

## Simulation-based inference on warm dark matter from HERA forecasts

Q. Decant <sup>a</sup>, A. Dimitriou <sup>b</sup>, L. Lopez-Honorez <sup>a,c</sup> and B. Zaldivar <sup>b</sup>

<sup>a</sup>*Service de Physique Théorique, C.P. 225, Université Libre de Bruxelles, Boulevard du Triomphe, B-1050 Brussels, Belgium*

<sup>b</sup>*Instituto de Física Corpuscular (IFIC), Consejo Superior de Investigaciones Científicas (CSIC) and Universitat de València, 46980, Valencia, Spain*

<sup>c</sup>*Theoretische Natuurkunde & The International Solvay Institutes, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium*

*E-mail:* [quentin.decant@ulb.be](mailto:quentin.decant@ulb.be), [androniki.dimitriou@ific.uv.es](mailto:androniki.dimitriou@ific.uv.es), [llopezho@ulb.be](mailto:llopezho@ulb.be), [bryan.zaldivar@ific.uv.es](mailto:bryan.zaldivar@ific.uv.es)

**ABSTRACT:** The redshifted 21cm signal from Cosmic Dawn promises to open a new window into the early history of our universe and enable the probing of an unprecedented comoving survey volume. In this work, we revisit the imprint of Warm Dark Matter (WDM) on the 21cm signal power spectrum using an updated implementation of the WDM effect in the public code `21cmFast` and considering a single population of cosmic dawn galaxies. By focusing on inferring the WDM mass, we analyze the degeneracies between the latter and the astrophysics parameters characterizing star formation and X-ray heating and we emphasize the role of the threshold mass for star-forming galaxies,  $M_{\text{turn}}$ . We study the capability of the recently built HERA telescope to reconstruct the WDM mass by adopting the statistical approach of simulation-based inference. We include a comparison of the per-parameter reconstruction quality for different number of simulations used in the training of the algorithm. Our results indicate that HERA could surpass current Lyman- $\alpha$  forest constraints if Cosmic Dawn galaxies exhibit a threshold mass  $M_{\text{turn}} \lesssim 10^8 M_{\odot}$ . The X-ray source properties considered in this study may also influence the strength of the WDM constraint for lower threshold masses.

**KEYWORDS:** cosmological simulations, dark matter theory, Machine learning, physics of the early universe

ARXIV EPRINT: [2412.10310](https://arxiv.org/abs/2412.10310)

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>21cm cosmology and astrophysics model</b>	<b>4</b>
2.1	The redshifted 21cm signal	4
2.2	Astrophysics parameters	5
2.3	Warm dark matter	6
2.4	WDM imprint on 21cm signal and degeneracies	9
<b>3</b>	<b>Analysis</b>	<b>11</b>
3.1	Priors from the UV luminosity data with MCMC	11
3.2	21 cm Mock data	13
3.3	Simulation-based inference	14
<b>4</b>	<b>Results</b>	<b>17</b>
4.1	Parameter reconstruction	17
4.2	Degeneracies in the posterior distributions	22
4.3	Impact of $M_{\text{turn}}$ on the $m_{\text{WDM}}$ lower bound	23
<b>5</b>	<b>Conclusion</b>	<b>27</b>
<b>A</b>	<b>Some extra information</b>	<b>28</b>
<b>B</b>	<b>Coverage tests</b>	<b>29</b>

---

## 1 Introduction

21cm Cosmology is expected to open a new window on our understanding of the history of the Universe in a redshift range beyond the reach of any existing experiments. The redshifted 21cm signal arising from the so-called Cosmic Dawn (CD) until the Epoch of Reionization (EoR) shall soon be probed to a high signal-to-noise ratio thanks to ground-based radio interferometers such as the Hydrogen Epoch of Reionization Array (HERA) [1] and the Square Kilometer Array (SKA) experiment [2, 3]. The former, though, is already fully built and running, having produced its first results in 2021 [4]. An interesting possibility is to study the dark matter hypothesis using the 21cm signal. In this work, we focus on the HERA probes of thermal warm dark matter (WDM) candidates, that suppress small scale structures due to WDM free-streaming, as a benchmark scenario for non-cold dark matter (NCDM) models [5–7]. These DM models can leave a measurable imprint on the 21cm signal delaying the appearance of peaks and troughs both in the sky-averaged signal and in the power spectrum of 21 cm fluctuations, see e.g. [8–20].

Modeling the evolution of the 21cm signal from EoR and CD is challenging due to the complex interplay of many astrophysical processes, as well as potential physics beyond the

Standard Model. Consequently, the simulations are computationally costly, even those based on semi-numerical methods as performed by the `21cmFast` public code [21, 22], which we have adapted for the present analysis. Clearly, the complexity of the simulations hinders the corresponding statistical analyses, typically relying on likelihood-based approaches such as Fisher matrix analysis or Markov Chain Monte Carlo (MCMC), see e.g. [16, 17, 19, 20] for NCDM-related studies. In our analysis, we use a likelihood-free approach instead, dubbed Simulation-based Inference (SBI), which has proven to computationally surpass standard MCMC methods for an increasing number of problems in areas such as cosmology [23–26], astrophysics, and particle physics.<sup>1</sup>

SBI is already being used in studies concerning 21cm Cosmology, see e.g. [18, 27–32]. The works [27–32] focus on the ability of SKA to reconstruct astrophysics parameters related to reionization and star formation, while considering the cold dark matter (CDM) scenario, unlike the present study which focuses on WDM instead. In particular, [27] focus as well on the HERA experiment and considers an SBI flavour (“DELFI”) designed to estimate the likelihood function itself. In [28, 31] the SBI flavour used in the analysis is the same as the one used here, called Neural Ratio Estimation (NRE, cfr. section 3.3). In [30], they compare different summary statistics of the signal in terms of the Fisher Information, while reporting their results based on a combination of fiducial values of the parameters, given that there are no physics-motivated values a priori. Here instead, we report our results for a selection of fiducial scenarios.

On the other hand, [18] does analyze a WDM scenario aiming at reconstructing the WDM mass. They do so with a different SBI method, so-called Neural Posterior Estimation (NPE), in the context of the SKA experiment. Note that the galaxy parametrization used in that work corresponds to an older implementation in `21cmFast` than the one used here. Furthermore, the default WDM implementation in `21cmFast` is based on [33] and makes use of a top-hat window function in the halo mass function (HMF), complemented with a sharp cut in the integrated number of ionizing photons, X-ray emissivity, and Lyman- $\alpha$  photon flux at a certain WDM Jeans mass computed from the WDM free-streaming length. Note that the top-hat window function is expected to give rise to a significant over-estimation of the number of low-mass haloes in a WDM cosmology, see [34, 35], while those haloes play a key role in estimating the WDM imprint on the 21cm signal. Here, in contrast, we modified `21cmFast` to use a sharp-k window function to evaluate the HMF. The latter has been shown to provide a better fit to numerical simulations up to redshifts  $z \sim 20$ , see [7] (and also [36] for a discussion). Furthermore, beyond a different implementation of the SFR and of the WDM HMF, the threshold mass for star formation,  $M_{\text{turn}}$ , is effectively kept in [18] within the upper end of the  $M_{\text{turn}}$  range considered here.<sup>2</sup> As we will show, this is expected to hinder the sensitivity to the WDM imprint. Other relevant works (based on Fisher matrix or MCMC instead of SBI) on WDM or related non-cold dark matter scenarios, using the brightness temperature fluctuations in the 21cm signal as a probe, have already been provided; see e.g. [15–17, 19]. Most recently, the analysis of [16], as well as in [19], were

<sup>1</sup>Given the high number of works in these areas, we have decided to cite here only the most recent ones.

<sup>2</sup>Our  $M_{\text{turn}}$  is equivalent to the  $T_{\text{vir}}$  (the minimum virial temperature parameter) parameter used in [18] but on a smaller range of values.

performed using other simulation codes than `21cmFast`. Those codes in particular allowed to perform MCMC analysis making use of the default assumption of a Gaussian likelihood.

In this work, we reassess the WDM imprint on the 21cm signal, as potentially measured by the fully deployed HERA experiment, starting with a description of the assumed WDM implementation. We first provide an MCMC analysis of the viable WDM scenario parameter space that is in agreement with existing UV luminosity data, and then analyze the expected degeneracies of the WDM imprint with the astrophysics parameters controlling the overall shape of the 21cm power spectrum. One of our main goals is the assessment of the role of the threshold mass for star-forming galaxies,  $M_{\text{turn}}$ , in probing thermal WDM. Note that this threshold value is intuitively expected to affect the 21cm signal in the same way as WDM, since it effectively prevents the smallest haloes from contributing to the signal. We evaluated under which conditions HERA is expected to surpass the existing sensitivity of the Lyman- $\alpha$  forest data to the WDM mass, see [37].<sup>3</sup> We also carefully report on the performance of our SBI configuration to reconstruct the astrophysics parameters.

In our analysis, we adopt a  $\Lambda$ CDM cosmology with parameters  $(\Omega_m, \Omega_b, \Omega_\Lambda, h, \sigma_8, n_s) = (0.31, 0.049, 0.69, 0.68, 0.81, 0.97)$  following the *TT, TE, EE + lowE + lensing + BAO* result of Planck 2018 [40]. For modelling the 21cm signal during the CD-EoR period, specifically its power spectrum (PS) summary statistics, we use our own modified version of the `21cmFast` public code [21, 22], with the flexible galaxy parameterization of [41]. We model the thermal noise by making use of the estimates from `21cmSense` [42, 43]. We consider a seven-dimensional parameter space, including the WDM mass, and six other astrophysics parameters expected to show some level of degeneracy with the former. For simplicity, and to minimize computational expense, the astrophysics model considered here involves a single population of galaxies. A similar approach was recently employed in an SBI analysis within the context of 21cm cosmology in CDM scenarios to assess the impact of the HMF parametrization on the reconstruction of the astrophysics parameters, see [32]. In our work, we vary the threshold mass  $M_{\text{turn}}$  between  $10^5$  and  $10^{10} M_\odot$ . Masses larger than  $10^8 M_\odot$  at  $z = 10$  could be representative of atomic-cooling galaxies (ACGs) that we have observed at late times [44]. The ACGs are expected to be dominated by Pop II stars. On the other hand, the lower end of our threshold mass range would include the very first galaxies, hosted by so-called minihalos (with a virial mass  $M \lesssim 10^8 M_\odot$  at  $z \sim 10$ ) and referred to as molecular-cooling galaxies (MCGs). They are expected to predominantly host Population III stars, whose detailed properties are still open to debate. For a discussion on this topic, see e.g. [33, 45]. The single population of galaxies approach followed here serves as a simple set-up to evaluate the sensitivity of 21cm Cosmology to a modification of the small-scale power spectrum and the halo mass function. It will highlight the important effect of the galaxy population assumption on the 21cm sensitivity reach.

We organize this paper as follows. In section 2, we briefly summarize the parametrization of the astrophysics and WDM imprint on the 21cm signal considered here, as well as the expected degeneracies. In section 3, we provide the details of our SBI analysis, present our results in section 4, and conclude in section 5.

---

<sup>3</sup>Other recent constraints on WDM from other probes can also be found in e.g. [38, 39] and references there-in. Here we use the exclusion of a 5.3 keV WDM mass at 95%CL from [37] as a reference.

## 2 21cm cosmology and astrophysics model

We start with a brief review of 21cm Cosmology and the astrophysics parameters that are expected to have a noticeable impact on the inference of the warm dark matter mass from the power spectrum.

### 2.1 The redshifted 21cm signal

The redshifted cosmic 21cm signal, arising from the hyperfine spin-flip transition of neutral hydrogen at different redshifts, is characterized by the differential brightness temperature of the 21cm signal with respect to a radio background source, which can be approximated by [46]

$$\delta T_b \approx 20\text{mK} \left(1 - \frac{T_{\text{CMB}}}{T_{\text{S}}}\right) x_{\text{HI}}(1 + \delta_b) \left(1 + \frac{1}{H} \frac{dv_r}{dr}\right)^{-1} \sqrt{\frac{1+z}{10} \frac{0.15}{\Omega_{\text{m}} h^2} \frac{\Omega_{\text{b}} h^2}{0.023}}, \quad (2.1)$$

when considering the Cosmic Microwave Background (CMB), with temperature  $T_{\text{CMB}}$ , as the radio background. In eq. (2.1),  $x_{\text{HI}}$  is the neutral hydrogen fraction,  $\delta_b$  is the baryon number density contrast,  $\Omega_{\text{m},\text{b}}$  are the matter and baryon energy densities relative to the critical energy density today,  $H$  is the Hubble rate and  $\frac{dv_r}{dr}$  is the gradient of the proper velocity along the line of sight. The 21cm signal is seen in emission or absorption when the CMB and spin temperature,  $T_{\text{S}}$ , differ. The spin temperature quantifies the relative occupancy of the two hyperfine levels of the ground state of neutral atomic hydrogen. For redshifts  $z$  below  $\sim 30$ , it is obtained from the equilibrium balance of absorption/emission of 21cm photons from/to the CMB background and resonant scattering of Lyman- $\alpha$  photons coupling the spin temperature to the gas temperature  $T_{\text{k}}$ . Spatial variations of IGM properties lead to fluctuations in the 21cm signal. In what follows, we refer to the 21cm global signal,  $\overline{\delta T_b}$ , as the sky-averaged brightness temperature and to the 21cm power spectrum in terms of the dimension-full quantity  $\overline{\delta T_b^2} \Delta_{21}^2$ . The latter is defined as:

$$\overline{\delta T_b^2} \Delta_{21}^2(k, z) = \overline{\delta T_b^2(z)} \times \frac{k^3}{2\pi^2} P_{21}(k, z) \quad (2.2)$$

with the two-point correlation function corresponding to

$$\langle \tilde{\delta}_{21}(\mathbf{k}, z) \tilde{\delta}_{21}(\mathbf{k}', z) \rangle = (2\pi)^3 \delta^D(\mathbf{k} - \mathbf{k}') P_{21}(k, z), \quad (2.3)$$

with  $\langle \cdot \rangle$ , the ensemble average,  $\mathbf{k}$  the comoving wave vector, and  $\tilde{\delta}_{21}(\mathbf{k}, z)$  the Fourier transform of  $\delta_{21}(\mathbf{x}, z) = \delta T_b(\mathbf{x}, z) / \overline{\delta T_b}(z) - 1$ , where  $\mathbf{x}$  denotes the position vector.

Specific features can be observed in the redshift evolution of the global 21cm signal and the power spectra at a fixed scale. In particular, the global signal is expected to feature an absorption ( $\overline{\delta T_b^2} < 0$ ) trough between the epoch of Lyman- $\alpha$  or Wouthuysen-Field (WF) coupling and the subsequent epoch of heating (EoH); while the signal shall be seen in emission ( $\overline{\delta T_b^2} > 0$ ) between the EoH and the EoR. The HERA telescope, on which we will focus here, is expected to probe the 21cm power spectrum. Depending on the scale considered, the latter can feature up to 3 peaks, see e.g. figure 2, corresponding to dominant contributions of the auto-power spectrum of fluctuations in the Lyman-alpha coupling, in the gas temperature, and in the neutral fraction during the WF, EoH and EoR periods, respectively. The detailed

parameter	units	range	scale
$M_{\text{turn}}$	$M_{\odot}$	$[10^5, 10^{10}]$	log
$f_{\star,10}$	—	$[10^{-3}, 1]$	log
$\alpha_{\star}$	—	$[0.0, 1.0]$	linear
$t_{\star}$	—	$[0.1, 1.0]$	linear
$L_X$	erg yr/(s $M_{\odot}$ )	$[10^{38}, 10^{42}]$	log
$E_0$	keV	$[0.2, 1.5]$	linear
$m_{\text{WDM}}$	keV	$[3, 30]$	linear

**Table 1.** Type (4th column) and prior ranges (3rd column) of the six astrophysics parameters (1st column) considered in our analysis, together with the DM mass ( $m_{\text{WDM}}$ ). See the text for the definition of the parameters.

shape of the power spectrum, and its testability, will depend on the IGM temperature and ionized fraction. These quantities, in turn, are influenced by the radiation emitted by stars in the late Universe. In the next section, we briefly introduce the modelisation of the astrophysics sources considered in this work, following its implementation in the public code `21cmFAST` from [41]. Concerning the testability, we focus on the HERA experiment, as it has completed deployment [4], and is currently analysing data from an extended observational campaign. A first observational result performed with 71 antennas (out of the total 331) and only 94 nights of measurement has already provided the most constraining upper bounds on the 21 cm power spectrum at redshifts  $z = 8$  and 10 [1].

## 2.2 Astrophysics parameters

In this paper, we focus on the astrophysics parameters that are expected to influence the inference of the WDM properties. In what follows, we provide a short description of those parameters. We consider a single population of galaxies and follow the astrophysics model implemented in `21cmFAST` as described in [41]. The parameters involved in our analysis, see section 3, are listed in table 1 together with the scale (linear or logarithmic) and the ranges of the prior values that have been explored in this work.<sup>4</sup>

The stellar mass of a galaxy, denoted as  $M_{\star}$ , is assumed to be related to the virial mass of the host halo,  $M$ , through the relation  $M_{\star} = f_{\star}M$ , where  $f_{\star}$  is the stellar-to-halo mass ratio. The latter is assumed to scale as:

$$f_{\star} = \frac{\Omega_{\text{b}}}{\Omega_{\text{m}}} \min \left\{ 1, f_{\star,10} \left( \frac{M}{10^{10} M_{\odot}} \right)^{\alpha_{\star}} \right\}, \quad (2.4)$$

where  $f_{\star,10}$  and  $\alpha_{\star}$  are free parameters in our study. This power-law relation between  $f_{\star}$  and  $M$  has been inferred in [47, 48]. On the other hand, the star formation rate (SFR) inside a halo of virial mass  $M$  is parameterised by [49]

$$\dot{M}_{\star} = \frac{M_{\star}}{t_{\star} H(z)^{-1}} \quad (2.5)$$

<sup>4</sup>The selection of parameters is similar to the one that was made in the MCMC analysis of [16], yet with a different astrophysics model and 21cm signal simulator.

where  $t_* H^{-1}$  is a characteristic star-formation time-scale. Note that in the parametrization of [41],  $\dot{M}_*$  appears in the computation of the X-ray luminosity and of the estimation of the non-ionizing UV photon flux, while  $M_*$  drives the number of ionizing photons per baryon. This allows for disentangling the effects of  $t_*$  and  $f_{*,10}$  on the 21cm signal.

On the other hand, the number density of galaxies per unit halo mass that contributes to the number of UV ionizing and non-ionizing photons, the X-ray emissivity and the Lyman- $\alpha$  flux is accounted for by the halo mass function (HMF),  $dn/dM$ , weighted by a duty cycle,  $f_{\text{duty}}$ . Here we adopt the Sheth, Mo and Tormen (SMT) [50] first crossing distribution in the HMF, and we make use of a sharp-k window function to evaluate the root-mean-square (rms) variance of density perturbations, see more details in section 2.3. The duty cycle accounts for inefficient star formation in low-mass galaxies due to inefficient cooling and/or feedback mechanisms. It is assumed to take the following form [49, 51–54]:

$$f_{\text{duty}} = \exp\left(-\frac{M_{\text{turn}}}{M}\right). \quad (2.6)$$

While e.g. [41] concluded that  $M_{\text{turn}} < 2 \times 10^9 M_\odot$  to account for UV luminosity data, molecular cooling would point to threshold masses as low as  $M_{\text{turn}} = 10^5 M_\odot$  [33, 45]. In our analysis, we consider a single population of galaxies whose threshold mass is fixed to a constant value  $M_{\text{turn}}$  lying in the range between  $10^5$  and  $10^{10} M_\odot$ .<sup>5</sup>

Finally, the X-ray luminosity, which brings the global signal in emission w.r.t. the CMB, is normalized with the integrated soft-band ( $E < 2 \text{ keV}$ ) luminosity per SFR (in units of  $\text{erg yr s}^{-1} M_\odot^{-1}$ ), that we denote:

$$L_X = \int_{E_0}^{2\text{keV}} dE \mathcal{L}_X, \quad (2.7)$$

where  $\mathcal{L}_X$  refers to the specific X-ray luminosity per unit star formation rate, and  $E_0$  is the X-ray energy threshold below which photons cannot escape the host galaxy.  $L_X$  and  $E_0$  are considered as free parameters in our analysis. The X-ray threshold  $E_0$  depends on the interstellar medium as well as the environment of the X-ray sources. Furthermore, the first HERA results [58], analyzed in the context of a CDM cosmology, tend to favor  $L_X \geq 10^{40} \text{ erg yr}/(\text{s } M_\odot)$ . Here, in the context of our WDM analysis, we consider the ranges reported in table 1.

### 2.3 Warm dark matter

Thermal WDM is assumed to be in kinetic and chemical equilibrium with the rest of the plasma before its decoupling, which settles its relic abundance through a relativistic freeze-out.

---

<sup>5</sup>The threshold mass  $M_{\text{turn}}$  depends on the minimum virial mass from which the gas can cool, and on feedback processes. Atomic cooling is expected to be efficient above minimum virial temperatures of  $T_{\text{vir}} \sim 10^4 \text{ K}$ , while molecular cooling provides cooling from  $T_{\text{vir}} \sim 300 \text{ K}$  [55]. At  $z = 20$ , making use of the virial theorem and neglecting feedback processes, this would correspond to minimal threshold masses of the order of  $M_{\text{mol}} = 10^5 M_\odot$  and  $M_{\text{atom}} = 2 \times 10^7 M_\odot$ , respectively. See also e.g. [41, 49, 56, 57] for a discussion on the impact of those feedback mechanisms on the threshold mass. An implementation of these effects is provided in **21cmFAST**, but running simulations taking into account these effects entails a significant computational cost.

As WDM decouples while being relativistic, it free-streams over a characteristic length scale, referred to as the free-streaming length, that is larger for smaller WDM mass. Note that it is well known that in order to account for all the DM in the form of *thermal* WDM more than  $10^3$  relativistic degrees of freedom must to be invoked at the time of WDM production for  $m_{\text{WDM}}$  larger than a few keV. This is certainly not very natural or well motivated. It can be shown, however, that many other non-cold dark matter (NCDM) scenarios, that are not necessarily thermal and involving a much reduced extension of the SM, give rise to the same or a very similar imprint on the observables, see e.g. [59–62]. The thermal WDM scenario considered here is only used as a well defined benchmark for free-streaming DM imprint on the linear matter power spectrum and halo mass function.

Thermal WDM, or any equivalent NCDM, induces a suppressed linear matter power spectrum at small scales, as well as a reduced number of low-mass halos contributing to astrophysics radiation sources. In practice, WDM implies a delay in the EoR, EoH and WF eras, which gets imprinted in the 21cm signal and related observables, see e.g. [8–16, 18, 19, 36]. The lower the WDM mass,  $m_{\text{WDM}}$ , the larger the WDM free-streaming scale, and as a result, the more suppressed the number of low-mass haloes. In table 1, the prior on the DM mass is chosen such that for its high-mass end, 30 keV, the DM scenario is equivalent to CDM in the context of our analysis. On the other hand, the low-mass bound of 3 keV is chosen to be such that the WDM candidate is marginally allowed by existing Lyman- $\alpha$  forest bounds on WDM, see [37, 63].

The WDM suppression of the linear matter power spectrum,  $P(k)$ , has been parameterized in terms of a transfer function,  $T_{\text{WDM}}(k)$ , which is defined as a function of the ratio of the WDM and the CDM power spectra:  $T_{\text{WDM}}^2(k) = P_{\text{WDM}}(k)/P_{\text{CDM}}(k)$ . Assuming that WDM makes up all DM, the transfer function has been fitted to [63, 64]:

$$T_{\text{WDM}}(k) = (1 + (\alpha_{\text{WDM}}k)^{2\nu})^{-5/\nu}, \quad (2.8)$$

where  $k$  denotes the amplitude of the wave vector, Fourier dual of a comoving distance. In this work, we use  $\nu = 1.12$  and a breaking scale

$$\alpha_{\text{WDM}} = 0.0445 \left( \frac{\text{keV}}{m_{\text{WDM}}} \right)^{1.11} \left( \frac{\Omega_{\text{WDM}}}{0.25} \right)^{0.11} \left( \frac{h}{0.7} \right)^{1.22} \text{Mpc}/h, \quad (2.9)$$

which indicates the typical scale from which the WDM power spectrum gets significantly suppressed and effectively provides an estimate of the WDM free-streaming length. Notice that in this analysis, we have set the prefactor to 0.0445 Mpc/ $h$  to account for the case of  $m_{\text{WDM}} = 5.3$  keV, which saturates the current Lyman- $\alpha$  data lower bound [37]. Indeed, as already emphasized in e.g. [62], previous fits from [63, 64] were originally designed for lower mass dark matter.

To simulate the 21cm signal, one has to evaluate the imprint of WDM on non-linearly evolved matter perturbations that have given rise to a distribution of haloes. The halo mass function, i.e., the number of halos per unit mass as a function of mass and redshift, is defined as [65]

$$\frac{dn(M, z)}{dM} = \frac{\rho_{\text{m}}}{M^2} \frac{d \ln \sigma^{-1}}{d \ln M} f(\sigma), \quad (2.10)$$

where  $\rho_m = \Omega_m \rho_c$  is today's matter energy density with  $\rho_c$  the critical energy density,  $\sigma^2 = \sigma^2(M, z)$  is the variance of density perturbations on a mass-scale  $M$  at redshift  $z$ , and  $f(\sigma)$  is the so-called first crossing distribution. Here we consider the Sheth-Thormen description given by [66]

$$f(\sigma) = A \sqrt{\frac{2q}{\pi}} \left( 1 + \left( \frac{\sigma^2}{q \delta_c^2} \right)^p \right) \left( \frac{\delta_c}{\sigma} \right) e^{-\frac{q \delta_c^2}{2\sigma^2}}, \quad (2.11)$$

where  $\delta_c = 1.686$  is the critical overdensity. In this work, we use  $p = 0.3$ ,  $q = 1$ , and  $A = 0.322$  [67]. The root-mean-square (rms) variance of density perturbations at a given radius  $R$ , is evaluated using

$$\sigma^2(M(R), z) = \left( \frac{D(z)}{D(0)} \right)^2 \int \frac{d^3k}{(2\pi)^3} P_{\text{WDM}}(k) |W(kR)|^2, \quad (2.12)$$

where the redshift dependence is driven by the linear growth function,  $D(z)$  and  $P_{\text{WDM}}(k)$  is the linear power spectrum at  $z = 0$  obtained with the transfer function of eq. (2.8). For this work, we implemented the Fourier transform of the filter function,  $W(kR)$ , referred to as a sharp-k window:

$$W(kR) = \Theta(1 - kR), \quad (2.13)$$

where  $\Theta$  denotes the Heaviside function. The latter has indeed been argued to account for the observed WDM cosmology behaviour in N-body simulations, namely the WDM suppression of the halo mass function at low halo masses [34, 35], see also e.g. [7, 17] in the context of 21cm cosmology.

When making use of the sharp-k (SK) filter, a free parameter  $c_{\text{SK}}$  has to be introduced to relate the halo mass to its radius as follows:

$$M_{\text{SK}}(R) = \frac{4\pi}{3} (c_{\text{SK}} R)^3. \quad (2.14)$$


This parameter has been set to  $c_{\text{SK}} = 2.5$  in our work, see [7, 35, 67] for references. Considering a radius of the order of the breaking scale,  $R = \alpha_{\text{WDM}}$ , in the SK halo mass of eq. (2.14), we get a corresponding halo mass of:

$$M_{\text{WDM}}^{\text{cut}}(m_{\text{WDM}}) = 2.6 \left( \frac{5.3 \text{keV}}{m_{\text{WDM}}} \right)^{3.33} \times 10^6 M_{\odot}, \quad (2.15)$$

below which the HMF is significantly suppressed due to WDM free streaming.

For the follow-up analysis we use our own modified version of **21cmFast** accounting for eqs. (2.8)–(2.14). More details on this code are provided in the appendix A. The imprint of WDM on the HMF is illustrated in figure 1 displaying  $f_{\text{duty}} \times dn/d \ln M$  as a function of the halo mass  $M$  at  $z = 9$ . The continuous blue line corresponds to a 30 keV WDM, which in practice, is equivalent to CDM in the context of our analysis,<sup>6</sup> while the red and green lines assume a WDM cosmology with  $m_{\text{WDM}} = 5.3$  and 2 keV, respectively. The lower

<sup>6</sup>Indeed, for  $m_{\text{WDM}} = 30 \text{ keV}$ ,  $M_{\text{WDM}}^{\text{cut}} \sim 10^4 M_{\odot}$ , with  $M_{\text{WDM}}^{\text{cut}}$  defined in eq. (2.15), which is below the minimum threshold mass for star formation considered in this work.




**Figure 1.** Halo mass function as a function of the halo mass at  $z = 9$ . With the continuous lines, we show the case of  $m_{\text{WDM}} = 2, 5.3,$  and  $30 \text{ keV}$  WDM mass with green, red and blue colours, assuming  $M_{\text{turn}} = 10^5 M_{\odot}$ . The  $m_{\text{WDM}} = 30 \text{ keV}$  case is shown as a representative of a CDM scenario in the context of our analysis. The vertical thin coloured lines indicate the cutoff mass scale,  $M_{\text{WDM}}^{\text{cut}}$ , see eq. (2.15). This scale provides a rule-of-thumb estimate of the threshold below which halo formation is significantly suppressed due to the free-streaming of WDM particles. The dashed and dot-dashed blue lines correspond to the  $30 \text{ keV}$  DM HMF with a threshold mass increased to  $M_{\text{turn}} = 10^7$  and  $10^8 M_{\odot}$ , respectively.

the DM mass, the longer the free streaming scale of eq. (2.9) and the larger is the halo mass at which the HMF is suppressed, as expected from eq. (2.15). In particular,  $M_{\text{WDM}}^{\text{cut}}$  for  $m_{\text{WDM}} = 5.3$  and  $2 \text{ keV}$  are indicated with thin colored vertical red and green lines in figure 1. This WDM halo mass cut-off estimate can conveniently be compared to the threshold mass for star formation  $M_{\text{turn}}$  entering in the duty cycle of eq. (2.6). Indeed, as mentioned above, the ionization, X-rays and Lyman- $\alpha$  fluxes from stars induce the 21 cm signal for a given halo mass is weighted by the product of  $f_{\text{duty}} \times dn/d \ln M$ . With the continuous, dashed and dot-dashed blue lines, we assume increasing threshold masses,  $M_{\text{turn}} = 10^5, 10^7$  and  $10^8 M_{\odot}$ , respectively, dictating the form of the duty cycle factor. The plot suggests that a threshold mass  $M_{\text{turn}} > 10^8 M_{\odot}$  may hinder the straightforward reconstruction of a WDM cosmology characterized by  $m_{\text{WDM}} = 5.3 \text{ keV}$ . In general, when  $M_{\text{turn}} \gg M_{\text{WDM}}^{\text{cut}}$ , the WDM free-streaming effect is expected to suppress haloes that are too small to host stars. The 21cm signal shall then not be sensitive to the WDM imprint. Equivalently, if  $M_{\text{WDM}}^{\text{cut}} \gg M_{\text{turn}}$ , WDM would prevent haloes from forming that would have otherwise significantly contributed to the onset of the EoR, EoH and WF eras. This would for example be the case of  $m_{\text{WDM}} = 2 \text{ keV}$  with  $M_{\text{turn}} \lesssim 10^7$ . Yet, such a low mass WDM is already excluded by Lyman- $\alpha$  forest observations from [37].

## 2.4 WDM imprint on 21cm signal and degeneracies

In this section, we briefly comment on the imprint of WDM on the 21cm power spectrum given the parametrization detailed in section 2.3 as well as on the expected degeneracies



**Figure 2.** 21cm power spectrum  $\overline{\delta T_b^2 \Delta_{21}^2}$  as a function of the redshift for fixed scale  $k = 0.15/\text{Mpc}$ . Left: effect of varying the WDM mass. Right: effect of varying the threshold mass for star formation. The astrophysics parameters that are not mentioned in the figure have been fixed to  $\log(f_{\star,10}) = -0.9$ ,  $\alpha_{\star} = 0.46$ ,  $E_0 = 0.5 \text{ keV}$  and  $L_X = 10^{40} \text{ erg yr}/(s M_{\odot})$ .

between  $m_{\text{WDM}}$  and the astrophysics parameters introduced in section 2.2. For that purpose the 21cm power spectrum  $\overline{\delta T_b^2 \Delta_{21}^2}$  is illustrated in figure 2 as a function of the redshift at a fixed scale  $k = 0.15 \text{ Mpc}^{-1}$  for a selection of parameter values. The continuous blue lines correspond to a fiducial scenario assuming  $M_{\text{turn}} = 10^5 M_{\odot}$  with  $m_{\text{WDM}} = 30 \text{ keV}$  in both plots. In the left plot, we illustrate the effect of the WDM mass, the red and green curves corresponding to a DM mass of  $m_{\text{WDM}} = 5.3$  and  $2 \text{ keV}$ , respectively. We see that the lower the WDM mass, the later the different features of the 21cm power spectrum appear. Indeed, as discussed in section 2.3, lower WDM masses induce larger DM-free streaming length at the time of structure formation, preventing the lowest mass halos from forming efficiently. This effectively reduces the ionizing, X-ray, and Lyman- $\alpha$  photons luminosity from cosmic dawn galaxies, delaying the EoR, EoH and WF eras. On the other hand, in the right plot, we illustrate the impact of different choices of threshold mass for star-forming galaxies in a CDM-like cosmology, the dashed and dot-dashed blue lines corresponding to  $M_{\text{turn}} = 10^7$  and  $10^8 M_{\odot}$ , respectively. Increasing  $M_{\text{turn}}$  prevents the lower mass haloes to host stars. Again, this gives rise to a shift of the power spectrum features to later redshifts. A higher value of the threshold mass  $M_{\text{turn}}$  has a qualitatively similar effect on the power spectrum as a lower  $m_{\text{WDM}}$ . We expect thus a positive correlation between  $M_{\text{turn}}$  and  $m_{\text{WDM}}$  in our analysis as an increase in  $m_{\text{WDM}}$  could potentially be compensated by an increase of  $M_{\text{turn}}$ . Also, relatively large values of  $M_{\text{turn}}$  shall prevent any sensitivity of the 21 cm signal to the WDM properties beyond minimal  $m_{\text{WDM}}$  as the suppressed haloes can not host stars. The effect on the 21cm power spectrum of the other parameters of table 1 is as follows. The parameters  $\alpha_{\star}$  and  $f_{\star,10}$  regulate the fraction of stars into haloes, while  $t_{\star}$  further tunes the amplitude of the SFR. Increasing  $\alpha_{\star}$  favours larger haloes in hosting stars, which can effectively be mimicked with decreasing  $m_{\text{WDM}}$ , thereby suppressing the number of low mass haloes. We expect the  $m_{\text{WDM}}$  and  $\alpha_{\star}$  parameters to be positively correlated e.g. an increase in  $\alpha_{\star}$  can be compensated by an increase of  $m_{\text{WDM}}$ . On the other hand,  $f_{\star,10}$  defines the overall amplitude of the fraction of stars into haloes (independently of the halo masses). Increasing  $f_{\star,10}$  enhances any source of radiation, leading to the emergence of earlier features

in the power spectrum. An increase of  $f_{\star,10}$  in a WDM scenario could thus be mimicked by an increase in  $m_{\text{WDM}}$ . Those parameters could therefore appear to be anticorrelated when the WDM imprint is not screened by a large value of  $M_{\text{turn}}$ . On the other hand,  $t_{\star}$  regulates the overall amplitude of the SFR that is relevant for non-ionizing UV and X-ray luminosities. Larger  $t_{\star}$  suppresses the SFR, which can be mimicked with lower  $m_{\text{WDM}}$ . As a result,  $t_{\star}$  shall display a positive correlation with  $m_{\text{WDM}}$ , as is the case of  $M_{\text{turn}}$  and  $\alpha_{\star}$ . The above degeneracies will be further studied in section 4.2.

The last two parameters mainly affect the EoH.  $L_X$  regulates the overall amplitude of the X-ray emissivity at fixed energy threshold  $E_0$ . Considering low  $E_0$ , we make the X-ray spectrum is softer, i.e. involving X-rays with shorter mean free path. The latter induces an inhomogeneous heating of the IGM [68] which induces a large EoH peak. On the other hand, e.g. for low  $E_0 = 0.5 \text{ keV}$ , one can reduce the amplitude of the power spectrum at a given redshift by increasing  $L_X$  as this might shift the EoH peak to an earlier time, where HERA has no sensitivity, see e.g. [69] for illustration. Changing  $E_0$  is thus expected to show some level of degeneracy with changing  $L_X$  and could indirectly affect the  $m_{\text{WDM}}$  reconstruction. This is in particular further illustrated in section 4.3.

### 3 Analysis

In this section, we describe the different steps of our analysis. We start by imposing constraints on our parameters of interest from UV luminosity data in section 3.1. We then describe the simulations of the power spectrum in section 3.2 before introducing our statistical analysis (the SBI implementation) in section 3.3.


#### 3.1 Priors from the UV luminosity data with MCMC

Some of the astrophysics parameters discussed in section 2.2 can already be constrained by existing UV luminosity data, see e.g. [4, 41, 70]. In contrast to ionizing UV photons, the soft UV responsible for coupling the gas and spin temperatures can have long mean free paths, even through the neutral IGM. The non-ionizing UV luminosity functions,  $\phi$ , are evaluated with 21cmFAST assuming:

$$\phi(M_{\text{UV}}) = f_{\text{duty}} \frac{dn}{dM} \times \frac{dM}{dM_{\text{UV}}} . \quad (3.1)$$

The last factor accounts for the conversion of the halo mass to UV magnitude. It is evaluated assuming that the SFR,  $\dot{M}_{\star}$  introduced in eq. (2.5), is proportional to the rest-frame UV luminosity of a galaxy which itself is a function of the UV magnitude  $M_{\text{UV}}$ , see [41] for details. In this paper, we use the data compiled by [71] and restrict the analysis to  $M_{\text{UV}} > -20$  and  $z \geq 6$ . As in previous works [4, 32, 41], the latter procedure is done to avoid uncertainties related to the estimation of dust extinction, which becomes important at low  $z$  and high brightness, see e.g. [72].

We compare the data from [71] to the UV luminosity functions obtained with 21cmFAST for a given choice of parameters in an MCMC analysis with a Gaussian likelihood. The resulting posterior distributions are shown in figure 3. Except for  $\alpha_{\star}$  and  $f_{\star,10}$ , constrained through the dependence of  $\dot{M}_{\star}$  in  $dM/dM_{\text{UV}}$ , the posteriors do not significantly constrain the range



**Figure 3.** Results of the MCMC analysis taking into account UV-luminosity data constraints. Darker (lighter) contours correspond to 68% (95%) credible intervals. The range of the axes is made to coincide with the prior ranges of table 1.

of the parameters considered here. The favoured range of the star formation rate parameters<sup>7</sup> are  $\log(f_{*,10}) \in \{-1.29, -0.77\}$  and  $\alpha_* \in \{0.35, 0.50\}$  at 68% CL. In our SBI analysis, we will make use of the posteriors depicted in figure 3 as priors for our parameters. Notice that the UV-luminosity function data considered here give rise to distributions compatible with those obtained in [41] for comparable ranges of the parameter values, even though we consider WDM cosmologies with a range of  $M_{\text{turn}}$  extended to lower values and a more recent UV luminosity data set.

Note that the optical depth to reionization measurements from Planck are also well known to impose constraints on the astrophysics parameters, in particular, on the parameters

<sup>7</sup>Notice that log refers to log 10 in our paper.


driving the fraction of ionizing photons escaping the galaxies. Yet, to estimate the ionized fraction, a lightcone has to be simulated which induces a significant time and computational cost compared to the evaluation of the UV luminosity functions. Furthermore, a dedicated analysis accounting correctly for the history of reionization resulting from the simulations, and not from a default tanh reionization assumed for Planck analysis, could be needed, see e.g. [45] for such study in a CDM context. To avoid such an extra non-negligible cost of computational resources, in our analysis, we fix the escape fraction parameters to the central values obtained in [58] from UV + optical depth + HERA I data, namely  $\log(f_{\text{esc},10}) = -1.1$  and  $\alpha_{\text{esc}} = 0.02$ . In the latter case, we checked in the course of our study that the optical depth measurement from Planck did not significantly affect the posteriors of the parameters presented in figure 3. All other parameters, not included in table 1, are set to their default values in 21cmFast.

### 3.2 21 cm Mock data

For our SBI analysis, we simulate 21 cm signal lightcones with 21cmFast within boxes of  $[300 \text{ Mpc}]^3$  comoving volume on a  $[200]^3$  cells grid. From those lightcones we obtain 21cm cosmological power spectra on fixed  $k$  and  $z$  grids reported in the appendix A, see table 3. Note that, to ensure the stochasticity of our mock data set, we picked a new *random seed* for each realization of the cosmological 21cm power spectrum obtained from 21cmFast, see e.g. [29]. We add to each simulated cosmological power spectra several realizations of the thermal noise contribution, see the discussion in section 3.3. The thermal noise is assumed to follow a zero-mean Gaussian distribution with variance given by the experimental sensitivity,  $\sigma_{\text{therm}}^2(k, z)$ , associated with the HERA telescope. We sample noisy data from such a Gaussian truncated at zero to enforce positive power spectrum values. To compute the sensitivities, we use the 21cmSense code [42, 43]. Within the latter, the HERA experiment is described as a hexagonal array of 331 antennas of 14 m dish size. In addition, we have specified a bandwidth of  $B = 8 \text{ MHz}$ , a spectral resolution of  $\delta\nu \sim 97.7 \text{ kHz}$ , and a total operating time of 1000 hrs ( $\sim 167$  nights with 6 hours of observation per day). The thermal noise depends on the system temperature,  $T_{\text{sys}}$ , as  $\sigma_{\text{therm}}(k, z) \propto T_{\text{sys}}^2(z)$  [42], where the frequency dependence is assumed to follow [73]:

$$T_{\text{sys}}(z) = 100 \text{ K} + 120 \text{ K} \left( \frac{\nu(z)}{150 \text{ MHz}} \right)^{-2.55}. \quad (3.2)$$

HERA will measure the redshifted 21cm signal in a frequency band of 50-250 MHz, which shall allow to probe the redshift range  $z \sim 4.6 - 27$ , with  $z = \nu_{21}/\nu - 1$  and  $\nu_{21} = 1420.4 \text{ MHz}$  is the rest frame frequency of the 21cm signal. The bandwidth  $B$  sets the size of frequency bins, which, expressed in terms of redshifts, implies increasing the size of redshift bins when increasing redshift. In this work we use a bandwidth  $B = 8 \text{ MHz}$  and 19 redshift bins in redshift range  $[5.9 - 25.3]$ , see table 3. The upper bound has been set to 25.3, as larger redshift power spectra get large noise contributions. Furthermore, given the HERA configuration, the accessible  $k$  modes are mostly dominated by the  $k$  component parallel to the line of sight (l.o.s.),  $k_{\parallel}$  [4]. Mapping angles in the sky and frequency to cosmological distances, one shall thus consider  $k$ -bins set by  $\Delta k_{\parallel}(z, \Delta\nu) \equiv 2\pi(\nu_{21}/\Delta\nu) \times H(z)/(1+z)^2$ , where  $\Delta\nu$  refers to a frequency step, implying that  $\Delta k_{\parallel}(z, \Delta\nu)$  decreases with increasing redshift. In this analysis,



**Figure 4.** Example of mock 21cm power spectrum  $\overline{\delta T_b}^2 \Delta_{21}^2$  as a function of the wave number  $k$  at different redshifts for  $m_{\text{WDM}} = 5.3 \text{ keV}$ ,  $M_{\text{turn}} = 10^7 M_{\odot}$ . The dashed red lines show the noiseless power spectrum; the blue starred lines are noisy realizations of the power spectrum, while the blue (purple) bands correspond to 68% (95%) CL uncertainties. See text for details. The rest of parameters have been fixed as in figure 2.

we use fixed values of the  $k$  bins width. We set the latter to the maximal value of  $\Delta k_{\parallel}$ , obtained at the minimum redshift, i.e.  $\Delta k = \Delta k_{\parallel}(z_{\text{min}}, B) = 0.06 \text{ Mpc}^{-1}$ . As a result, we use 15  $k$ -bins, with our minimum  $k$ -bin center set to  $0.15 \text{ Mpc}^{-1}$ , and the maximum  $k$ -bin center value is set to  $0.99 \text{ Mpc}^{-1}$  to avoid large thermal noise, see table 3.

Figure 4 illustrates the case of  $m_{\text{WDM}} = 5.3 \text{ keV}$  and  $M_{\text{turn}} = 10^7 M_{\odot}$  with the other parameters set to the same fiducial values than in figure 2. The red dashed lines depict the noiseless power spectra as a function of the wave mode  $k$  for the entire  $k$  range considered in our analysis. Note that a selection of 10 values of the redshifts, out of the 19 considered redshift bins, is displayed. The blue and purple areas indicate the 68% and 95% CL uncertainty bands around the noiseless power spectra. In particular, the blue lines with stars, indicating the  $k$ -bin centres, provide one realization of the noisy simulated power spectra. As expected from eq. (3.2), we see in figure 4 that the thermal noise becomes more important for the larger redshift values considered, dominating completely the signal for  $z \gtrsim 12$  in all  $k$  range. On the other hand, for the fiducial values of the parameters considered here, the thermal noise becomes negligible for  $6 \lesssim z \lesssim 8$ , while it rises again for  $z \lesssim 6$ . Notice that the variation of the relative importance of the thermal noise contribution, as a function of the redshift, depends on the fiducial parameter values considered.

### 3.3 Simulation-based inference

The general working principle of SBI is to compare an observed dataset with a sufficiently large set of simulated datasets (provided we have a simulator at hand), each of which has known values of the input parameters. As such, by selecting somehow the simulated datasets closest to the observed one, we can deduce the distribution of values of the parameters behind the latter. Of course, this process entirely relies on the assumption that the simulator captures reasonably well the physics responsible for the phenomena that we are observing. If this was not the case, the whole procedure would surely give misleading results. While

conceptually the SBI approach is not new,<sup>8</sup> the exponential advance in computing power allowing to work with largely parameterized Machine Learning (ML) models as deep neural networks and others have unlocked the potential of SBI to solve modern problems in science requiring complex simulations. This is e.g. the case in areas like cosmology, astrophysics, particle physics, and others [26, 29, 75–82]. Modern SBI flavors use these ML models to approximate the posterior distribution of the input parameters.

SBI differs from the MCMC approach in three essential ways. First, it is an approximate inference method that provides samples from approximations of the true posterior. Second, it is a likelihood-free method, referring to the fact that there is no need to evaluate a likelihood function. Third, it can be designed as a fully “pre-trained” strategy, meaning that the inference on new datasets takes essentially no computational time after the model has been fitted/trained to a large number of simulated datasets. Precisely the latter point is behind the computational advantages of SBI over MCMC seen in many analyses. At this point, it is important to note that none of the two approaches can guarantee to give, in practice, unbiased results. Indeed, while MCMC is surely unbiased provided the true likelihood is available (which is hardly the case for data coming from complex simulators<sup>9</sup>), SBI, on the other hand, in any of its flavours, approximate at the end of the day the posterior distribution in one way or another, thus being inherently biased. For sufficiently complex and large parametrisations (as in the case of modern Machine Learning models), such approximations can however become increasingly better with respect to the true unknown posterior distribution. On the other hand, the scaling of the computational time with the number of parameters entering the simulation has been shown to be much better with SBI as compared to typical flavours of MCMC [83]. It is worth noting that the above discussion on biased results refer to the statistical analysis, and not to potential simulation biases, a.k.a. simulation mis-modeling or “model mis-specification”. Clearly, the reliability of both MCMC and SBI methods rely on the quality of the simulation process, and assumes it gives a very good description of the generative process behind real data. On the other hand, the model mis-specification problem is well identified in the SBI community, which nowadays is working actively in trying to address it, see e.g. [84]. While addressing the mis-modeling of 21cmFAST and 21cmSENSE is beyond the scope of this work, we note that we have adopted a simplified simulation setup, including only one type of stars. Considering instead two or three types of stars would surely be more realistic, and the number of astrophysical parameters would double or triple accordingly. While this is still manageable for SBI algorithms, given the expected level of degeneracy between astrophysical parameters and the WDM mass, and the methodological scope of this paper, we have deliberately simplified the astrophysical setup, at the cost of having a certain level of model mis-specification.

As in most cases, the simulator we use in this work, contains, on the one hand, physical parameters of interest  $\theta$  and, on the other hand, it contains other latent (internal) stochastic

---

<sup>8</sup>As far as we know, likelihood-free methods appeared already in the 1980s, with the Approximate Bayesian Computation (ABC) method [74].

<sup>9</sup>Modulo observables for which it makes sense to apply the Central Limit Theorem, thus being Gaussian in practice. However, typically, complex simulators produce data with intractable likelihoods. Note that simulating data is formally equivalent to sampling the corresponding likelihood. This is where the correspondence between simulators and likelihoods arise.

variables  $\boldsymbol{\eta}$ . In our case, the simulator is the combination of `21cmFast` and `21cmSense`, giving rise to multiple realizations of the noisy 21cm power spectra, as detailed in section 3.2. The physical parameters  $\boldsymbol{\theta}$  are given in table 1 while the  $\boldsymbol{\eta}$  latent variables are given by the randomly sampled initial conditions of the density and velocity fields. Clearly, the simulator output depends on both  $\boldsymbol{\theta}$  and  $\boldsymbol{\eta}$ , and typically one is only interested in making inferences on the former by estimating the posterior  $p(\boldsymbol{\theta}|\mathcal{D})$  via the Bayes theorem.

Had we had access to the true likelihood  $p(\mathcal{D}|\boldsymbol{\theta})$ , we could confidently use MCMC, even if slower, but providing unbiased inference. However, this likelihood is typically intractable, since it is the result of marginalizing over the latent variables:  $p(\mathcal{D}|\boldsymbol{\theta}) = \int d\boldsymbol{\eta} p(\mathcal{D}|\boldsymbol{\eta}, \boldsymbol{\theta})p(\boldsymbol{\eta}|\boldsymbol{\theta})$ . While it is common in the literature to approximate this likelihood somehow (for example, as a Gaussian), the resulting posteriors can show noticeable deviations with respect to likelihood-free approaches [29]. That being said, for the 21cm power spectrum data the assumption of a Gaussian likelihood is reasonable, at least for the smallest physical scales, while for the highest scales, the Gaussian approximation is expected to break down.<sup>10</sup> Still, the use of SBI is motivated here, since it does not rely on evaluating the (poorly or adequately approximated) likelihood, leave apart its computational advantages. Note that while in SBI there is no need to evaluate the likelihood, the simulated power spectra coming from `21cmFAST` that we provide to the algorithm are formally samples from the true (unknown) likelihood, so we do use information of the latter.

In this work, we use a flavour of SBI called Neural-Ratio Estimation (NRE), through the Python library `swyft` [85]. It aims at estimating the posterior probability  $p(\boldsymbol{\theta}|\mathcal{D}_{\text{obs}})$ , given an observed dataset  $\mathcal{D}_{\text{obs}}$ , by approximating the likelihood-to-evidence ratio  $r(\mathcal{D}_{\text{obs}}, \boldsymbol{\theta}) \equiv p(\mathcal{D}_{\text{obs}}|\boldsymbol{\theta})/p(\mathcal{D}_{\text{obs}})$  with a neural network model (here  $p(\mathcal{D}_{\text{obs}})$  is the evidence), and subsequent use of the Bayes theorem to get the posterior via multiplication by the prior. The neural network model is used as a binary classifier, trained to discriminate matching sample pairs  $\{\mathcal{D}_i, \boldsymbol{\theta}_i\}$  obtained through the simulation process (i.e.,  $\mathcal{D}_i$  is the simulated dataset corresponding to the input value  $\boldsymbol{\theta}_i$ ), from those unmatching pairs  $\{\mathcal{D}_i, \boldsymbol{\theta}'_i\}$ , where  $\boldsymbol{\theta}'_i$  is unrelated to  $\mathcal{D}_i$  according to the simulator. This makes sense based on the well-known fact that probabilistic binary classifiers (as neural networks) are formally an approximation of the likelihood-ratio between the two concerning hypotheses [86].

The pipeline of our SBI procedure is the following:

1. We generate  $\bar{n} = 20000$  samples of  $\boldsymbol{\theta}$  from  $p(\boldsymbol{\theta}|\mathcal{D}_{\text{UV}})$  via the MCMC analysis performed in section 3.1 consistent with UV luminosity data  $\mathcal{D}_{\text{UV}}$ . These constitute samples from the prior as far as SBI is concerned and are expected to represent an exhaustive enough coverage of the whole parameter space, for the sake of good reconstruction capabilities (see section 4).
2. For each of these values of  $\boldsymbol{\theta}$ , we use the simulator to generate the corresponding (noiseless) simulated dataset  $\tilde{\mathcal{D}}$ . We represent the latter as a 1-dimensional array of power spectrum as a function of  $k$ , concatenating sub-arrays corresponding to different redshifts, see [28] using the same procedure. The training data contains  $n = 18000$  noiseless simulations, while we

<sup>10</sup>This is all depending on the extent at which the conditions for the Central Limit Theorem applies since the power spectrum is a sum of a large (small) number of contributions at small (large) physical scales.

use the remaining 2000 simulations as validation ( $n_{\text{val}} = 1900$ ) plus test ( $n_{\text{test}} = 100$ ) data. We corrupt the training datasets  $\tilde{\mathcal{D}}$  with ten realisations of thermal noise, thus resulting in a total of  $10 \times n = 180000$  noisy datasets  $\mathcal{D}_i$ ,  $i = 1, \dots, 10 \times n$ . We thus build tuples  $\{\mathcal{D}_i, \theta_i\}$  of “good samples” (labelled `class 1`).

3. We obtain the same amount of “bad samples” (labelled `class 0`) by building tuples  $\{\mathcal{D}_i, \theta'_i\}$ , where the previously obtained datasets are now paired with random values of the parameters instead, different from that  $\theta_i$  used to generate them.
4. These two sets of tuples will be the input of a classification model, subsequently trained to distinguish `class 1` samples from `class 0` samples. The training consists of minimizing a standard binary “cross-entropy” cost function, with the `Adam` stochastic gradient descent algorithm [87]. The classifier is given by a feed-forward neural network with three hidden layers, each containing 256 units. The network is trained with a batch size of 64, and we decay the initial learning rate of  $10^{-3}$  by 0.95 after every epoch.
5. Once trained, the classifier is used to build the function  $r(\mathcal{D}_{\text{obs}}, \theta)$ , which together with the prior  $p(\theta|\mathcal{D}_{\text{UV}})$  can be evaluated in order to obtain the desired posterior  $p(\theta|\mathcal{D}_{\text{obs}})$ . Consistency checks of the resulting posteriors are presented in the form of coverage plots, see appendix B for the results. The interested reader can find more details in [88].

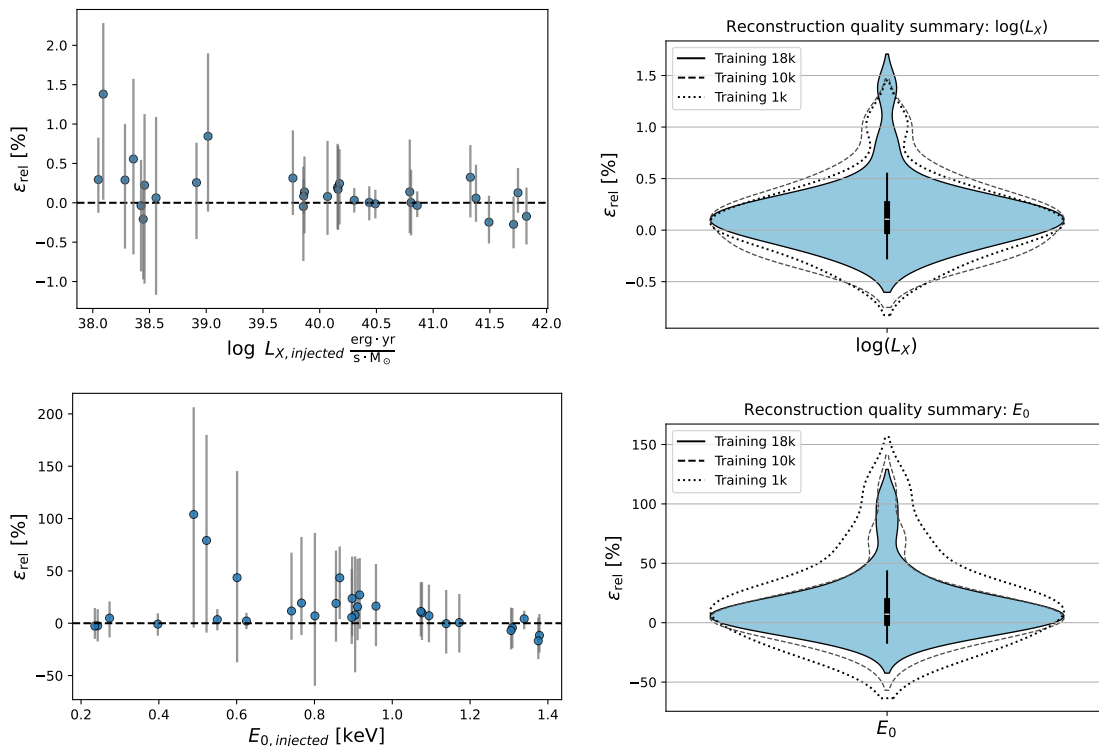
## 4 Results

In this section, we present the results of our study. In section 4.1, we show the generic reconstruction capabilities of our SBI algorithm in terms of marginal (single parameter) posteriors. Given the high computational cost of a single `21cmFast` simulation, we include as well a comparison of the reconstruction capabilities when the number of simulations,  $n'$ , used for training decreases, i.e. when  $n' < n$ . In section 4.2, we discuss the reconstruction of the DM mass in the presence of degeneracies or correlations with the other astrophysics parameters analyzing the corresponding 2D posterior distributions. We close the discussion presenting the  $m_{\text{WDM}}$  bound resulting from our analysis as a function of  $M_{\text{turn}}$  in section 4.3.

### 4.1 Parameter reconstruction

We first discuss the performance of our procedure in reconstructing the different parameters as follows:

1. For each parameter, we compute the marginal posteriors corresponding to each of  $N_{\text{sim}} = 30$  simulations. The latter has been chosen randomly from the set of  $n_{\text{test}}$  simulations that were kept apart from the ones used to train our algorithms. The posterior of every parameter  $\theta$  gives rise to an expected value,  $\langle \theta \rangle$ , which we compare to the true value  $\theta_{\text{true}}$  of the corresponding simulation to construct the relative error,  $\varepsilon_{\text{rel}} = 100(\langle \theta \rangle - \theta_{\text{true}})/\theta_{\text{true}}$ , as a function of  $\theta_{\text{true}}$ . The results are shown in the left panels of figures 5, 6, and 7 where the error bars correspond to the 95% CL of the posteriors.
2. From the above procedure, we have  $N_{\text{sim}}$  values of  $\varepsilon_{\text{rel}}$  for each parameter  $\theta$ , whose distributions are shown in violins plots in the right panels of the same figures. Blue violins

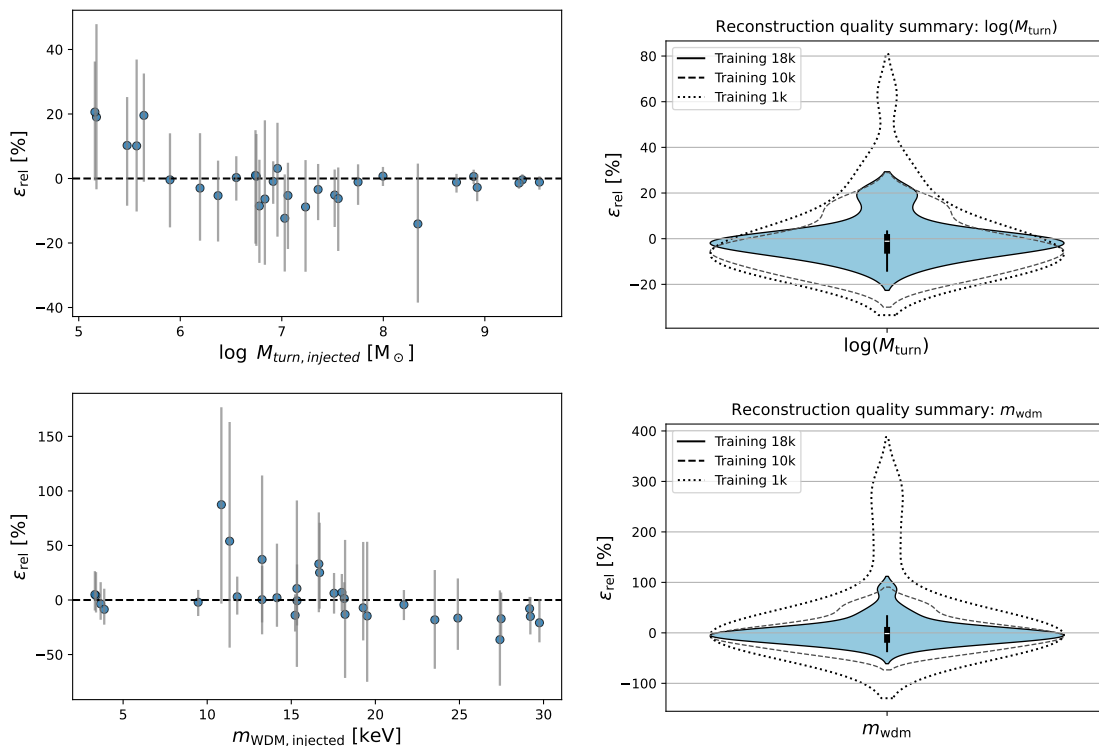


**Figure 5.** Reconstruction quality in terms of relative error  $\epsilon_{\text{rel}}$  for the X-ray parameters used in this work. Left panel: relative error in the reconstruction as a function of the injected value of the parameter for  $N_{\text{sim}} = 30$  random simulations. The error bars represent the 95% CI of the posterior distribution in each case. Right panels: corresponding distribution of the relative error for a training setup with  $n = 18\text{k}$  simulations (blue violins). The transparent violins are analogous to the blue one, but using instead  $n' = 10\text{k}$  simulations (dashed) and  $n' = 1\text{k}$  simulations (dotted). See text for details.

are shown together with their corresponding “candle” in the centre. The latter represents the mean (central white marker), the Interquartile range (IQR, spanning from the 25% to the 75% quartile) at the extremes of the thick bar, as well as the points extending up to 1.5 times the IQR (thin inner vertical segment). All those violins are reported in figure 8 to ease their one-to-one comparison.

Given the high computational cost of running 21cmFAST in our machines ( $\sim 20'$  per simulation to obtain the power spectrum), it is natural to ask how the reconstruction capabilities degrade as we decrease the number of simulations used for training. While so far the results have used all the  $n = 18\text{k}$  available simulations (blue violins), we repeat the exercise with  $n' = 10\text{k}$  simulations (dashed transparent violins) and  $n' = 1\text{k}$  simulations (dotted transparent violins).

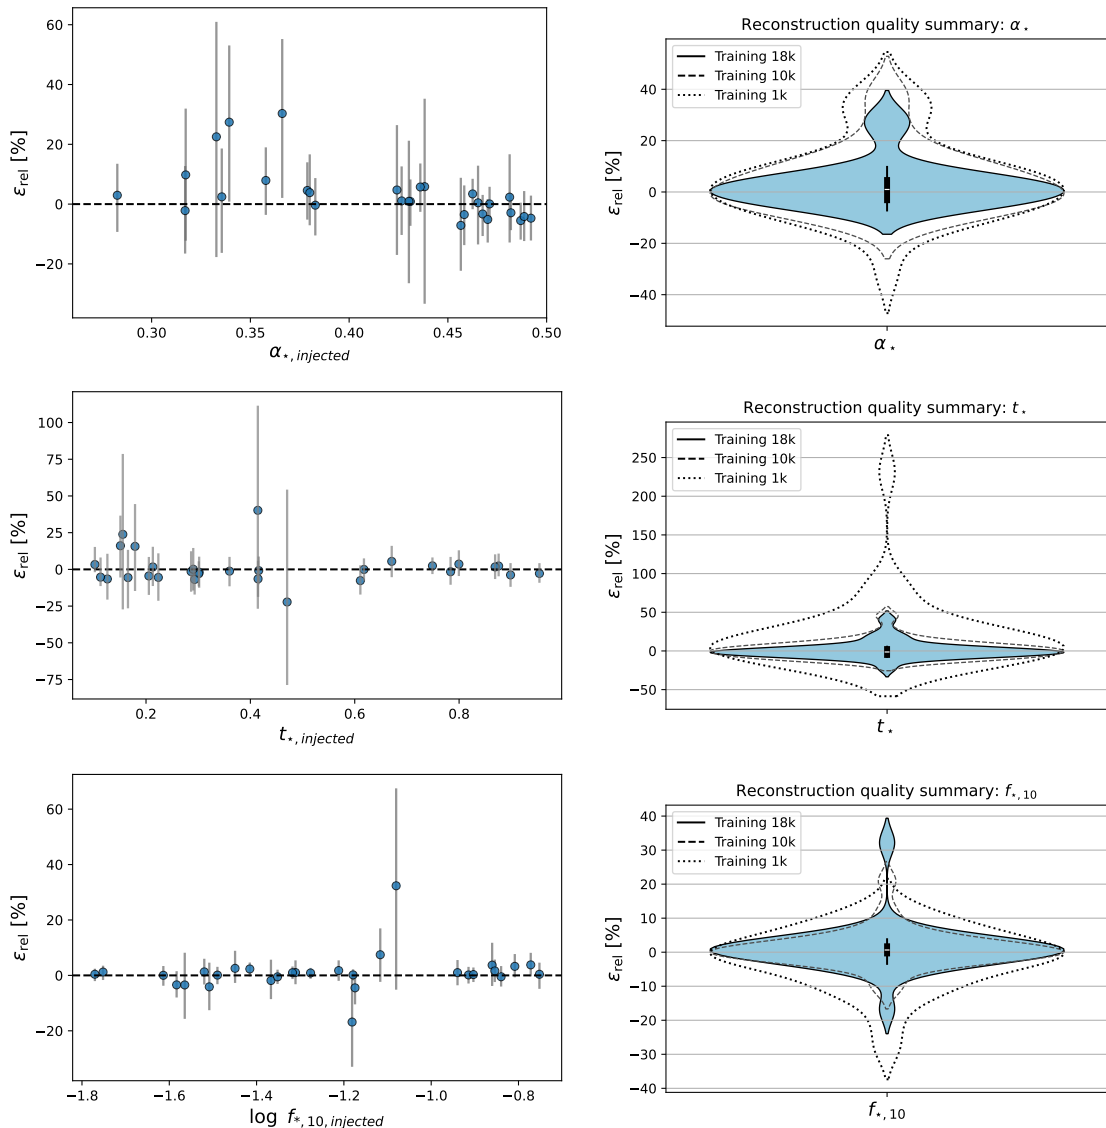
Figure 8 highlights different reconstruction capabilities across the parameters. The best reconstruction is achieved for the parameter  $L_X$ , see the upper panels of figure 5 for more details. We see that the  $\epsilon_{\text{rel}}$  is indeed below the per cent level in most cases. Note that actually, the first results of HERA [58] were already able to constrain this parameter. In the case of  $E_0$ , the reconstruction quality worsens for the true values of  $E_0 \gtrsim 0.6\text{ keV}$ , as



**Figure 6.** Analogous to figure 5 but for the  $m_{\text{WDM}}$  and  $M_{\text{turn}}$  parameters. The 30 points in each plot correspond to the same 30 random simulations used in figure 7.

visible in the bottom panels of figure 5. Similar behaviour has been observed in [18] for a different parametrization of the 21cm signal, implementation of the WDM imprint and a different range of the corresponding parameters, see the discussion in section 1. Here in particular we consider a wider range of  $m_{\text{WDM}}$  and  $M_{\text{turn}}$  values. Note that larger values of the energy threshold for X-rays escaping galaxies,  $E_0$ , make the X-ray photon spectrum heating the IGM harder. Harder X-rays have longer mean free path,  $\lambda_X \sim E_X^{2.6}$  where  $E_X$  denotes the X-ray energy [89]. This implies that for larger  $E_0$  the heating is happening in a more homogeneous manner leading to less contrast in the 21cm signal at EoH when other parameters are fixed. Consequently, we have a suppressed power spectrum at EoH and thus a more difficult reconstruction of  $E_0$  for these higher values.

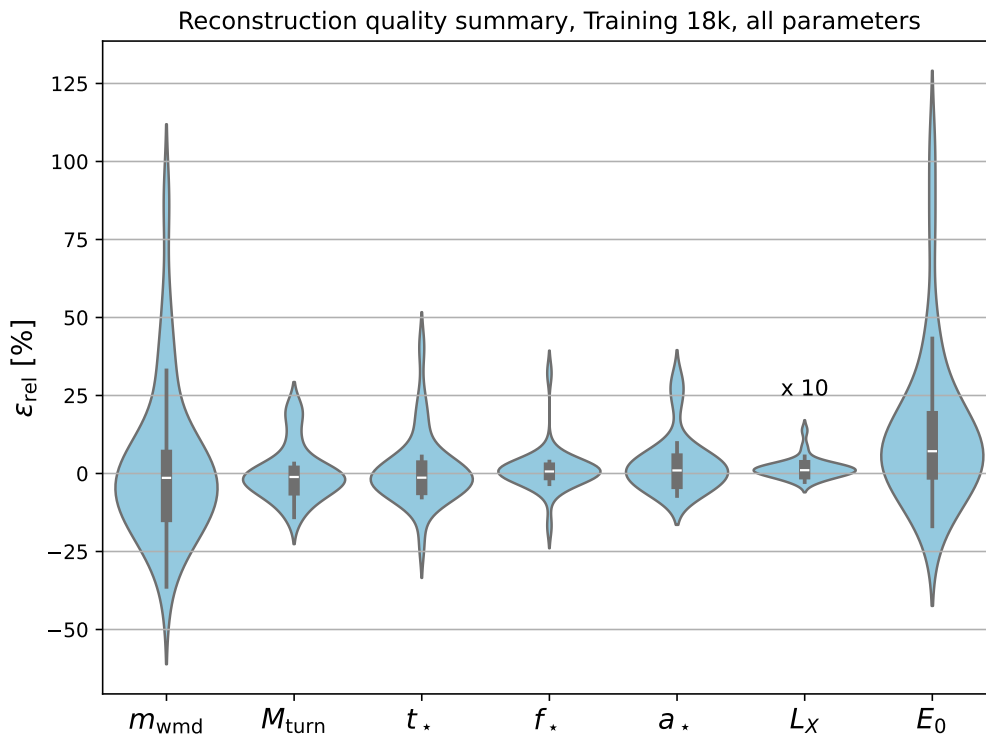
On the other hand, focusing now on the  $M_{\text{turn}}$  parameter in the upper panels of figure 6 we observe that the reconstruction quality is very high for the largest true values of this parameters while becoming more challenging for  $M_{\text{turn}} \lesssim 10^{7.5} M_{\odot}$ . This is expected, since  $M_{\text{turn}} > 10^{7.5} M_{\odot}$  is expected to drive the overall shift of the power spectrum to later redshifts as our largest considered WDM effect induces an  $M_{\text{WDM}}^{\text{cut}} = 2 \times 10^7 M_{\odot}$  for  $m_{\text{WDM}} = 3$  keV. Lower  $M_{\text{turn}}$  values have an imprint that is more difficult to disentangle from other parameters e.g. the free-streaming effect of WDM could prevent a good reconstruction of the former. We have checked that this is the case for the two left-most points in the bottom-left panel of figure 6, corresponding to a very low  $m_{\text{WDM}}$ . Overall, the reconstruction error of  $M_{\text{turn}}$  is always better than  $\sim 20\%$  for the set of simulations that we have used. On the other hand,



**Figure 7.** Analogous to figure 5 but for the SFR parameters considered in this work. The 30 points in each plot correspond to the same 30 random simulations used in figure 5.

the reconstruction of the DM mass is much better at the lowest  $m_{\text{WDM}}$ , see the bottom panels of figure 6. This is for reasons analogous to the better reconstruction of large  $M_{\text{turn}}$ , i.e. the impact of a sufficiently low  $m_{\text{WDM}}$  on the power spectrum can be more easily distinguished from the effect of the other parameters. Note that, at the highest DM masses considered in this example, the  $\epsilon_{\text{rel}}$  is also reasonably good. We have checked that this is due to relatively low true values of  $M_{\text{turn}}$  for our particular choice of simulations.

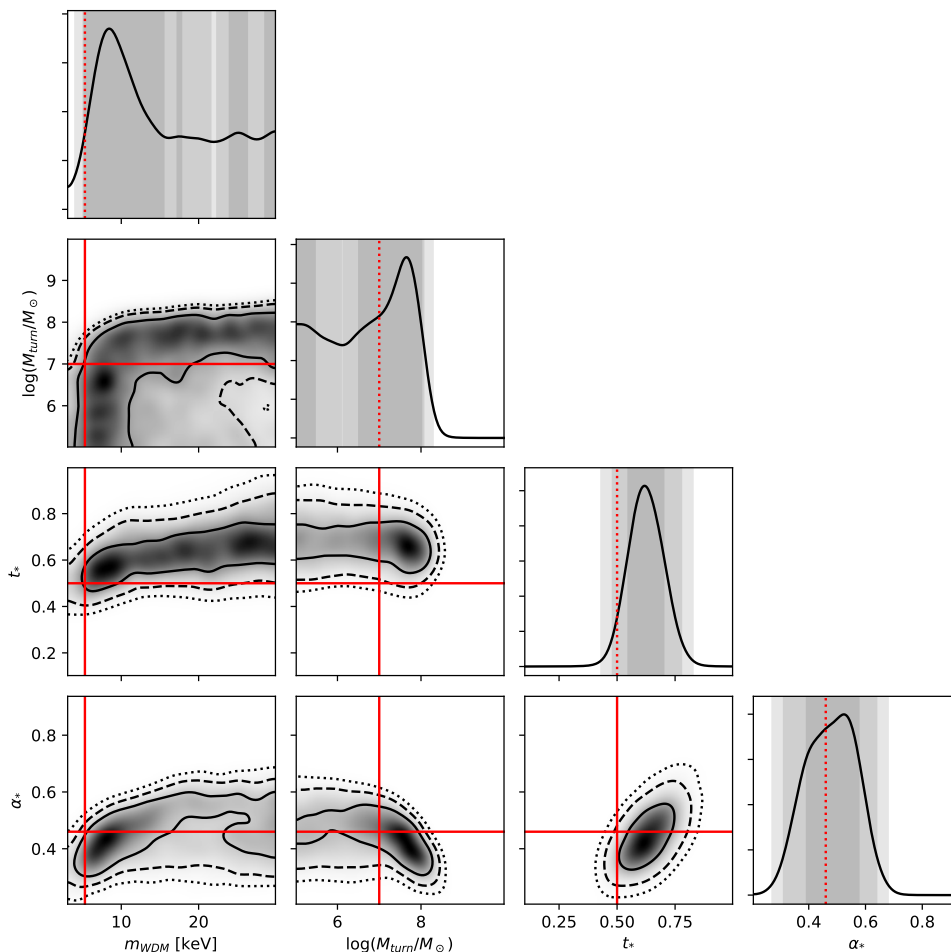
On the other hand, we see in figure 7 that the SFR parameters are reconstructed at the  $\mathcal{O}(10 - 50\%)$  level. Yet, their  $\epsilon_{\text{rel}}$  distributions clearly peak around zero, for the other parameters. We observe a better reconstruction in general for the largest true values of the SFR parameters. This is expected as, for example, in the case of  $\alpha_*$  (upper panels), sufficiently large values of this parameter will imply larger star-forming halos, which in turn



**Figure 8.** Summary of the violin plots representing the distribution of the relative error  $\varepsilon_{\text{rel}}$  when considering  $n = 18000$  simulations in the reconstruction of the parameters reported at the bottom of the figure. Note that the violin for  $L_X$  has been rescaled by a factor 0.1 for visibility purposes. A zoom-in version of those violins and their degradation reducing the number of simulations were shown in the previous figures 5, 6 and 7.

becomes the dominating cause for the delay in the appearance of the power spectrum features, compared to the effect of the other parameters. Similar reasoning can be made for  $t_*$  (central panels), controlling the SFR time scale. Larger  $t_*$  values imply a suppression of the SFR, which causes a delay in the power spectrum features.

Concerning the degradation of the reconstruction when training with a smaller number of simulations, we can observe in the right panels of figures 5, 6 and 7 that the lowest-budget training of  $n' = 1\text{k}$  noiseless simulations is not enough to converge on the reconstruction, thus producing a much wider distribution of  $\varepsilon_{\text{rel}}$  for all the parameters. In contrast, the intermediate-budget training, using  $n' = 10\text{k}$  simulations, produces in general a reconstruction with similar quality to the one with the highest-budget training ( $n = 18\text{k}$ ). As a result, we can conclude that using the 21cmFast simulator and the SBI technique that we have adopted for our setup, the minimum number of simulations to produce good quality results is 10k.



**Figure 9.** Reduced set of recovered 1D and 2D marginalized posteriors on our astrophysics and WDM parameters assuming a mock 1000h observation of the 21-cm power spectrum with HERA for  $m_{\text{WDM}} = 5.3 \text{ keV}$  and  $M_{\text{turn}} = 10^7 M_{\odot}$ . The other parameters have been fixed to the same fiducial values as in figure 2. The entire triangle plot, including all the free parameters considered in our SBI analysis, is shown in figure 12.

## 4.2 Degeneracies in the posterior distributions

Having presented the generic reconstruction capabilities of our algorithm, here we address the degeneracies between the DM mass and the other astrophysics parameters discussed in section 2.4. We do so by considering the 2D posterior distributions for one illustrative fiducial scenario and identifying the shape of the different contours.

In figure 9, we focus on the correlation between  $m_{\text{WDM}}$  and  $M_{\text{turn}}, \alpha_*, t_*$  to ease the reading of the plot. Indeed, it appears that the HERA sensitivity is such that no strong

correlation appears between  $m_{\text{WDM}}$  and  $f_{\star,10}, L_X, E_0$ , see the full triangle plot in figure 12 in the appendix. For illustrative purposes, we have fixed  $m_{\text{WDM}} = 5.3 \text{ keV}$  and  $M_{\text{turn}} = 10^7 M_{\odot}$  for the fiducial scenario and we have fixed all the other parameters to the same values as in figures 2 and 4. Since our choice of  $M_{\text{turn}}$  is such that  $M_{\text{turn}} \sim M_{\text{WDM}}^{\text{cut}}(5.3\text{keV})$ , we expect that degeneracies shall appear between the  $m_{\text{WDM}}$  and  $M_{\text{turn}}$ . Furthermore, from the discussion in the previous section, we expect a relatively good recovery of low dark matter mass with HERA sensitivity. All parameters shall thus be relatively well reconstructed for our choice of fiducial scenario.

These expected trends can indeed be observed in the posterior distributions of figure 9. The top panels correspond to the marginalised 1D posteriors whereas the other panels show the 2D posteriors. The continuous, dashed and dotted contours in the 2D panels correspond to the 68%, 95% and 99% confidence levels (CL) contours. The same percentiles are represented as darker, medium and lighter grey regions in the marginalised 1D posteriors. The vertical and horizontal red lines indicate our fiducial parameter values. As visible here and in the full triangle plot, the true values of the parameters are always lying within the 95% CL intervals confirming the relatively good reconstruction of the parameters. Furthermore, a positive correlation between  $m_{\text{WDM}}$  and  $M_{\text{turn}}$  is visible, showing the expected behaviour, see section 2.4. Indeed, to compensate for an increase in  $M_{\text{turn}}$ , which shifts the power spectrum features to later times, one could consider colder dark matter by increasing the WDM mass (and vice versa). For sufficiently low WDM mass, low  $M_{\text{turn}}$  cannot be discarded, since the power spectrum would be insensitive to the latter; resulting in a high tail at low  $M_{\text{turn}}$  values as seen in the corresponding 1D distribution. The same logic applies to the WDM mass: for sufficiently high  $M_{\text{turn}}$ , high values of  $m_{\text{WDM}}$  will have no impact on the power spectrum, thus resulting in a high tail at high  $m_{\text{WDM}}$ , shown in its 1D posterior.

An analogous reasoning can be made when looking at the degeneracy between  $m_{\text{WDM}}$  and the other two parameters shown in figure 9. To compensate for an increase in  $\alpha_{\star}$ , leading to a larger fraction of stars into heavier haloes, we should increase  $m_{\text{WDM}}$  so as not to suppress low mass haloes population. On the other hand, when  $t_{\star}$  increases the SFR decreases, which is again compensated by a colder DM, i.e. larger  $m_{\text{WDM}}$ , thus featuring a similar positive correlation among these two parameters.

### 4.3 Impact of $M_{\text{turn}}$ on the $m_{\text{WDM}}$ lower bound

We now turn our attention to the potential for 21cm Cosmology, in the particular case of the HERA telescope, to surpass existing probes of the warmness of dark matter. More concretely, we aim to put lower bounds on the WDM mass as a function of the injected threshold mass  $M_{\text{turn}}$ , while keeping fixed the rest of the astrophysical parameters. The injected  $m_{\text{WDM}}$  mass has been fixed to 28 keV, as a representative of a CDM scenario, while not being right at the extreme of our prior range to avoid border effects. Following the standard procedure in Bayesian statistics, for every value of the injected  $M_{\text{turn}}$  we obtain the 95% CL bounds,  $m_{\text{WDM}}^{\text{min},95}$ , by solving for the latter in the equation

$$0.95 = \int_{m_{\text{WDM}}^{\text{min},95}}^{\infty} dm_{\text{WDM}} p(m_{\text{WDM}}|\mathcal{D}),$$

benchmark	$\log(f_{\star,10})$	$\alpha_{\star}$	$t_{\star}$	$L_X$ [erg yr/(s $M_{\odot}$ )]	$E_0$ [keV]	$m_{\text{WDM}}$ [keV]
1	-0.9	0.46	0.5	$10^{40}$	0.5	28
2	-0.9	0.46	0.5	$10^{41}$	1.1	28

**Table 2.** Benchmark models submitted to the trained NRE to extract the 95%CL bounds on  $m_{\text{WDM}}$  as a function of  $M_{\text{turn}}$  in figures 10 and 11.

where  $p(m_{\text{WDM}}|\mathcal{D})$  is the  $m_{\text{WDM}}$  posterior distribution given the data  $\mathcal{D}$ . Such data consists of 5 simulations  $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_5\}$ , each with different initial conditions but generated from identical fiducial parameter values, and including thermal noise. The latter was sampled from HERA sensitivities using `21cmSense` as in the training procedure. We make inference on each of these simulated datasets using our NRE algorithm, previously trained on the maximum budget of  $n = 18\text{k}$  simulations, see section 4.1. The combined posterior can thus be simply estimated as:<sup>11</sup>

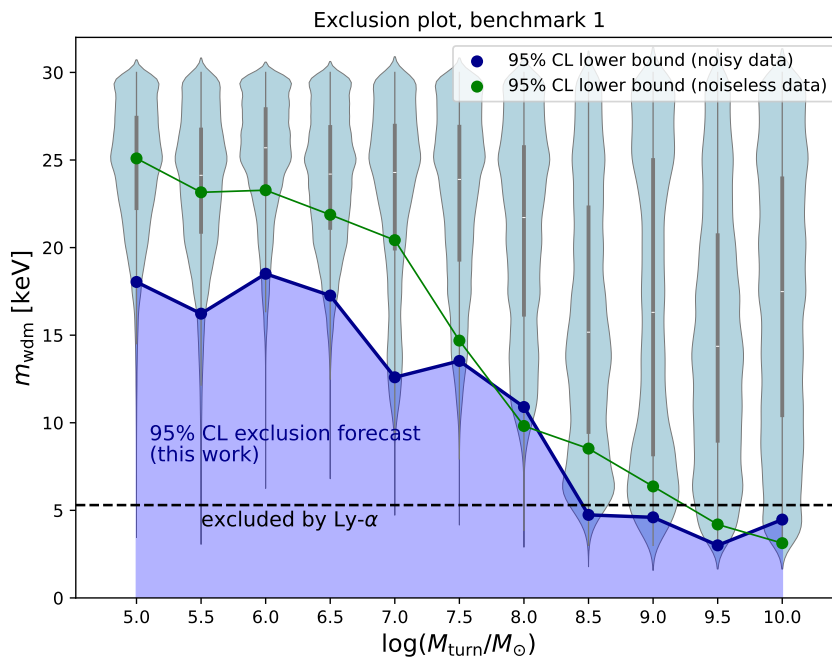
$$p(m_{\text{WDM}}|\mathcal{D}) = p(m_{\text{WDM}}) \prod_{i=1}^5 r_i(m_{\text{WDM}}; \mathcal{D}_i), \quad (4.1)$$

where  $p(m_{\text{WDM}})$  is the DM mass prior, and  $r_i(m_{\text{WDM}}; \mathcal{D}_i)$  is the output of the algorithm for the dataset  $\mathcal{D}_i$ , noting that this is nothing but the estimation of the likelihood-to-evidence ratio.

We compute the bound for a regular grid of  $M_{\text{turn}}$  masses in the range  $[10^5 M_{\odot}, 10^{10} M_{\odot}]$ . The results are shown in figures 10 and 11 for two benchmarks, specified in table 2. We consider those two benchmarks to illustrate the dependence on the bound both on  $M_{\text{turn}}$  and on the X-ray parameters. Benchmark 2, in particular, exhibits higher  $E_0$  and  $L_X$  values, which shall result in a more homogeneous heating of the IGM. Consequently, a more suppressed noiseless power spectrum is expected within the redshift range probed by HERA. With all other parameters fixed, this makes benchmark 2 potentially more challenging to test. In both figures, each violin represents the posterior distribution  $p(m_{\text{WDM}}|\mathcal{D})$  for a given injected  $M_{\text{turn}}$ , reported on the x-axis, while the blue dots represent the corresponding 95% lower bound on  $m_{\text{WDM}}$ . On the other hand, for both benchmarks we repeat the exercise by using instead *noiseless* data, i.e. not adding thermal noise, intuitively leading to stronger WDM bounds with respect to the analogous inference on noisy data. The results, represented with the green dots, are just for academic purposes. They indicate how much the thermal noise degrades the bounds.

As expected from the discussion in section 2.3, larger threshold masses would prevent a good reconstruction of the DM mass, resulting in weaker bounds on the latter. Nonetheless, for both benchmarks, we find that the lower bound on  $m_{\text{WDM}}$  can surpass the Lyman- $\alpha$  forest constraint of 5.3 keV (horizontal black dashed line in both figures) if  $M_{\text{turn}}$  is below  $10^8 M_{\odot}$ . Furthermore, if all CD galaxies have  $M_{\text{turn}} = 10^5 M_{\odot}$ , or equivalently, if all CD galaxies are assumed to be in the form of MCGs, thermal DM masses up to  $m_{\text{WDM}} \simeq 17$  keV and  $m_{\text{WDM}} \simeq 10$  keV could be excluded with the X-ray parameters of benchmark 1 and 2, respectively. Even if it might be unrealistic to consider that all CD galaxies could be part of


<sup>11</sup>We assume that all the  $\mathcal{D}_{1,\dots,5}$  are independent.



**Figure 10.** Forecast of the HERA sensitivity to  $m_{\text{WDM}}$  as a function of the injected  $M_{\text{turn}}$  value. For every  $M_{\text{turn}}$  we show in blue violins the posterior distribution of  $m_{\text{WDM}}$  from 5 noisy simulations of benchmark 1, see table 2. The corresponding 95% CL lower bound on  $m_{\text{WDM}}$  is shown with the dark blue bullets, while the analogous lower bound using noiseless data is shown with green bullets. The lower bound from Ly- $\alpha$  forest data is included as a dashed line for comparison.

one single population of galaxies, the latter number provides us with a quantitative *lower limit* on the WDM mass that could be probed with HERA for our choices of fiducial scenario. As expected, in the noiseless case (green dots), the resulting lower bounds on  $m_{\text{WDM}}$  are stronger than in the noisy case (blue dots).

One could also provide an analysis involving multiple galaxy populations. An instance of the latter is available in **21cmFast** combining MGC and AGC while accounting for feedback effects, see [49]. Yet, this implies much slower simulation and a much larger set of parameters to account for the two population galaxy properties or, equivalently, a larger number of simulations to run in order for the training procedure to converge. Furthermore, this would induce extra degeneracies between those parameters and the WDM mass. Such an analysis is beyond the scope of this paper. Our main objective was to emphasize that the choice of threshold mass for star-forming galaxies, or equivalently, the choice of galaxy populations plays a major role in the determination of the sensitivity of 21cm Cosmology experiments on the NCDM imprint. Similar conclusions were reached in [16] providing a sensitivity estimate of SKA on a discrete set of source modeling below virial masses of  $10^8 M_{\odot}$ . In addition, comparing figures 10 and 11, for the selected X-ray parameters considered, we have a rather mild effect on the maximal threshold mass that would help to improve on existing bounds on  $m_{\text{WDM}}$ . In contrast, the lower bound  $m_{\text{WDM}}$  at the lowest threshold masses considered in this



**Figure 11.** Same as figure 10 but for benchmark 2 displaying higher values of the X-ray parameters, see table 2.

analysis could become less stringent for the higher explored values of the X-ray parameters giving rise to a more suppressed power spectrum at the redshifts of interest.

Note that it is expected that the results presented here for WDM can actually be easily generalized to other free-streaming NCDM scenarios that would give rise to an unimodal DM momentum distribution times the momentum squared (i.e. with the product  $p^2 f(p)$  presenting a single peak). This has been discussed in e.g. [59–62].<sup>12</sup> For example, for a NCDM candidate produced through the freeze-in mechanism associated to the decay of a heavy bath particle, a 5.3 keV bound for WDM can be reinterpreted into a 16 keV constraint on the DM mass from FI (existing Lyman- $\alpha$  bound). On the other hand the 10 keV WDM bound, that could be obtained with HERA if  $M_{\text{turn}} = 10^5 M_{\odot}$  in our benchmark model 2 of figure 11, would be reinterpreted into a 37 keV bound on DM from freeze-in. For other NCDM scenarios, involving e.g. dark acoustic oscillations, a further analysis would be needed to determine if telescope as HERA could distinguish WDM from such NCDM, see e.g. [9, 19, 90] in that direction.

<sup>12</sup>Equating the rms velocity of unbounded DM today in a given NCDM scenario, to the one of thermal WDM today, one can derive a constraint on the NCDM DM mass  $m_{\chi}$  as a function of the WDM mass bound  $m_{\text{WDM}}^{\text{Lim}}$  as follows:  $m_{\chi} \gtrsim 1.75 \text{ keV} \times (m_{\text{WDM}}^{\text{Lim}}/5.3 \text{ keV})^{4/3} \times \langle q_{\star}^2 \rangle^{1/2} T_{\text{ncdm}}$  when the NCDM is produced through a single production mechanism and the mass bound  $m_{\text{WDM}}^{\text{Lim}}$  is obtained from the imputing the NCDM power spectrum that is suppressed at high wave numbers. The dimensionless constant parameter  $T_{\text{ncdm}}$  is scaling as the ratio of the temperatures of the bath at DM production and today divided by the scale factors at the same periods while  $\langle q_{\star}^2 \rangle^{1/2} = \langle p^2 \rangle^{1/2} / T_{\text{ncdm}} \times a(t) / T_{\gamma}(t_0)$  is the root means squared of a dimensionless comoving momentum, see [62] for details. In the case of freeze-in from decays,  $\langle q_{\star}^2 \rangle = 35/4$  and, for freeze-in above electroweak phase transition,  $T_{\text{ncdm}} \simeq 3$ .

## 5 Conclusion

The redshifted 21cm signal from CD and EoR will soon be probed to a high signal-to-noise ratio thanks to the development of radio interferometers, such as HERA and SKA. These are expected to probe the distribution of IGM fluctuations at those epochs, inaccessible as of today. The 21cm signal will depend on -yet unknown- astrophysics phenomena and cosmic dawn galaxy evolution together with potential distinctive imprints beyond the standard model (BSM) phenomena. In this paper, we go beyond the  $\Lambda$ CDM model considering the possibility for dark matter to be non-cold. As is well known, a plethora of BSM scenarios can give rise to an NCDM behaviour, yet, for concreteness we focus on thermal Warm Dark Matter as a benchmark for free-streaming dark matter.

Thermal WDM is expected to significantly suppress the number density of low halo masses, giving rise to a delay in the 21cm signal features. In this work, we parametrize the WDM imprint on the linear matter power spectrum by making use of a transfer function parametrized as in eq. (2.8) involving a free-streaming scale described in eq. (2.9). Furthermore, for the non-linear evolution of the matter density perturbations, we have updated the WDM effect on the halo mass function in the public code `21cmFast` by considering a sharp-k cutoff window function instead of the default top-hat one, see section 2.4. Our astrophysics framework is the cosmic dawn galaxy parametrization introduced in [41] considering one single galaxy population characterized by a threshold mass for star formation  $M_{\text{turn}}$ , a star formation rate amplitude,  $f_{*,10}$ , and tilt  $\alpha_*$ . The X-ray heating from stars is driven by the soft X-ray luminosity amplitude,  $L_X$ , and the energy threshold for X-rays to escape galaxies,  $E_0$ . All these astrophysics parameters shape the 21cm signal and can be expected to display some level of degeneracy with the WDM mass,  $m_{\text{WDM}}$ , that controls the WDM free streaming length. In this paper in particular, we emphasize the key role played by the  $M_{\text{turn}}$  parameter in the sensitivity of the 21cm signal to the WDM imprint.

Traditionally, likelihood-based inference methods, and especially MCMC, have been used in cosmology, in general. Apart from relying on a good likelihood specification (something which is not possible in general), MCMC methods have other features which render them computationally inconvenient. First, they require to produce joint samples of the full parameter space (even beyond those parameters of direct interest). Second, a full MCMC from scratch has to be performed for every newly observed dataset. In our particular case of the 21cm power spectrum, the likelihood may be Gaussian to a good approximation. Yet, the computational disadvantages mentioned above still apply. Consequently, here we address these problems by making use of a Simulation-based Inference (SBI) analysis, which does not rely on a likelihood specification, while being on the other hand advantageous from the computational point of view with respect to MCMC methods.

We generate mock data with `21cmFast` while varying the initial conditions from a random seed to account for the cosmic variance. Before our Bayesian inference on such data, we constrain the parameters of interest from the UV luminosity data from [71], as a result of which we obtain priors to our SBI analysis. Our MCMC results reported, in figure 3, provide qualitatively similar results to [41], but making use of a WDM cosmology and of more recent data sets. We finally discuss the results of our SBI analysis in section 4. We begin by commenting on the parameter reconstruction capabilities of our procedure, before

discussing their degeneracies. In particular, from figure 9 we can corroborate our expectations that low thresholds for star formation galaxies are a necessary ingredient to probe the WDM scenario in the low mass regime. This is furthermore quantified in our conclusion plots of figures 10 and 11, where we display the sensitivity of a fully built radio interferometer, namely HERA, to  $m_{\text{WDM}}$  as a function of  $M_{\text{turn}}$ , assuming a single population of galaxies and 1000h of observations. From these plots, we see that the lower bound prospects on the WDM mass obtained in this work are stronger than the existing bound coming from Lyman- $\alpha$  forest data, provided that the threshold mass  $M_{\text{turn}} \lesssim 10^8 M_{\odot}$ . Although this result is fairly robust with respect to the X-ray source properties considered in this study, these properties may influence the strength of the WDM constraint for lower threshold masses.

As a final comment on the astrophysics modelling, the galaxy model considered here is rather simple and certainly does not account for multiple cosmic dawn galaxy populations that are expected to be present in the early universe. Yet, accounting for the latter will introduce a larger number of parameters that will display extra degeneracies with the WDM mass that should carefully be treated. From our analysis, it is clear that the assumed properties of molecular cooling galaxies, that are expected to be hosted in minihaloes with lower  $M_{\text{turn}}$ , will play a decisive role in the 21cm sensitivity to NCDM properties. Similar conclusions were already drawn in e.g. [16] considering a different analysis and 21cm signal implementation. Our work thus advocates for clearly specifying both the NCDM and astrophysics implementations when reporting prospective constraints on NCDM models with 21cm Cosmology.

## Acknowledgments

We thank G. Facchinetti and S. Palomares-Ruiz for useful discussions. LLH and QD have been supported by the Fonds de la Recherche Scientifique F.R.S.-FNRS through a research associate position and a FRIA PhD grant, and acknowledge support of the FNRS research grant number J.0134.24, the ARC program of the Federation Wallonie-Bruxelles and the IISN convention No. 4.4503.15. Computational resources have been provided by the Consortium des Equipements de Calcul Intensif (CECI), funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under Grant No. 2.5020.11 and by the Walloon Region of Belgium. AD and BZ acknowledge the support from CIDEAGENT/2020/055. This work was supported by Generalitat Valenciana grant PROMETEO/2021/083, and by Spanish Ministerio de Ciencia e Innovacion grant PID2020- 113644GB-I00. We also thank the computer resources at IFIC; in particular the SOM cluster, nowadays funded through the grants EUR2022-134028 and ASFAE/2022/020, as well as Artemisa, funded by the European Union ERDF and Comunitat Valenciana as well as the technical support provided by the Instituto de Física Corpuscular, IFIC (CSIC-UV), Generalitat Valenciana, Project ‘‘ARTEMISA’’, ref. ASFAE/2022/024.

## A Some extra information

In this appendix, we first briefly describe the modifications done to the `21cmFast` code for our analysis and then precise the tabulated values of  $k$  and  $z$  considered in this work.

The main modifications done to the default 21cmFast WDM implementation has been to:

- implement the sharp-k cutoff window function introduced in eq. (2.13) to account for WDM halo mass function as originally proposed in [34, 64].
- cancel/comment the default implementation of WDM which was accounted for as a WDM Jeans mass, computed from some estimate of the WDM free streaming length, and considered as the relevant threshold mass when the WDM Jeans mass was larger than the threshold mass for star formation.

Most of the changes have been done at the level of the `ps.c` source file in `21cmFast`, where we have modified existing functions, and in the header files, to properly take into account the numerical values of the parameters considered here in the Seth-Thormen first crossing distribution and the sharp-k cutoff Mass relation to its radius. The modifications to `ps.c` include adding the sharp-k (eq. (2.13)) filter function and the WDM transfer function as defined in eq. (2.8) within all the relevant functions (mainly `dsigma_dk`, `power_in_k`, `dsigma_sqdm_z0` and `dNdM_st`). Note in particular that we have used the analytic expression of the scale derivative of the root-mean-square variance of density perturbations that give rise to a simple expression of the halo mass function as a function of the linear power spectrum and the scale  $R_{\text{SK}}$  related to a given mass  $M_{\text{SK}}(R_{\text{SK}})$  through the relation given in eq. (2.14) (see [34, 64] and also [36]):

$$\frac{dn}{dM_{\text{SK}}} = \frac{1}{2} \frac{\rho_{\text{m}}}{M_{\text{SK}}^2} f(\sigma) \frac{1}{2\pi^2 \sigma^2} \frac{P(1/R_{\text{SK}})}{R_{\text{SK}}^3} \frac{d \ln R_{\text{SK}}}{d \ln M_{\text{SK}}} . \quad (\text{A.1})$$

where  $P$  denotes the linear power spectrum and  $\sigma$  is the root-mean-square (rms) variance of density perturbations as defined in eq. (2.12).

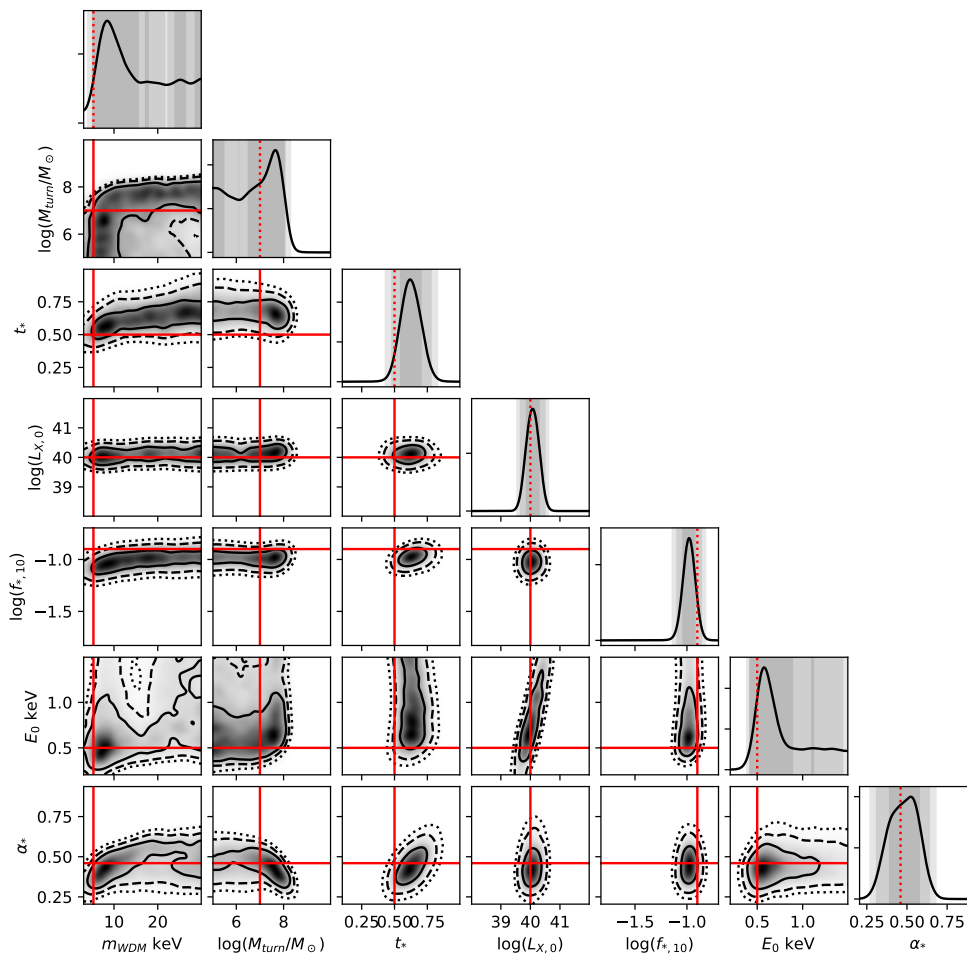
The modified version of `21cmFast` used here can be shared upon request to the authors of this paper. Yet, note that G. Facchinetti, collaborating with one of us (LLH) on a follow-up project, has recently made available a more advanced modified version of `21cmFast` that allows to account for all the above features for non-cold dark matter in a more user friendly way, see [91].<sup>13</sup>

Let us also provide the tabulated values of the considered  $z$  and  $k$  values in our analysis as described in section 3.2, see table 3. We also provide in figure 12 the full set of recovered 1D and 2D marginalized posteriors on our astrophysics and WDM parameters. We have assumed a mock 1000h observation of the 21-cm power spectrum with HERA for  $m_{\text{WDM}} = 5.3$  keV and  $M_{\text{turn}} = 10^7 M_{\odot}$  and all other parameters set to the same values as in figure 2. As further discussed in section 4.2, for this choice of fiducial, we see that except for  $M_{\text{turn}}$ ,  $\alpha_{\star}$ ,  $t_{\star}$  and  $f_{\star,10}$ , no other astrophysics parameter shows significant degeneracies with  $m_{\text{WDM}}$ .

## B Coverage tests

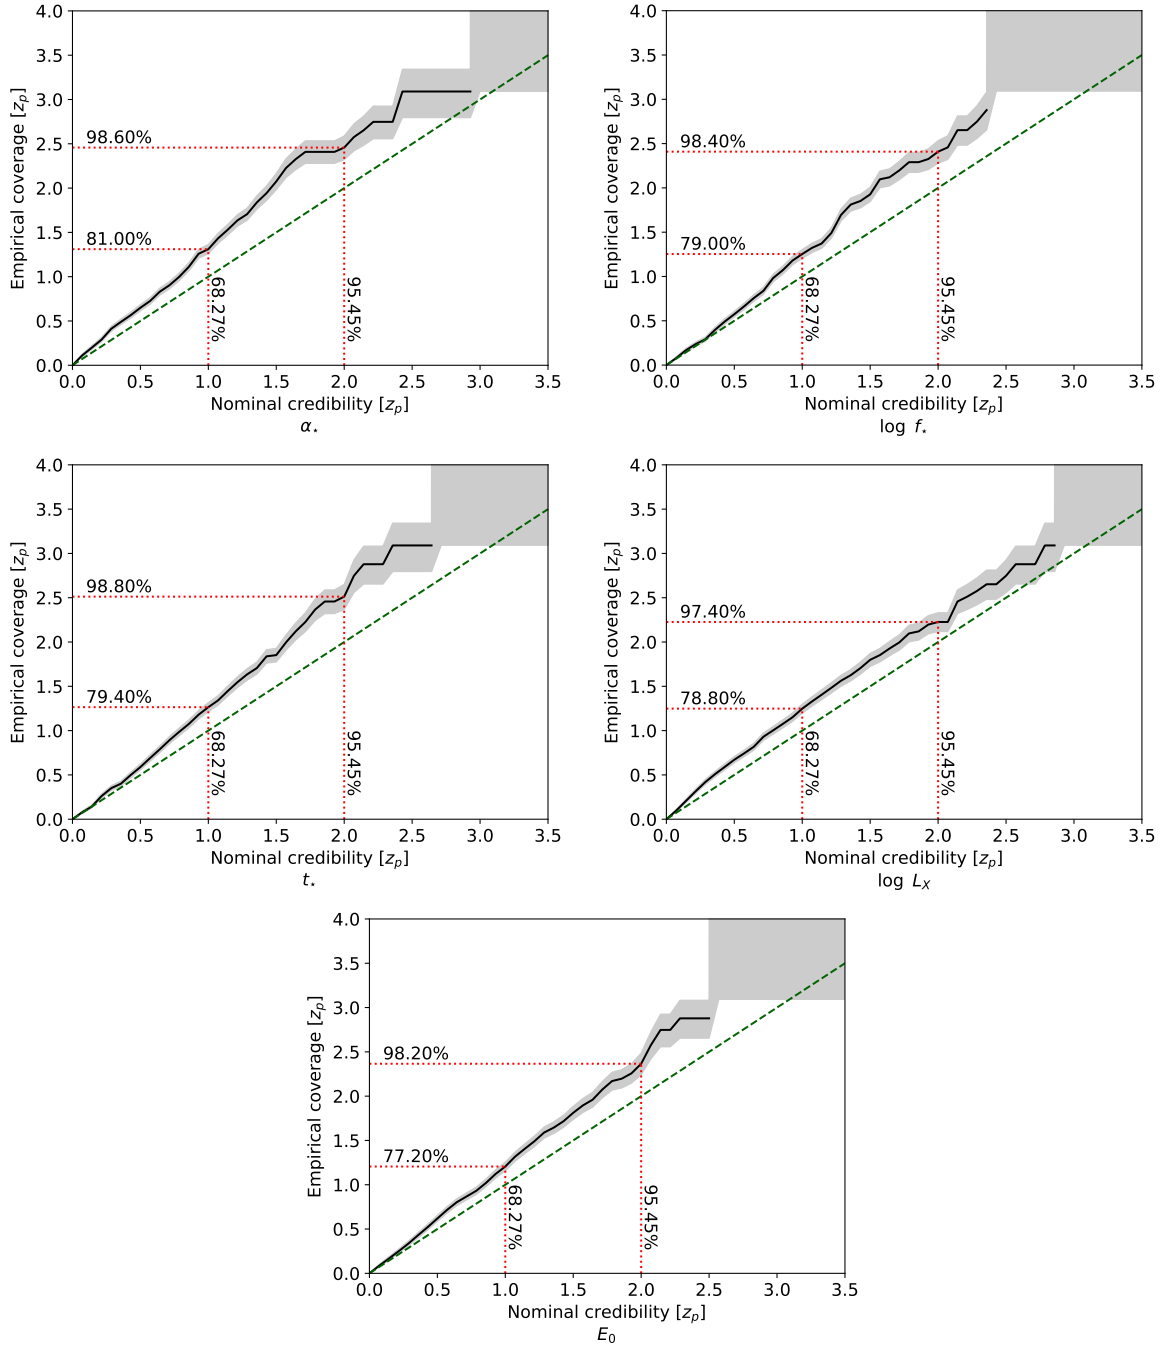
In this appendix, we explore methods to evaluate the quality of the approximate posterior distributions generated by our approach, without relying on MCMC comparison.

<sup>13</sup>The modified version of `21cmFast` developed by G. Facchinetti is available at <https://github.com/gaeta/facchinetti/21cmCLAST/tree/ncdm>.



**Figure 12.** Recovered 1D and 2D marginalized posteriors on all astrophysics and WDM parameters considered here assuming a mock 1000h observation of the 21-cm power spectrum with HERA for  $m_{\text{WDM}} = 5.3 \text{ keV}$  and  $M_{\text{turn}} = 10^7 M_{\odot}$ . The other parameters have been fixed to the same fiducial values as in figure 2. A reduced version of this plot can be found in figure 9.


Assessing the reliability of posterior inference in simulation-based methods is an active field of research [86]. In this work, we focus on one of the most widely used techniques, commonly referred to as coverage tests or “percentile-percentile” plots. The underlying concept is straightforward: for a given parameter, the  $x\%$  credible interval (CI) of the estimated posterior distribution should encompass the true parameter value in approximately  $x\%$  of the cases (empirical coverage, EC), when inference is repeated across multiple simulated datasets. Further discussion of this idea can be found in [86]. A well-calibrated posterior would produce an empirical coverage curve that aligns closely with the diagonal line  $\text{EC}=\text{CI}$  in a EC versus CI plot.



**Figure 13.** Here we follow [92] displaying the empirical expected coverage probability as a function of the confidence level,  $1 - \alpha = 0.6827, 0.9545, 0.9997$ , or equivalently  $z = 1, 2, 3$ , for all the astrophysical parameters (except  $M_{\text{turn}}$ ). The vertical dashed red lines represent the confidence levels, while the horizontal lines depict the empirical estimates. Reproduced from [92]. © 2022 IOP Publishing Ltd and Sissa Medialab. All rights reserved.

Bin	$z_{\min}$	$z_{\max}$	$z_{\text{center}}$	$k_{\min}$	$k_{\max}$	$k_{\text{center}}$
1	5.9	6.2	6.0	0.12	0.18	0.15
2	6.2	6.5	6.3	0.18	0.24	0.21
3	6.5	6.8	6.6	0.24	0.30	0.27
4	6.8	7.2	7.0	0.30	0.36	0.33
5	7.2	7.6	7.4	0.36	0.42	0.39
6	7.6	8.0	7.8	0.42	0.48	0.45
7	8.0	8.5	8.2	0.48	0.54	0.51
8	8.5	9.0	8.7	0.54	0.60	0.57
9	9.0	9.6	9.3	0.60	0.66	0.63
10	9.6	10.3	9.9	0.66	0.72	0.69
11	10.3	11.0	10.6	0.72	0.78	0.75
12	11.0	11.9	11.5	0.78	0.84	0.81
13	11.9	12.9	12.4	0.84	0.90	0.87
14	12.9	14.1	13.5	0.90	0.96	0.93
15	14.1	15.5	14.8	0.96	1.02	0.99
16	15.5	17.2	16.3	—	—	—
17	17.2	19.3	18.2	—	—	—
18	19.3	21.9	20.5	—	—	—
19	21.9	25.3	23.5	—	—	—

**Table 3.** Table of  $z$  and  $k$  (in units of  $[\text{Mpc}^{-1}]$ ) bins edges with their centers.



**Figure 14.** Same as figure 13, but for  $M_{\text{turn}}$  and  $m_{\text{WDM}}$  parameters. Reproduced from [92]. © 2022 IOP Publishing Ltd and Sissa Medialab. All rights reserved.

Figures 13 and 14 present our coverage plots for all parameters considered in our analysis. The gray error bands correspond to the range of the empirical coverage variable, which follows a beta distribution, derived from the probability-point function (PPF). Ideally, the curves in these plots should lie near the diagonal. If the empirical coverage is larger than the nominal credibility (thus, being above the diagonal), the credible intervals are conservative, meaning they capture the true value more often than expected. On the contrary, if the empirical coverage is smaller than the nominal credibility (thus being below the diagonal), the intervals are overconfident and capture the true value less frequently than expected. According to [86], a posterior estimator is deemed acceptable when the empirical coverage probability meets or exceeds the specified confidence level.

The results in figures 13 and 14 indicate that the network has converged with reasonable coverage, as the empirical expected coverage probabilities are close to (but above) the corresponding confidence levels, so obtaining relatively conservative posteriors.

## References

- [1] HERA collaboration, *Improved constraints on the 21 cm EoR power spectrum and the X-ray heating of the IGM with HERA phase I observations*, *Astrophys. J.* **945** (2023) 124 [[arXiv:2210.04912](#)] [[INSPIRE](#)].
- [2] G. Mellema et al., *Reionization and the cosmic dawn with the Square Kilometre Array*, *Exper. Astron.* **36** (2013) 235 [[arXiv:1210.0197](#)] [[INSPIRE](#)].
- [3] L.V.E. Koopmans et al., *The cosmic dawn and epoch of reionization with the Square Kilometre Array*, *PoS AASKA14* (2015) 001 [[arXiv:1505.07568](#)] [[INSPIRE](#)].
- [4] HERA collaboration, *First results from HERA phase I: upper limits on the epoch of reionization 21 cm power spectrum*, *Astrophys. J.* **925** (2022) 221 [[arXiv:2108.02263](#)] [[INSPIRE](#)].
- [5] M. Archidiacono et al., *Constraining dark matter-dark radiation interactions with CMB, BAO, and Lyman- $\alpha$* , *JCAP* **10** (2019) 055 [[arXiv:1907.01496](#)] [[INSPIRE](#)].
- [6] D.C. Hooper et al., *One likelihood to bind them all: Lyman- $\alpha$  constraints on non-standard dark matter*, *JCAP* **10** (2022) 032 [[arXiv:2206.08188](#)] [[INSPIRE](#)].
- [7] A. Schneider, *Constraining noncold dark matter models with the global 21-cm signal*, *Phys. Rev. D* **98** (2018) 063021 [[arXiv:1805.00021](#)] [[INSPIRE](#)].
- [8] M. Sitwell, A. Mesinger, Y.-Z. Ma and K. Sigurdson, *The imprint of warm dark matter on the cosmological 21-cm signal*, *Mon. Not. Roy. Astron. Soc.* **438** (2014) 2664 [[arXiv:1310.0029](#)] [[INSPIRE](#)].
- [9] M. Escudero et al., *A fresh look into the interacting dark matter scenario*, *JCAP* **06** (2018) 007 [[arXiv:1803.08427](#)] [[INSPIRE](#)].
- [10] T. Schaeffer, S.K. Giri and A. Schneider, *BEoRN: a fast and flexible framework to simulate the epoch of reionization and cosmic dawn*, *Mon. Not. Roy. Astron. Soc.* **526** (2023) 2942 [[arXiv:2305.15466](#)] [[INSPIRE](#)].
- [11] M. Safarzadeh, E. Scannapieco and A. Babul, *A limit on the warm dark matter particle mass from the redshifted 21 cm absorption line*, *Astrophys. J. Lett.* **859** (2018) L18 [[arXiv:1803.08039](#)] [[INSPIRE](#)].
- [12] A. Lidz and L. Hui, *Implications of a prereionization 21-cm absorption signal for fuzzy dark matter*, *Phys. Rev. D* **98** (2018) 023011 [[arXiv:1805.01253](#)] [[INSPIRE](#)].

- [13] L. Lopez-Honorez, O. Mena and P. Villanueva-Domingo, *Dark matter microphysics and 21 cm observations*, *Phys. Rev. D* **99** (2019) 023522 [[arXiv:1811.02716](#)] [[INSPIRE](#)].
- [14] J.B. Muñoz et al., *ETHOS — an effective theory of structure formation: impact of dark acoustic oscillations on cosmic dawn*, *Phys. Rev. D* **103** (2021) 043512 [[arXiv:2011.05333](#)] [[INSPIRE](#)].
- [15] D. Jones et al., *Fuzzy dark matter and the 21 cm power spectrum*, *Astrophys. J.* **913** (2021) 7 [[arXiv:2101.07177](#)] [[INSPIRE](#)].
- [16] S.K. Giri and A. Schneider, *Imprints of fermionic and bosonic mixed dark matter on the 21-cm signal at cosmic dawn*, *Phys. Rev. D* **105** (2022) 083011 [[arXiv:2201.02210](#)] [[INSPIRE](#)].
- [17] J.J. Hibbard et al., *Constraining warm dark matter and population III stars with the global 21 cm signal*, *Astrophys. J.* **929** (2022) 151 [[arXiv:2201.02638](#)] [[INSPIRE](#)].
- [18] B. Schosser, C. Heneka and T. Plehn, *Optimal, fast, and robust inference of reionization-era cosmology with the 21cmPIE-INN*, *SciPost Phys. Core* **8** (2025) 037 [[arXiv:2401.04174](#)] [[INSPIRE](#)].
- [19] J. Verwohlt et al., *Separating dark acoustic oscillations from astrophysics at cosmic dawn*, *Phys. Rev. D* **110** (2024) 103533 [[arXiv:2404.17640](#)] [[INSPIRE](#)].
- [20] J. Flitter and E.D. Kovetz, *New tool for 21-cm cosmology. I. Probing  $\Lambda$ CDM and beyond*, *Phys. Rev. D* **109** (2024) 043512 [[arXiv:2309.03942](#)] [[INSPIRE](#)].
- [21] A. Mesinger, S. Furlanetto and R. Cen, *21cmFAST: a fast, semi-numerical simulation of the high-redshift 21-cm signal*, *Mon. Not. Roy. Astron. Soc.* **411** (2011) 955 [[arXiv:1003.3878](#)] [[INSPIRE](#)].
- [22] S.G. Murray et al., *21cmFAST v3: a python-integrated C code for generating 3D realizations of the cosmic 21 cm signal*, *J. Open Source Softw.* **5** (2020) 2582 [[arXiv:2010.15121](#)] [[INSPIRE](#)].
- [23] A. Spurio Mancini, K. Lin and J.D. McEwen, *Field-level cosmological model selection: field-level simulation-based inference for stage IV cosmic shear can distinguish dynamical dark energy*, [arXiv:2410.10616](#) [[INSPIRE](#)].
- [24] F. List, N. Anau Montel and C. Weniger, *Bayesian simulation-based inference for cosmological initial conditions*, in the proceedings of the 37th conference on neural information processing systems, (2023) [[arXiv:2310.19910](#)] [[INSPIRE](#)].
- [25] C. Modi et al., *Sensitivity analysis of simulation-based inference for galaxy clustering*, *Mon. Not. Roy. Astron. Soc.* **536** (2024) 254 [[arXiv:2309.15071](#)] [[INSPIRE](#)].
- [26] A. Dimitriou, D.G. Figueroa and B. Zaldívar, *Fast likelihood-free reconstruction of gravitational wave backgrounds*, *JCAP* **09** (2024) 032 [[arXiv:2309.08430](#)] [[INSPIRE](#)].
- [27] X. Zhao, Y. Mao and B.D. Wandelt, *Implicit likelihood inference of reionization parameters from the 21 cm power spectrum*, *Astrophys. J.* **933** (2022) 236 [[arXiv:2203.15734](#)] [[INSPIRE](#)].
- [28] A. Saxena et al., *Constraining the X-ray heating and reionization using 21-cm power spectra with marginal neural ratio estimation*, *Mon. Not. Roy. Astron. Soc.* **525** (2023) 6097 [[arXiv:2303.07339](#)] [[INSPIRE](#)].
- [29] D. Prelogović and A. Mesinger, *Exploring the likelihood of the 21-cm power spectrum with simulation-based inference*, *Mon. Not. Roy. Astron. Soc.* **524** (2023) 4239 [[arXiv:2305.03074](#)] [[INSPIRE](#)].
- [30] D. Prelogović and A. Mesinger, *How informative are summaries of the cosmic 21 cm signal?*, *Astron. Astrophys.* **688** (2024) A199 [[arXiv:2401.12277](#)] [[INSPIRE](#)].

- [31] B. Greig et al., *Inferring astrophysical parameters using the 2D cylindrical power spectrum from reionization*, *Mon. Not. Roy. Astron. Soc.* **533** (2024) 2530 [[arXiv:2403.14060](#)] [[INSPIRE](#)].
- [32] B. Greig et al., *Exploring the role of the halo-mass function for inferring astrophysical parameters during reionization*, *Mon. Not. Roy. Astron. Soc.* **533** (2024) 2502 [[arXiv:2403.14061](#)] [[INSPIRE](#)].
- [33] R. Barkana, Z. Haiman and J.P. Ostriker, *Constraints on warm dark matter from cosmological reionization*, *Astrophys. J.* **558** (2001) 482 [[astro-ph/0102304](#)] [[INSPIRE](#)].
- [34] A. Schneider, R.E. Smith and D. Reed, *Halo mass function and the free streaming scale*, *Mon. Not. Roy. Astron. Soc.* **433** (2013) 1573 [[arXiv:1303.0839](#)] [[INSPIRE](#)].
- [35] A.J. Benson et al., *Dark matter halo merger histories beyond cold dark matter: I — methods and application to warm dark matter*, *Mon. Not. Roy. Astron. Soc.* **428** (2013) 1774 [[arXiv:1209.3018](#)] [[INSPIRE](#)].
- [36] L. Lopez-Honorez, O. Mena, S. Palomares-Ruiz and P. Villanueva-Domingo, *Warm dark matter and the ionization history of the universe*, *Phys. Rev. D* **96** (2017) 103539 [[arXiv:1703.02302](#)] [[INSPIRE](#)].
- [37] N. Palanque-Delabrouille et al., *Hints, neutrino bounds and WDM constraints from SDSS DR14 Lyman- $\alpha$  and Planck full-survey data*, *JCAP* **04** (2020) 038 [[arXiv:1911.09073](#)] [[INSPIRE](#)].
- [38] R.E. Keeley et al., *JWST lensed quasar dark matter survey — II. Strongest gravitational lensing limit on the dark matter free streaming length to date*, *Mon. Not. Roy. Astron. Soc.* **535** (2024) 1652 [[arXiv:2405.01620](#)] [[INSPIRE](#)].
- [39] C.Y. Tan, A. Dekker and A. Drlica-Wagner, *Mixed warm dark matter constraints using Milky Way satellite galaxy counts*, *Phys. Rev. D* **111** (2025) 063079 [[arXiv:2409.18917](#)] [[INSPIRE](#)].
- [40] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](#)] [[INSPIRE](#)].
- [41] J. Park, A. Mesinger, B. Greig and N. Gillet, *Inferring the astrophysics of reionization and cosmic dawn from galaxy luminosity functions and the 21-cm signal*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 933 [[arXiv:1809.08995](#)] [[INSPIRE](#)].
- [42] J.C. Pober et al., *The baryon acoustic oscillation broadband and broad-beam array: design overview and sensitivity forecasts*, *Astron. J.* **145** (2013) 65 [[arXiv:1210.2413](#)] [[INSPIRE](#)].
- [43] J.C. Pober et al., *What next-generation 21 cm power spectrum measurements can teach us about the epoch of reionization*, *Astrophys. J.* **782** (2014) 66 [[arXiv:1310.7031](#)] [[INSPIRE](#)].
- [44] R. Barkana and A. Loeb, *In the beginning: the first sources of light and the reionization of the universe*, *Phys. Rept.* **349** (2001) 125 [[astro-ph/0010468](#)] [[INSPIRE](#)].
- [45] Y. Qin et al., *Reionization inference from the CMB optical depth and E-mode polarization power spectra*, *Mon. Not. Roy. Astron. Soc.* **499** (2020) 550 [[arXiv:2006.16828](#)] [[INSPIRE](#)].
- [46] S. Furlanetto, S.P. Oh and F. Briggs, *Cosmology at low frequencies: the 21 cm transition and the high-redshift universe*, *Phys. Rept.* **433** (2006) 181 [[astro-ph/0608032](#)] [[INSPIRE](#)].
- [47] M. Stefanon et al., *Galaxy stellar mass functions from  $z \sim 10$  to  $z \sim 6$  using the deepest Spitzer/IRAC data: no significant evolution in the stellar-to-halo mass ratio of galaxies in the first Gyr of cosmic time*, [arXiv:2103.16571](#) [[DOI:10.3847/1538-4357/ac1bb6](#)].

- [48] M. Shuntov et al., *COSMOS2020: cosmic evolution of the stellar-to-halo mass relation for central and satellite galaxies up to  $z \sim 5$* , *Astron. Astrophys.* **664** (2022) A61 [[arXiv:2203.10895](#)] [[INSPIRE](#)].
- [49] Y. Qin et al., *A tale of two sites — I. Inferring the properties of minihalo-hosted galaxies from current observations*, *Mon. Not. Roy. Astron. Soc.* **495** (2020) 123 [[arXiv:2003.04442](#)] [[INSPIRE](#)].
- [50] R.K. Sheth, H.J. Mo and G. Tormen, *Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes*, *Mon. Not. Roy. Astron. Soc.* **323** (2001) 1 [[astro-ph/9907024](#)] [[INSPIRE](#)].
- [51] H. Xu, M.L. Norman, B.W. O’Shea and J.H. Wise, *Late pop III star formation during the epoch of reionization: results from the renaissance simulations*, *The Astrophysical Journal* **823** (2016) 140 [[arXiv:1604.03586](#)].
- [52] E. Sobacchi and A. Mesinger, *The depletion of gas in high-redshift dwarf galaxies from an inhomogeneous reionization*, *Mon. Not. Roy. Astron. Soc.* **432** (2013) 51 [[arXiv:1301.6776](#)] [[INSPIRE](#)].
- [53] E. Sobacchi and A. Mesinger, *How does radiative feedback from a UV background impact reionization?*, *Mon. Not. Roy. Astron. Soc.* **432** (2013) 3340 [[arXiv:1301.6781](#)] [[INSPIRE](#)].
- [54] B.W. O’Shea, J.H. Wise, H. Xu and M.L. Norman, *Probing the ultraviolet luminosity function of the earliest galaxies with the renaissance simulations*, *Astrophys. J.* **807** (2015) L12 [[arXiv:1503.01110](#)].
- [55] R. Barkana and A. Loeb, *In the beginning: the first sources of light and the reionization of the universe*, *Phys. Rept.* **349** (2001) 125 [[astro-ph/0010468](#)] [[INSPIRE](#)].
- [56] E. Sobacchi and A. Mesinger, *How does radiative feedback from a UV background impact reionization?*, *Mon. Not. Roy. Astron. Soc.* **432** (2013) 3340 [[arXiv:1301.6781](#)] [[INSPIRE](#)].
- [57] J.B. Muñoz et al., *The impact of the first galaxies on cosmic dawn and reionization*, *Mon. Not. Roy. Astron. Soc.* **511** (2022) 3657 [[arXiv:2110.13919](#)] [[INSPIRE](#)].
- [58] HERA collaboration, *HERA phase I limits on the cosmic 21 cm signal: constraints on astrophysics and cosmology during the epoch of reionization*, *Astrophys. J.* **924** (2022) 51 [[arXiv:2108.07282](#)] [[INSPIRE](#)].
- [59] K.J. Bae, A. Kamada, S.P. Liew and K. Yanagi, *Light axinos from freeze-in: production processes, phase space distributions, and Ly- $\alpha$  forest constraints*, *JCAP* **01** (2018) 054 [[arXiv:1707.06418](#)] [[INSPIRE](#)].
- [60] G. Ballesteros, M.A.G. Garcia and M. Pierre, *How warm are non-thermal relics? Lyman- $\alpha$  bounds on out-of-equilibrium dark matter*, *JCAP* **03** (2021) 101 [[arXiv:2011.13458](#)] [[INSPIRE](#)].
- [61] F. D’Eramo and A. Lenoci, *Lower mass bounds on FIMP dark matter produced via freeze-in*, *JCAP* **10** (2021) 045 [[arXiv:2012.01446](#)] [[INSPIRE](#)].
- [62] Q. Decant, J. Heisig, D.C. Hooper and L. Lopez-Honorez, *Lyman- $\alpha$  constraints on freeze-in and superWIMPs*, *JCAP* **03** (2022) 041 [[arXiv:2111.09321](#)] [[INSPIRE](#)].
- [63] M. Viel et al., *Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman- $\alpha$  forest*, *Phys. Rev. D* **71** (2005) 063534 [[astro-ph/0501562](#)] [[INSPIRE](#)].
- [64] P. Bode, J.P. Ostriker and N. Turok, *Halo formation in warm dark matter models*, *Astrophys. J.* **556** (2001) 93 [[astro-ph/0010389](#)] [[INSPIRE](#)].

- [65] A. Jenkins et al., *The mass function of dark matter halos*, *Mon. Not. Roy. Astron. Soc.* **321** (2001) 372 [[astro-ph/0005260](#)] [[INSPIRE](#)].
- [66] R.K. Sheth and G. Tormen, *Large scale bias and the peak background split*, *Mon. Not. Roy. Astron. Soc.* **308** (1999) 119 [[astro-ph/9901122](#)] [[INSPIRE](#)].
- [67] A. Schneider, *Structure formation with suppressed small-scale perturbations*, *Mon. Not. Roy. Astron. Soc.* **451** (2015) 3117 [[arXiv:1412.2133](#)] [[INSPIRE](#)].
- [68] F. Pacucci, A. Mesinger, S. Mineo and A. Ferrara, *The X-ray spectra of the first galaxies: 21 cm signatures*, *Mon. Not. Roy. Astron. Soc.* **443** (2014) 678 [[arXiv:1403.6125](#)] [[INSPIRE](#)].
- [69] B. Greig and A. Mesinger, *Simultaneously constraining the astrophysics of reionization and the epoch of heating with 21CMMC*, *Mon. Not. Roy. Astron. Soc.* **472** (2017) 2651 [[arXiv:1705.03471](#)] [[INSPIRE](#)].
- [70] J. Mirocha, S.R. Furlanetto and G. Sun, *The global 21-cm signal in the context of the high- $z$  galaxy luminosity function*, *Mon. Not. Roy. Astron. Soc.* **464** (2016) 1365.
- [71] R.J. Bouwens et al., *New determinations of the UV luminosity functions from  $z \sim 9$  to 2 show a remarkable consistency with halo growth and a constant star formation efficiency*, *Astron. J.* **162** (2021) 47 [[arXiv:2102.07775](#)].
- [72] N. Sabti, J.B. Muñoz and D. Blas, *First constraints on small-scale non-Gaussianity from UV galaxy luminosity functions*, *JCAP* **01** (2021) 010 [[arXiv:2009.01245](#)] [[INSPIRE](#)].
- [73] D.R. DeBoer et al., *Hydrogen Epoch of Reionization Array (HERA)*, *Publ. Astron. Soc. Pac.* **129** (2017) 045001 [[arXiv:1606.07473](#)] [[INSPIRE](#)].
- [74] D.B. Rubin, *Bayesianly justifiable and relevant frequency calculations for the applied statistician*, *Annals Statist.* **12** (1984) 1151.
- [75] SIMBIG collaboration, *Field-level simulation-based inference of galaxy clustering with convolutional neural networks*, *Phys. Rev. D* **109** (2024) 083536 [[arXiv:2310.15256](#)] [[INSPIRE](#)].
- [76] U. Bhardwaj et al., *Sequential simulation-based inference for gravitational wave signals*, *Phys. Rev. D* **108** (2023) 042004 [[arXiv:2304.02035](#)] [[INSPIRE](#)].
- [77] J. Alvey et al., *Simulation-based inference for stochastic gravitational wave background data analysis*, *Phys. Rev. D* **109** (2024) 083008 [[arXiv:2309.07954](#)] [[INSPIRE](#)].
- [78] A. Saxena et al., *Simulation-based inference of the sky-averaged 21-cm signal from CD-EoR with REACH*, *RAS Tech. Instrum.* **3** (2024) 724 [[arXiv:2403.14618](#)] [[INSPIRE](#)].
- [79] C. Pacilio, S. Bhagwat and R. Cotesta, *Simulation-based inference of black hole ringdowns in the time domain*, *Phys. Rev. D* **110** (2024) 083010 [[arXiv:2404.11373](#)] [[INSPIRE](#)].
- [80] Y. Zheng et al., *Angular power spectrum of gravitational-wave transient sources as a probe of the large-scale structure*, *Phys. Rev. Lett.* **131** (2023) 171403 [[arXiv:2305.02652](#)] [[INSPIRE](#)].
- [81] F. Stachurski, C. Messenger and M. Hendry, *Cosmological inference using gravitational waves and normalizing flows*, *Phys. Rev. D* **109** (2024) 123547 [[arXiv:2310.13405](#)] [[INSPIRE](#)].
- [82] C.H. Hahn et al., *Cosmological constraints from non-Gaussian and nonlinear galaxy clustering using the SimBIG inference framework*, *Nature Astron.* **8** (2024) 1457 [[arXiv:2310.15246](#)] [[INSPIRE](#)].
- [83] B.K. Miller, A. Cole, G. Louppe and C. Weniger, *Simulation-efficient marginal posterior estimation with swyft: stop wasting your precious time*, [arXiv:2011.13951](#) [[INSPIRE](#)].

- [84] A. Wehenkel et al., *Addressing misspecification in simulation-based inference through data-driven calibration*, [arXiv:2405.08719](#).
- [85] B.K. Miller et al., *swyft: truncated marginal neural ratio estimation in python*, *J. Open Source Softw.* **7** (2022) 4205 [[INSPIRE](#)].
- [86] J. Hermans et al., *A trust crisis in simulation-based inference? Your posterior approximations can be unfaithful*, [arXiv:2110.06581](#) [[INSPIRE](#)].
- [87] D.P. Kingma and J. Ba, *Adam: a method for stochastic optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [88] B.K. Miller et al., *Truncated marginal neural ratio estimation*, in the proceedings of the *35th conference on neural information processing systems*, (2021) [[DOI:10.5281/zenodo.5043706](#)] [[arXiv:2107.01214](#)] [[INSPIRE](#)].
- [89] A. Mesinger, A. Ferrara and D.S. Spiegel, *Signatures of X-rays in the early universe*, *Mon. Not. Roy. Astron. Soc.* **431** (2013) 621 [[arXiv:1210.7319](#)] [[INSPIRE](#)].
- [90] M.R. Mosbech, C. Boehm and Y.Y.Y. Wong, *Probing dark matter interactions with 21 cm observations*, *JCAP* **03** (2023) 047 [[arXiv:2207.03107](#)] [[INSPIRE](#)].
- [91] G. Facchinetti, *Neural network emulation of reionization to constrain new physics with early- and late-time probes*, [arXiv:2503.11261](#) [[INSPIRE](#)].
- [92] A. Cole et al., *Fast and credible likelihood-free cosmology with truncated marginal neural ratio estimation*, *JCAP* **09** (2022) 004 [[arXiv:2111.08030](#)] [[INSPIRE](#)].