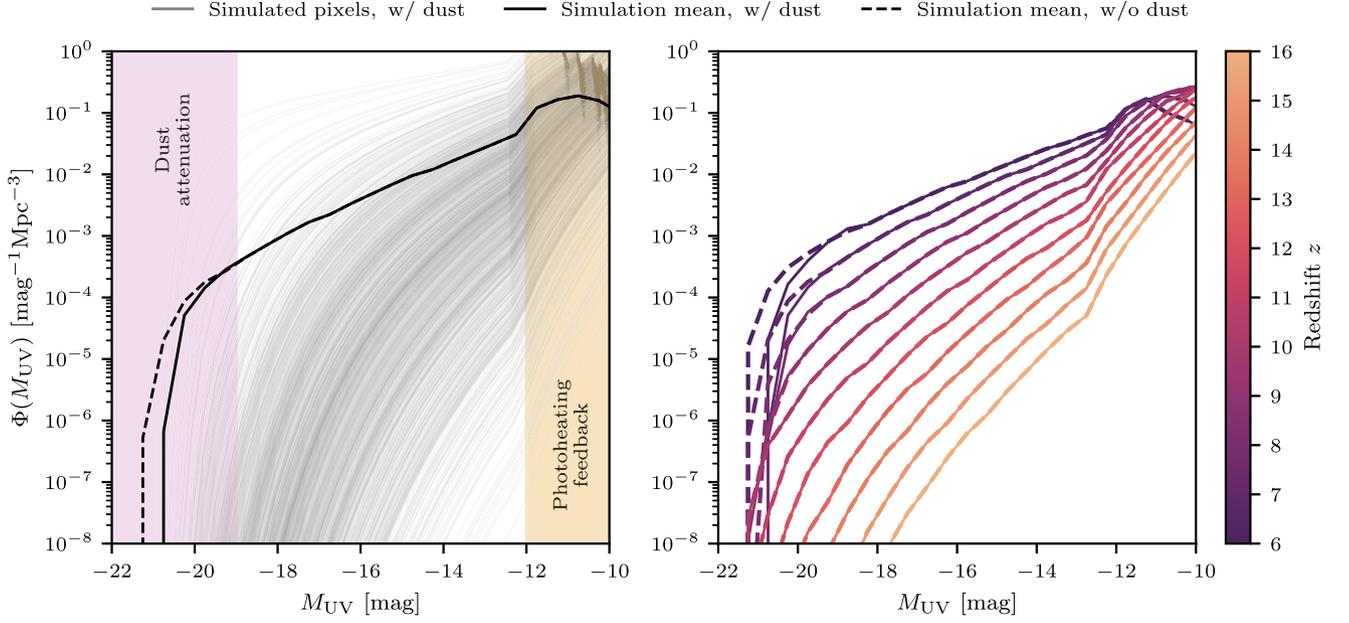
Figure 2. UVLFs in the simulation for a fiducial parameter set: $V_c = 4.2 \, \text{km s}^{-1}$, $\tilde{f}_{\star,\text{II}} = 10^{-3}$, $M_0 = 5 \times 10^9 \, \text{M}_\odot$, $\alpha_\star = 1.0$, $\beta_\star = 0.5$. Left: UVLF at $z = 8$ for a random sample of ~ 1000 pixels (grey), averaged across the simulation volume with and without dust effects (black solid and black dash–dotted lines, respectively). The suppression at the bright end is a result of dust attenuation and a finite simulation volume (left, purple shaded region), while the suppression at the faint end is due to photoheating feedback during the EoR (right, orange shaded region). The finite volume and photoheating feedback are simulation features not captured by simple analytic models. Right: Redshift evolution of the UVLF from $z = 6$ to $z = 15$ for the same parameter set.

galaxies at $z \approx 3$, Meurer, Heckman & Calzetti (1999) derive the widely used infrared-excess (IRX) and β slope relation:

$$A_{\text{UV}} = C_0 + C_1 \beta \qquad (25)$$

with best-fitting parameters $C_0 = 4.33$ and $C_1 = 1.99$. Note that to avoid unphysical attenuation, the above relation is only valid for $\beta > -2.33$, and galaxies with larger UV slopes are assumed to have negligible dust effects.

The UV spectral slope is not constant and can strongly vary with UV magnitude as well as redshift of the galaxy (e.g. Bouwens et al. 2014). We use the parametrization of Zhao & Furlanetto (2024),

$$\langle\beta\rangle(M_{\text{UV}}^{\text{dust}}, z) = \frac{\text{d}\beta}{\text{d}M_{\text{UV}}^{\text{dust}}}(z) \left[M_{\text{UV}}^{\text{dust}} - M_{\text{UV},0}\right] + \beta_{M_{\text{UV},0}}(z), \qquad (26)$$

where the observations from *HST* in Bouwens et al. (2014) and *JWST* JADES in Topping et al. (2024) are combined to give the best-fit values

$$\text{d}\beta/\text{d}M_{\text{UV}}^{\text{dust}} = -0.012z - 0.216,$$
$$\beta_{M_{\text{UV},0}}(z) = -0.081z - 1.58$$

for the reference magnitude $M_{\text{UV},0} = -19.5$ mag. Using an aggregate fit such as equation (26) avoids the issue of abrupt jumps or incompatible results from different *HST* and *JWST* observations and survey sizes.

Combining equations (24)–(26), we get an analytic dust correction model that can be applied across the whole simulation:

$$M_{\text{UV}}^{\text{dust}} = \frac{1}{1 - C_1 \frac{\text{d}\beta}{\text{d}M_{\text{UV}}^{\text{dust}}}(z)} \left[M_{\text{UV}}^{\text{dust-free}} + C_0 \right.$$
$$\left. + C_1 \beta_{M_{\text{UV},0}} - C_1 \frac{\text{d}\beta}{\text{d}M_{\text{UV}}^{\text{dust}}}(z) M_{\text{UV},0}\right]. \qquad (27)$$

Physically, this leads to a suppression of the brightest galaxies at redshifts $z \lesssim 10$, but does not significantly affect galaxies at $z \gtrsim 10$. This is because galaxies are generally not large enough to reach the high magnitudes where dust effects become important, and there is an 'expected' absence of dust in this calibration. A potential mechanism for lack of dust in galaxies at high-z (the so-called blue monsters) is by radiation driven outflows (Fiore et al. 2023; Ziparo et al. 2023; Ferrara, Pallottini & Sommovigo 2025).

2.5 Simulation parameters

As the early Universe remains uncertain, the cosmological and astrophysical processes described in Section 2.1 can be turned on/off or scaled through various parameters in 21cmSPACE. In this work, we focus on exploring the effect of different star formation models. Hence, our free parameter space consists of eight parameters: $\{V_c, \tilde{f}_{\star,\text{II}}, M_0, \alpha_\star, \beta_\star, \tau, f_X, f_r\}$. The method of obtaining constraints on these parameters using Bayesian analysis closely follows that of previous works (Gessey-Jones et al. 2023; Bevins et al. 2024; Pochinda et al. 2024), but with the updated SFE and inclusion of the UVLF observable. We briefly list the parameters and their prior ranges below.

(i) V_c – Minimum circular velocity for star formation in haloes in the absence of feedback. As described in Section 2.2, the feedback processes increase the M_{crit} value for each pixel, but V_c sets the minimum threshold for star formation in the absence of feedback processes (M_{vir}). We vary it in the range $V_c \in [4.2, 100] \, \text{km s}^{-1}$, corresponding to $M_{\text{vir}} \in [3 \times 10^6, 4 \times 10^{10}] \, \text{M}_\odot$ at $z = 6$ (see equation 3), and $T_{\text{vir}} \in [6.35 \times 10^2, 3.6 \times 10^5]$ K (see equation 6).

(ii) $\tilde{f}_{\star,\text{II}}$ – Level of the Pop II SFE. Previous works used a constant SFE in the atomic cooling regime denoted by $f_\star$ or $f_{\star,\text{II}}$ (e.g. Bevins et al. 2022a; Pochinda et al. 2024). In this work, we

introduce $f_{\star,\mathrm{II}}(M_h, z)$ which can be halo-mass and redshift dependent (equation 7), and define $\tilde{f}_{\star,\mathrm{II}}$ as the level of efficiency in this model. We vary this parameter in the wide range $\tilde{f}_{\star,\mathrm{II}} \in \left[10^{-4}, 10^{-0.3}\right]$, chosen to allow for both low and high SFE below the power-law transition. This is to ensure that we do not make any assumptions regarding faint galaxy populations that lie outside of observational constraints.

(iii) M_0 – Normalization of the turning point, M_{high} (anchored at $z = 6$), at which the dependence of SFE on M_h transitions from flat to a power law. We vary this parameter in the range $M_0 \in \left[2 \times 10^8, 10^{11}\right]$ M$_\odot$. The lower limit is chosen to be just above the atomic cooling threshold $M_{\mathrm{atm}}(z = 6) \approx 1.5 \times 10^8$ M$_\odot$, and the upper limit is close to the maximum halo mass in the simulation volume at low redshifts.

(iv) $\alpha_\star$ – The power-law index of the high-mass end of the SFE, varied in the range $\alpha_\star \in [0, 2]$. We choose positive values to 'boost' the SFE at the high-mass end in agreement with observations (e.g. Behroozi & Silk 2015; Behroozi et al. 2019). A value of 0 corresponds to a flat SFE in the atomic cooling regime, while a value of 2 corresponds to a steep power-law (i.e. efficiency rapidly increasing with halo mass).

(v) $\beta_\star$ – Redshift evolution of M_{high}, varied in the range $\beta_\star \in [0, 5]$. Higher values of $\beta_\star$ correspond to a faster evolution of the SFE with redshift, while a value of 0 corresponds to no evolution (i.e. redshift-independence). This parameter is introduced to capture an effective boost in SFE in the early Universe for small halo masses, in order to explain the observed high-z galaxy populations from *JWST*.

(vi) τ – In 21CMSPACE, the effective ionization efficiency of galaxies is captured by ξ_{ion} as defined in equation (2). Since the optical depth of CMB photons to reionization τ monotonically scales with ξ_{ion} when V_c, $f_{\star,\mathrm{II}}$ and other parameters are fixed (Cohen et al. 2020), we use τ instead of ξ_{ion} in our analysis as the free parameter. We adopt a uniform prior of 3σ around the measured *Planck 2018* value of $\tau = 0.054 \pm 0.07$ (Planck Collaboration VI 2020), i.e. $\tau \in [0.033, 0.075]$.[7] By indirectly using this reionization era constraint, we ensure that no unphysical reionization histories are realized in our models. We do not explore more complex parametrizations here, instead focusing on the SFE, due to uncertainties in the galaxy ionizing emissivities, escape fraction of ionizing sources, and IGM clumping which contribute to ζ_{ion} (e.g. Mitra & Chatterjee 2023; Davies, Bosman & Furlanetto 2024; Asthana et al. 2025).

(vii) f_X – The X-ray emission efficiency of high-redshift galaxies, relative to the theoretical predictions of low metallicity HMXBs of Fragos et al. (2013b), as defined in equation (16). We vary this in the wide range $f_X \in \left[10^{-3}, 10^3\right]$. Since this parameter is degenerate with the SFR of galaxies (and thus the SFE), the wide range aptly allows for extreme parameter combinations such as very low SFE and very high f_X.

(viii) f_r – The radio emission efficiency of high-redshift galaxies, relative to present-day galaxies as defined in equation (19). We vary this parameter in the range $f_r \in \left[10^{-1}, 10^5\right]$, chosen to be wide for the same reason as f_X.

The above described parameter priors are summarized in Table 1. Using these, we run a suite of 30 000 simulations across the eight-dimensional astrophysical parameter space. Each simulation takes a couple of hours to run, and the entire suite takes a few weeks

Table 1. Parameter priors used in our Bayesian analysis. The astrophysical parameters consist of those concerning star formation (V_c, $\tilde{f}_{\star,\mathrm{II}}$, M_0, $\alpha_\star$, $\beta_\star$) and others concerning radiation from stellar/post-stellar sources (τ, f_X, f_r). The SARAS parameters consist of a 6th order polynomial foreground model and a thermal noise term, and are only included in the constraints that use SARAS 3 data (see Section 3.1 for data description and 4.2 for the Bayesian likelihood that includes these nuisance parameters). The prior ranges for the SARAS parameters are chosen to be centred around the best-fit values in Bevins et al. (2022a).

Type	Parameter	Prior	Minimum	Maximum
Astrophysical				
	V_c	Log-uniform	4.2 km s^{-1}	100 km s^{-1}
	$\tilde{f}_{\star,\mathrm{II}}$	Log-uniform	10^{-4}	$10^{-0.3}$
	M_0	Log-uniform	2×10^8 M$_\odot$	10^{11} M$_\odot$
	$\alpha_\star$	Uniform	0	2
	$\beta_\star$	Uniform	0	5
	τ	Uniform	0.033	0.075
	f_X	Log-uniform	10^{-3}	10^3
	f_r	Log-uniform	10^{-1}	10^5
SARAS	σ_{S3}	Log-uniform	0.01 K	1 K
	a_0	Uniform	3.54	3.55
	a_1	Uniform	-0.23	-0.21
	a_2	Uniform	0	0.01
	a_3	Uniform	-0.01	0
	a_4	Uniform	0	0.01
	a_5	Uniform	-0.01	0.01
	a_6	Uniform	-0.01	0.01

to complete on the DiRAC high-performance computing facility COSMA8.[8]

3 OBSERVATIONAL DATA

Models created with 21CMSPACE are often utilized to constrain astrophysical parameters using various observational data sets, as summarized in Table A1 (listed for ease of parsing previous literature). We build upon these works by including the UVLF as an additional observable, and by exploring the effect of the time-evolving enhancement of SFE. We describe all the data sets used in our analysis below.

3.1 SARAS 3

The Shaped Antenna measurement of the background RAdio Spectrum (SARAS) experiments are a series of radio experiments designed to detect the global 21-cm signal. We use measurements of the global sky temperature from the latest experiment, SARAS 3, in the band $55 - 85$ MHz ($z \approx 15 - 25$; Jishnu Nambissan et al. 2021).[9] The data consists of 15 h of observation from a 14-d period during its deployment in Dandiganahalli Lake and Sharavati backwaters in Southern India (Singh et al. 2022). The data is calibrated to remove any radio-frequency interference (RFI), receiver systematics, and thermal emission from the lake. It is thus expected to consist of the global 21-cm signal, the time-averaged foreground from both Galactic and extragalactic sources, and any residual noise.

---

[7]For late reionization scenarios, beyond the end of our simulation ($z \lesssim 6$), we calculate the τ by extrapolating the ionization history as a function of redshift in log-log space.

[8]See https://cosma.readthedocs.io/en/latest/cosma8 for more details on CPU specifications.

[9]We do not include SARAS 2 measurements (Singh et al. 2017) in this work due to uncertainty in the systematics modelling in the data set (see Bevins et al. 2024, for a joint analysis of both SARAS 2 and 3).

For our analysis, since the foregrounds are expected to be smooth, we fit the data with a 6th order log-log polynomial (coefficients a_i) and also include a thermal noise term (σ_{S3}) for the residuals, as done in previous works (Singh et al. 2022; Bevins et al. 2022a, 2024). These 7 + 1 parameters are varied in the prior range shown in Table 1.

3.2 HERA

The Hydrogen Epoch of Reionization Array (HERA) is a radio interferometer with cross-dipole feeds and the PAPER correlator (in its Phase 1), designed to detect the 21-cm power spectrum in the frequency range $100 - 200$ MHz (DeBoer et al. 2017). We use the upper limits on the 21-cm power spectrum from HERA Phase 1 observations (HERA Collaboration 2023), which improve upon the previous limits from HERA Collaboration (2022b). This data set contains 94 nights of observations using $35 - 41$ antennas across the experiment's entire frequency range. We use the two frequency bands least contaminated by radio frequency interference, Band 1 in Field D ($6.25 - 9.25$ h LST) and Band 2 in Field C ($4.0 - 6.25$ h LST) in the range $117.19 - 133.11$ MHz and $152.25 - 167.97$ MHz respectively. The lowest limits in the two bands are:

$$\Delta^2_{21,\text{Band 1}}(k = 0.36\,h\,\text{Mpc}^{-1}, z \approx 10.35) \leq 3496\,\text{mK}^2,$$
$$\Delta^2_{21,\text{Band 2}}(k = 0.34\,h\,\text{Mpc}^{-1}, z \approx 7.87) \leq 457\,\text{mK}^2.$$

We reduce the data to only include every other k-bin to ensure neighbouring data points are uncorrelated in our analysis, as done in Pochinda et al. (2024). As HERA employs a foreground avoidance technique, no foreground modelling is required and we marginalize residual systematics above the thermal noise using the likelihood function described in Section 4.1, as done previously in HERA Collaboration (2023).

3.3 Cosmic X-ray background

The CXB is the unresolved X-ray flux from sources outside of the Milky Way Galaxy (see Bambi & Sangangelo 2022, for a recent review). CXB includes contribution from point sources such as active galactic nucleus (AGN) and galaxies that are too faint to be resolved by telescopes, diffuse emission from the hot IGM, and redshifted emission from X-ray binaries both local and distant (barring unknown instrumental systematics). Hence, measurements of the CXB should provide upper limits for the X-ray emissivity of XRBs in the early Universe (e.g. Fialkov et al. 2017) and complement the inference on SFE (equation 16).

For our analysis, we use the known point-source subtracted CXB flux from Chandra X-ray Observatory (Hickox & Markevitch 2006):

$$S_X(1 - 2\,\text{keV}) \leq (1.04 \pm 0.14) \times 10^{-12}\,\text{erg cm}^{-2}\,\text{s}^{-1}\,\text{sr}^{-1},$$
$$S_X(2 - 8\,\text{keV}) \leq (3.4 \pm 1.7) \times 10^{-12}\,\text{erg cm}^{-2}\,\text{s}^{-1}\,\text{sr}^{-1},$$

and the collated data from table 1 of Harrison et al. (2016):

$$S_X(8 - 24\,\text{keV}) \leq (1.832 \pm 0.042) \times 10^{-11}\,\text{erg cm}^{-2}\,\text{s}^{-1}\,\text{sr}^{-1},$$
$$S_X(20 - 50\,\text{keV}) \leq (2.0 \pm 0.083) \times 10^{-11}\,\text{erg cm}^{-2}\,\text{s}^{-1}\,\text{sr}^{-1},$$

which includes measurements from HEAO (Marshall et al. 1980; Gruber et al. 1999), BeppoSAX (Frontera et al. 2007), INTEGRAL (Churazov et al. 2007) and Swift BAT (Ajello et al. 2008).

3.4 Cosmic radio background

The CRB, akin to the CXB, is the unresolved radio flux from extragalactic sources, after the removal of Galactic foreground emission and the CMB monopole. This includes contribution from high-redshift radio galaxies (potentially AGN or Pop III SNe), and provides upper limits for the radio emissivity of galaxies in the early Universe, alongside the joint inference on SFE (equation 19).

For our analysis, we use the CRB data collated in table 2 of Dowell & Taylor (2018), which includes measurements from LWA1 Low Frequency Sky Survey (Dowell et al. 2017) in the range $40 - 80$ MHz, ARACDE2 (Fixsen et al. 2011) in the range $3 - 11$ GHz, and other single frequency experiments at 22 MHz (Roger et al. 1999), 46 MHz (Alvarez et al. 1997; Maeda et al. 1999), 408 MHz (Haslam et al. 1982; Remazeilles et al. 2015), and 1420 MHz (Reich, Testori & Reich 2001).

3.5 Ultraviolet luminosity function

In this work, in addition to the aforementioned data sets which have been previously used in Pochinda et al. (2024), we include UVLF determinations from Hubble and *JWST* fields. These are:

(i) Bouwens et al. (2021a, B21): This data set combines several *HST* surveys to provides UVLFs in the range $z = 2 - 10$. This includes the HUDF/XDF and parallel fields, BORG/HIPPIES, five CANDELS fields, HFF and parallel fields, and observations from the ERS program (see their table 1 for details). For our analysis, we use the step-wise UVLF estimates from their table 4 in the redshift range $z \approx 6 - 9$ containing 1066, 601, 246 and 33 sources respectively, and omit the 9 sources at $z \approx 10$ from Oesch et al. (2018) in favour of the larger sample size in *JWST* data at this limiting redshift.

(ii) Donnan et al. (2024, D24): This multifield data set combines several *JWST* surveys, including PRIMER, NGDEEP and JADES to provide UVLFs in the redshift range $z \approx 9 - 15$. We pick this data set due to its inclusion of the latest *JWST* data, its large survey area of $\simeq 370$ arcmin^2, and its consistent findings with other data sets at both $z = 11$ (Leung et al. 2023; Pérez-González et al. 2023; McLeod et al. 2024) and $z = 12.5$ (Adams et al. 2024; Robertson et al. 2024). For our analysis, we use UVLF estimates from the entire data set given in their table 2, consisting of a total of 2548 sources, in the redshift bins $8.5 < z < 9.5$, $9.5 < z < 10.5$, $10.5 < z < 11.5$, $11.5 < z < 13.5$ and a tentative $13.5 < z < 15.5$ bin based on an equivalent of ~ 1.3 sources.[10] However, note that due to the large error bars of the highest-z data point, it does not affect our constraints strongly.

Fig. 3 shows the collated UVLF data used in this work. By eye, one can see that the UVLF evolution with redshift slows down at $z \gtrsim 10$.

4 BAYESIAN INFERENCE

In order to perform our parameter inference, we use a Bayesian approach similar to that of previous works (e.g. Bevins et al. 2022a; Pochinda et al. 2024). In this statistical paradigm, given a set of parameters θ and observed data $\mathcal{D}$, the *a priori* probability of the parameters being true $\pi(\theta)$ is updated by the likelihood of the data given the parameters $\mathcal{L}(\theta) = P(\mathcal{D}|\theta)$, to give the *a posteriori* probability distribution $\mathcal{P}(\theta) = P(\theta|D)$:

$$\mathcal{P}(\theta) = \frac{\mathcal{L}(\theta)\pi(\theta)}{\mathcal{Z}}, \tag{28}$$

---

[10]This is the total contribution from summed probability $p(z)$ of all photometric solutions; for more details on the sample selection process, see Section 3 of their work.
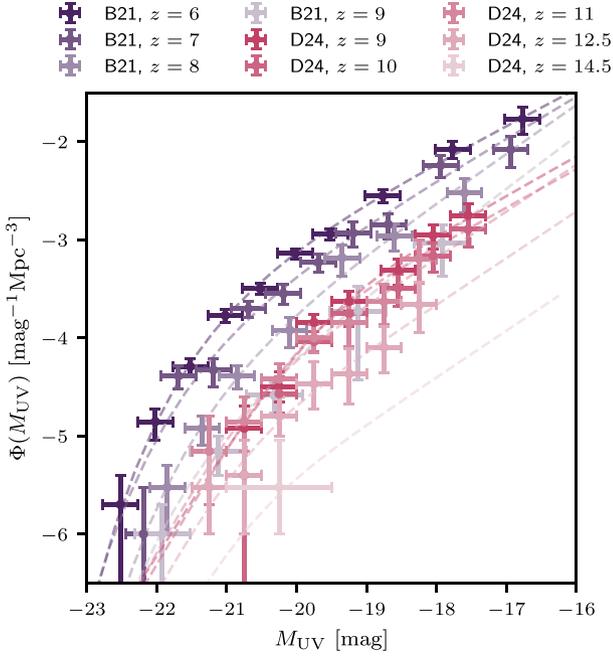
Figure 3. UVLF data used for constraints in this work. The data points are from Bouwens et al. (2021a, B21, in purple) and Donnan et al. (2024, D24, in pink), and represent some of the largest compilations of galaxies with *HST* and *JWST*, respectively. The dashed lines are analytic Schechter and double power-law fits from the respective works.

where $\mathcal{Z} = P(\mathcal{D}) = \int \mathcal{L}(\theta)\pi(\theta)d\theta$ is the Bayesian evidence (the integral of the likelihood over the parameter space). A uniform prior is uninformed and suitable for parameters which have an expected range, while a log-uniform prior works better for parameters whose order of magnitude is not known.

The Bayesian approach allows us to marginalize over the parameter space to obtain the posterior distribution of each parameter (or indeed subset of parameters) θ_i as

$$\mathcal{P}(\theta_i) = \frac{1}{\mathcal{Z}} \int \mathcal{L}(\theta)\pi(\theta)\mathrm{d}\theta_{j \neq i}. \tag{29}$$

Thus, we can treat the SARAS 3 foreground parameters θ_{fg} described in Section 3.1 as nuisance parameters, and marginalize over them to obtain the posterior distribution of the astrophysical parameters θ_{ast}.

We use the slice-sampling-based nested-sampling package POLY-CHORD (Handley, Hobson & Lasenby 2015a, b) to perform our Bayesian inference. A single run of POLYCHORD can require millions of likelihood evaluations, which is computationally unfeasible to do by running the full 21CMSPACE simulation for each set of parameters. To make the inference tractable, we train neural network (NN) emulators that can rapidly evaluate the observables of interest within tens of milliseconds. We describe the emulator training and likelihood evaluation in the following sections.

4.1 Emulator training

We train emulators for our five observables of interest: the 21-cm global signal $T_{21}(z)$, the 21-cm power spectrum $\Delta_{21}^2(k, z)$, the CXB $S_X(E)$, CRB $T_r(\nu)$, and UVLF $\Phi(M_{\mathrm{UV}}, z)$. We use the TENSORFLOW (Abadi et al. 2016) based package GLOBALEMU (Bevins et al. 2021) to train the one-dimensional (1D) observables: 21-cm global signal, CRB and CXB, and SCIKIT-LEARN (Pedregosa et al. 2011) to

train the two-dimensional (2D) observables: 21-cm power spectrum and UVLF. All emulators are trained with the Adam optimization algorithm using a tanh activation function, a mean squared error loss function, and a test/train split of 1/9. The input pre-processing and training details for each emulator are as follows:

(i) 21-cm global signal: We train the emulator on the 21-cm global signal T_{21} output from the simulations in the range $z = 6 - 30$. We pre-process the training data by first resampling the default simulation redshift step of $\Delta z = 1$ to $\Delta z = 0.1$ ($N_{\mathrm{samp}} = 250$) via simple interpolation to increase the training data set size and resolution. We then down-scale and normalize the signals (subtracting the mean and dividing by the standard deviation of training data), as is a standard feature in GLOBALEMU. The emulator is trained on a 4-layer NN of 16 nodes each, a batch-size of 250, and with early-stopping based on validation loss to prevent over-fitting. The resulting emulator has an accuracy of < 11 per cent error in the SARAS 3 band at 68 per cent confidence, and an RMSE of ~ 70 mK across the entire redshift range at the 95th percentile of test data. For a more detailed analysis of emulator accuracy (and its dependence on the input-space z), see Appendix B.

(ii) 21-cm power spectrum: We use a similar methodology here as described in Pochinda et al. (2024) using SCIKIT-LEARN. The simulation outputs the 21-cm power spectrum Δ_{21}^2 at redshifts $z = 7 - 11$ and wavevector range $k \approx 0.1 - 2.5\, h\, \mathrm{Mpc}^{-1}$. We pre-process the training data by log-transforming both z and k, as well as Δ_{21}^2 (truncating values $< 10^{-8}\, \mathrm{mK}^2$) to improve the performance of the emulator over the large dynamic range. We then resample and interpolate the power spectrum to a finer input grid by drawing $N_{\mathrm{samp}} = 20000$ samples from uniform pairs of z and k values. The emulator is trained using an MLPRegressor with 4 hidden layers of 100 nodes each and a batch-size of 10 000. The resulting emulator has an accuracy of < 13 per cent in linear space in the k and z regime of HERA data at 68 per cent confidence.

(iii) CXB/CRB: We train the CXB S_X and radio background T_r in a similar manner to the 21-cm global signal. In case of the former, the CXB (equation 18) is output by the simulation in the energy range $E = 10^{-1} - 10^3$ keV with $N_{\mathrm{samp}} = 400$. The CRB on the other hand is calculated in post-processing from star formation rates output by the simulation (equation 20) in $\nu = 10 - 10^3$ MHz with $N_{\mathrm{samp}} = 400$ to cover the CRB data comfortably. The data is first log-transformed in both cases, and then normalized as described above for the 21-cm global signal emulator. The emulator training parameters remain the same, except for a batch size of 400 for the CXB emulator and 100 for the CRB emulator. The resulting emulators have an accuracy of < 5 per cent and 4 per cent error in the observational energy bands and frequencies, respectively, in linear space at 68 per cent confidence.

(iv) UVLF: We train the UVLF $\Phi(M_{\mathrm{UV}}, z)$ output from the simulation in a similar manner to the 21-cm power spectrum described above. The simulation code outputs the UVLF (in log space) at integer redshifts in the range $z = 6 - 16$ and magnitude $M_{\mathrm{UV}} = [-23, -13]$ mag with a bin size of $\Delta M_{\mathrm{UV}} = 0.5$ to match most observations. The training data is pre-processed by drawing $N_{\mathrm{samp}} = 15000$ from uniform pairs of M_{UV} and z values, and the emulator training follows that of the 21-cm power spectrum with a batch size of 20000. The resulting emulator has an accuracy of < 2 per cent in log space at the observed median z and magnitude bins M_{UV} at 68 per cent confidence, corresponding to an RMSE of 0.8 dex across the full magnitude range at 95th percentile of the test data.

The architecture and training parameters for all the emulators above have been summarized in Table 2.

Table 2. Summary of neural network emulator training parameters for each observable: 21-cm global signal, 21-cm power spectrum, CXB, CRB, and UVLF. All emulators are trained with the Adam optimization algorithm using a tanh activation function, a mean squared error loss function, and a test/train split of 1/9. The N_{samp} parameter denotes the number of sample points per simulation used in the emulator training (e.g. $T_{21}(z)$ consists of 250 redshift points per simulation for the global 21-cm signal emulator). The emulator accuracy quoted here is the maximum 1σ error in the part of the input space occupied by data (e.g. the SARAS 3 band $z = 15 - 25$ for global 21-cm signal). For a complete error analysis of the emulators, see Appendix B.

Emulator	Framework	Network architecture	Batch size	N_{samp}	Emulator accuracy ($\mathbf{1\sigma}$)
21-cm global signal	GLOBALEMU	16-16-16-16	250	250	< 11%
21-cm power spectrum	SCIKIT-LEARN (MLPRegressor)	100-100-100-100	10000	20000	< 13%
CXB	GLOBALEMU	16-16-16-16	400	400	< 5%
CRB	GLOBALEMU	16-16-16-16	100	100	< 4%
UVLF	SCIKIT-LEARN (MLPRegressor)	100-100-100-100	20000	15000	< 2%

4.2 Likelihood evaluation

For the Bayesian inference described in Section 4, we evaluate the likelihood of the model given the data using the emulators described in Section 4.1 and the likelihood functions for each data set we shall describe below. In this work, we introduce a new way to characterize the 'model error' as the incoherent sum of the theory/simulation error and the emulator error:

$$\text{model error} = \sqrt{\text{theory error}^2 + \text{emulator error}^2}. \tag{30}$$

The former accounts for the assumed error on the semi-numerical model used in the simulation, fixed to an ad-hoc 20 per cent (e.g. Zahn et al. 2011; Hutter 2018), while the latter is derived by the performance of the emulator on test data. For a target function $f(x)$, the emulator performance may vary as a function of x. We evaluate this for all emulators in Appendix B, which then feeds into our likelihood functions described below.

(i) SARAS 3: In order to evaluate the likelihood of the model given the SARAS 3 data, we first calculate the foreground model T_{fg} using the polynomial fit described in Section 3.1:

$$\log_{10}\left(T_{\text{fg}}(\nu)/\,\text{K}\right) = \sum_{i=0}^{6} a_i \left[f_{\text{N}}\left(\log_{10}\left(\nu/\,\text{MHz}\right)\right) \right]^i, \tag{31}$$

where function f_{N} normalizes the log-frequency to be in the range $[-1, 1]$. Next, we use the emulator to predict the 21-cm global signal at the frequencies of the SARAS 3 data and evaluate the likelihood function:

$$\mathcal{L}_{\text{SARAS}}(\theta) = \prod_{j}^{N_{\text{freq}}} \frac{1}{\sqrt{2\pi\left[\sigma_{\text{S3}}^2 + \sigma_{\text{model}}^2(\nu_j)\right]}} \exp\Bigg[$$
$$- \frac{1}{2}\left(\frac{T_{\text{SARAS3},j} - T_{\text{fg}}(\nu_j; \theta_{\text{fg}}) - T_{21,\text{model}}(\nu_j; \theta_{\text{ast}})}{\sigma_{\text{S3}}^2 + \sigma_{\text{model}}^2(\nu_j)}\right)^2 \Bigg], \tag{32}$$

where j indexes the observational frequency bins, T_{SARAS3} is the observed SARAS 3 sky temperature corrected for known systematics, T_{fg} is the foreground model, σ_{S3} is the thermal noise nuisance parameter as mentioned in Section 3.1, $T_{21,\text{model}}$ is the theory model signal, and $\sigma_{\text{model}}(\nu_j)$ is the model error at the frequency ν_j.

(ii) HERA: We use the emulator to predict the 21-cm power spectrum at the k and z values of the HERA data and adopt the same likelihood function as in HERA Collaboration (2022a):

$$\mathcal{L}_{\text{HERA}}(\theta_{\text{ast}}) = $$
$$\prod_{j}^{N_{\text{bins}}} \frac{1}{2}\left[1 + \text{erf}\left(\frac{\Delta_{21,\text{HERA},j}^2 - \Delta_{21,\text{model}}^2(k_j, z_j; \theta_{\text{ast}})}{\sqrt{2\left[\sigma_{\text{HERA},j}^2 + \sigma_{\text{model}}^2(k_j, z_j)\right]}}\right)\right], \tag{33}$$

where j indexes the observational k and z bins, $\Delta_{21,\text{HERA}}^2$ is the observed HERA upper limits, $\Delta_{21,\text{model}}^2$ is the theory model predicted 21-cm power spectrum, σ_{HERA} is the observational error, and σ_{model} is the model error at those k_j and z_j values.

(iii) CXB/CRB: For the CXB and CRB inference, we use a similar likelihood to HERA since we use the data in the form of upper limits. In case of the former, the observable is the CXB flux $S_X(E_{\text{band}})$ and we use the emulator to predict this at a high resolution within each energy band of the CXB data and then integrate the flux over the bin. In case of CRB, the observable is the temperature $T_r(\nu)$, and we use the emulator to predict the CRB flux at the precise frequencies of the CRB data. The likelihood function is then evaluated as in equation (33) for both data sets.

(iv) UVLF: To evaluate the likelihood of the model given the UVLF data, we use the emulator to predict $\log_{10}(\Phi)$ at the M_{UV} bins and median z values of the UVLF observations, and use a two-piece Gaussian likelihood function to account for asymmetric data errors (as in Fig. 3):

$$\mathcal{L}_{\text{UVLF}}(\theta_{\text{ast}}) = \prod_{i}^{z-\text{medians}} \prod_{j}^{M_{\text{UV}}-\text{bins}} \sqrt{2/\pi}\,(\sigma_+ + !\sigma_-)^{-1}$$
$$\exp\left[-\frac{1}{2}\left(\frac{\log_{10}\left(\Phi_{\text{obs},i,j}\right) - \log_{10}\left(\Phi_{\text{model}}(M_{\text{UV},i}, z_j; \theta_{\text{ast}})\right)}{\sigma_{+/-}}\right)^2\right], \tag{34}$$

where $\sigma_{+/-} = \sqrt{\left(\sigma_{+/-,\text{obs},i,j}\right)^2 + \left(\sigma_{\text{model}}(M_{\text{UV},i}, z_j)\right)^2}$ is the asymmetric error on the UVLF data point, and the denominator of the exponential is σ_+ if $\log_{10}(\Phi_{\text{th}}) > \log_{10}(\Phi_{\text{obs}})$, and σ_- otherwise.

Equipped with the likelihood function for each data set, we can now perform our Bayesian analysis to infer the astrophysical parameters θ_{ast} given the data.

4.3 Joint inference

As mentioned previously, we use POLYCHORD to perform our Bayesian inference. We sample the parameter space using 1000 live points, and use the standard stopping criterion to ensure the sampling has converged. For the joint constraints, the final likelihood is just the product of the individual likelihoods,

$$\mathcal{L}_{\text{Joint}} = \mathcal{L}_{\text{SARAS 3}} \times \mathcal{L}_{\text{HERA}} \times \mathcal{L}_{\text{CXB}} \times \mathcal{L}_{\text{CRB}} \times \mathcal{L}_{\text{UVLF}}. \tag{35}$$

In order to read the POLYCHORD chains, generate corner-plots and perform statistics, we use the package ANESTHETIC (Handley 2019). We also quantify how much constraining power exists in each data set by calculating the contraction of the astrophysical prior volume to the posterior volume. In order to do this, we use the MARGARINE package (Bevins et al. 2023) that trains masked

autoregressive flows to learn marginal posterior probability density functions (PDFs; marginalizing over nuisance parameters like the SARAS 3 foregrounds). We can then calculate the Kullback–Liebler (KL) divergence statistic on the marginal posteriors to calculate:

$$\text{per cent prior volume consistent with data} = 100 \times \frac{V_{\text{posterior}}}{V_{\text{prior}}} \quad (36)$$
$$\approx 100 \times \exp(-\mathcal{D}_{\text{KL}}).$$

Expectedly, the KL divergence is zero when the prior and posterior are the same, and increases as the posterior diverges from the prior. We would also intuitively expect the most informative data set to be the joint constraints, which we quantify in the results section below.

5 RESULTS AND DISCUSSION

Fig. 4 shows the posterior on the astrophysical parameters recovered for the different data sets independently and when they are combined in a joint analysis. A detailed comparison of constraints between the data sets, excluding UVLF observations, is done in Pochinda et al. (2024) using an SFE model that is constant in the atomic cooling regime (dashed orange line in Fig. 1). Their results are retrieved in our analysis as a special case of our more general SFE model—i.e. the posterior PDF of M_0 saturates to the upper prior limit $10^{11}\,\mathrm{M_\odot}$, rendering $\alpha_\star$ and $\beta_\star$ parameters redundant (since most haloes are $< 10^{11}\,\mathrm{M_\odot}$ at $z \geq 6$). Furthermore, the $\tilde{f}_{\star,\mathrm{II}}$ posterior PDF saturates to the lower prior limit at 10^{-4}. This is shown in Fig. C1, where we compare the joint constraints with and without (shortened to w/o) the contribution of the UVLF data set.

The individual data sets provide different constraining power on the astrophysical parameters; the UVLF data is the most constraining, shrinking the posterior volume to 1.2 per cent of the prior volume, followed by CXB, CRB, and HERA at ~ 50 per cent and lastly SARAS 3 at ~ 88 per cent. Note that these are individual constraints from each data set on the full set of astrophysical parameters. Different data sets can provide more or less constraining power for specific parameters. The combination of SARAS 3, HERA, CXB and CRB show great synergies with the posterior volume corresponding to ~ 29 per cent of the prior. With the inclusion of UVLF data, the complete joint constraints are the strongest, as expected, bringing the posterior volume further down to 0.3 per cent of the prior.

We shall present our results in the order shown in the Fig. 4. We first discuss our results for the minimum circular virial velocity for star formation in haloes V_c in Section 5.1. We then discuss the constraints on the SFE of galaxies in the early Universe in Section 5.2, followed by the constraints on the radiative efficiencies of galaxies f_X and f_r in Section 5.3.

5.1 Constraints on V_c

The minimum circular virial velocity V_c for star formation in haloes is largely informed by the combination of SARAS 3 and UVLF data. The CXB and CRB data sets return flat 1D posterior PDFs, while HERA data has a weak (statistically insignificant) preference for lower values of V_c.

The dominant constraint on V_c comes from the UVLFs. They directly probe faint magnitude galaxies at $M_{\mathrm{UV}} \lesssim -17$ mag, providing a hard upper limit on how high V_c can be. The data disfavour[11]:

$$\log_{10}\left(V_c/\,\mathrm{km\,s^{-1}}\right) \lesssim 1.41\,(1.79) \text{ at } 68\%\,(95\%)\,\mathrm{CL}$$
$$\Rightarrow V_c \lesssim 25.7\,(62.1)\,\mathrm{km\,s^{-1}} \text{ at } 68\%\,(95\%)\,\mathrm{CL}.$$

Furthermore, $V_c > 95\,\mathrm{km\,s^{-1}}$ lies outside the 5σ confidence region, which corresponds to a DM halo mass of $M_h \approx 3 \times 10^{10}\,\mathrm{M_\odot}$ at $z = 6$, and stellar masses of $M_\star \sim 10^7\,\mathrm{M_\odot}$ in galaxies (assuming per cent-level SFE) which is consistent with the faintest galaxies observed at $z \sim 6$ (see e.g. fig. 5 of Simmonds et al. 2024). This constraint is robust across different redshift data, as seen in Fig. 5 (first column). In order to further tighten this derived constraint, future works will include lensed luminosity functions such as the ones in Bouwens et al. (2022) that go down to $M_\star \sim 10^{5-6}\,\mathrm{M_\odot}$ at $z \sim 6$ (Bouwens et al. 2021b). Since our model has the added benefit of including photoheating feedback at the faint end, these lensed UVLFs can shed light on the suppression of star formation in small haloes.

Recently, Sipple & Lidz (2024) performed an analysis of lensed UVLF data using truncated SFE models and found that non-lensed UVLFs disfavour a truncation of star formation above $M_h \sim 3 \times 10^{10}\,\mathrm{M_\odot}$ at $z = 6$ (see their fig. 11), consistent with our results. Furthermore, they suggest that lensed UVLFs disfavour a truncation above $2 \times 10^9\,\mathrm{M_\odot}$ at $z = 6$ (i.e. strong photoheating feedback disfavoured) at 2σ confidence. This would correspond to a $V_c \lesssim 36.5\,\mathrm{km\,s^{-1}}$ being favoured at 2σ, indeed stronger than our $V_c \lesssim 62.1\,\mathrm{km\,s^{-1}}$ constraint. Another recent multiwavelength analysis by Zvi Katz, Redigolo & Volansky (2025) similar to ours also finds that UVLFs play a crucial role in providing an upper bound on the minimum halo mass required for star formation.

We extend the aforementioned analyses by incorporating SARAS 3 data in our work, which provides a weak lower bound on the V_c (or equivalently, M_{vir}) parameter space. The SARAS 3 data residual has an RMS of order $200\,\mathrm{mK}$ after foreground modelling and removal, essentially setting an upper limit on the amplitude of the 21-cm absorption signal in the $z = 15 - 25$ redshift band. Hence, it would favour 21-cm signals that either have shallower absorption troughs or have the deepest absorption at redshifts lower than $z \lesssim 15$. Physically, the latter implies a delayed contribution of first stars and galaxies to the Lyα coupling, along with heating of the IGM from XRBs. In other words, star formation is suppressed in small haloes which corresponds to a higher V_c threshold. This is reflected in the 1D posterior PDF, which gives us:

$$\log_{10}\left(V_c/\,\mathrm{km\,s^{-1}}\right) \gtrsim 1.16\,(0.72) \text{ at } 68\%\,(95\%)\,\mathrm{CL}$$
$$\Rightarrow V_c \gtrsim 14.4\,(5.3)\,\mathrm{km\,s^{-1}} \text{ at } 68\%\,(95\%)\,\mathrm{CL}.$$

At 1σ, this V_c constraint alone would imply that star formation is suppressed in most molecular-cooling haloes ($V_c < 16.5\,\mathrm{km\,s^{-1}}$). At 2σ however, the data only weakly disfavours small haloes, allowing for $V_c \gtrsim 5.3\,\mathrm{km\,s^{-1}}$, and 1D posterior PDF does not go to zero probability at the lower end. Although this has potential to be a powerful constraint, it has two added complexities: (i) we find that V_c is degenerate with the SARAS 3 foreground parameters. In particular, the lowest order polynomial coefficients a_0 and a_1 (which have the largest effect due to the nature of the log polynomial function used) show a negative correlation with V_c. This is because (for all other astrophysical parameters fixed) lower values of V_c correspond to deeper 21-cm absorption troughs in the SARAS 3 band. The

---

[11]The confidence levels (CLs) here, and for the rest of the results, are defined as the smallest intervals in the cumulative distribution function (CDF) containing the required probability volume.
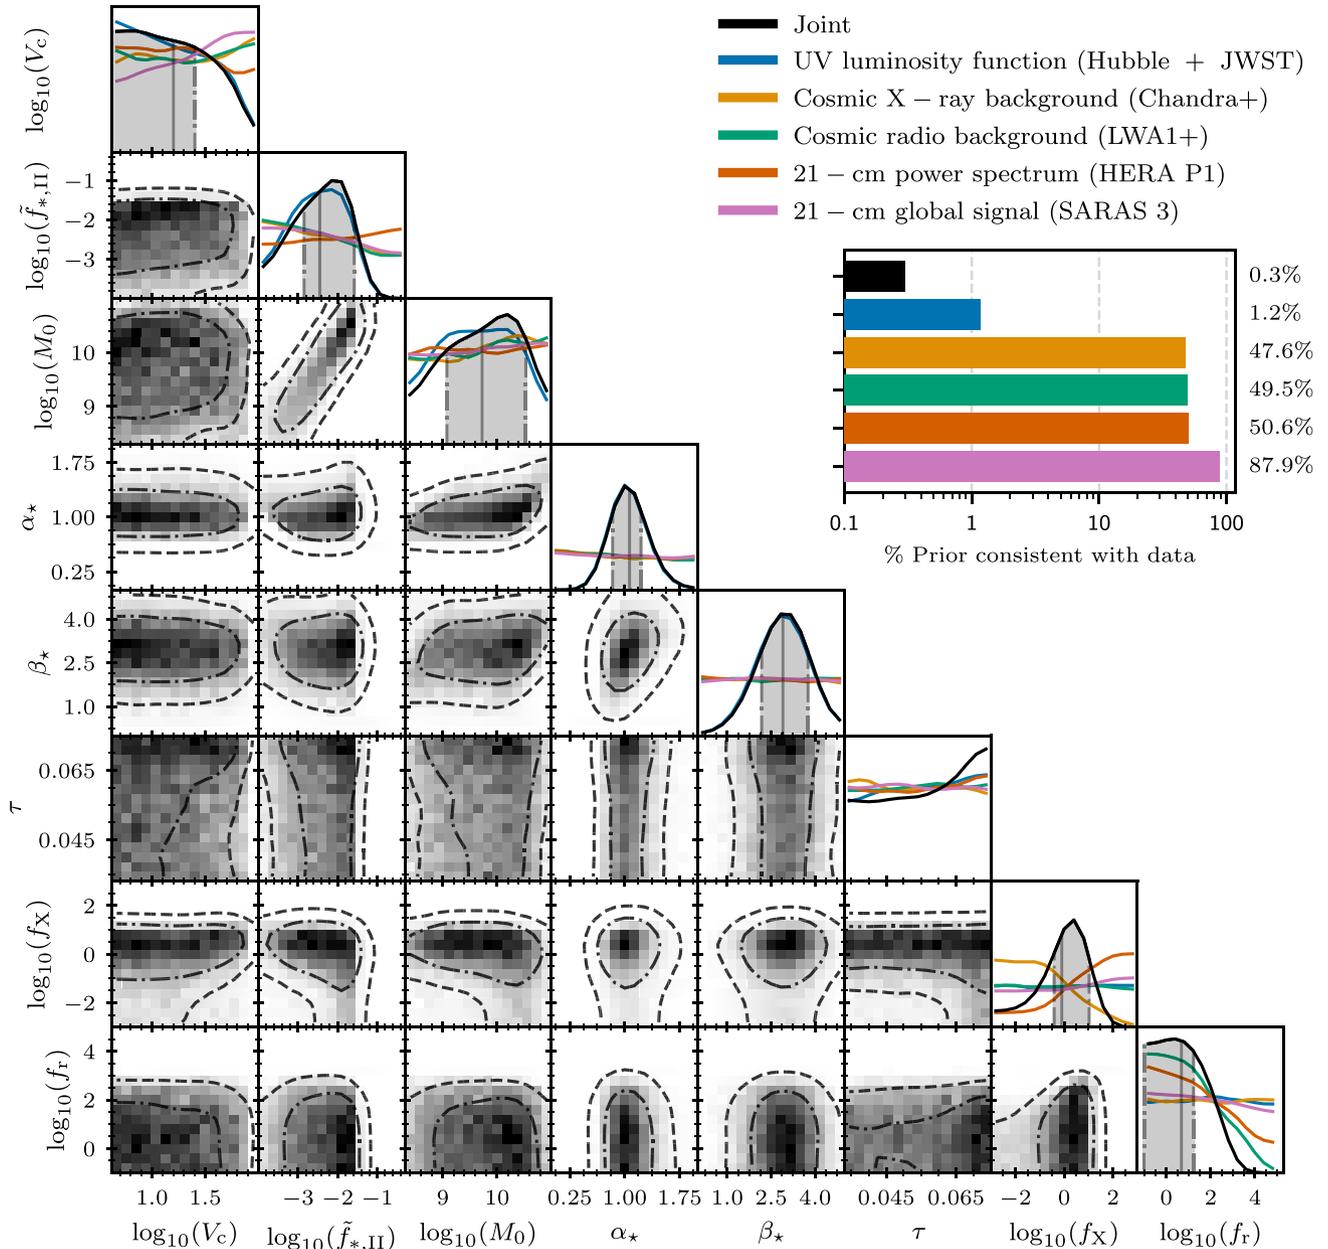
Figure 4. Marginal 1D and 2D posteriors of astrophysical parameters from the joint analysis and individual analysis of the data sets: UVLF, CXB, CRB, HERA, and SARAS 3 (colour-coded as shown in the legend). The 2D posterior contours show the 68 per cent (dash–dotted) and 95 per cent (dashed) credible regions emerging from the joint analysis. The 1D marginal posterior PDFs show the weighted mean (solid line) and 68 per cent credible region (shaded with dash–dotted outline; calculated as the smallest CDF intervals containing the required probability volume) for the joint fit (black), as well as the 1D PDFs for each individual observable (indicated in the legend). Top right panel shows the posterior to prior volume contraction for each data set, calculated using MARGARINE. For a similar plot comparing joint constraints with and without the UVLF data set, see Fig. C1, and for joint constraints with and without the SARAS 3 data set, see Fig. C2. For the full parameter space, including SARAS 3 nuisance parameters, see Fig. C3.

foreground model compensates for this by assuming a larger a_0 and a_1, which boosts the foreground by tens of mK, especially at the low redshift edge of the SARAS 3 band ($z \sim 15$). This is seen in the 2D posterior PDFs in Fig. C3. (ii) The SARAS 3 foreground model likely subsumes some of the 21-cm signal. Thus, in future works, it would be good to check if the choice of the form of foreground model biases our results.

In the joint analysis combining all data sets, the inclusion of SARAS 3 data weakens the upper limit on V_c set by UVLFs as

shown in Fig. C2. The joint constraints yield a V_c value of:

$$V_c \lesssim 21.3\,(60.4)\,\mathrm{km\,s^{-1}} \text{ at 68\% (95\%) CL, without SARAS 3, and}$$

$$V_c \lesssim 25.2\,(62.0)\,\mathrm{km\,s^{-1}} \text{ at 68\% (95\%) CL, with SARAS 3.}$$

Table 3 summarizes these constraints at both 68 per cent and 95 per cent confidence levels, converted to T_{vir} and $M_{\mathrm{vir}}(z = 6)$ for ease of comparison with works that use alternate parametrizations. Given the SARAS 3 constraint's sensitivity to the foreground fit,
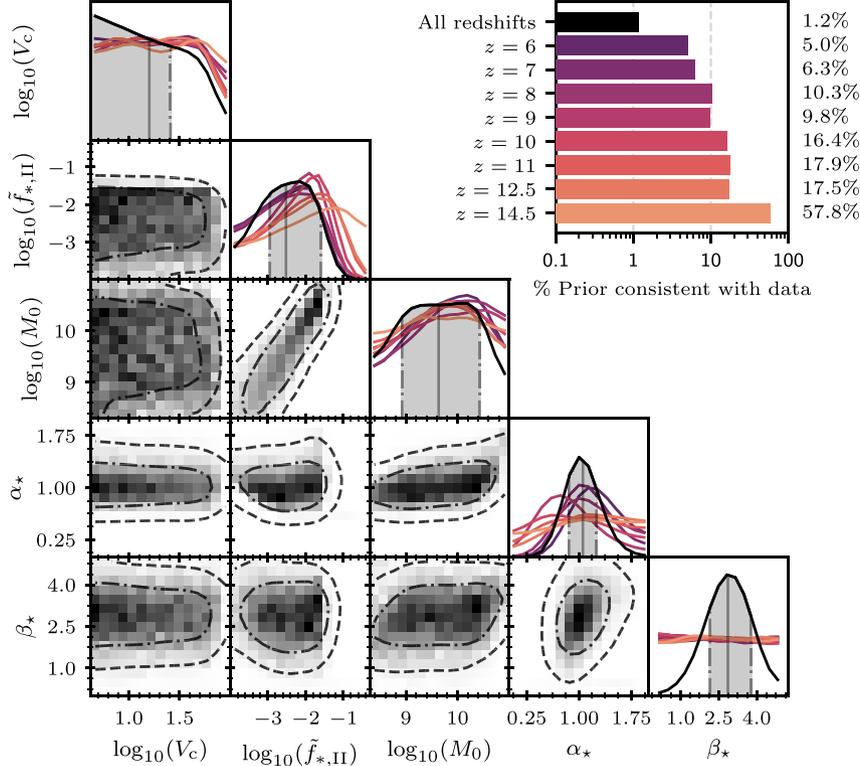
Figure 5. Marginal 1D and 2D posteriors of SFE parameters from the full UVLF data set and individual redshift bins. As expected, the lowest redshift bins are most constraining. Furthermore, since a single redshift bin contains no redshift-evolution information, the PDF of the $\beta_\star$ parameter remains flat, and the $\tilde{f}_{\star,\mathrm{II}}$ parameter compensates by shifting the peak of the 1D posterior PDF to higher values with increasing redshift in order to explain the high-z *JWST* observations.

Table 3. Marginal 1D constraints on the parameter V_c, and related quantities M_{vir} (equation 3) and T_{vir} (equation 6). The quoted errors are at 68% (95%) confidence levels around the weighted mean. We list SARAS 3 and UVLF individually since they are the most constraining data sets, providing lower and upper limits, respectively. The joint analysis combines the two, with SARAS 3 weakening the upper bound on V_c set by UVLFs as seen in Fig. C2 (or providing a lower bound depending on the foreground fit, as in Fig. 6).

Parameter	Prior	SARAS 3	UVLF	Joint w/o SARAS 3	Joint
$\log_{10}\left(V_c/\,\mathrm{km\,s^{-1}}\right)$	$\mathcal{U}\,(0.62,2)$ or	$\gtrsim 1.16\,(0.72)$	$\lesssim 1.41\,(1.79)$	$\lesssim 1.33\,(1.78)$	$\lesssim 1.40\,(1.79)$
$\Rightarrow V_c\,(\mathrm{km\,s^{-1}})$	$\log\mathcal{U}(4.2,100)$	$\gtrsim 14.4\,(5.3)$	$\lesssim 25.7\,(62.1)$	$\lesssim 21.3\,(60.4)$	$\lesssim 25.2\,(62.0)$
$\log_{10}(M_{\mathrm{vir}}(z=6)/\,\mathrm{M_\odot})$	$\mathcal{U}\,(6.5,10.6)$	$\gtrsim 8.09\,(6.77)$	$\lesssim 8.84\,(9.99)$	$\lesssim 8.60\,(9.95)$	$\lesssim 8.82\,(9.99)$
$\log_{10}(T_{\mathrm{vir}}/\,\mathrm{K})$	$\mathcal{U}\,(2.8,5.6)$	$\gtrsim 3.87\,(3.00)$	$\lesssim 4.38\,(5.14)$	$\lesssim 4.21\,(5.12)$	$\lesssim 4.36\,(5.14)$

even a marginal change in the foreground model can lead to stronger exclusion of low V_c values. This is seen in Fig. 6 (first panel) in the case of joint constraints without high-z UVLF data. In the absence of a more physically motivated foreground model, we refrain from over-interpreting this result but emphasize the potential for future 21-cm signal measurements in providing constraints on V_c.

Our results demonstrate the synergy that exists between the 21-cm signal data and high-z galaxy surveys from *HST/JWST*, even when the former is a non-detection, in constraining the star-forming properties of small haloes.

5.2 Constraints on the star formation efficiency

The 21-cm signal data is not informative enough yet to constrain the SFE parameters independently and, in the absence of the UVLF data, the 1D posterior PDF for the scaling of the SFE $\tilde{f}_{\star,\mathrm{II}}$ saturates to the lower limit of the prior at 10^{-4}, while the probability of the turning point for the SFE power-law slope M_0 saturates to the upper limit

$\sim 10^{11}\,\mathrm{M_\odot}$ (see Fig. C1). This is consistent with predictions from a fixed constant SFE model as adopted in Pochinda et al. (2024). Thus, constraints on the SFE parameters $\{\tilde{f}_{\star,\mathrm{II}}, M_0, \alpha_\star, \beta_\star\}$ are almost entirely driven by UVLF data. Table 4 summarizes our results at 68 per cent and 95 per cent confidence levels. As seen in the 2D posterior plot for $\tilde{f}_{\star,\mathrm{II}}$ and M_0 in Fig. 4, there is a strong degeneracy below $M_0 \lesssim 10^{10.5}\,\mathrm{M_\odot}$ which is well-described by a simple linear fit to the 2D histogram:

$$\log_{10}(M_0/\,\mathrm{M_\odot}) \approx 0.9 \times \log_{10}\left(\tilde{f}_{\star,\mathrm{II}}\right) + 12, \qquad (37)$$

before saturating at $\tilde{f}_{\star,\mathrm{II}} < 0.1$. This is as expected from the observational limit on faint galaxies, as discussed in the previous section. To better understand the constraining power of the UVLF data, we now perform different redshift cuts and analyse the SFE parameters inferred.
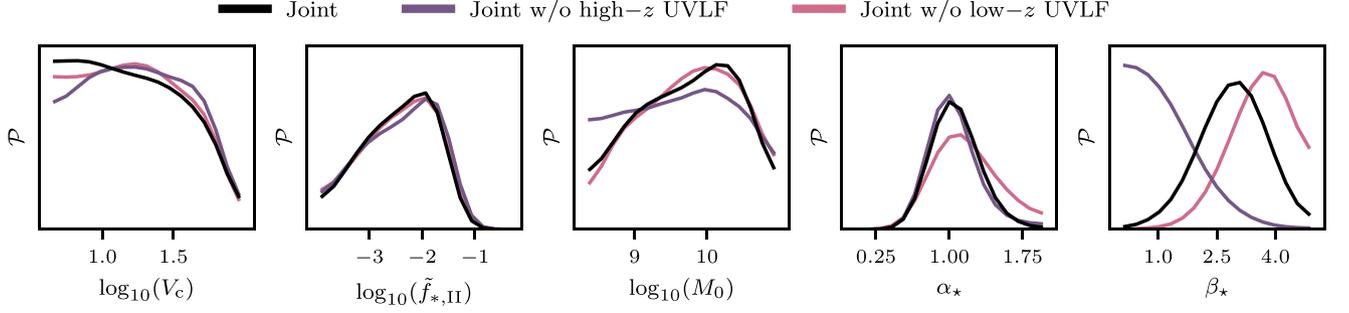
Figure 6. Marginal 1D posteriors of the SFE parameters $\tilde{f}_{\star,\mathrm{II}}$, M_0, $\alpha_\star$, and $\beta_\star$ from all observational data sets, with the UVLFs split into two redshift bins: $z = 7 - 10$ as the 'low-z' (purple) and $z = 11 - 14.5$ as the 'high-z' (pink) fits, with the $z = 6$ data points from B21 included in both (see Footnote 12). The main takeaway here is that the redshift-evolution term $\beta_\star$ showing little to no evolution at low-z (i.e. consistent with 0) and a rapid evolution at high-z. The V_c constraints, discussed in Section 5.1, show slight variation depending on the SARAS 3 foreground parameter, with the low-z fit exhibiting the strongest synergy between the UVLF and SARAS 3 data.

Table 4. Marginal 1D posteriors of the SFE parameters as seen in Fig. 6. The quoted errors are at 68% (95%) confidence levels around the weighted mean.

Parameter	Prior	Joint w/o high-zUVLF (low-z fit)	Joint w/o low-zUVLF (high-z fit)	Joint (joint fit)
$\log_{10}\left(\tilde{f}_{\star,\mathrm{II}}\right)$	$\mathcal{U}(-4, -0.3)$	$-2.4^{+0.9}_{-0.4}\left(^{+1.0}_{-1.2}\right)$	$-2.5^{+0.9}_{-0.4}\left(^{+1.0}_{-1.2}\right)$	$-2.4^{+0.9}_{-0.4}\left(^{+1.0}_{-1.2}\right)$
$\log_{10}\left(M_0/\mathrm{M}_\odot\right)$	$\mathcal{U}(8.3, 11)$	$9.6^{+0.8}_{-0.8}\left(^{+1.2}_{-1.3}\right)$	$9.8^{+0.9}_{-0.5}\left(^{+1.1}_{-1.2}\right)$	$9.7^{+0.8}_{-0.6}\left(^{+1.1}_{-1.2}\right)$
$\alpha_\star$	$\mathcal{U}(0, 2)$	$1.0^{+0.1}_{-0.2}\left(^{+0.5}_{-0.4}\right)$	$1.2^{+0.2}_{-0.4}\left(^{+0.7}_{-0.5}\right)$	$1.1^{+0.2}_{-0.2}\left(^{+0.4}_{-0.4}\right)$
$\beta_\star$	$\mathcal{U}(0, 5)$	$\lesssim 1.4\ (2.8)$	$3.6^{+0.7}_{-0.7}\left(^{+1.4}_{-1.3}\right)$	$2.9^{+0.9}_{-0.7}\left(^{+1.5}_{-1.6}\right)$

5.2.1 How much do individual redshifts contribute?

We first quantify the constraining power of the UVLF data from different redshift bins individually. This is shown in Fig. 5. The SFE parameter space is constrained to $\lesssim 10$ per cent of the prior volume by data at $z \approx 6-9$, $\lesssim 20$ per cent at $z \approx 10-13$, and ~ 60 per cent at $z \approx 14.5$ as the data becomes increasingly uncertain at high redshifts.

The constraint on V_c is consistent across the redshift bins. In particular, observations of faint galaxies at $z = 12.5$ provide just as good constraints as those at $z = 6$. Deeper observations in fainter magnitude bins, including lensed objects with *JWST*, are expected to provide more stringent upper limits on V_c in the near future.

Similarly, the constraints on the halo-mass dependence of the SFE, $\alpha_\star$, also show consistency between redshift bins. The scatter around the mean value inferred from the combined data sets is well within errors. Furthermore, since there is no redshift-evolution information encoded in individual redshift bins, the PDF of $\beta_\star$ is flat across all redshifts. To compensate for this lack of evolution, the peak of the PDF for $\tilde{f}_{\star,\mathrm{II}}$ shifts to higher values with increasing redshift, thus hinting at an enhanced SFE at early times.

Combining all redshifts gives us a ~ 1 per cent level volume consistency of prior and data, demonstrating the strength of using UVLF data from *HST* and *JWST* across a wide redshift range. This not only provides a tighter 1D posterior on $\alpha_\star$, but also encodes important information about the evolution of the SFE. We explore this evolution in the next section.

5.2.2 How does the SFE evolve with redshift?

In order to properly quantify the evolution of the SFE with redshift, $f_{\star,\mathrm{II}} \propto [(1 + z)/7]^{-\beta_\star}$, we now take a closer look at the $\beta_\star$ parameter.

We split the UVLF data set into two redshift bins, now including the rest of the observational data sets (SARAS 3, HERA, CXB and CRB data) for both, to compare to the joint constraint which uses the complete UVLF data set across all redshifts. We denote $z = 7-10$ as the 'low-z' cut and $z = 11-14.5$ as the 'high-z' cut, including the $z = 6$ data point from B21 in both.[12] The choice of redshift bins defines the high-redshift frontier before and after the launch of the *JWST*, hence informing us on how the new instrument has changed our understanding of the SFE. Fig. 6 shows the result of this analysis. The $\tilde{f}_{\star,\mathrm{II}}$ and $\alpha_\star$ parameters are consistent across the two redshift ranges and M_0 remains degenerate with $\tilde{f}_{\star,\mathrm{II}}$, as suggested by equation (37). However, $\beta_\star$ shows a stark contrast:

(i) Joint w/o high-z UVLF (or low-z fit for short): At $z \leq 10$, the UVLF data favours a very slow/non-evolving SFE. The 1D posterior PDF of $\beta_\star$ saturates at the lower edge of the prior, favouring $\lesssim 1.4$ at 68 per cent CL, and remains consistent with ~ 0. This behaviour aligns with pre-*JWST* inferences using *HST* data (Tacchella, Trenti & Carollo 2013; Mason, Trenti & Treu 2015; Bouwens et al. 2015, 2021a; Tacchella et al. 2018; Harikane et al. 2018; Stefanon et al. 2021), where a simple redshift-independent parametrization calibrated at $z \lesssim 4$ can reproduce the observed UVLFs at higher redshifts. This trend of slow evolution is also consistent up to $z \approx 10$ with the FirstLight numerical simulation suite (Ceverino et al. 2024).

---

[12]This 'anchoring' at $z = 6$ is done for two reasons. First, the SFE as defined in equation (7) and (8), is normalized to this redshift which means the $\beta_\star$ parameter has no effect at $z = 6$. Secondly, by fixing $\tilde{f}_{\star,\mathrm{II}}$, $\alpha_\star$ and M_0 to the best-fitting values at $z = 6$, we can isolate the effect of SFE evolution via $\beta_\star$.

(ii) Joint w/o low-z UVLF (or high-z fit for short): At $z > 10$, the UVLF data favours a rapidly evolving SFE with the 1D posterior PDF of $\beta_\star$ having a weighted mean of $\langle\beta_\star\rangle = 3.6$. This deviation from the low-z behaviour is an interesting outcome of the abundance of bright galaxies observed in the *JWST* data. The bright excess has been alternatively attributed in the literature to either a top-heavy IMF of the early populations of stars (Inayoshi et al. 2022; Wang et al. 2023; Trinca et al. 2024; Hutter et al. 2025), or stochasticity/variability in star-forming galaxies (Mason et al. 2023; Mirocha & Furlanetto 2023; Sun et al. 2023a, b; Pallottini & Ferrara 2023; Shen et al. 2023).[13] However, both these solutions give rise to other problems – for example, Cueto et al. (2024) find that a top-heavy IMF alone leads to strong stellar feedback which inhibits stellar mass growth. Similarly, Nikolić et al. (2024) and Yung et al. (2024) conclude that larger than physical scatter would be needed to explain the bright galaxies. Although a redshift dependence in the SFE has been previously proposed (e.g. Sun & Furlanetto 2016), we constrain this SFE evolution for the first time at high-z using a subgrid analytic prescription. The redshift evolution is quantitatively similar to the feedback free burst (FFB) galaxies explored in Dekel et al. (2023), Li et al. (2024), as well as the DMSFEs in Somerville et al. (2025). Indeed, we shall see how our inferred SFE posteriors align with the former in the next section. A similar transition from low SFE to high SFE at $z \sim 10$ has also been inferred in Chakraborty & Choudhury (2025, see their fig. 10).

(iii) Joint (or joint fit): The joint constraints, combining both low-z and high-z UVLF data, give a weighted mean of $\langle\beta_\star\rangle = 2.9$. This fit is effectively a compromise between the two redshift bins, showing a moderate evolution of the SFE. Although not an ideal fit at the lower redshifts due to small error-bars on the data, it provides a simple parametrization that can be applied across the entire redshift range.

Table 4 summarizes the constraints on the SFE parameters at 68 per cent and 95 per cent confidence levels for the different UVLF cuts. Equipped with this information, we can now look at the inferred SFE in the two redshift bins.

5.2.3 SFE functional posteriors and best fit

As suggested in the previous section, the SFE appears to have a complex evolution with time. Namely, the $\beta_\star$ parameter alone is not enough to capture the evolution at both low-z and high-z. Thus, our inference across the full redshift range $z = 6 - 14.5$ is essentially a composite of the low-z fit at $z \leq 10$ and the high-z fit at $z > 10$. In order to visualize the results from our Bayesian analysis in the two redshift bins, we use FGIVENX (Handley 2018) to convert our parameter posterior samples into functional posteriors of the SFE and its associated UVLF (as a cross-check to see how our model fits the data). This is shown in Fig. 7.

We also plot the posterior-mean model in the two redshift bins (Table 4 using equation 7) for comparison, alongside other empirical SFE models from literature. Note that other works may use different cosmologies and IMFs, with potentially different halo mass definitions and mass accretion models which could lead to biases compared to our results. Hence, renormalizing their results to our choice of cosmology and SFR–luminosity conversion is non-trivial and we show the SFE curves as presented in their works.

---

[13]For completeness, we note that other solutions have also been proposed. This includes the presence of a hidden population of AGN (Hegde et al. 2024), or simply that a formation epoch of $z \approx 15$ implies young, bright stellar population (requiring no new physics to explain current observations; Donnan et al. 2025).

At a glance, one can see that the 1σ SFE constraints are generally informative up to $M_h \gtrsim 10^{10}\,\mathrm{M_\odot}$ at $z = 6$ moving down to $M_h \gtrsim 10^9\,\mathrm{M_\odot}$ at $z = 14.5$. At lower halo masses, the SFE is unconstrained and hence, the degeneracy between $\tilde{f}_{\star,\mathrm{II}}$ and M_0 allows for a wide range of curves that flatten below the observed mass threshold. The shrinking of SFE priors (grey contours in Fig. 7) to posteriors (coloured contours in the same figure) at the low-mass end is due to our implicit assertion that the SFE is continuous across halo masses, which is not a priori guaranteed, but a reasonable assumption (but see also the 'bursty' or 'high scatter' model of Sipple & Lidz 2024, in their fig. 11). Furthermore, the steep contours for SFE and the drop in UVLF at the faint magnitude end (which is redshift dependent) are a physical feature – this is the direct result of the constraint on high values of minimum circular velocity V_c (equivalently T_{vir} or M_{vir}) for star formation in DM haloes.

At the lowest redshift, $z = 6$, we find strong agreement with the SFE magnitude and slope of Tacchella et al. (2018, T + 18) and Sun & Furlanetto (2016, SF16) to $\lesssim 2\sigma$, both of which predict $f_{\star,\mathrm{II}} \propto M_h$ (or equivalently, $M_\star \propto M_h^2$). The former is an empirical redshift-independent model based on informed accretion histories from N-body simulations, while the latter uses halo-abundance matching assuming smooth gas accretion. The results are also within 2σ of the deepest Spitzer and *HST*-calibrated results from Stefanon et al. (2021, S + 21). Compared to the recent work of Sipple & Lidz (2024, SL24, in particular their 'B21 + B22' model), although roughly consistent at 3σ, we find a steeper slope than their $f_{\star,\mathrm{II}} \propto M_h^{0.56}$ and lower efficiency at the faint end. The feedback-free burst model from Dekel et al. (2023); Li et al. (2024, L + 24, their $\epsilon_{\mathrm{max}} = 0.2$ model) is a factor of few lower than our prediction, which is expected because FFB galaxies are less abundant compared to non-FFB galaxies at this redshift.

At $z = 8$, we find an almost exact match with the SFE of T + 18 and SF16. We are also $\lesssim 2\sigma$ consistent with L + 24 across the halo mass range, and with SL24 at the high-mass end (deviating by a factor of few at the low-mass end). The UVLF best-fitting model at this redshift is in fact more consistent with latest spectroscopic measurements from *JWST* (Rojas-Ruiz et al. 2025), lying slightly above the B21 data points at the bright end.

At $z = 10$, we are consistent with most works at the low-mass end, while predicting an intermediate SFE at the high-mass end, between the FFB model of L + 24 and others. At $z > 10$, we predict an order of magnitude higher SFEs than T+18 and strongly favour the FFB model of L + 24 quantitatively.

This result recovers the slow evolution of the SFE inferred in pre-*JWST HST*-calibrated results at $z \lesssim 10$, and enhanced SFE from a rapid evolution at $z \gtrsim 10$ in post-*JWST* models (a transition consistent with the findings of Chakraborty & Choudhury 2025, although with a different parameterization). Across the entire redshift range, the SFE slope $\langle\alpha_\star\rangle \approx 1.1$ is a reasonable fit to data (see Table 4 for exact values and error-bars), which is consistent with T + 18 within errors.

In Fig. 8, we show the evolution of the SFE of DM haloes at $M_h = 10^{10}\,\mathrm{M_\odot}$ with redshift, and the 1σ upper limits for $M_h = 10^9\,\mathrm{M_\odot}$. We find,

$$f_{\star,\mathrm{II}}(M_h = 10^{10}\,\mathrm{M_\odot}) = \begin{cases} 1.3^{+0.5}_{-0.8} \text{ per cent at } z = 6 \text{ (low-}z\text{fit)}, \\ 1.7^{+0.4}_{-0.9} \text{ per cent at } z = 9 \text{ (low-}z\text{fit)}, \\ 8.2^{+2.1}_{-4.0} \text{ per cent at } z = 12 \text{ (high-}z\text{fit)}, \\ 20.8^{+3.3}_{-14.5} \text{ per cent at } z = 15 \text{ (high-}z\text{fit)}, \end{cases}$$
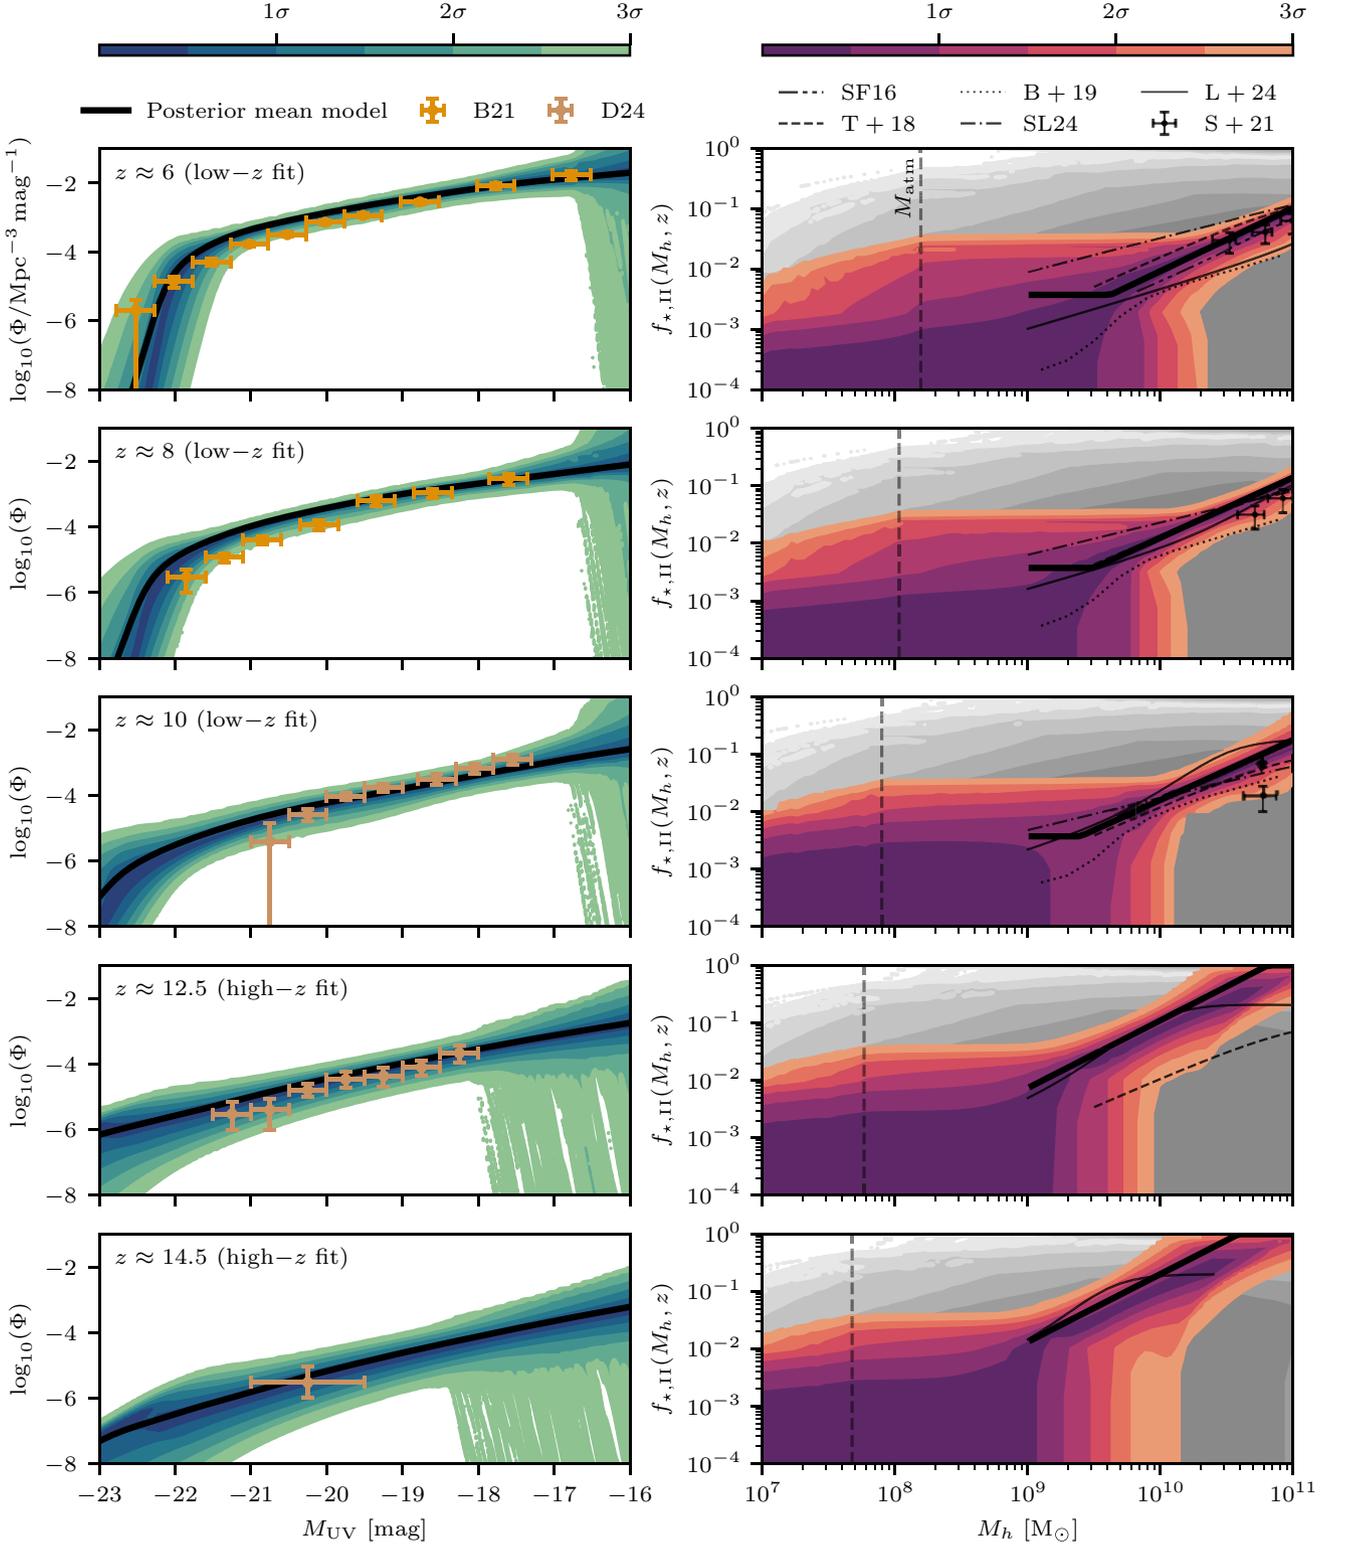
Figure 7. Functional posteriors of the UVLF (left) and constrained SFE (right) for the fit using UVLFs at low and high redshift (including all other data sets). In the upper three panels, $z = 6, 8, 10$ we use the low-z fit, while in the lower two panels, $z = 12.5, 14.5$ we use the high-z fit (see Fig. 6 and Table 4). In all the panels, the filled coloured contours show 1, 2, 3σ confidence of the functional posterior (also showing functional SFE priors in the right panels in grey contours), while the thick black line shows the posterior-mean model (i.e. the weighted means listed in Table 4). Note that the black line and contours need not align since $\langle f(x|\theta_{ast}) \rangle \neq f(x|\langle \theta_{ast} \rangle)$. The UVLF data (scatter points, left panels) used for the constraints are shown as a cross-check to see whether the model fits the data. In the right column, we compare the SFE posteriors to various models in the literature: halo-abundance matching results from Sun & Furlanetto (2016, SF16), empirical fits from Tacchella et al. (2018, T + 18), UniverseMachine (Behroozi et al. 2019, B + 19), semi-empirical model of Sipple & Lidz (2024, SL24), feedback-free burst model of Dekel et al. (2023); Li et al. (2024, L + 24), and *HST/Spitzer* based results from Stefanon et al. (2021, S + 21).
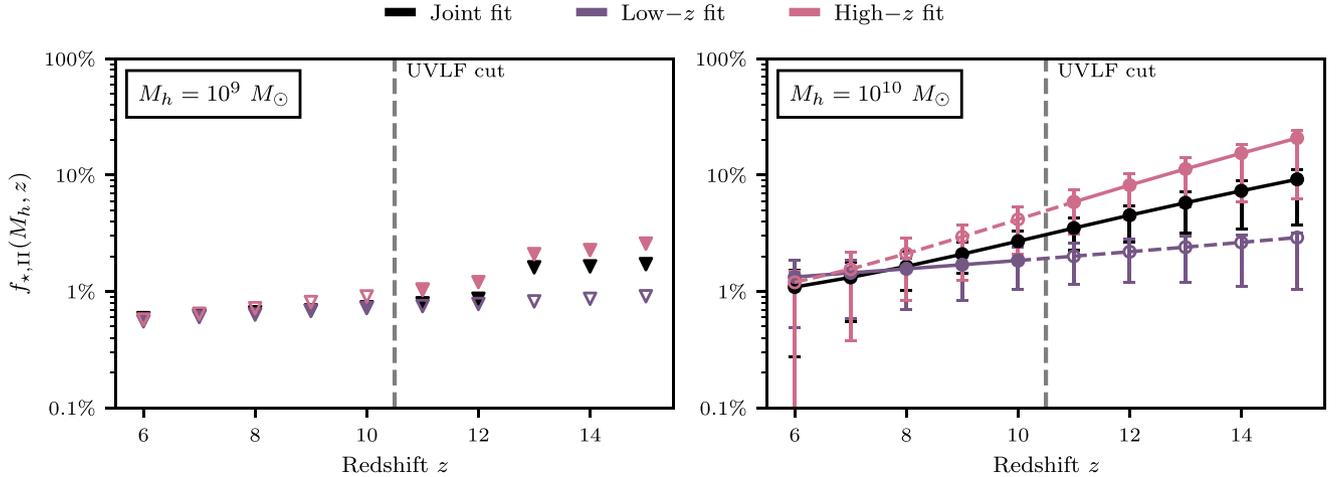
Figure 8. Evolution of SFE at fixed halo masses $M_h = 10^9 \, \mathrm{M_\odot}$ (left, triangles) and $M_h = 10^{10} \, \mathrm{M_\odot}$ (right, circles) with redshift, for the case of the three joint analyses: low-z UVLFs (purple) and high-z UVLFs (pink), separated by a cut at $z \approx 10.5$ (dashed vertical line), and all UVLFs (black). Unfilled markers and dashed coloured lines denote extrapolations outside of the redshift range of the fits. In the case of $M_h = 10^{10} \, \mathrm{M_\odot}$, the error bars are 1σ confidence levels while for $M_h = 10^9 \, \mathrm{M_\odot}$, the SFE are 1σ upper limits and relatively unconstrained due to the lack of UVLF data at the faint magnitude end. See Fig. 7 for the full SFE posteriors.

while that of $f_{\star,\mathrm{II}}(M_h = 10^9 \, \mathrm{M_\odot})$ remains at a few sub- per cent level, reaching $\approx 1{-}2$ per cent at the highest redshifts, and is weakly constrained due to the lack of UVLF data at the faint end. Such an increase of $\times 10$ in SFE at $M_h = 10^{10} \, \mathrm{M_\odot}$ from the global $1{-}3$ per cent (at low redshifts) to $10{-}30$ per cent or higher (at high redshifts) is consistent with observations of compact, super star clusters in nuclear starburst galaxies and giant molecular clouds (Kennicutt 1998; Keto, Ho & Lo 2005; Murray, Quataert & Thompson 2010; Leroy et al. 2018; Emig et al. 2020; Rico-Villas et al. 2020; Costa et al. 2021; Sun et al. 2024).[14] The increased SFE also quantitatively matches previous results in the literature (Inayoshi et al. 2022; Yung, Somerville & Iyer 2025) which conclude that high-z galaxy observations are not yet inconsistent with ΛCDM.

Thus, the UVLF is a powerful data set providing the strongest current constraints on early galaxies via the SFE. The data favours a constant SFE at $z \lesssim 10$ and a rapidly evolving SFE at $z \gtrsim 10$. We conduct the first multiwavelength analysis utilizing high-z galaxy observations out to $z \approx 15$.

5.2.4 SFRD evolution

We show the evolution of the SFRD in Fig. 9. In order to compare our results directly with the literature, we compute the SFRD by integrating the UVLFs down from a magnitude of $M_{\mathrm{UV}} = -23$ mag to $M_{\mathrm{UV}} = -17$ mag, instead of the simulation volume-averaged SFRD defined in equation (1). The latter includes the contribution from faint galaxies hosted in haloes above V_c, which can be as low as $\sim 10^6 \, \mathrm{M_\odot}$ at $z = 15$ and $\sim 10^5 \, \mathrm{M_\odot}$ at $z = 30$ for low V_c values (equation 3), and contribute significantly to the SFR due to a flattening SFE slope. The luminosities are converted to SFR

---

[14]These high efficiencies are also achieved in numerical simulations of compact giant molecular clouds with high surface densities (Kim, Kim & Ostriker 2018; Fukushima & Yajima 2021; see also Wang et al. 2025 for interesting discussions on GMC to galaxy scale SFE).
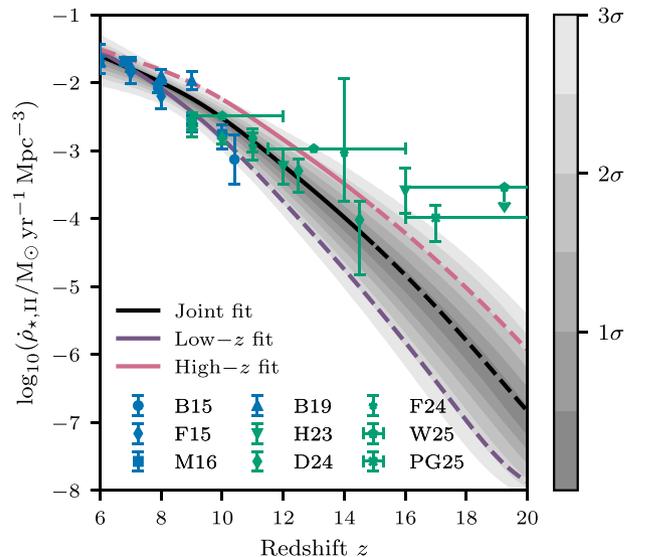
Figure 9. Evolution of the SFRD, calculated by integrating the UVLF ($\dot{\rho}_{\star,\mathrm{II}} = \int \Phi \dot{M}_{\star,\mathrm{II}} \, dM_{\mathrm{UV}}$) in the magnitude range $M_{\mathrm{UV}} \in [-23, -17]$. The filled contours show functional posteriors of the joint fit, while the lines show the posterior-mean model (see Table 4) for the joint fit (black), low-z fit (purple), and high-z fit (pink). The lines are solid in the redshift range of the fit and dashed outside where the model is extrapolated. Various literature values are shown for comparison: results from *HST* in blue (Bouwens et al. 2015; Finkelstein et al. 2015; McLeod, McLure & Dunlop 2016; Bhatawdekar et al. 2019) and *JWST* in green (Harikane et al. 2023a; Donnan et al. 2024; Finkelstein et al. 2024; Whitler et al. 2025; Pérez-González et al. 2025).

using the conversion factor $\kappa_{\mathrm{UV}} = 1.15 \times 10^{-28} \, \mathrm{M_\odot \, yr^{-1} \, erg^{-1} \, s \, Hz}$, consistent with the UVLF modelling described in Section 2.4.

By construction in equation (8), the SFRDs for the low-z and high-z fits are consistent at $z = 6$, but diverge at higher redshifts. The low-z fit aligns with literature values based on *HST* observations at $z \lesssim 10$ (e.g. Bouwens et al. 2015; Finkelstein et al. 2015; McLeod et al.

Table 5. Marginal 1D constraints on the radiative efficiencies f_X and f_r as seen in Figs 4 and C1. The quoted errors are at 68% (95%) confidence levels around the weighted mean. We find that the constraints are stronger and more accurate when UVLF data is included.

Parameter	Prior	Joint w/o UVLF	Joint w/ UVLF
$\log_{10}(f_X)$	$\mathcal{U}(-3, 3)$	$0.05^{+1.76}_{-1.07}\left(^{+2.24}_{-2.82}\right)$	$-0.10^{+1.12}_{-0.31}\left(^{+1.33}_{-2.22}\right)$
$\Rightarrow f_X$	$\log\mathcal{U}(10^{-3}, 10^3)$	$1.1^{+63.1}_{-1.0}\left(^{+192.1}_{-1.1}\right)$	$0.8^{+9.7}_{-0.4}\left(^{+16.3}_{-0.8}\right)$
$\log_{10}(f_r)$	$\mathcal{U}(-1, 5)$	$\lesssim 1.45\,(3.36)$	$\lesssim 1.23\,(2.39)$
$\Rightarrow f_r$	$\log\mathcal{U}(10^{-1}, 10^5)$	$\lesssim 28.0\,(2282.6)$	$\lesssim 16.9\,(243.7)$

2016), while the high-z fit is consistent with *JWST* observations at $z \gtrsim 10-15$ (e.g. Harikane et al. 2023a; Donnan et al. 2024; Whitler et al. 2025). At $z \gtrsim 15$ however, outside the range of our fits, the tentative observations of Pérez-González et al. (2025) suggest even higher SFRDs. The difference between FFB and non-FFB models of Libanore et al. (2024, see their fig. 1 where the curves differ by an order of magnitude) is quantitatively similar to the low-z and high-z fits in our analysis. Furthermore, since the SFRD converges at the lowest redshifts – which contribute most to the diffuse X-ray and radio backgrounds – the constraints on the radiative efficiencies f_X and f_r, described in the next section, are robust to the choice of the $\beta_\star$ parameter.

5.3 Constraints on f_X and f_r

We find that the constraints on the X-ray and radio efficiencies of early galaxies, f_X and f_r as defined in equation (16) and equation (19), respectively, become tighter with the inclusion of UVLF data by breaking their degeneracy with the SFE of galaxies. Fig. C1 shows the 2D joint posteriors for $\{\tilde{f}_{\star,\mathrm{II}}, f_X, f_r\}$ when UVLF data is included and excluded. In the absence of UVLF data, the parameter $\tilde{f}_{\star,\mathrm{II}}$ saturates to $\sim 10^{-4}$, allowing for larger values of f_X and f_r (and indeed this matches the inference of Pochinda et al. 2024).

Several different data sets contribute to constraints on the X-ray efficiency of early galaxies f_X. CXB data (i.e. Chandra and other satellites) sets upper limits on the diffuse X-ray background disfavouring high values of f_X; on the other hand, HERA favours high values as efficient heating would suppress the 21-cm signal (as found previously in HERA Collaboration 2023; Pochinda et al. 2024). This squeezes the prior from either ends of its range resulting in a measurement of f_X.[15] Furthermore, since the diffuse background scales as $f_X \times$ SFR, the inclusion of UVLF data breaks the degeneracy between $f_{\star,\mathrm{II}}$ and f_X, anchoring the former to observed galaxies as discussed earlier. The joint analysis then yields a measurement of f_X:

$$\log_{10}(f_X) = -0.10^{+1.12}_{-0.31}\left(^{+1.33}_{-2.22}\right) \text{ at } 68\% \ (95\%) \text{ CL}$$
$$\Rightarrow f_X = 0.8^{+9.7}_{-0.4}\left(^{+16.3}_{-0.8}\right) \text{ at } 68\% \ (95\%) \text{ CL}.$$

which is strongly consistent with theoretical predictions of low metallicity HMXBs at high-z ($f_X = 1$; Fragos et al. 2013b). This is indeed a tighter constraint than the one obtained without UVLF data (see Table 5). Hence, $f_X \gtrsim 11$ is ruled out at 1σ confidence, $f_X \gtrsim 17$ is ruled out at 2σ confidence, and $f_X \gtrsim 49$ is completely

ruled out (5σ). At the low f_X end, the constraints are weaker. Although disfavouring low f_X values, we do not rule them out. This is on account of the degeneracy between τ and f_X (see 2D posterior in Fig. 4), where an advanced reionization at $z \approx 8$ (i.e. higher τ) suppresses the 21-cm power spectrum enough to allow for weak X-ray heating. Inclusion of reionization-era data sets (e.g. Qin et al. 2021) can potentially tighten this constraint further, and shed light on whether weak X-ray efficiencies in early galaxies are indeed allowed.

The upper limits from CRB data (i.e. LWA1 and other telescopes) directly constrain the radio efficiency of early galaxies f_r by disfavouring high values. The upper limits from HERA data provide a secondary constraint disfavouring high values of f_r which would result in bright 21-cm signals.[16] The joint analysis gives us an upper limit on f_r:

$$\log_{10}(f_r) < 1.23\,(2.39) \text{ at } 68\% \ (95\%) \text{ CL}$$
$$\Rightarrow f_r < 16.9\,(243.7) \text{ at } 68\% \ (95\%) \text{ CL}.$$

As with the f_X parameter, the UVLF data strengthens the constraints by $\times 10$ at 2σ (see Table 5). Thus, in this model, $f_r \gtrsim 17$ is ruled out at 1σ confidence, $f_r \gtrsim 244$ is ruled out at 2σ confidence, and $f_r \gtrsim 1800$ is completely ruled out (5σ). The 2D joint posterior in the $f_X - f_r$ plane, shown in Fig. 4, has an interesting feature ruling out even lower values of $f_r \gtrsim 100$ at 2σ confidence for $f_X \lesssim 0.1$. This constraint comes from the HERA data disfavouring large 21-cm signals – a low heating efficiency (i.e. f_X) of galaxies would imply that the matter temperature can cool down for longer, leading to a large 21-cm signal. Hence, f_r cannot be too high as it would further increase the contrast between the matter temperature and radio background temperature, leading to an even larger 21-cm signal. Recently, Sikder et al. (2024b) derived constraints on f_r from clustering of radio sources (not included here) using the anisotropy upper limits from VLA at 4.9 GHz and ATCA at 8.7 GHz (Holder 2014), which amount to $f_r(z = 7) \lesssim 20$ (although assuming a fixed SFE of 10 per cent). Our results here are as stringent as this at 1σ, showing the power of UVLF data. Future works will combine the analysis here with radio source clustering data to further tighten the constraint on f_r.

The main caveat of our results here is the assumption of a continuous model for X-ray and radio emission across redshifts. The diffuse X-ray and radio backgrounds are integrated observables, building up over cosmic time, with the largest contribution coming from galaxies at $z \approx 6$ where the SFRD constraints are tightest.[17] Thus, our inference of f_X and f_r parameters is most robust close to $z = 6$ and becomes increasingly model-dependent at higher redshifts. Just as there is evidence of SFE varying with z (via $\beta_\star$), the f_X and f_r could indeed be very different for early galaxy populations (e.g. Pop III galaxies as in Gessey-Jones et al. 2025 or for metallicity-dependent XRBs as in Kaur et al. 2022). This is a limitation of our analysis, and we leave the exploration of such models to future works.

By jointly constraining the CXB and CRB data with UVLFs, we have significantly tightened the constraints on the X-ray and

---

[15]These constraints build upon previous works where only one of the two data set combinations are used: UVLF + HERA in HERA Collaboration 2023, and UVLF + CXB in Zvi Katz et al. 2025, both of which give strong constraint on one end of the f_X prior range.

[16]For both parameters, f_X and f_r, SARAS 3 constraints are weak with only a marginal preference for higher f_X and lower f_r values. This is because the contribution of faint high-z galaxies (driving the signal in the SARAS 3 band) to the diffuse X-ray and radio backgrounds is small. Compared to Pochinda et al. (2024), we probe an order of magnitude lower SFEs, especially for the low-mass haloes abundant at high-z, leading to weaker constraints.

[17]Since this is also where the SFRDs converge between the low-z and high-z fit, the f_X and f_r constraints are insensitive to $\beta_\star$.

radio properties of galaxies in the early Universe. This highlights the importance of a well-calibrated SFE model, which enables flexibility and supports broad priors in poorly probed regimes, such as low-mass and high-redshift galaxies.

6 CONCLUSIONS

In this work, we constrain astrophysical parameters of galaxies in the early Universe combining a suite of 30 000 simulations performed using the code 21CMSPACE with multiwavelength observations. The simulations span the redshift range $z = 6 - 50$, and we vary eight astrophysical parameters: the minimum circular velocity for star formation in DM haloes V_c, four SFE parameters which describe a halo mass and redshift dependent SFE of Pop II galaxies $\left(\tilde{f}_{\star,\mathrm{II}}, M_0, \alpha_\star, \beta_\star \right)$, the ionization efficiency of galaxies ζ_{ion}, X-ray emission efficiency of galaxies f_X, and the radio emission efficiency of galaxies f_r.

In order to constrain these parameters with recent data from *HST/JWST*, we implement a pixel-level UVLF model as a new observable in 21CMSPACE, including an analytic prescription for dust attenuation. We then train neural networks to emulate five key observables as functions of the astrophysical parameters: the 21-cm global signal and power spectrum, the present-day CXB, the CRB, and UVLFs. We perform Bayesian analysis to infer our astrophysical parameters using the following observational data sets: 21-cm global signal non-detection from SARAS 3 (Singh et al. 2022), 21-cm power spectrum upper limits from HERA Phase 1 observations (HERA Collaboration 2023), CXB measurements from Chandra and other X-ray telescopes (Hickox & Markevitch 2006), CRB measurements from LWA1 and other telescopes (Dowell & Taylor 2018), UVLF measurements from *HST* at redshifts $z = 6 - 9$ (Bouwens et al. 2021a) and *JWST* at redshifts $z = 9 - 14.5$ (Donnan et al. 2024). We also indirectly include the *Planck 2018* measurement of optical depth of CMB photons to reionization $\tau = 0.054 \pm 0.07$ (Planck Collaboration VI 2020) with a 3σ prior on it, as a proxy of excluding ζ_{ion} values that lead to very late or very early reionization.

Our joint analysis thus builds on the work of Pochinda et al. (2024) and improves upon it by incorporating the flexible SFE model and inclusion of UVLF data for constraints. Our main findings include the following:

(i) The minimum circular velocity for star formation in DM haloes V_c is constrained by a combination of SARAS 3 data and UVLF data. The UVLF data provides the dominant constraint, disfavouring large V_c values via direction observation of faint magnitude galaxies. This upper limit is weakened by the SARAS 3 data which either disfavours or begins to rule out (depending on the foreground fit) low V_c values due to the absence of a deep 21-cm absorption signal in their observed redshift band $z = 15 - 25$. The joint analysis with and without the SARAS 3 data yields $V_c \lesssim 21.3 \, \mathrm{km \, s^{-1}}$ and $V_c \lesssim 25.2 \, \mathrm{km \, s^{-1}}$ at 68 per cent confidence, respectively. We thus demonstrate the synergy that exists between 21-cm signal (even non-detections) and high-z galaxy observations in tightening the constraints on this parameter, improving upon works that only use either UVLF data (Sipple & Lidz 2024; Zvi Katz et al. 2025) or 21-cm signal data (Pochinda et al. 2024). The main challenge in the analysis is the degeneracy between V_c and the log-log polynomial foreground model used to fit SARAS 3 data. Using physically motivated foregrounds, future works can tighten this constraint further.

(ii) The SFE model, as illustrated in Fig. 1, consists of four parameters: $\tilde{f}_{\star,\mathrm{II}}$ scales the magnitude of the SFE, M_0 sets the halo mass below which the curve flattens, $\alpha_\star$ is the power-law slope

of the M_h-dependence, and $\beta_\star$ controls the redshift-dependence anchored to $z = 6$. These parameters are almost entirely constrained by the UVLF data with other data sets providing much weaker constraints. The data favours a non-monotonic evolution of the SFE with redshift: slow evolution at $z \approx 6 - 10$ with $\beta_\star \lesssim 1.4$ (and the posterior saturating to zero, consistent with no evolution), and then a rapid evolution at $z \approx 10 - 15$ with $\beta_\star = 3.6^{+0.7}_{-0.7}$ at 68 per cent confidence. The joint fit combining both redshift bins gives a maximum posterior estimate $\beta_\star = 2.9^{+0.9}_{-0.7}$. Thus, we find consistency with pre-*JWST HST*-calibrated results at low redshifts (Sun & Furlanetto 2016; Tacchella et al. 2018), and post-*JWST* enhanced SFE models at high redshifts (in particular the feedback-free burst model of Dekel et al. 2023; Li et al. 2024, and the density-modulated SFE of Somerville et al. 2025). The slope of the SFE with halo mass $\alpha_\star = 1.1^{+0.2}_{-0.2}$ is consistent across the two redshift bins for $M_h \gtrsim 10^{10} \, \mathrm{M_\odot}$, in agreement with the findings of Tacchella et al. (2018). Furthermore, under our composite model, we infer that DM haloes of mass $M_h = 10^{10} \, \mathrm{M_\odot}$ have a SFE of $1 - 2$ per cent at $z \lesssim 10$, 8 per cent at $z = 12$ and 21 per cent at $z = 15$ (for the full functional posteriors of the SFE in the low and high-redshift regime, see Fig. 7).

(iii) The X-ray efficiency of early galaxies f_X, defined relative to the theoretical predictions of low-metallicity, high-mass X-ray binaries (HMXBs; Fragos et al. 2013b) in equation (16), is predominantly constrained by a combination of CXB and HERA data. CXB data disfavour very efficient sources by their direct contribution to the diffuse background seen today. HERA disfavours very inefficient X-ray sources, as late heating would lead to stronger fluctuations in the 21-cm brightness temperature at $z = 7 - 11$. The joint analysis gives us a measurement of $f_X = 0.8^{+9.7}_{-0.4} \left(^{16.3}_{-0.8} \right)$ at 68 per cent (95 per cent) confidence levels. We rule out $f_X \gtrsim 17$ at 2σ confidence and $f_X \gtrsim 49$ at 5σ confidence level. Our findings imply that the X-ray properties of early Pop II galaxies at $z \gtrsim 6$ are consistent with predictions of low-metallicity HMXBs from Fragos et al. (2013a). Note that this constraint becomes increasingly model-dependent at higher redshifts, since the CXB is an integrated observable (with largest contribution from galaxies at $z \sim 6$) and we assume a constant model for X-ray emission across redshifts (see e.g. Kaur et al. 2022, for an alternate redshift-evolving model).

(iv) The radio efficiency of early galaxies f_r, defined relative to those in the local Universe in equation (19), is constrained by a combination of CRB and HERA data (with a weak contribution from SARAS 3). CRB data disfavours high radio efficiency of galaxies by their direct contribution to the diffuse background seen today. HERA further disfavours strong radio sources since they would boost the amplitude of 21-cm signal fluctuations beyond observed upper limits. The joint analysis results in an upper limit $f_r \lesssim 16.9 \, (243.7)$ at 68 per cent (95 per cent) confidence levels. We rule out $f_r \gtrsim 1800$ at 5σ confidence level. At 1σ, this is comparable to the constraints obtained by Sikder et al. (2024b) using radio source clustering data from VLA and ATCA (not included here).

(v) The joint analysis without UVLF data provides weaker constraints on f_X and f_r due their degeneracy with an unconstrained SFE in the calculation of the diffuse backgrounds (e.g. disfavouring $f_X \gtrsim 193$ and $f_r \gtrsim 2280$ at 2σ). By including UVLF data and thus informing the model of SFE in observed galaxies, we break this degeneracy and tighten the constraints by a factor of ~ 10 (disfavouring $f_X \gtrsim 17$ and $f_r \gtrsim 244$ at 2σ; see also Zvi Katz et al. 2025, whose result supports this conclusion).

In conclusion, we have shown the critical role of UVLFs in our understanding of the Cosmic Dawn and EoR: they set an upper limit

on the minimum mass of star-forming haloes, constrain the SFE of galaxies across cosmic epochs, and thus also break degeneracies with other astrophysical parameters constrained by different wavelength probes.

Exciting prospects remain for future works on multiwavelength inferences. The analysis done here can be extended to include other reionization era data sets such as Lyman-line, Lyα forest and quasar observations to constrain ionizing properties of galaxies (Qin et al. 2021, 2025; Sims et al. 2025), flexible Pop III models to constrain their role during the Cosmic Dawn (Gessey-Jones et al. 2024; Pochinda et al. 2024; Zvi Katz et al. 2025), as well as deeper *JWST* data (e.g. $z = 15 - 25$, Pérez-González et al. 2025, published during the final stages of preparation of this manuscript). Furthermore, future synergies using cross-correlations with CMB (La Plante, Sipple & Lidz 2022), line intensity mapping experiments (see, e.g. Bernal & Kovetz 2022, for a recent review) and the cosmic near-infrared background (CNIRB; Mao 2014; Sun et al. 2021) will enable a comprehensive multiwavelength multiscale picture of the early Universe IGM and ISM in the coming years.

DATA AVAILABILITY

The 30 000 simulations performed using 21CMSPACE, the trained emulators, the POLYCHORD chains, and other analysis scripts used in this work can all be made available upon reasonable request to the corresponding author.

REFERENCES

Abadi M. et al., 2016, preprint (arXiv:1603.04467)
Acharya S. K., Dhandha J., Chluba J., 2022, MNRAS, 517, 2454
Adams N. J. et al., 2024, ApJ, 965, 169
Ajello M. et al., 2008, ApJ, 689, 666
Alvarez H., Aparici J., May J., Olmos F., 1997, A&AS, 124, 205
Asthana S., Kulkarni G., Haehnelt M. G., Bolton J. S., Keating L. C., Simmonds C., 2025, MNRAS, 539, L18
Austin D. et al., 2023, ApJ, 952, L7
Bambi C., Sangangelo A., 2022, Handbook of X-ray and Gamma-ray Astrophysics. Springer Nature
Barkana R., 2016, Phys. Rep., 645, 1
Barkana R., Loeb A., 2001, Phys. Rep., 349, 125
Barkana R., Loeb A., 2004, ApJ, 609, 474
Barkana R., Loeb A., 2005, ApJ, 626, 1
Behroozi P. S., Silk J., 2015, ApJ, 799, 32
Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, MNRAS, 488, 3143
Bernal J. L., Kovetz E. D., 2022, A&AR, 30, 5
Bevins H. T. J., Handley W. J., Fialkov A., de Lera Acedo E., Javid K., 2021, MNRAS, 508, 2923
Bevins H. T. J., Fialkov A., de Lera Acedo E., Handley W. J., Singh S., Subrahmanyan R., Barkana R., 2022a, Nature Astron., 6, 1473
Bevins H. T. J., de Lera Acedo E., Fialkov A., Handley W. J., Singh S., Subrahmanyan R., Barkana R., 2022b, MNRAS, 513, 4507
Bevins H. T. J., Handley W. J., Lemos P., Sims P. H., de Lera Acedo E., Fialkov A., Alsing J., 2023, MNRAS, 526, 4613
Bevins H. T. J., Heimersheim S., Abril-Cabezas I., Fialkov A., de Lera Acedo E., Handley W., Singh S., Barkana R., 2024, MNRAS, 527, 813
Bhatawdekar R., Conselice C. J., Margalef-Bentabol B., Duncan K., 2019, MNRAS, 486, 3805
Bogdán Á. et al., 2024, Nature Astron., 8, 126
Bolliet B., Spurio Mancini A., Hill J. C., Madhavacheril M., Jense H. T., Calabrese E., Dunkley J., 2024, MNRAS, 531, 1351
Bouwens R. J. et al., 2014, ApJ, 793, 115
Bouwens R. J. et al., 2015, ApJ, 803, 34
Bouwens R. J. et al., 2021a, AJ, 162, 47
Bouwens R. J., Illingworth G. D., van Dokkum P. G., Ribeiro B., Oesch P. A., Stefanon M., 2021b, AJ, 162, 255
Bouwens R. J., Illingworth G., Ellis R. S., Oesch P., Stefanon M., 2022, ApJ, 940, 55
Bowman J. D., Rogers A. E. E., Monsalve R. A., Mozdzen T. J., Mahesh N., 2018, Nature, 555, 67
Boylan-Kolchin M., 2025, MNRAS, 538, 3210
Bradley R. F., Tauscher K., Rapetti D., Burns J. O., 2019, ApJ, 874, 153
Brandenberger R., Cyr B., Shi R., 2019, J. Cosmology Astropart. Phys., 2019, 009
Breitman D., Mesinger A., Murray S. G., Prelogović D., Qin Y., Trotta R., 2024, MNRAS, 527, 9833
Bull P. et al., 2024, preprint (arXiv:2410.00076)
Calzetti D., 2001, PASP, 113, 1449
Carniani S. et al., 2024, Nature, 633, 318
Castellano M. et al., 2023, ApJ, 948, L14
Ceverino D., Klessen R. S., Glover S. C. O., 2018, MNRAS, 480, 4842
Ceverino D., Nakazato Y., Yoshida N., Klessen R. S., Glover S. C. O., 2024, A&A, 689, A244
Chakraborty A., Choudhury T. R., 2025, preprint (arXiv:2503.07590)
Chen X., Miralda-Escudé J., 2004, ApJ, 602, 1
Chiaki G., Susa H., Hirano S., 2018, MNRAS, 475, 4378
Churazov E. et al., 2007, A&A, 467, 529
Ciardi B., Madau P., 2003, ApJ, 596, 1
Ciardi B., Salvaterra R., Di Matteo T., 2010, MNRAS, 401, 2635
Cohen A., Fialkov A., Barkana R., 2016, MNRAS, 459, L90
Cohen A., Fialkov A., Barkana R., Monsalve R. A., 2020, MNRAS, 495, 4845
Costa A. H., Johnson K. E., Indebetouw R., Finn M. K., Brogan C. L., Reines A., 2021, ApJ, 918, 76
Cueto E. R., Hutter A., Dayal P., Gottlöber S., Heintz K. E., Mason C., Trebitsch M., Yepes G., 2024, A&A, 686, A138
Cyr B., Acharya S. K., Chluba J., 2024, MNRAS, 534, 738

Davé R., Finlator K., Oppenheimer B. D., 2012, MNRAS, 421, 98

Davies F. B., Bosman S. E. I., Furlanetto S. R., 2024, preprint (arXiv:2406.18186)

Dawson K. S. et al., 2013, AJ, 145, 10

Dayal P., Ferrara A., Dunlop J. S., Pacucci F., 2014, MNRAS, 445, 2545

de Lera Acedo E. et al., 2022, Nature Astron., 6, 1332

DeBoer D. R. et al., 2017, PASP, 129, 045001

Dekel A., Mandelker N., 2014, MNRAS, 444, 2071

Dekel A., Sarkar K. C., Birnboim Y., Mandelker N., Li Z., 2023, MNRAS, 523, 3201

DESI Collaboration, 2024, AJ, 168, 58

Donnan C. T. et al., 2024, MNRAS, 533, 3222

Donnan C. T., Dunlop J. S., McLure R. J., McLeod D. J., Cullen F., 2025, MNRAS, 539, 2409

Dorigo Jones J., Rapetti D., Mirocha J., Hibbard J. J., Burns J. O., Bassett N., 2023, ApJ, 959, 49

Dowell J., Taylor G. B., 2018, ApJ, 858, L9

Dowell J., Taylor G. B., Schinzel F. K., Kassim N. E., Stovall K., 2017, MNRAS, 469, 4537

Emig K. L. et al., 2020, ApJ, 903, 50

Ewall-Wice A. et al., 2016, MNRAS, 460, 4320

Ewall-Wice A., Chang T. C., Lazio J., Doré O., Seiffert M., Monsalve R. A., 2018, ApJ, 868, 63

Ewall-Wice A., Chang T.-C., Lazio T. J. W., 2020, MNRAS, 492, 6086

Feng C., Holder G., 2018, ApJ, 858, L17

Ferrara A., Pallottini A., Sommovigo L., 2025, A&A, 694, A286

Fialkov A., Barkana R., 2014, MNRAS, 445, 213

Fialkov A., Barkana R., 2019, MNRAS, 486, 1763

Fialkov A., Barkana R., Tseliakovich D., Hirata C. M., 2012, MNRAS, 424, 1335

Fialkov A., Barkana R., Visbal E., Tseliakovich D., Hirata C. M., 2013, MNRAS, 432, 2909

Fialkov A., Barkana R., Visbal E., 2014, Nature, 506, 197

Fialkov A., Cohen A., Barkana R., Silk J., 2017, MNRAS, 464, 3498

Fialkov A., Gessey-Jones T., Dhandha J., 2024, Phil. Trans. R. Soc. London Ser. A, 382, 20230068

Field G. B., 1958, Proc. IRE, 46, 240

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

Finkelstein S. L. et al., 2024, ApJ, 969, L2

Fiore F., Ferrara A., Bischetti M., Feruglio C., Travascio A., 2023, ApJ, 943, L27

Fixsen D. J. et al., 2011, ApJ, 734, 5

Fragos T. et al., 2013a, ApJ, 764, 41

Fragos T., Lehmer B. D., Naoz S., Zezas A., Basu-Zych A., 2013b, ApJ, 776, L31

Fraser S. et al., 2018, Phys. Lett. B, 785, 159

Frontera F. et al., 2007, ApJ, 666, 86

Fujimoto S. et al., 2024, ApJ, 977, 250

Fukushima H., Yajima H., 2021, MNRAS, 506, 5512

Furlanetto S. R., Mirocha J., 2022, MNRAS, 511, 3895

Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004, ApJ, 613, 1

Furlanetto S. R., Oh S. P., Briggs F. H., 2006, Phys. Rep., 433, 181

Furlanetto S. R., Mirocha J., Mebane R. H., Sun G., 2017, MNRAS, 472, 1576

Furtak L. J. et al., 2024, Nature, 628, 57

Gardner J. P. et al., 2006, Space Sci. Rev., 123, 485

Gehrz R., 1989, in Allamandola L. J., Tielens A. G. G. M., eds, Proc. IAU Symp. Vol. 135, Interstellar Dust. Kluwer Academic Publishers. p. 445

Gelli V., Mason C., Hayward C. C., 2024, ApJ, 975, 192

Gessey-Jones T. et al., 2022, MNRAS, 516, 841

Gessey-Jones T., Fialkov A., de Lera Acedo E., Handley W. J., Barkana R., 2023, MNRAS, 526, 4262

Gessey-Jones T., Pochinda S., Bevins H. T. J., Fialkov A., Handley W. J., de Lera Acedo E., Singh S., Barkana R., 2024, MNRAS, 529, 519

Gessey-Jones T. et al., 2025, Nature Astron., 9, 1268

Gnedin N. Y., 2000, ApJ, 542, 535

Grimm H. J., Gilfanov M., Sunyaev R., 2003, MNRAS, 339, 793

Gruber D. E., Matteson J. L., Peterson L. E., Jung G. V., 1999, ApJ, 520, 124

Guo Q., White S., Li C., Boylan-Kolchin M., 2010, MNRAS, 404, 1111

Gürkan G. et al., 2018, MNRAS, 475, 3010

Hainline K. N. et al., 2024, ApJ, 964, 71

Handley W., 2018, J. Open Source Soft., 3, 849

Handley W., 2019, J. Open Source Soft., 4, 1414

Handley W. J., Hobson M. P., Lasenby A. N., 2015a, MNRAS, 450, L61

Handley W. J., Hobson M. P., Lasenby A. N., 2015b, MNRAS, 453, 4384

Hardcastle M. J. et al., 2016, MNRAS, 462, 1910

Harikane Y. et al., 2018, PASJ, 70, S11

Harikane Y. et al., 2023a, ApJS, 265, 5

Harikane Y. et al., 2023b, ApJ, 959, 39

Harrison F. A. et al., 2016, ApJ, 831, 185

Hartwig T. et al., 2022, ApJ, 936, 45

Haslam C. G. T., Salter C. J., Stoffel H., Wilson W. E., 1982, A&AS, 47, 1

Hegde S., Wyatt M. M., Furlanetto S. R., 2024, J. Cosmology Astropart. Phys., 2024, 025

HERA Collaboration, 2022a, ApJ, 924, 51

HERA Collaboration, 2022b, ApJ, 925, 221

HERA Collaboration, 2023, ApJ, 945, 124

Hickox R. C., Markevitch M., 2006, ApJ, 645, 95

Hills R., Kulkarni G., Meerburg P. D., Puchwein E., 2018, Nature, 564, E32

Holder G. P., 2014, ApJ, 780, 112

Hutter A., 2018, MNRAS, 477, 1549

Hutter A., Cueto E. R., Dayal P., Gottlöber S., Trebitsch M., Yepes G., 2025, A&A, 694, A254

Inayoshi K., Harikane Y., Inoue A. K., Li W., Ho L. C., 2022, ApJ, 938, L10

Jaacks J., Finkelstein S. L., Bromm V., 2019, MNRAS, 488, 2202

Jana R., Nath B. B., Biermann P. L., 2019, MNRAS, 483, 5329

Jin X. et al., 2023, ApJ, 942, 59

Jishnu Nambissan T. et al., 2021, Exp. Astron., 51, 193

Kannan R., Garaldi E., Smith A., Pakmor R., Springel V., Vogelsberger M., Hernquist L., 2022, MNRAS, 511, 4005

Kaur H. D., Qin Y., Mesinger A., Pallottini A., Fragos T., Basu-Zych A., 2022, MNRAS, 513, 5097

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

Keto E., Ho L. C., Lo K. Y., 2005, ApJ, 635, 1062

Kim J.-G., Kim W.-T., Ostriker E. C., 2018, ApJ, 859, 68

Kravtsov A., Belokurov V., 2024, preprint (arXiv:2405.04578)

Kuhlen M., Madau P., Montgomery R., 2006, ApJ, 637, L1

Kulkarni G., Worseck G., Hennawi J. F., 2019, MNRAS, 488, 1035

La Plante P., Sipple J., Lidz A., 2022, ApJ, 928, 162

Larson R. L. et al., 2023, ApJ, 953, L29

Lehmer B. D., Alexander D. M., Bauer F. E., Brandt W. N., Goulding A. D., Jenkins L. P., Ptak A., Roberts T. P., 2010, ApJ, 724, 559

Lehmer B. D. et al., 2016, ApJ, 825, 7

Lehmer B. D. et al., 2021, ApJ, 907, 17

Leitherer C. et al., 1999, ApJS, 123, 3

Leroy A. K. et al., 2018, ApJ, 869, 126

Leung G. C. K. et al., 2023, ApJ, 954, L46

Lewis A., Challinor A., 2011, Astrophysics Source Code Library, record ascl:1102.026

Li Z., Dekel A., Sarkar K. C., Aung H., Giavalisco M., Mandelker N., Tacchella S., 2024, A&A, 690, A108

Libanore S., Flitter J., Kovetz E. D., Li Z., Dekel A., 2024, MNRAS, 532, 149

Lovell C. C., Vijayan A. P., Thomas P. A., Wilkins S. M., Barnes D. J., Irodotou D., Roper W., 2021, MNRAS, 500, 2127

Ma X. et al., 2018, MNRAS, 478, 1694

Machacek M. E., Bryan G. L., Abel T., 2001, ApJ, 548, 509

Madau P., Dickinson M., 2014, ARA&A, 52, 415

Madau P., Meiksin A., Rees M. J., 1997, ApJ, 475, 429

Maeda K., Alvarez H., Aparici J., May J., Reich P., 1999, A&AS, 140, 145

Magg M. et al., 2022, MNRAS, 514, 4433
Maiolino R. et al., 2024, A&A, 691, A145
Mao X.-C., 2014, ApJ, 790, 148
Marshall F. E., Boldt E. A., Holt S. S., Miller R. B., Mushotzky R. F., Rose L. A., Rothschild R. E., Serlemitsos P. J., 1980, ApJ, 235, 4
Mason C. A., Trenti M., Treu T., 2015, ApJ, 813, 21
Mason C. A., Trenti M., Treu T., 2023, MNRAS, 521, 497
Mauerhofer V., Dayal P., Haehnelt M. G., Kimm T., Rosdahl J., Teyssier R., 2025, A&A, 696, A157
McClymont W. et al., 2025, preprint (arXiv:2503.00106)
McGreer I. D., Mesinger A., D'Odorico V., 2015, MNRAS, 447, 499
McLeod D. J., McLure R. J., Dunlop J. S., 2016, MNRAS, 459, 3812
McLeod D. J. et al., 2024, MNRAS, 527, 5004
Mebane R. H., Mirocha J., Furlanetto S. R., 2018, MNRAS, 479, 4544
Mebane R. H., Mirocha J., Furlanetto S. R., 2020, MNRAS, 493, 1217
Mertens F. G. et al., 2020, MNRAS, 493, 1662
Mesinger A., 2019, The Cosmic 21-cm Revolution; Charting the First Billion Years of Our Universe. IOP Publishing
Mesinger A., Furlanetto S., Cen R., 2011, MNRAS, 411, 955
Meurer G. R., Heckman T. M., Calzetti D., 1999, ApJ, 521, 64
Mineo S., Gilfanov M., Sunyaev R., 2012, MNRAS, 419, 2095
Mineo S., Gilfanov M., Lehmer B. D., Morrison G. E., Sunyaev R., 2014, MNRAS, 437, 1698
Mirabel I. F., Dijkstra M., Laurent P., Loeb A., Pritchard J. R., 2011, A&A, 528, A149
Mirocha J., Furlanetto S. R., 2019, MNRAS, 483, 1980
Mirocha J., Furlanetto S. R., 2023, MNRAS, 519, 843
Mitra S., Chatterjee A., 2023, MNRAS, 523, L35
Mittal S., Kulkarni G., 2022, MNRAS, 510, 4992
Mondal R. et al., 2020, MNRAS, 498, 4178
Monsalve R. A., Fialkov A., Bowman J. D., Rogers A. E. E., Mozdzen T. J., Cohen A., Barkana R., Mahesh N., 2019, ApJ, 875, 67
Monsalve R. A. et al., 2024, MNRAS, 530, 4125
Moster B. P., Naab T., White S. D. M., 2018, MNRAS, 477, 1822
Muñoz J. B., 2023, MNRAS, 523, 2587
Muñoz J. B., Qin Y., Mesinger A., Murray S. G., Greig B., Mason C., 2022, MNRAS, 511, 3657
Murray N., Quataert E., Thompson T. A., 2010, ApJ, 709, 191
Naoz S., Barkana R., Mesinger A., 2009, MNRAS, 399, 369
Nikolić I., Mesinger A., Davies J. E., Prelogović D., 2024, A&A, 692, A142
Oesch P. A., Bouwens R. J., Illingworth G. D., Labbé I., Stefanon M., 2018, ApJ, 855, 105
Oh S. P., 2001, ApJ, 553, 499
Oke J. B., Gunn J. E., 1983, ApJ, 266, 713
Pacucci F., Mesinger A., Mineo S., Ferrara A., 2014, MNRAS, 443, 678
Pallottini A., Ferrara A., 2023, A&A, 677, L4
Pallottini A. et al., 2022, MNRAS, 513, 5621
Park J., Mesinger A., Greig B., Gillet N., 2019, MNRAS, 484, 933
Pedregosa F. et al., 2011, J. Mach. Learn. Res., 12, 2825
Pérez-González P. G. et al., 2023, ApJ, 951, L1
Pérez-González P. G. et al., 2025, preprint (arXiv:2503.15594)
Philip L. et al., 2019, J. Astron. Instr., 8, 1950004
Planck Collaboration XVI, 2014, A&A, 571, A16
Planck Collaboration VI, 2020, A&A, 641, A6
Pochinda S. et al., 2024, MNRAS, 531, 1113
Pritchard J. R., Furlanetto S. R., 2007, MNRAS, 376, 1680
Pritchard J. R., Loeb A., 2012, Rep. Prog. Phys., 75, 086901
Purcell E. M., Field G. B., 1956, ApJ, 124, 542
Qin Y., Mesinger A., Bosman S. E. I., Viel M., 2021, MNRAS, 506, 2390
Qin Y. et al., 2025, PASA, 42, e049
Reich P., Testori J. C., Reich W., 2001, A&A, 376, 861
Reis I., Fialkov A., Barkana R., 2020, MNRAS, 499, 5993
Reis I., Fialkov A., Barkana R., 2021, MNRAS, 506, 5479
Reis I., Barkana R., Fialkov A., 2022a, MNRAS, 511, 5265
Reis I., Barkana R., Fialkov A., 2022b, ApJ, 933, 51
Remazeilles M., Dickinson C., Banday A. J., Bigot-Sazy M. A., Ghosh T., 2015, MNRAS, 451, 4311

Rico-Villas F., Martín-Pintado J., González-Alfonso E., Martín S., Rivilla V. M., 2020, MNRAS, 491, 4573
Ricotti M., Ostriker J. P., 2004, MNRAS, 352, 547
Robertson B. et al., 2024, ApJ, 970, 31
Roger R. S., Costain C. H., Landecker T. L., Swerdlyk C. M., 1999, A&AS, 137, 7
Rojas-Ruiz S. et al., 2025, ApJ, 985, 80
Schauer A. T. P., Liu B., Bromm V., 2019, ApJ, 877, L5
Schmidt M., 1959, ApJ, 129, 243
Scott D., Rees M. J., 1990, MNRAS, 247, 510
Seager S., Sasselov D. D., Scott D., 2011, Astrophysics Source Code Library, record ascl:1106.026
Shen X., Vogelsberger M., Boylan-Kolchin M., Tacchella S., Kannan R., 2023, MNRAS, 525, 3254
Sikder S., Barkana R., Fialkov A., Reis I., 2024a, MNRAS, 527, 10975
Sikder S., Barkana R., Fialkov A., 2024b, ApJ, 970, L25
Silk J., 1997, ApJ, 481, 703
Silk J., Begelman M. C., Norman C., Nusser A., Wyse R. F. G., 2024, ApJ, 961, L39
Simmonds C. et al., 2024, MNRAS, 535, 2998
Sims P. H., Pober J. C., 2020, MNRAS, 492, 22
Sims P. H. et al., 2025, preprint (arXiv:2504.09725)
Singh S., Subrahmanyan R., 2019, ApJ, 880, 26
Singh S. et al., 2017, ApJ, 845, L12
Singh S. et al., 2018, ApJ, 858, 54
Singh S. et al., 2022, Nature Astron., 6, 607
Sipple J., Lidz A., 2024, ApJ, 961, 50
Sobacchi E., Mesinger A., 2013, MNRAS, 432, 3340
Somerville R. S., Yung L. Y. A., Lancaster L., Menon S., Sommovigo L., Finkelstein S. L., 2025, preprint (arXiv:2505.05442)
Stefanon M., Bouwens R. J., Labbé I., Illingworth G. D., Gonzalez V., Oesch P. A., 2021, ApJ, 922, 29
Sun G., Furlanetto S. R., 2016, MNRAS, 460, 417
Sun G., Mirocha J., Mebane R. H., Furlanetto S. R., 2021, MNRAS, 508, 1954
Sun G., Faucher-Giguère C.-A., Hayward C. C., Shen X., 2023a, MNRAS, 526, 2665
Sun G., Faucher-Giguère C.-A., Hayward C. C., Shen X., Wetzel A., Cochrane R. K., 2023b, ApJ, 955, L35
Sun J. et al., 2024, ApJ, 967, 133
Tacchella S., Trenti M., Carollo C. M., 2013, ApJ, 768, L37
Tacchella S., Bose S., Conroy C., Eisenstein D. J., Johnson B. D., 2018, ApJ, 868, 92
Tingay S. J. et al., 2013, PASA, 30, e007
Topping M. W. et al., 2024, MNRAS, 529, 4087
Trinca A., Schneider R., Valiante R., Graziani L., Ferrotti A., Omukai K., Chon S., 2024, MNRAS, 529, 3563
Tseliakhovich D., Barkana R., Hirata C. M., 2011, MNRAS, 418, 906
Vale A., Ostriker J. P., 2004, MNRAS, 353, 189
van Haarlem M. P. et al., 2013, A&A, 556, A2
Venkatesan A., Giroux M. L., Shull J. M., 2001, ApJ, 563, 1
Venumadhav T., Dai L., Kaurov A., Zaldarriaga M., 2018, Phys. Rev. D, 98, 103513
Verner D. A., Ferland G. J., Korista K. T., Yakovlev D. G., 1996, ApJ, 465, 487
Visbal E., Barkana R., Fialkov A., Tseliakhovich D., Hirata C. M., 2012, Nature, 487, 70
Visbal E., Bryan G. L., Haiman Z., 2020, ApJ, 897, 95
Wang Y.-Y., Lei L., Yuan G.-W., Fan Y.-Z., 2023, ApJ, 954, L48
Wang Z. et al., 2025, preprint (arXiv:2505.05554)
Weingartner J. C., Draine B. T., 2001, ApJ, 548, 296
White M., 2001, A&A, 367, 27
Whitler L. et al., 2025, preprint (arXiv:2501.00984)
Wouthuysen S. A., 1952, AJ, 57, 31
Wyithe J. S. B., Loeb A., 2004, Nature, 432, 194
Yajima H., Khochfar S., 2015, MNRAS, 448, 654
Yue M. et al., 2024, ApJ, 966, 176

Yung L. Y. A., Somerville R. S., Finkelstein S. L., Wilkins S. M., Gardner J. P., 2024, MNRAS, 527, 5929

Yung L. Y. A., Somerville R. S., Iyer K. G., 2025, preprint (arXiv:2504.18618)

Zahn O., Mesinger A., McQuinn M., Trac H., Cen R., Hernquist L. E., 2011, MNRAS, 414, 727

Zhang Z., Wang H., Luo W., Zhang J., Mo H. J., Jing Y., Yang X., Li H., 2022, A&A, 663, A85

Zhao R. J., Furlanetto S. R., 2024, J. Cosmology Astropart. Phys., 2024, 018

Zhu Y. et al., 2023, ApJ, 955, 115

Ziparo F., Ferrara A., Sommovigo L., Kohandel M., 2023, MNRAS, 520, 2445

Zvi Katz O., Redigolo D., Volansky T., 2025, preprint (arXiv:2502.03525)

APPENDIX A: IMPACT OF FIXED UVLF PARAMETERS

Here, we briefly assess the impact of the fixed but uncertain parameters such as $t_{\star,II}$, δ_{max} and κ_{UV} on the UVLF.

(i) Impact of $t_{\star,II}$: The effect of the averaging time-scale $t_{\star,II}$ on the UVLF is to shift the entire distribution by a fixed $\Delta M_{UV} = 2.5 \times \Delta \log(t_{\star,II})$. This is because the UVLF is a direct tracer of the SFR, which is scaled by $t_{\star,II}$. In Fig. A1, we see the effect of varying $t_{\star,II} = 0.2 \to 0.5$ resulting in a shift in magnitude $\Delta M_{UV} = 1$ mag. The parameter is completely degenerate with $\tilde{f}_{\star,II}$, introduced in equation (7), and can be subsumed into its interpretation (see Footnote 4). Since we typically account for an error magnitude of 20 per cent from semi-numerical modelling in our Bayesian analysis, the effect of $t_{\star,II}$ (varied by a factor of $2-3$) is within the error budget of our inference.

(ii) Impact of δ_{max}: Next, we find that the UVLF is sensitive to the overdensity cap δ_{max}, which determines the largest overdensity allowed in the simulation, in a notable way. It directly affects the abundance of the brightest galaxies in the simulation, and we attempt to set this parameter in an informed manner for our adopted hybrid HMF (as described in Footnote 3) to $\delta_{max} = 1.3$. The effect is apparent at low redshifts and becomes less pronounced at $z \gtrsim 10$ as overdensities are less evolved. We leave the development of structure-formation in the non-linear regime within the framework of our semi-numerical simulations to a future work.

(iii) Impact of κ_{UV}: Finally, the effect of assuming a fixed metallicity Salpeter IMF or a Chabrier IMF to calculate the SFR-L_{UV} conversion factor κ_{UV}, compared to the Madau & Dickinson (2014) value, proves to be sub-dominant compared to the effect of δ_{max} or $t_{\star,II}$. Note that we do not compare to Pop III top-heavy IMF, which can indeed have an elevated κ_{UV}, instead choosing to focus our inference on Pop-II type star formation which is generally dominant at the redshifts of interest ($z < 15$, e.g. Mebane, Mirocha & Furlanetto 2018; Jaacks, Finkelstein & Bromm 2019; Visbal, Bryan & Haiman 2020; Hartwig et al. 2022). Furthermore, Cueto et al. (2024) suggested that a top-heavy IMF alone is insufficient to explain the observed UVLFs on account of the resulting strong stellar feedback.
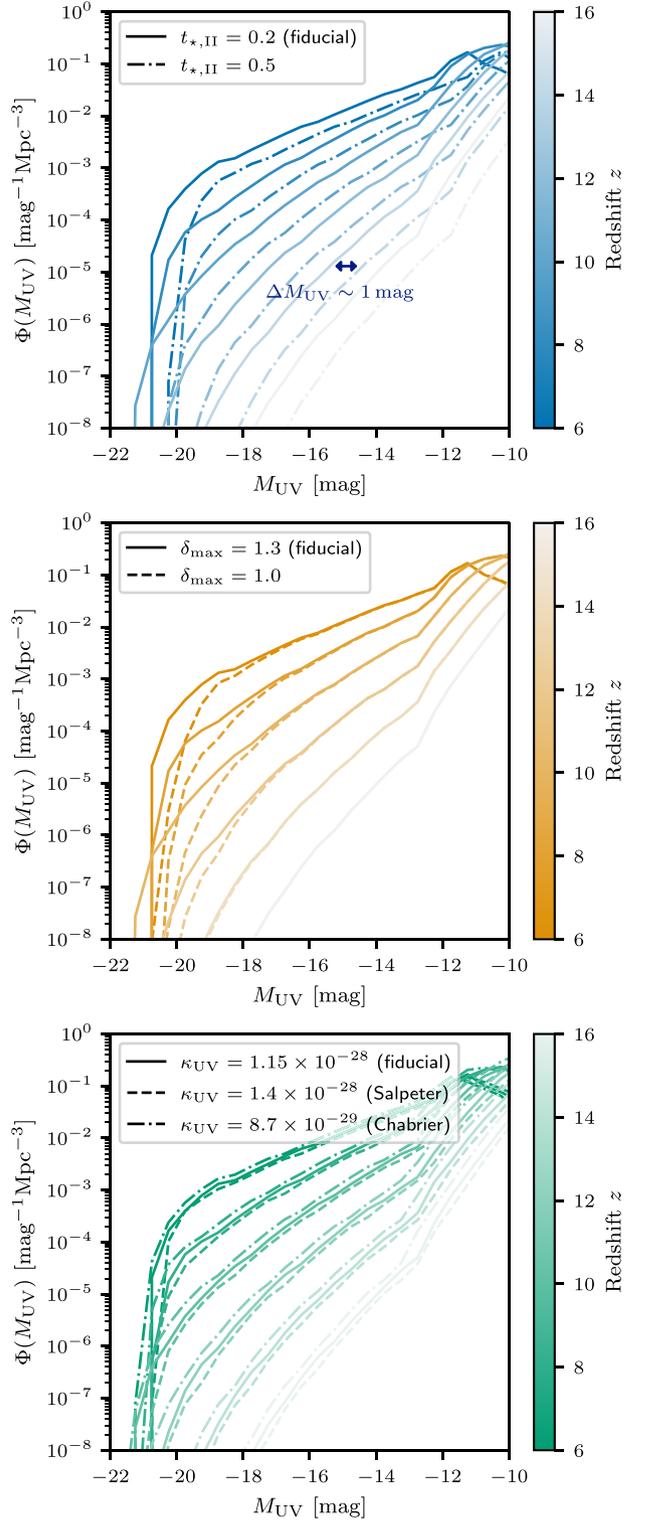
Figure A1. Assessing the impact of the fixed parameters in our UVLF modelling. Top: Changing the averaging time-scale $t_{\star,II}$ from $0.2H(z)^{-1}$ (fiducial) to $0.5H(z)^{-1}$ shifts the UVLF by $\Delta M_{UV} = 1$ mag. This parameter is completely degenerate with $\tilde{f}_{\star,II}$. Middle: Varying the threshold for linear overdensities δ_{max} in the simulation affects the abundance of bright galaxies. The effect is most pronounced at low redshifts. Bottom: The effect of varying the SFR-L_{UV} conversion factor κ_{UV}, from the fiducial value based on Madau & Dickinson (2014), to a fixed-metallicity Salpeter and Chabrier IMF. The effect is sub-dominant to the impact of $t_{\star,II}$ or δ_{max}.

Table A1. Summary of works on astrophysical constraints using the code 21CMSPACE. The astrophysical parameters used in each work are listed, along with the observational data used. For details on the parameters, please refer to the cited works – the model changes between works, but the qualitative interpretation of parameters is similar. Previous works use a constant Pop II SFE in the atomic cooling regime ($f_\star$ or $f_{\star,\mathrm{II}}$). In this work, we modify and introduce halo mass and redshift-dependent SFE with the efficiency level normalized to $\tilde{f}_{\star,\mathrm{II}}$. We also introduce new parameters $\alpha_\star$, M_0, and $\beta_\star$ (see equation 7). The † indicates improved upper limits from HERA Phase 1 (HERA Collaboration 2023), as opposed to the earlier limits (HERA Collaboration 2022b).

Work	Focus of the work	Astrophysical parameters	Observational data
Fialkov et al. (2017)	CXB constraint	$V_c, \tau, \alpha_X, E_{\mathrm{min,X}}, f_X$	CXB
Singh et al. (2017)	SARAS 2 first analysis	$V_c, f_\star, \tau, R_{\mathrm{mfp}}, \alpha_X, E_{\mathrm{min,X}}, f_X$	SARAS 2
Singh et al. (2018)	SARAS 2 reanalysis	$V_c, f_\star, \tau, R_{\mathrm{mfp}}, \alpha_X, E_{\mathrm{min,X}}, f_X$	SARAS 2
Monsalve et al. (2019)	EDGES High-Band analysis	$V_c, f_\star, \tau, R_{\mathrm{mfp}}, \alpha_X, E_{\mathrm{min,X}}, f_X$	EDGES
Mondal et al. (2020)	LOFAR excess radio constraints	$V_c, f_\star, \tau, R_{\mathrm{mfp}}, A_r$	LOFAR
Bevins et al. (2022b)	SARAS 2 Bayesian analysis	$V_c, f_\star, \tau, R_{\mathrm{mfp}}, \alpha_X, E_{\mathrm{min,X}}, f_X, f_r$	SARAS 2
HERA Collaboration (2022a)	HERA excess radio constraints	$V_c, f_\star, \tau, f_X, f_r, A_r$	HERA P1
Bevins et al. (2022a)	SARAS 3 Bayesian analysis	$V_c, f_\star, \tau, R_{\mathrm{mfp}}, \alpha_X, E_{\mathrm{min,X}}, f_X, f_r, A_r$	SARAS 3
HERA Collaboration (2023)	HERA excess radio constraints	$V_c, f_\star, \tau, R_{\mathrm{mfp}}, \alpha_X, \nu_{\mathrm{min,X}}, f_X, f_r$	HERA P1†
Bevins et al. (2024)	21-cm joint constraint	$V_c, f_\star, \tau, f_X, f_r$	HERA P1, SARAS 2/3, LOFAR, MWA
Gessey-Jones et al. (2024)	Cosmic string constraints	$V_c, f_{\star,\mathrm{II}}, f_{\star,\mathrm{III}}, t_{\mathrm{delay}}, \tau, \alpha_X, E_{\mathrm{min,X}}, f_X, A_r$	HERA P1†, SARAS 3, CXB
Pochinda et al. (2024)	Multiwavelength synergies	$V_c, f_{\star,\mathrm{II}}, f_{\star,\mathrm{III}}, t_{\mathrm{delay}}, \tau, \alpha_X, E_{\mathrm{min,X}}, f_X, f_r$	HERA P1†, SARAS 3, CXB, CRB
Sikder, Barkana & Fialkov (2024b)	Radio galaxy clustering constraints	f_r, f_X	CRB, ATCA, VLA
This work	Enhanced SFE and ULVFs	$V_c, \tilde{f}_{\star,\mathrm{II}}, \tau, M_0, \alpha_\star, \beta_\star, f_X, f_r$	HERA P1†, SARAS 3, CXB, CRB, UVLF

APPENDIX B: EMULATOR ACCURACY

Here, we analyse the accuracy of our emulators in predicting the observables of interest, and describe how we incorporate this into our likelihoods. Often, the test data is used as a validation set to calculate the loss of an emulator (whether it is root-mean squared error, or mean absolute error). For a 1D observable target function $f(x|\theta)$ such as the 21-cm global signal $T_{21}(z|\theta_{\mathrm{ast}})$, where x is the 'input space' and θ is the parameter space, these metrics are calculated by averaging all test data points, and across the input space. However, we can extend this idea to characterize the error as a function of the input space x, instead of a fixed average value:

$$\epsilon(x|\theta_i) = \left| \frac{f_{\mathrm{theory}}(x|\theta_i) - f_{\mathrm{emu}}(x|\theta_i)}{f_{\mathrm{theory}}(x|\theta_i)} \right|. \tag{B1}$$

where θ_i is the ith parameter set in the test data, $f_{\mathrm{theory/emu}}$ are the theory (simulated) and emulated observables respectively. Instead of averaging, we can sort the errors (over θ_i, keeping x bins separate) in ascending order and define 68 per cent and 95 per cent confidence intervals which yields $\epsilon_{1\sigma}(x)$ and $\epsilon_{2\sigma}(x)$, respectively. This gives us an idea of where in the x-space the emulator is performing well, and provides a way to accurately account for this error in final analysis. This method has been previously explored in the context of posterior validation for 21-cm emulators (Dorigo Jones et al. 2023;

Breitman et al. 2024) and for CMB emulators (Bolliet et al. 2024). In particular, its utility in our analysis is to check whether the emulator is performing well in the region of the input space we care about the most – i.e. where the observational data lies. For example, in case of the 21-cm signal, we are interested in the SARAS 3 band of $z = 15 - 25$. The emulator accuracy outside this range is less important. This can be easily extended to 2D functions like 21-cm power spectrum $\Delta_{21}^2(k, z|\theta_{\mathrm{ast}})$ or UVLFs $\Phi(M_{\mathrm{UV}}, z|\theta_{\mathrm{ast}})$.

Fig. B1 shows the errors on all emulators at the 68 per cent and 95 per cent confidence levels, alongside the observational data. Note that, since the global signal potentially crosses zero, we define the error metric in equation (B1) with $\max(T_{21,\mathrm{true}}(z|\theta_{\mathrm{ast}}))$ instead of $T_{21,\mathrm{true}}(z|\theta_{\mathrm{ast}})$ in the denominator. In principle, this metric can be defined in a variety of ways (using squared error instead of absolute, for example).

Once the emulator error is characterized in this way, we calculate the total 'model error' as the incoherent sum of both simulation error and emulator error:

$$\sigma_{\mathrm{model}}(x) = f(x) \times \sqrt{\epsilon_{\mathrm{theory}}^2 + \epsilon_{\mathrm{emu},1\sigma}^2(x)}, \tag{B2}$$

where $\epsilon_{\mathrm{theory}} = 20$ per cent is the assumed error on the semi-numerical model implemented in 21CMSPACE. This is then used to calculate the likelihood of the data given the model, as described in Section 4.2.
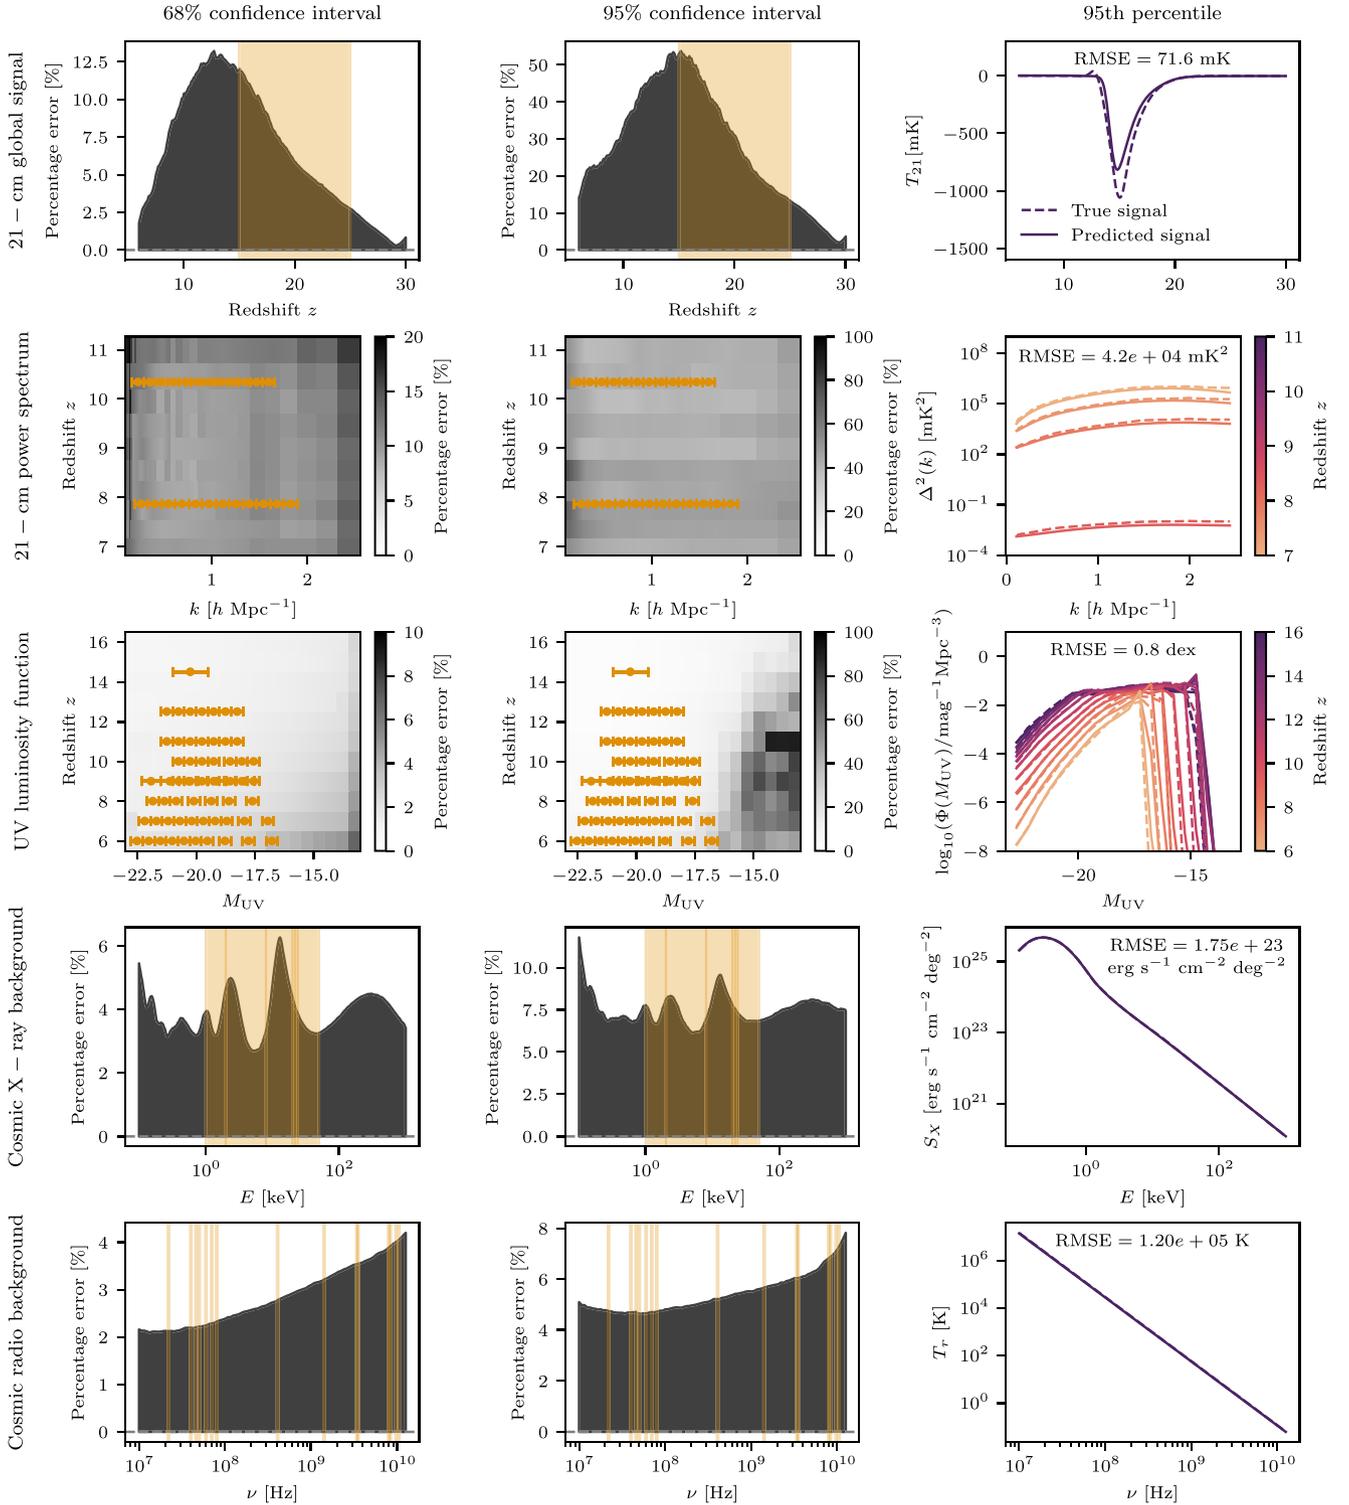
Figure B1. Errors on the 1D emulators for the 21-cm global signal, CXB, CRB, and 2D emulators for the 21-cm power spectrum and UVLFs, as a function of the input space (e.g. z for the global signal) shown at the 68 per cent and 95 per cent confidence levels calculated using the test data. The third column shows the test data alongside its emulated counterpart, at the 95th percentile of RMSE. The data and shaded regions in yellow show the observational data used in the constraints, for comparison.

APPENDIX C: EXTRA TABLES AND FIGURES

Table A1 shows a summary of past works constraining astrophysics using 21CMSPACE listed for ease of parsing literature and comparing results.

Fig. C1 shows the joint constraints on the astrophysical parameter space with and without the UVLF data set, for comparison with the joint analysis in Fig. 4. Although the constraints are model dependent, the case without UVLF data is consistent with the inference of Pochinda et al. (2024) who assume a constant SFE in the atomic cooling regime. This is a special case of our more general model, where the SFE depends on both halo-mass and redshift.

Fig. C2 shows the joint constraints on the astrophysical parameter space with and without the SARAS 3 data set, for comparison with the joint analysis in Fig. 4. We demonstrate the synergy between SARAS 3 (or more generally, 21-cm global signal limits) and UVLF data in constraining the minimum circular velocity V_c for star formation in DM haloes.
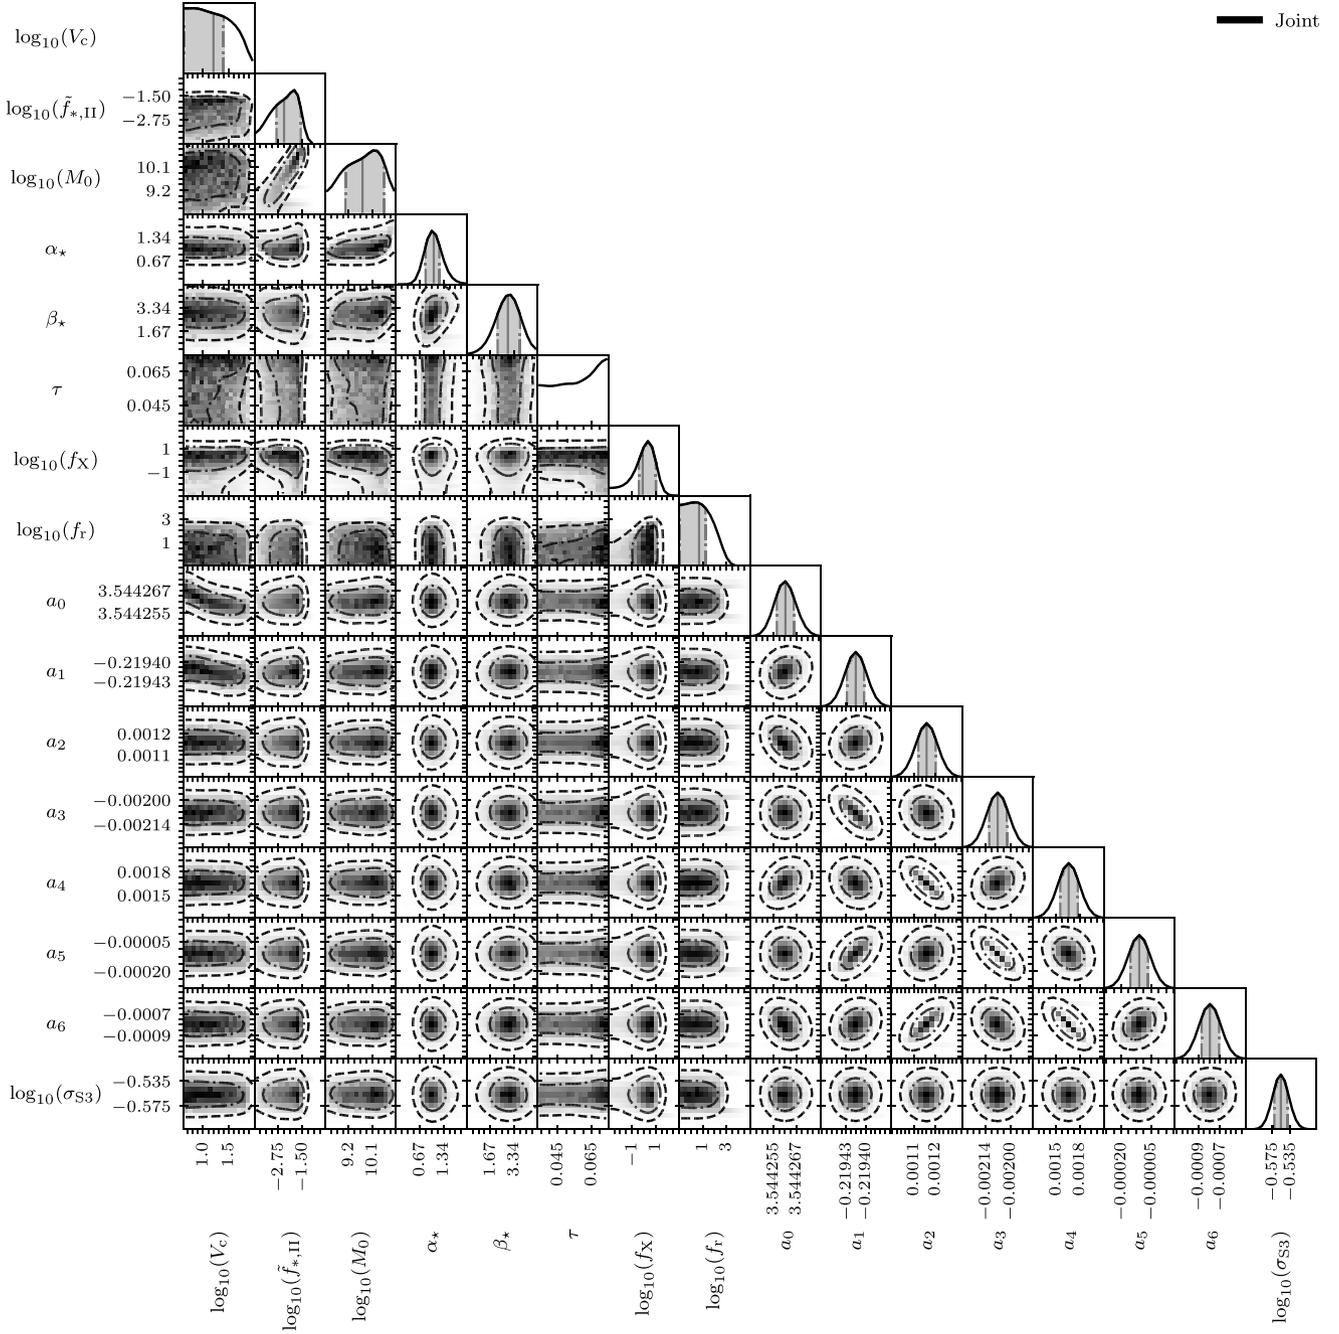
Fig. C3 shows the constraints on the full parameter space (i.e. astrophysical parameters + SARAS 3 nuisance parameters) for the joint analysis. We draw attention particularly to the degeneracy between the minimum circular velocity V_c and the lowest order polynomial coefficients of the SARAS 3 foreground polynomial, a_0 and a_1, which is explained in more detail in Section 5.1.
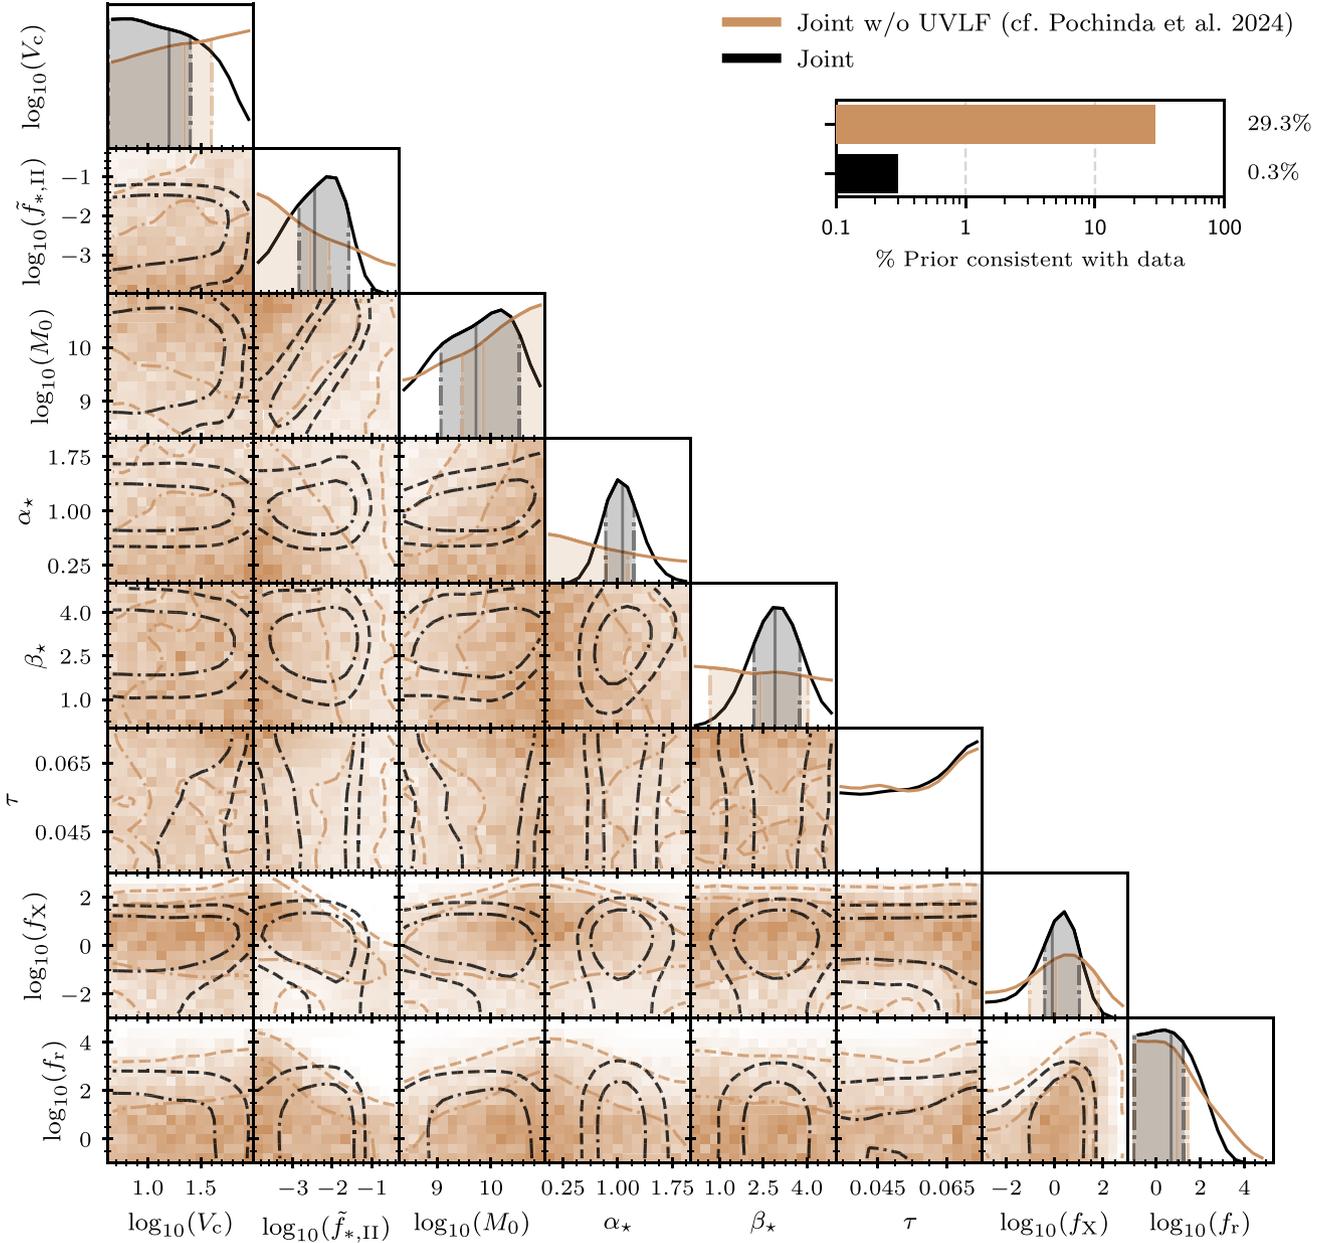
Figure C1. Similar to Fig. 4, but showing joint constraints with and without the UVLF data set. The 2D posterior contours show the 68 per cent (dash–dotted) and 95 per cent credible regions (dashed), with filled contours for the case without UVLF data. The 1D marginal posterior PDFs show the weighted mean (solid line) and 68 per cent credible region (shaded with dash–dotted outline). The top right panel shows the prior to posterior volume contraction for each data set. Note in particular, without the UVLF data set: V_c saturates to upper prior limit, the four SFE parameters are consistent with a constant, fixed value at the lower limit of 0.01 per cent (recovering the model used in Pochinda et al. 2024, and indeed their results), and f_X and f_r are weakened by their degeneracy with SFE.
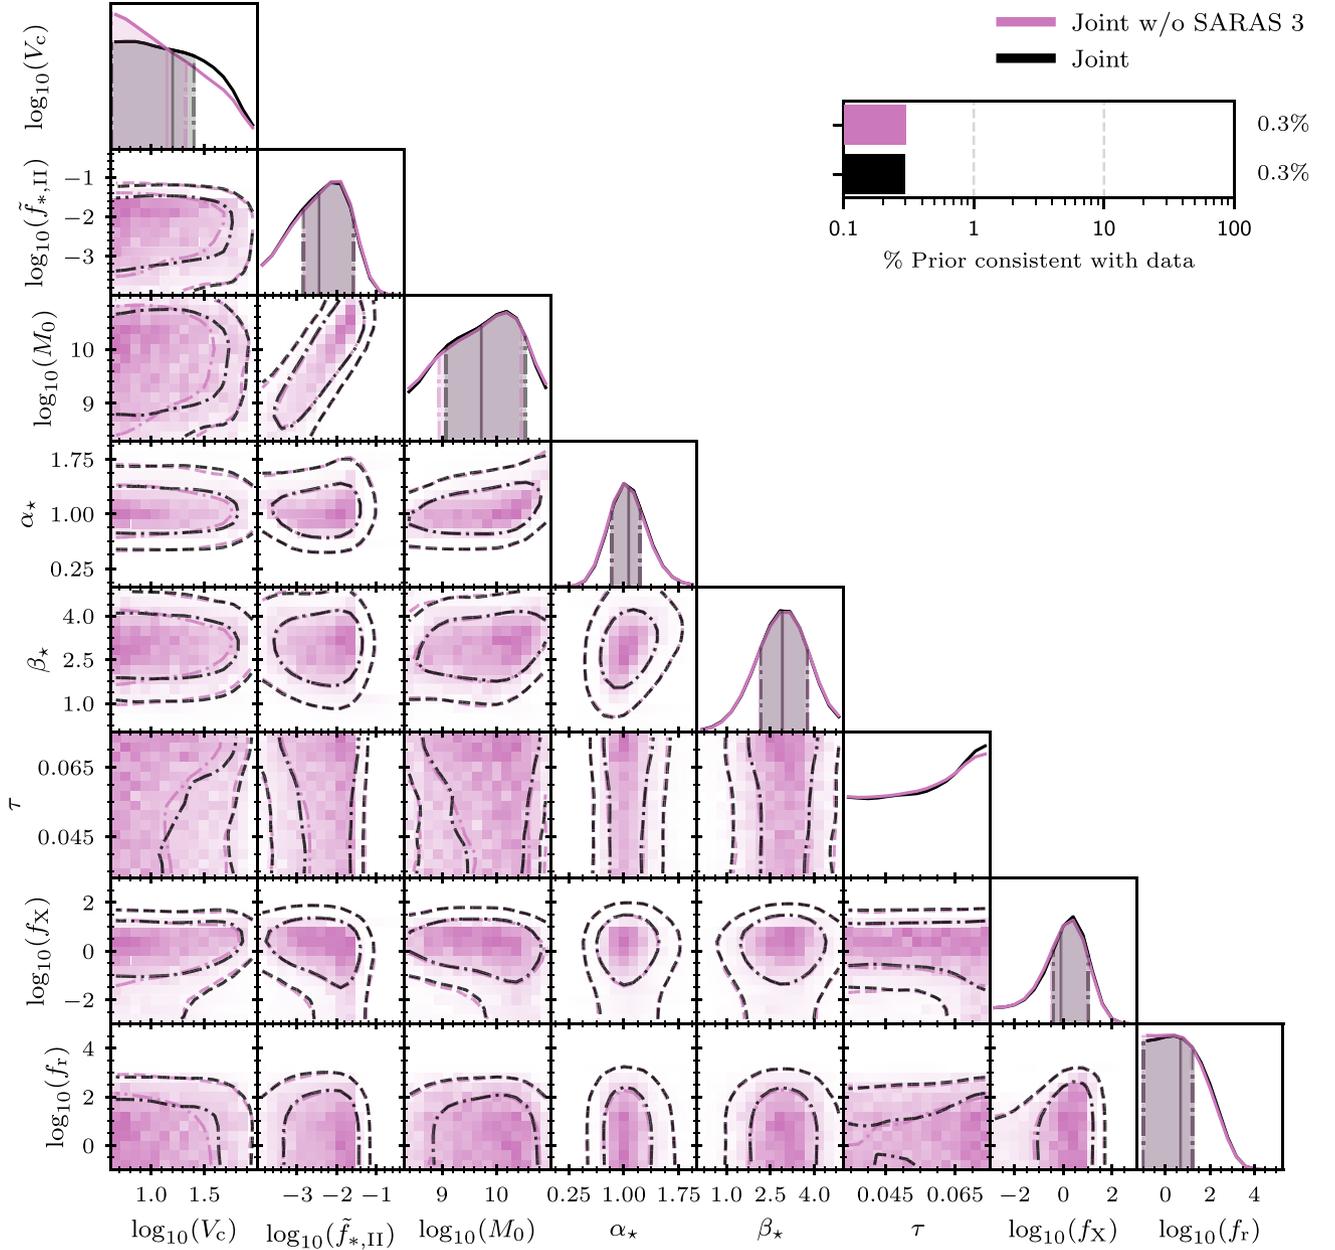
Figure C2. Similar to Fig. 4, but showing joint constraints with and without the SARAS 3 data set. The 2D posterior contours show the 68 per cent (dash–dotted) and 95 per cent credible regions (dashed), with filled contours for the case without SARAS 3 data. The 1D marginal posterior PDFs show the weighted mean (solid line) and 68 per cent credible region (shaded with dash–dotted outline). The top right panel shows the prior to posterior volume contraction for each data set. The SARAS 3 data weakly constrains the minimum circular velocity V_c at the lower end of the prior (as opposed to the UVLF data which disfavours the higher end). The constraint is highly sensitive to the SARAS 3 foreground fit as discussed in Section 5.1 (see e.g. Fig. 6 where the constraint is strongest, and the synergy between UVLFs and SARAS 3 is most visible).

Figure C3. Similar to Fig. 4, but showing the constraints on the full parameter space (i.e. astrophysical parameters + SARAS 3 nuisance parameters) for the joint analysis. The 2D contours show the 68 per cent (dash–dotted) and 95 per cent (dashed) credible regions, while the 1D marginal posterior PDFs show the weighted mean (solid line) and 68 per cent credible region (shaded with dash–dotted outline). Note in particular, the constraint on the minimum circular velocity V_c has some degeneracy with the lowest order coefficients of the SARAS 3 foreground polynomial, a_0 and a_1 (which have the largest effect).

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