21cmfish: Fisher-matrix framework for fast parameter forecasts from the cosmic 21-cm signal

Charlotte A. Mason[1,2]⋆, Julian B. Muñoz[3,4], Bradley Greig[5,6], Andrei Mesinger[7], and Jaehong Park[8]

[1]*Cosmic Dawn Center (DAWN)*
[2]*Niels Bohr Institute, University of Copenhagen, Jagtvej 128, 2200 København N, Denmark*
[3]*Center for Astrophysics | Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA*
[4]*The University of Texas at Austin, Department of Astronomy, Austin, TX 78712*
[5]*School of Physics, University of Melbourne, Parkville, VIC 3010, Australia*
[6]*ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D)*
[7]*Scuola Normale Superiore, 56126 Pisa, PI, Italy*
[8]*School of Physics, Korea Institute for Advanced Study (KIAS), 85 Hoegiro, Dongdaemun-gu, Seoul 02455, Republic of Korea*

ABSTRACT

The 21-cm signal from neutral hydrogen in the early universe will provide unprecedented information about the first stars and galaxies. Extracting this information, however, requires accounting for many unknown astrophysical processes. Semi-numerical simulations are key for exploring the vast parameter space of said processes. These simulations use approximate techniques such as excursion-set and perturbation theory to model the 3D evolution of the intergalactic medium, at a fraction of the computational cost of hydrodynamic and/or radiative transfer simulations. However, exploring the enormous parameter space of the first galaxies can still be computationally expensive. Here we introduce 21cmfish, a Fisher-matrix wrapper for the semi-numerical simulation 21cmFAST. 21cmfish facilitates efficient parameter forecasts, scaling to significantly higher dimensionalities than MCMC approaches, assuming a multivariate Gaussian posterior. Our method produces comparable parameter uncertainty forecasts to previous MCMC analyses but requires $\sim 10^4\times$ fewer simulations. This enables a rapid way to prototype analyses adding new physics and/or additional parameters. We carry out a forecast for HERA using the largest astrophysical parameter space to-date, with 10 free parameters, spanning both population II and III star formation. We find X-ray parameters for the first galaxies could be measured to sub-percent precision, and, though they are highly degenerate, the stellar-to-halo mass relation and ionizing photon escape fraction for population II and III galaxies can be constrained to $\sim 10\%$ precision (logarithmic quantities). Using a principal component analysis we find HERA is most sensitive to the product of the ionizing escape fraction and the stellar-to-halo mass fraction for population II galaxies.

Key words: dark ages, reionisation, first stars – galaxies: high-redshift

1 INTRODUCTION

The cosmic 21-cm signal will soon open a new window on the early universe. The signal is a net absorption or emission of 21-cm photons relative to the cosmic microwave background (CMB) by neutral hydrogen gas in the intergalactic medium (IGM) that depends sensitively on the formation redshift and properties of the universe's first stars and galaxies, see e.g., reviews by Furlanetto (2006); Morales & Wyithe (2010); Pritchard & Loeb (2012); Mesinger (2019). When the first stars form ($z \sim 30$), their ultra-violet (UV) radiation couples the spin temperature of neutral hydrogen to the cooler kinetic

temperature of the gas, via the Wouthuysen-Field effect (e.g., Hirata 2006), driving a net absorption of CMB photons at 21-cm wavelengths. Later, as the galaxy population builds up, X-ray emission from the first galaxies is expected to heat the IGM, driving the 21-cm signal into emission ($z \lesssim 15$), and eventually, as hydrogen reionizes, the signal decays to zero ($z \lesssim 10$). As the formation redshift and abundance of the first stars is highly sensitive to the build up of low mass dark-matter halos ($M_h \sim 10^{6-8} M_\odot$) and astrophysical feedback effects which govern star formation, the redshift evolution of the 21-cm signal can be used to constrain astrophysical and cosmological properties in the early universe. This will be an exciting way to see beyond the limits of optical and near-infrared telescopes.

In the coming decade, a number of experiments will reach the

expected sensitivity required to measure the 21-cm power spectrum at $z > 6$, including e.g. the Hydrogen Epoch of Reionization Array (HERA, DeBoer et al. 2017; The HERA Collaboration et al. 2022c) and the Square Kilometre Array (SKA, Koopmans et al. 2015), while current instruments have set strong upper limits: the Murchison Widefield Array (MWA Tingay et al. 2013; Trott et al. 2020) and the Low Frequency Array (LOFAR, van Haarlem et al. 2013; Gehlot et al. 2019; Mertens et al. 2020), and a claimed first measurement of the 21-cm global signal has been made at $z \sim 18$ (Bowman et al. 2018 though c.f. Hills et al. 2018; Singh et al. 2022). To interpret these upcoming data, it is imperative to compare observations to models which encompass as much of the unknown physics of the early universe as possible. In recent years there has been a particular effort to create parametric models which describe the key astrophysical and cosmological mechanisms which could affect the 21-cm signal (e.g., 21cmFAST[1]: Mesinger & Furlanetto 2007; Mesinger et al. 2011; Murray et al. 2020, zreion: Battaglia et al. 2013; ares: Mirocha 2014; GRIZZLY: Ghara et al. 2015; SCRIPT: Choudhury & Paranjape 2018; ASTRAEUS: Hutter et al. 2021; DRAGONS: Mutch et al. 2016; SimFAST21: Santos et al. 2010; Visbal & Loeb 2010).

Unknown physics in the early universe determines the formation of the first stars and galaxies and the heating and ionization of the IGM which in turn sets the amplitude and spatial structure of the 21-cm signal. For example: the strength of star formation feedback in low mass halos; Lyman-Werner feedback; the impact of streaming velocities between baryons and dark matter; the typical X-ray emission of early galaxies; and the escape fraction of hydrogen ionizing photons from galaxies all contribute to the strength of the 21-cm signal via their impact on the formation redshifts and typical bias of the first galaxies, the level of inhomogeneous heating of the IGM during cosmic dawn, and the rate at which hydrogen is reionized (e.g., Muñoz et al. 2022). The degeneracies between some of these effects have been explored (e.g., Park et al. 2019; Qin et al. 2020, 2021a) but due to computational inefficiencies there has not been a thorough investigation of the degeneracies between effects that govern population II and population III star formation. It is furthermore unclear how much the astrophysical effects will hamper our efforts to understand the underlying cosmology: for example, Sitwell et al. (2014); Jones et al. (2021); Muñoz et al. (2020) have explored the impact of warm, fuzzy, and self-interacting (ETHOS) dark-matter models on the 21-cm signal, respectively, but did not carry out a full exploration of the degeneracies with astrophysical parameters. Only a few key observables, like the shape of the velocity-induced acoustic oscillations (VAOs, Muñoz 2019b) are immune to astrophysical uncertainties.

Once 21-cm detections are available, one would ideally map out the multi-dimensional posterior of astrophysics and cosmology using Bayesian inference techniques such as MCMC or Nested sampling (e.g. 21CMMC[2] Greig & Mesinger 2015, 2017). However, mapping out the full posterior is extremely expensive in the high dimensional parameter space of generalized galaxy and cosmological models. Before having a detection, however, decisions on observational strategies, telescope design, synergies, etc. can be guided by much faster (albeit more approximate) forecasting techniques, such as Fisher-matrix analyses[3].

In this paper, we use a Fisher-matrix analysis to explore an unprecedentedly large 21cmFAST astrophysical parameter space. We demonstrate that the Fisher-matrix analysis produces comparable parameter uncertainties to an MCMC with 21CMMC (which was restricted to population II galaxy parameters only), recovering the same degeneracies and very comparable error-bars. Having validated our approach we carry out a Fisher-matrix analysis for the largest parameter set to-date in 21cmFAST: 10 free parameters, including both population II and population III galaxies over $z \sim 5-30$. Our method reduces the computational cost significantly compared to a typical 21CMMC run: requiring a factor $\sim 10^4$ fewer individual simulations. We carry out a principal component analysis of our Fisher matrix to determine which combinations of parameters will be most easily constrained by observations. We provide a public release of our code 21cmfish[4], which is a python wrapper for 21cmFAST to perform Fisher-matrix analyses.

This approach will enable us to easily add new parameters, e.g. cosmological and dark-matter parameters, to quickly assess how well upcoming 21-cm experiments will be able to constrain them, which will be the topic of a future work (Mason et al. in prep). The Fisher matrix could also be used for prototyping instrument designs and instrumental noise estimates.

This paper is structured as follows. Section 2 describes our method for generating the 21-cm signal with 21cmFAST, the astrophysical model and our Fisher-matrix formalism in 21cmfish. Section 3 describes our results of the comparison between the Fisher-matrix analysis and MCMC, and our exploration of the extended model parameter space. We discuss our results in Section 4 and summarize our conclusions in Section 5.

In this work we fix the cosmological parameters to the best fit from Planck 2018 data (TT,TE,EE+lowE+lensing+BAO from Planck Collaboration et al. 2020), and all distances are comoving unless specified otherwise.

2   METHODS

Here we provide a brief overview of the astrophysical model (Section 2.1) and the 21-cm signal (Section 2.2), our Fisher-matrix analysis approach (Section 2.3), and the generation of simulated observations (Section 2.4). The parametric 21cmFAST astrophysical model we use is described by Park et al. (2019); Qin et al. (2020, 2021a,b); Muñoz et al. (2022) and we refer readers there for more details.

2.1   The astrophysical model

We use the public code 21cmFAST (Mesinger & Furlanetto 2007; Mesinger et al. 2011; Murray et al. 2020) to model the 21-cm signal. We use version 3.1.0[5] which includes a parametric model for galaxy formation in both atomic cooling galaxies and molecular cooling galaxies, accounting for all major feedback mechanisms on these galaxies, as well as detailed IGM physics such as inhomogeneous recombinations. The model is fully described by Park et al. (2019);

point. These are reasonable approximations, given that we already have a decent estimate of a "fiducial" model using current data, and previous inference results suggest that the posterior is likely only weakly multi-modal/non-Gaussian (e.g. Qin et al. 2021a,b; The HERA Collaboration et al. 2022b).

---

[1] https://github.com/21cmfast/21cmFAST
[2] https://github.com/21cmfast/21CMMC
[3] Fisher-matrix analyses assume a "fiducial" (i.e. maximum likelihood) parameter set, and computes a multi-variate Gaussian posterior around that

[4] https://github.com/charlottenosam/21cmfish
[5] https://github.com/21cmfast/21cmFAST/releases/tag/3.1.0

Qin et al. (2020) and Muñoz et al. (2022) and we refer the reader there for more details. Park et al. (2019) describes the halo-mass dependent star formation prescription for population II galaxies, Qin et al. (2020) added a prescription for population III galaxies, and Muñoz et al. (2022) self-consistently included feedback effects on population III star formation – Lyman Werner feedback and relative velocities between dark matter and baryons. Here we provide a brief overview of the astrophysical model.

As halos form hierarchically, we assume that the first generation of stars (population III) form out of pristine gas in molecular cooling galaxies, where gas can cool via transitions in H_2 ($T_{vir} \gtrsim 10^3$ K, $M_h \gtrsim 10^6 M_\odot$, e.g., Barkana & Loeb 2001), whereas later generations of stars (population II) form in more massive atomic cooling galaxies, where gas cools through atomic transitions ($T_{vir} \gtrsim 10^4$ K, corresponding to $M_h \gtrsim 10^8 M_\odot$), as halos grow and the galaxies' ISM is enriched. `21cmFAST` models these two galaxy populations separately, as described below. Throughout this work we use PopII or superscript "II" to refer to atomic cooling galaxies with population II stars, and PopIII or superscript "III" to refer to molecular cooling galaxies with population III stars.

The number densities of these galaxies are given by:

$$\frac{dn_g}{dM_h} = \frac{dn}{dM_h} \times f_{duty}(M_h) \tag{1}$$

where $\frac{dn}{dM_h}$ is the halo mass function[6], and f_{duty} accounts for the suppression of galaxy formation in halos below a certain mass scale due to a combination of cooling and feedback processes:

$$f_{duty} = \begin{cases} \exp\left(-\frac{M_{turn}^{II}}{M_h}\right) & \text{PopII} \\ \exp\left(-\frac{M_{turn}^{III}}{M_h}\right)\exp\left(-\frac{M_h}{M_{atom}}\right) & \text{PopIII} \end{cases} \tag{2}$$

Here, the second term for population III galaxies produces a smooth transition between the two galaxy populations. The characteristic mass scales where galaxy formation is suppressed, $M_{turn}^{II}, M_{turn}^{III}$, may be set as a free parameter to be inferred (e.g., Park et al. 2019) or modelled physically as a function of redshift and local radiative background strength (e.g., Muñoz et al. 2022).

For population II atomic cooling galaxies the suppression is determined by atomic cooling and photoionization feedback, and has a characteristic mass scale:

$$M_{turn}^{II} = \max\left[M_{atom}, M_{ion}\right] \tag{3}$$

where M_{atom} is the minimum mass scale of atomic cooling (corresponding to a virial temperature of $\sim 10^4$ K, e.g. Barkana & Loeb 2001; Oh & Haiman 2002), M_{ion} is the mass scale where photoheating feedback from inhomogeneous reionization becomes important (e.g., Sobacchi & Mesinger 2014).

For molecular cooling population III galaxies, in addition to the photoionization feedback, we follow Muñoz et al. (2022) (see also Fialkov et al. 2013; Muñoz 2019a) and include the two dominant sources of feedback that reduce the efficiency of star formation in molecular cooling halos: (i) Lyman-Werner feedback, due to stellar emission at $11.2 - 13.6$ eV which effectively photo-dissociates molecular hydrogen, stalling the cooling of gas in population III galaxies and thus impeding star formation; and (ii) dark matter - baryon streaming velocities which reduce the accretion of gas into dark-matter halos, which both slows the formation of halos, due to

the lower gravitational potential, and reduces the gas available to form stars.

The strength of these two effects depends on the Lyman-Werner background flux (J_{21}, in units 10^{-21} erg s^{-1} cm^{-2} Hz^{-1} sr^{-1}) and the relative velocity between dark matter and baryons (v_{cb} in km s^{-1}) as a function of redshift. Therefore,

$$M_{turn}^{III} = \max\left[M_{mol}, M_{ion}\right] \tag{4}$$

where we use the parameterisation of Muñoz et al. (2022):

$$M_{mol}(z, v_{cb}, J_{21}) = \frac{M_{mol,crit}}{(1+z)^{3/2}} f_{v_{cb}}(v_{cb}) f_{LW}(J_{21}) \tag{5}$$

where $M_{mol,crit} = 3.3 \times 10^7 M_\odot$ is the (no feedback) molecular cooling threshold, corresponding to $T_{vir} \approx 10^3$ K (e.g., Barkana & Loeb 2001), and f_{LW} and $f_{v_{cb}}$ are the strength of Lyman-Werner and relative-velocity feedback respectively.

The strength of the Lyman-Werner feedback is parameterised as (Machacek et al. 2001; Visbal et al. 2014):

$$f_{LW}(J_{21}) = 1 + A_{LW} \times (J_{21})^{\beta_{LW}} \tag{6}$$

where A_{LW} and β_{LW} are free parameters. In the following we set $\beta_{LW} = 0.6$, as simulations predict a small range for this parameter (e.g., Kulkarni et al. 2021; Schauer et al. 2021; Skinner & Wise 2020) and leave A_{LW} free, and study whether it can be determined from observations.

The strength of the relative-velocity feedback is modelled as:

$$f_{v_{cb}}(v_{cb}) = \left(1 + A_{v_{cb}} \frac{v_{cb}}{\sigma_{cb}}\right)^{\beta_{v_{cb}}} \tag{7}$$

where $\sigma_{cb} \approx 30$ km s^{-1} is the rms velocity (Tseliakhovich & Hirata 2010). We follow Muñoz et al. (2022) and use $A_{v_{cb}} = 1$ and $\beta_{v_{cb}} = 1.8$, as these values have been recovered consistently in independent simulations (Kulkarni et al. 2021; Schauer et al. 2021).

The stellar mass ($M_\star$) to halo mass ratio is described by a power law (e.g., Moster et al. 2013; Behroozi & Silk 2015; Sun & Furlanetto 2016; Mutch et al. 2016; Tacchella et al. 2018; Yung et al. 2020; Sabti et al. 2022b,a):

$$\frac{M_\star}{M_h} = \frac{\Omega_b}{\Omega_m} \times \min\left[1, \begin{cases} f_{\star,10}^{II}\left(\frac{M_h}{10^{10}M_\odot}\right)^{\alpha_\star^{II}} & \text{PopII} \\ f_{\star,7}^{III}\left(\frac{M_h}{10^7 M_\odot}\right)^{\alpha_\star^{III}} & \text{PopIII} \end{cases}\right] \tag{8}$$

with four free parameters ($f_{\star,10}^{II}, \alpha_\star^{II}, f_{\star,7}^{III}, \alpha_\star^{III}$) which set the normalization at the pivot mass and low mass slope for galaxies forming population II and population III stars respectively. We note that although the stellar mass – halo mass relation turns over at high masses ($M_h \sim 10^{12} M_\odot$, e.g., Behroozi et al. 2013; Mason et al. 2015), the radiation fields at these high redshifts are determined by the vast majority of halos which have much lower masses; thus a single power-law slope suffices for our purposes. This power-law slope encapsulates additional mass-dependent feedback processes, such as supernovae feedback.

We assume the star formation rate in halos is related to their dynamical time, which during matter domination scales as H^{-1}, to give us

$$\frac{dM_\star}{dt} = \frac{M_\star}{t_\star H(z)^{-1}} \tag{9}$$

where $t_\star$ can be a free parameter, but as SFR has the ratio $f_\star/t_\star$ there is a strong degeneracy between these parameters (Park et al. 2019), and $t_\star$ can typically be set to a constant.

---

[6] Halo mass functions are created using the Extended Press-Schechter formalism (e.g., Sheth et al. 2001) using a top-hat window function.

Galaxies hosting population II and population III stars can also have different hydrogen ionizing photon escape fractions:

$$f_{\rm esc} = \begin{cases} f_{\rm esc,10}^{\rm II} \left(\frac{M_h}{10^{10}M_\odot}\right)^{\alpha_{\rm esc}} & \text{PopII} \\ f_{\rm esc,7}^{\rm III} \left(\frac{M_h}{10^{7}M_\odot}\right)^{\alpha_{\rm esc}} & \text{PopIII} \end{cases} \tag{10}$$

where $f_{\rm esc,10}^{\rm II}$ and $f_{\rm esc,7}^{\rm III}$ set the normalization of the escape fractions and $\alpha_{\rm esc}$ their power-law slopes. Note that for simplicity, we take $\alpha_{\rm esc}$ to be the same for each population of galaxies as we assume it is a function of halo mass only.

These early populations of galaxies are also expected to produce X-rays, most likely from high mass X-ray binaries (e.g. Furlanetto 2006; Mineo et al. 2011; McQuinn 2012; Fragos et al. 2013). We assume the specific X-ray luminosity is proportional to the galaxy SFR and the X-ray SED is a power law with energy spectral index $= -1$, such that the specific X-ray luminosity per unit SFR is:

$$\frac{{\rm d}L_X/\dot{M}_\star}{{\rm d}E} = \frac{E^{-1}}{\int_{E_0}^{2\,\rm keV} {\rm d}E\, E^{-1}} \times \begin{cases} L_{X,<2\rm keV}^{\rm II}/\dot{M}_\star & \text{PopII} \\ L_{X,<2\rm keV}^{\rm III}/\dot{M}_\star & \text{PopIII} \end{cases} \tag{11}$$

where $L_{X,<2\rm keV}^{\rm II}/\dot{M}_\star$ and $L_{X,<2\rm keV}^{\rm III}/\dot{M}_\star$ are the total soft-band X-ray luminosity per unit SFR in atomic cooling and molecular cooling galaxies respectively. In the following we refer to these as $L_X^{\rm II}/\dot{M}_\star$ and $L_X^{\rm III}/\dot{M}_\star$ and assume they have the same value $= L_X/\dot{M}_\star$, motivated by simulations of high mass X-ray binaries in metal-poor environments in the early universe (Fragos et al. 2013) – though it is possible to vary them independently. E_0 is the cutoff in X-ray energies which can escape galaxies (e.g. X-rays with energy $< E_0$ are absorbed by the ISM of the host galaxies and thus do not interact with the IGM (e.g., Das et al. 2017). X-rays at higher energies ($E > 2$ keV) have a mean free path greater than the Hubble length and thus also do not efficiently heat the IGM.

The astrophysical model described above has a total of 14 free parameters. A summary of these parameters and which ones we vary in our analyses is given in Table 1.

2.2 Modelling the 21-cm signal

Hydrogen atoms have heavily disallowed hyperfine transitions, key amongst them is the 21-cm line between the singlet and triplet states of the ^{1}S orbital. The rate of these 'spin-flip' transitions is determined by the relative population of atoms in either states:

$$\frac{n_1}{n_0} = \frac{g_1}{g_0} e^{-\frac{T_*}{T_S}} \tag{12}$$

where g_i are the spin degeneracy factors of each state (with $g_1/g_0 = 3$ for 21-cm transitions), $T_* = E_{10}/k_B = 68$ mK corresponding to the transition energy E_{10}, and $T_S \gg T_*$ is the *spin temperature*. In the early universe, T_S is coupled to the kinetic temperature of the IGM through collisional excitations (Loeb & Zaldarriaga 2004), but by $z \sim 30$ collisions become rarer and interactions with CMB photons dominate, driving a thermal equilibrium between T_S and $T_{\rm CMB}$, resulting in no net absorption or emission of 21-cm photons relative to the CMB black-body. However, once the first stars form ($z \sim 20 - 30$) this thermal equilibrium is broken due to the injection of Lyα photons (the Wouthuysen–Field effect, e.g., Hirata 2006) and the spin temperature couples to the lower temperature of the IGM gas, resulting in a net absorption of CMB photons at 21-cm wavelengths. Later, as the IGM is heated by X-rays from the first galaxies ($z \sim 10 - 15$), the spin temperature exceeds the CMB temperature resulting in a net emission of 21-cm photons. Once hydrogen is reionized ($z \sim 6 - 10$) the density of neutral hydrogen

drops dramatically and the cosmological 21-cm signal is negligible (see e.g., Furlanetto et al. 2006; Morales & Wyithe 2010; Pritchard & Loeb 2012, for reviews).

The goal of Cosmic Dawn 21-cm experiments is to measure the 21-cm brightness temperature, T_{21}, the offset of the 21-cm signal relative to the temperature of the cosmic microwave background, $T_{\rm CMB}$ (e.g., Furlanetto et al. 2006):

$$\begin{aligned} T_{21} &= \frac{T_S - T_{\rm CMB}}{1+z}(1 - e^{-\tau_{21}}) \tag{13} \\ &\approx 27 x_{\rm HI}(1+\delta)\frac{H(z)}{{\rm d}v_t/{\rm d}r + H(z)}\left(1 - \frac{T_{\rm CMB}}{T_S}\right) \\ &\quad \times \left(\frac{1+z}{10}\frac{0.15}{\Omega_m h^2}\right)^{1/2}\left(\frac{\Omega_b h^2}{0.023}\right) {\rm mK}. \tag{14} \end{aligned}$$

where the second line is the $\tau_{21} \ll 1$ Taylor expansion. Here $x_{\rm HI}$ is the neutral hydrogen fraction, T_S is the spin temperature of the gas, $\delta = \rho/\overline{\rho} - 1$ is the overdensity of the gas, and ${\rm d}v_t/{\rm d}r$ is the gradient of the line-of-sight velocity component. All quantities are computed at $z = \nu_0/\nu - 1$, where $\nu_0 = 1420$ MHz is the 21-cm rest-frame frequency. This Taylor expansion is useful for physical intuition, though we note that 21cmFAST computes the full $\exp(-\tau_{21})$ term.

To create the neutral fraction, spin temperature, density and velocity fields required to find T_{21} we use the public code 21cmFAST (Mesinger et al. 2011; Park et al. 2019; Murray et al. 2020) using the astrophysical model described above in Section 2.1 and we refer the reader to these works and references therein. In Section 2.4 we describe the parameters and setup for our 21cmFAST simulations.

2.3 Fisher-matrix formalism

Given a posterior distribution $P(\theta|{\rm data})$ for model parameters θ, the Fisher-matrix components are given by:

$$\mathcal{F}_{ij} = -\left\langle \frac{\partial^2 \ln P}{\partial\theta_i \partial\theta_j} \right\rangle \tag{15}$$

The Cramér-Rao theorem states that any unbiased estimator for the parameters will produce a covariance matrix that is no more accurate than $\mathcal{F}^{-1}$: thus the Fisher matrix can be used to estimate the minimum uncertainties of parameters given observations (e.g., Albrecht et al. 2009).

Here we focus on the 21-cm power spectrum but the same analysis can be done with the 21-cm global signal (e.g., Liu et al. 2013; Muñoz et al. 2020), or any other observable. Previous works have used Fisher-matrix analyses of simulated 21-cm observations with 21cmFAST to forecast constraints, for example, on reionization and cosmological parameters, and small-scale dark-matter structure (e.g., Pober et al. 2014; Ewall-Wice et al. 2016; Liu & Parsons 2016; Shimabukuro et al. 2017; Muñoz et al. 2020; Jones et al. 2021; Greig et al. 2022). However, these works all used a restricted parameter space, with most using the 3 parameter astrophysical model of 21cmFAST v1. Here, we will use the extended 21cmFAST v3 parameter space (with 14 total degrees of freedom), which characterizes the UV and X-ray emission from both population II and III star formation, as detailed in Section 2.1.

Given a list of parameters θ_i, and assuming the posterior distribution can be described as Gaussian (i.e. $\ln P \propto -\chi^2$), we define the Fisher matrix for the 21-cm power spectrum $\Delta_{21}^2(k,z)$ as

$$\mathcal{F}_{ij} = \frac{\partial^2 \chi^2}{\partial\theta_i \partial\theta_j} = \sum_{i_k, i_z} \frac{\partial\Delta_{21}^2(k,z)}{\partial\theta_i}\frac{\partial\Delta_{21}^2(k,z)}{\partial\theta_j}\frac{1}{\sigma_{\Delta^2}^2(k,z)} \tag{16}$$

where $\sigma^2_{\Delta^2}$ is the measurement error in the power spectrum with k bin i_k and z bin i_z, and we have assumed uncorrelated errors between k and z bins (in principle we can compute the full measurement covariance matrix but we follow the assumption of uncorrelated errors to more easily compare to previously published MCMC results, e.g., Park et al. 2019). Our code `21cmfish` calculates the derivatives with respect to a set of input parameters to produce the Fisher matrix (Equation 16).

The inverse of the Fisher matrix is the covariance matrix $C = \mathcal{F}^{-1}$, and with this definition it is easy to see that the forecasted uncertainty in the i-th parameter is simply $\sigma(\theta_i) = \sqrt{C_{ii}}$.

2.4 Simulated 21-cm observations and uncertainties

As described in Section 2.3 above, we can use the Fisher matrix to derive the lowest possible estimate of the parameter uncertainties given observations. In this work we generate parameter constraints from two sets of mock observations, each serving a specific purpose:

(i) We validate our Fisher-matrix approach by comparing to the MCMC posterior of Park et al. (2019). We use their fiducial parameters, which only includes population II star formation in atomic-cooling galaxies. Simulation boxes have a comoving volume of $(250\,\mathrm{Mpc})^3$ on a 128^3 grid from $z = 5.9 - 28$ following Park et al. (2019) to match the HERA observing bandwidth.

(ii) An exploration of the updated 'Evolution of Structure' (EOS21) simulation presented by Muñoz et al. (2022)[7]. The EOS21 parameters include star formation in minihalos, and the impact of Lyman-Werner feedback and relative velocities between baryons and dark matter. The PopII parameters are based on those inferred by Qin et al. (2021b), who used the Planck Collaboration et al. (2016) electron scattering optical depth, rest-frame UV luminosity functions (LFs, Bouwens et al. 2015), dark pixel fraction (McGreer et al. 2015), and Lyα forest optical depths (Bosman et al. 2018) in their likelihood. On the other hand, the PopIII parameters are difficult to constrain with current observations, and so were chosen fairly arbitrarily. We note PopIII star formation in their fiducial model only dominates the total SFRD at $z \gtrsim 15$. Thus, the EOS21 model produces UV LFs, a reionization history and optical depths that are consistent with current observations at $z \sim 6 - 10$. The simulation boxes we use have a comoving volume of $(400\,\mathrm{Mpc})^3$ on a 200^3 grid from $z = 5 - 30$.

The fiducial parameters for these runs are given in Table 1. 21-cm power spectra are generated by dividing lightcones into 'chunks' following Greig & Mesinger (2018). For our comparisons with Park et al. (2019) we follow their approach and divide the lightcones into 12 chunks of equal comoving volume. For our EOS21 run we divide the lightcone into chunks corresponding to 8 MHz bandwidths, which more closely resembles what would be done with data. The 21-cm global signal, reionization history and power spectra at a few redshifts for the two models are plotted in Figure 1.

To calculate the Fisher matrices we vary each parameter $\pm 3\%$ of its fiducial value and calculate two-sided derivatives[8], finding the variation step-size such that the derivatives were converged. We tested convergence of the derivatives and found $\pm 3\%$ worked in all cases with the exception of $\log_{10} L_X/\dot{M}_\star$ which we varied by

$\pm 0.1\%$ and $\log_{10} M_{\mathrm{turn}}$ which we vary by $\pm 1\%$, as these parameters have narrower likelihoods in the 21-cm signal. We verified that almost one-sided derivatives produced almost identical values.

Thus, to create a derivative $\partial \Delta^2_{21}(k,z)/\partial \theta_i$ for each parameter we use 3 simulated power spectra: the fiducial, one setting the parameter to $\theta_i(1 + X)$ and one with $\theta_i(1 - X)$, where X is the variation (usually 3% as described). As we only need to create one fiducial simulation, for N parameters we create a total of $2N + 1$ simulations to calculate the Fisher matrix (Equation 16).

To make forecasts for upcoming observations with HERA we use the python package `21cmSense`[9] (Pober et al. 2014). We run three different setups, which we now describe. For the comparison with Park et al. (2019) we use the same noise as their study, which is essentially identical to the 'pessimistic' case described below. The noise was calculated assuming 1000 hours of observation using 331 antennae.

For our EOS21 forecasts (with PopIII stars in molecular cooling galaxies) we find the noise in two cases, with moderate and pessimistic foregrounds, respectively. The dominant sources of noise for 21-cm observations are the instrumental noise of the antennae and large foreground contamination, from both our Galaxy and the atmosphere. The combination of these noise sources can be parameterised by a "system temperature", $T_{\mathrm{sys}}(\nu)$ which is the total observed 21-cm brightness temperature (e.g. Morales & Wyithe 2010). The foregrounds substantially contaminate the signal in the Fourier plane, particularly in the so-called "wedge" (e.g., Liu et al. 2009; Pober et al. 2013; Liu & Shaw 2020), within wavenumbers $k_\parallel \leq a + bk_\perp$. Here $k_\parallel$ and $k_\perp$ are the wavenumbers in the line-of-sight and perpendicular direction, and a (called the super-horizon buffer in `21cmSense`) and b are constants that parameterise the severity of the foregrounds and thus which regions of k-space will be discarded in the analysis.

In the moderate foregrounds case we use the 'moderate' foregrounds model in `21cmsense`, with a wedge super-horizon buffer $a = 0.1\,h\,\mathrm{Mpc}^{-1}$ and a HERA system temperature (DeBoer et al. 2017) of

$$T_{\mathrm{sys}}(\nu) = 100\,\mathrm{K} + 120\,\mathrm{K}\,(\nu/150\mathrm{MHz})^{-2.55}\,. \tag{17}$$

In the pessimistic case we increase the wedge horizon buffer to $a = 0.15\,h\,\mathrm{Mpc}^{-1}$, and use the default system temperature in `21cmSense`, which is $\sim 3\times$ larger than that of Equation (17) (Pober et al. 2014). In both cases, the antennae temperature is the same but this larger system temperature is motivated by the sky temperature measured by LOFAR (van Haarlem et al. 2013). In the pessimistic case the signal-to-noise ratio of the power spectrum is lower by nearly an order of magnitude, and consequently increases the error bars on each parameter by approximately a factor of three, as we will see (Figures 3 and 5). This pessimistic case is comparable to that used in the forecasts by Park et al. (2019). In both moderate and pessimistic cases for EOS21 we assume 1 year (1080 hours) of data using 331 antennae across the 50-250 MHz range (corresponding to the redshifts $z \sim 5 - 28$), divided in equal-width bands of 8 MHz.

In all cases, we follow Park et al. (2019) and add Poisson noise (from the finite-size simulations) and a 20% modelling error to the power spectra in quadrature in addition to the Poisson noise for the mock observations. We calculate the likelihood only in the k-space window $0.1 - 1\,\mathrm{Mpc}^{-1}$, corresponding roughly to limits on the foreground noise and the shot noise, respectively. For the Park et al. (2019) comparison we calculate the likelihood over the

---

[7] Simulation data available at: https://www.dropbox.com/sh/dqh9r6wb0s68jfo/AACc9ZCqsNOSQ_JJN7GRVuqDa?dl=0

[8] The variation is in either linear or logarithmic quantities corresponding to the parameters in Table 1, i.e. we vary $\log_{10} f^{\mathrm{II}}_{\star,10} \pm 3\%$ of its fiducial value.

[9] https://github.com/jpober/21cmSense

Table 1. 21cmFAST parameters and their fiducial values for our two runs. For a detailed description see Sec. 2.1.

Parameter	Description	Population	Park et al. (2019) Fiducial	Vary?	EOS21 (Muñoz et al. 2022) Fiducial	Vary?
$\alpha_\star^{\rm II}$	Stellar – halo mass power law slope	II	0.50	✓	0.50	✓
$\log_{10} f_{\star,10}^{\rm II}$	Stellar – halo mass normalization (at $M_h = 10^{10} M_\odot$)	II	-1.30	✓	-1.25	✓
$t_\star$	SFR timescale as a fraction of the Hubble time	both	0.50	✓	0.50	
$\log_{10} M_{\rm turn}$	Halo mass turnover for atomic cooling halos [$M_\odot$]	II	8.7	✓	–	
$\alpha_\star^{\rm III}$	Stellar – halo mass power law slope	III	–		0.0	✓
$f_{\star,7}^{\rm III}$	Stellar – halo mass normalization (at $M_h = 10^7 M_\odot$)	III	–		-2.50	✓
$\alpha_{\rm esc}$	Ionizing escape fraction – halo mass power law slope	both	-0.50	✓	-0.30	✓
$\log_{10} f_{\rm esc,10}^{\rm II}$	Ionizing escape fraction normalization (at $M_h = 10^{10} M_\odot$)	II	-1.00	✓	-1.35	✓
$\log_{10} f_{\rm esc,7}^{\rm II}$	Ionizing escape fraction normalization (at $M_h = 10^7 M_\odot$)	III	–		-1.35	✓
$\log_{10} L_X^{\rm II}/\dot{M}_\star$	X-ray luminosity per SFR [erg s^{-1} M$_\odot^{-1}$ yr]	II	40.5	✓	40.5	✓
$\log_{10} L_X^{\rm III}/\dot{M}_\star$	X-ray luminosity per SFR [erg s^{-1} M$_\odot^{-1}$ yr]	III	–		40.5	$= \log_{10} L_X^{\rm II}$
E_0	minimum X-ray energy which escapes galaxies [eV]	both	500	✓	500	✓
$A_{v_{\rm cb}}$	Amplitude of DM-baryon relative velocity feedback	III	–		1.00	
A_{LW}	Amplitude of Lyman-Werner feedback	III	–		2.00	✓

redshift range $z = 5.9 - 28.0$ to match their analysis, and for EOS21 we calculate the likelihood over $z = 5.0 - 28.0$

3 RESULTS

Section 3.1 describes the validation of our Fisher-matrix approach by comparison to an MCMC by Park et al. (2019). Section 3.2 presents forecasted constraints on the population II and population III galaxy formation parameter space expected with HERA 1 year observations, assuming the EOS21 model from Muñoz et al. (2022). Section 3.3 presents the principal components of the Fisher matrix, providing insight into the best constrained linear combinations of parameters. In Section 3.4 we evaluate whether the Fisher matrix can be used for inference.

3.1 Comparison with Park et al. (2019) population II-only MCMC

To verify our Fisher-matrix approach we compare with the MCMC analysis by Park et al. (2019). Park et al. (2019) carried out an inference on mock HERA observations to infer properties of population II galaxies. The parameters varied in this analysis are listed in Table 1.

To compare with Park et al. (2019) we construct simulations using the same dimensions and redshift range, as described in Section 2.2, and with the same fiducial parameters, though with different initial conditions. We then divide lightcones into 12 chunks of equal comoving volume to create power spectra, replicating their approach, and calculate the derivatives of the power spectrum with respect to the parameters of interest as described in Section 2.3. To create the Fisher matrix (Equation 16) we use the same HERA noise estimate $\sigma[\Delta^2(k,z)]$ as Park et al. (2019), which was obtained using `21cmsense`. To obtain the total uncertainty on the power spectrum, we add in quadrature the HERA thermal noise plus the cosmic variance plus a 20% 'modelling uncertainty' on the power spectrum, again following Park et al. (2019).

The resulting $1, 2\sigma$ error ellipses obtained using our Fisher

matrix are shown in Figure 2, along with the 21cm-only[10] MCMC posteriors by Park et al. (2019). We note that we plot the MCMC chains produced by Park et al. (2019), but do not exactly reproduce their published posteriors and 68% credible intervals due to some small differences in burn-in steps and removal of local minima in the posterior between that work and ours. The chains have 229 steps and 35 walkers. We remove the first 50% of steps as burn-in and also remove $\sim 24\%$ of the chains which we identify as falling into local minima. The median maximum likelihood for the MCMC was $\mathcal{L}_{\rm max} \sim 10^{-4}$ and we remove chains in local minima – where the likelihood in the final step is $< 10^{-10}$.

Our Fisher-matrix approach clearly recovers the inferred degeneracies between parameters in all cases, and produces comparable uncertainties to those estimated using the MCMC. The marginalised 68% credible intervals on each parameter obtained with the Fisher matrix are within 40% of those obtained with the MCMC (with the exception of $L_X/\dot{M}_\star$ where the Fisher error is $\sim 50\%$ of the MCMC value). As noted in Section 2.3, the Fisher matrix will give the minimum uncertainties, so we do expect our technique to slightly underestimate the errors compared to a full mapping of the posterior. Nevertheless, this remarkable agreement validates our Fisher matrix approach as an efficient tool to explore 21-cm observational constraints on cosmology and astrophysics, providing the parameter degeneracies remain approximately Gaussian, dramatically reducing the need for expensive MCMC runs. For comparison, the MCMC run by Park et al. (2019) required 70,000 individual simulations, whereas our Fisher-matrix approach required only 17.

The parameter where the Fisher-matrix constraints are most different from the MCMC is $\alpha_\star^{\rm II}$. This is because Park et al. (2019) set a uniform prior on $\alpha_\star^{\rm II} = [-0.5, 1]$, whereas in the Fisher-matrix approach we do not set any bounding conditions and make the assumption of a Gaussian likelihood, so the confidence ellipse can extend to $\alpha_\star^{\rm II} > 1$, making it broader than the MCMC posterior.

Due to using the 'pessimistic' foreground model assumed (see Section 2.4), these forecasts find most parameters will be measured

---

[10] Park et al. (2019) also shows posteriors obtained by combining existing UV LF observations with the mock 21cm power spectra in the likelihood.
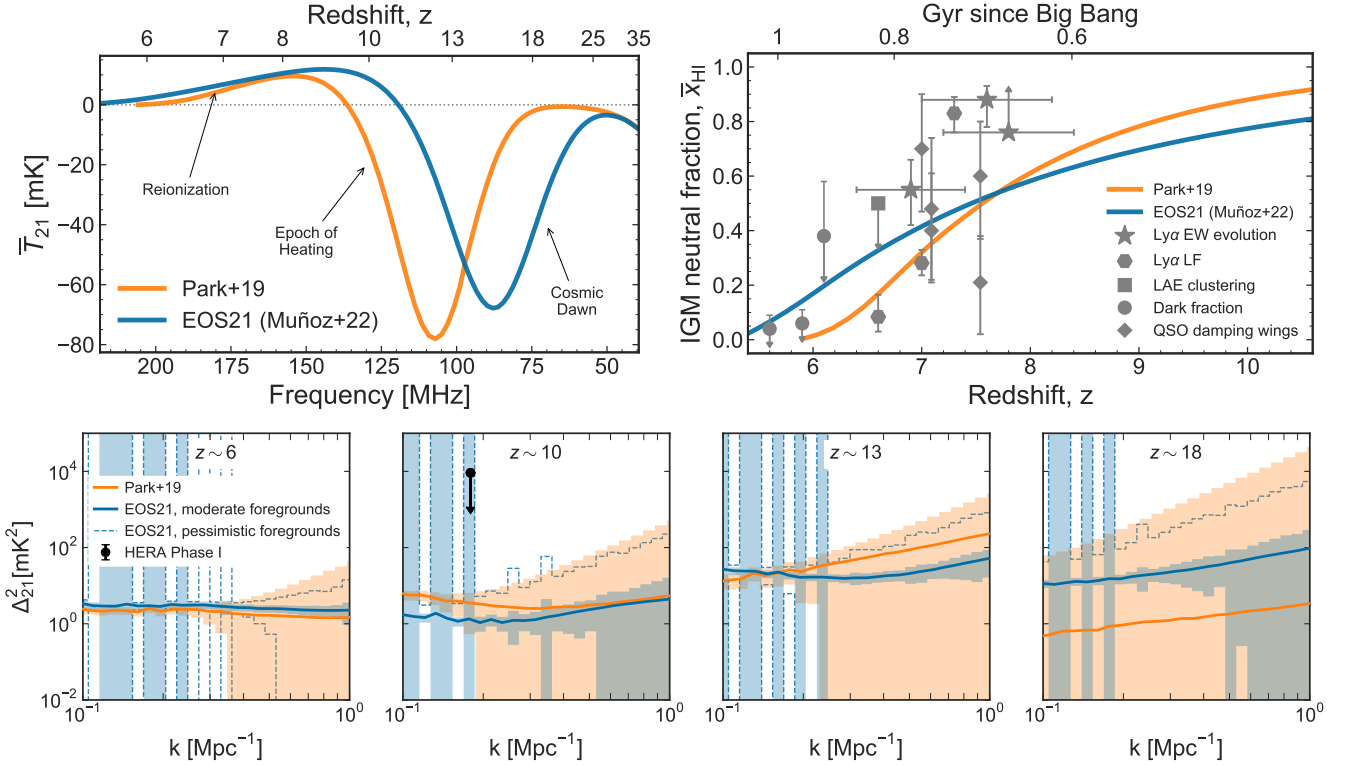
Figure 1. Comparison of our two fiducial models, with Park et al. (2019) in orange, EOS21 in blue. **Top left:** Global 21-cm signal. **Top right:** Reionization history. We also plot reionization history constraints derived from observations of the Lyα equivalent width distribution (star, Mason et al. 2018, 2019; Whitler et al. 2020; Hoag et al. 2019); the Lyα luminosity function (hexagon, Morales et al. 2021) the clustering of Lyα emitting galaxies (square, Ouchi et al. 2010; Sobacchi & Mesinger 2015); Lyα and Lyβ forest dark pixel fraction (circle; McGreer et al. 2015); and QSO damping wings (diamond, Greig et al. 2017; Davies et al. 2018; Greig et al. 2019; Wang et al. 2020). **Lower:** Example power spectra of our two fiducial models at four redshifts. We show our fiducial 'moderate' HERA noise estimates for each case as the filled regions, and the 'pessimistic' noise as dashed blue line, including Poisson noise and a 20% modelling error added in quadrature. We show the errors only in the k-space window where we calculate the likelihood ($0.1 - 1 \, \mathrm{Mpc}^{-1}$). Note that due to a spline interpolation of the power spectra errors by Park et al. (2019), the low k noise estimates are smooth, unlike those in the EOS21 model. For comparison, we also plot the 95% confidence upper limit at $z = 10.4$ from HERA Phase I observations (The HERA Collaboration et al. 2022c).

at the $20 - 30\%$ error level, with the exception of X-ray parameters and M_{turn} which can be measured to $< 10\%$ precision, whereas $\alpha_\star^{\mathrm{II}}$ and α_{esc} are more difficult to measure precisely. As discussed by Park et al. (2019), the inclusion of independent observations that can constrain galaxy formation properties, for example the UV luminosity function, can break degeneracies and reduce parameter uncertainties to $\lesssim 10\%$ precision.

3.2 Extended parameter space forecast

Given that our Fisher-matrix approach recovers the same degeneracies and very similar confidence intervals as an MCMC, we now make a forecast in an unexplored regime: 10 free astrophysical parameters governing both population II and population III galaxy formation, and the impact of streaming velocities and Lyman-Werner feedback on the formation of minihalos, over the full redshift range of the Cosmic Dawn 21-cm signal ($z \sim 5 - 30$). Due to computational expense and current instrumental sensitivity, previous works have performed MCMC or Nested sampling in a more limited parameter space, varying only PopII or PopIII parameters at a time (e.g., Qin et al. 2021a; The HERA Collaboration et al. 2022b, explored 7 and 9 parameters respectively, roughly requiring $\sim 10^5$ individual simulations and ~100k CPUh). Our `21cmfish` calculation, by contrast, required only 21 individual simulations (as we

fixed $L_X^{\mathrm{II}}/\dot{M}_\star = L_X^{\mathrm{III}}/\dot{M}_\star$). Although the Fisher matrix assumes a multi-variate Gaussian posterior and the setting of physical priors is less straightforward, these are reasonable approximations when the posterior is uni-modal and narrowly peaked around the maximum likelihood value (e.g. Trotta 2008). Indeed, we expect future HERA and SKA observations to be highly constraining, making Fisher forecasts useful approximations of the true posterior for futuristic data sets.

Figure 3 shows the forecasted parameter constraints from our Fisher-matrix approach for 10 EOS21 fiducial free parameters (Muñoz et al. 2022). The contours show $1, 2\sigma$ confidence intervals for the parameters. For this forecast, we assume one year (i.e., 1080 hours) of HERA data, and show the forecasts for moderate and pessimistic foreground noise levels, as described in Section 2.4.

The fractional error estimated for each parameter is plotted in Figure 4. With moderate foreground noise it would be possible to infer nearly all parameters to $\lesssim 10 - 30\%$ precision, with the exception of the amplitude of Lyman-Werner feedback, A_{LW}, which is much more poorly constrained – due to its smaller impact on the amplitude of the 21-cm signal (Muñoz et al. 2022), and X-ray parameters which can be constrained to $< 10\%$. This demonstrates 21-cm observations will be a powerful tool to learn about the formation of the first galaxies. As expected, parameter constraints would be substantially worse in the case of 'pessimistic' foreground noise,
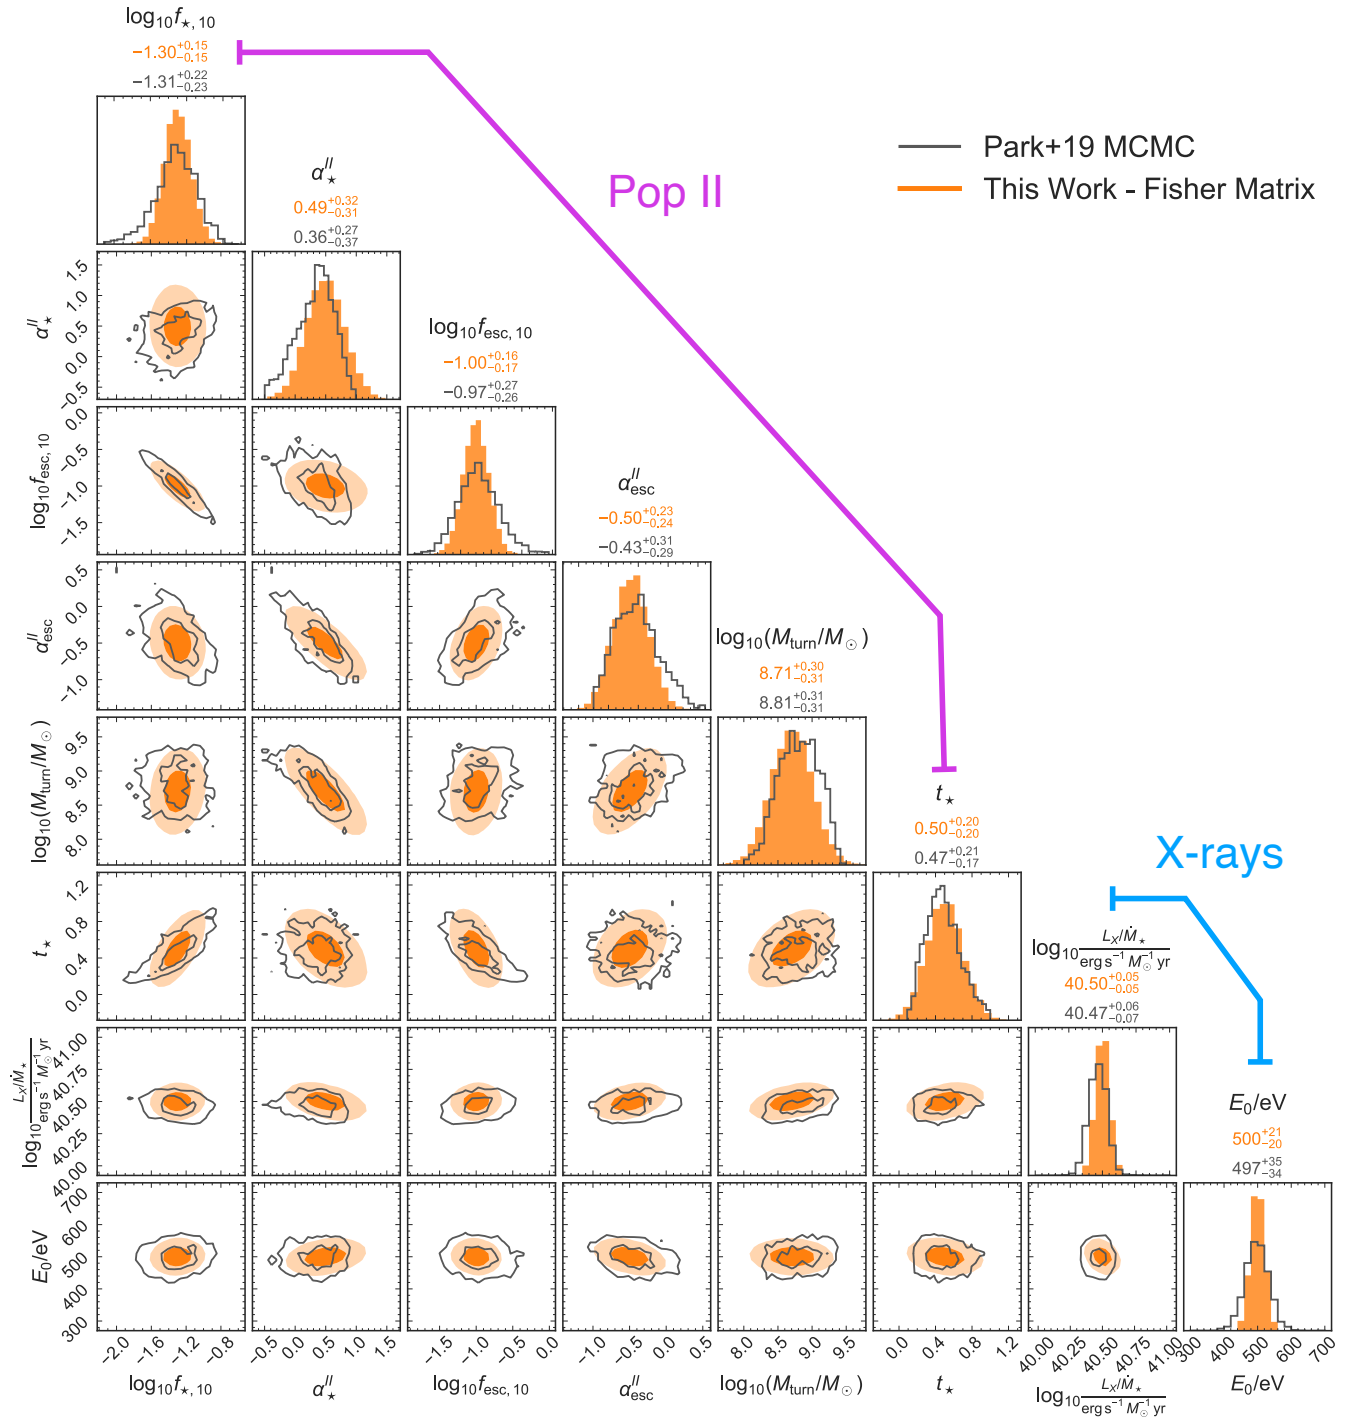
Figure 2. Validation of our Fisher-matrix technique by comparison with the MCMC runs of Park et al. (2019). We plot parameter constraints from Park et al. (2019) (grey line) and our results (orange filled). The top right plots show 1D marginalized posteriors and the lower triangle plots show the 2D marginalized posteriors. The contours show $1, 2\sigma$ confidence intervals for the posteriors. Our Fisher-matrix approach recovers the same degeneracies as the MCMC and comparable uncertainties, while requiring $\sim 10^4$ fewer simulations.

particularly on population III parameters. This is because the foreground noise is a strong function of frequency and thus redshift, with the lowest noise in the $z < 10$ – Epoch of Reionization – window, which is dominated by population II-hosting galaxies. This can be seen clearly in the signal-to-noise ratio of our fiducial 21-cm power spectrum as a function of redshift for the two noise models, which we plot in Figure 5.

We find that most parameters for *both* population II and III galaxies could be measured at $< 10\%$ precision under the assumption of moderate instrumental noise. This is due to the relative redshift dependence of the parameters and the wide frequency/redshift range of the HERA observations. In Figure 6 we plot the redshift dependence of the fractional error on each parameter, calculated in rolling redshift bins corresponding to two adjacent HERA fre-
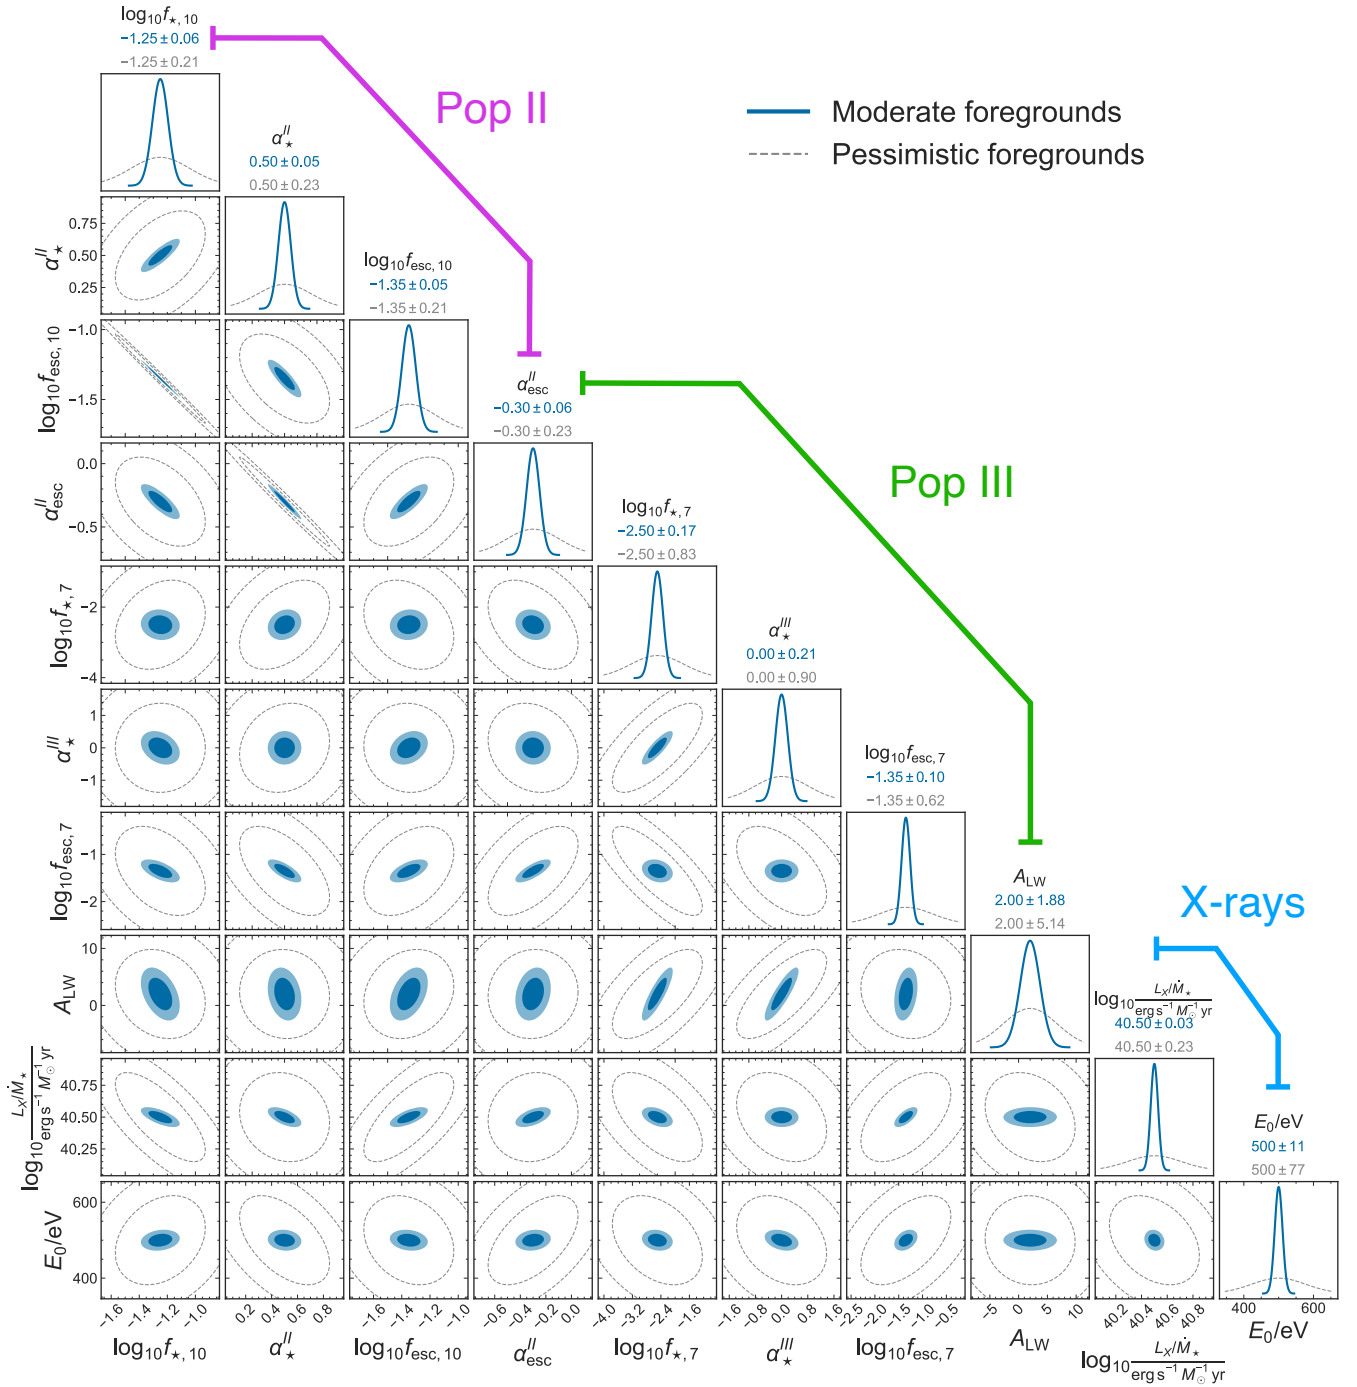
Figure 3. Parameter constraints for the fiducial EOS21 run, using the parameters defined by Muñoz et al. (2022) for moderate (blue solid) and pessimistic (grey dashed) foregrounds in HERA 1-year observations, as described in Section 2.4. The contours show 1 and 2σ confidence intervals and the quoted confidence interval for the 1D constraint is 1σ.

quency 8 MHz bandwiths. From this figure we see the lowest errors on population II parameters are obtained during the later stages of the Epoch of Reionization, $z \sim 5 - 7$ (see Figures 6 and A1), when they strongly affect the 21-cm signal. At these redshifts the observational noise is lowest (Figure 5) and the redshift resolution due to our sampling of the lightcone in fixed frequency bins is highest, which will naturally increase the signal-to-noise for population II parameters. But by $z \gtrsim 10$ the errors on population II parameters become comparable to the errors on population III parameters,

though the measurements are obtained in sparser redshift bins, and by $z \gtrsim 15$ the errors on population III parameters are lower than the population II parameters – as lower mass galaxies start to dominate the star formation rate density and thus the 21-cm signal (e.g., see Figure 6 by Muñoz et al. 2022). Due to the wide frequency range of HERA, the redshift dependence of the uncertainties of the parameters is captured, thus the population III parameters can be well-constrained. We can also see from Figure 6 that, as expected,
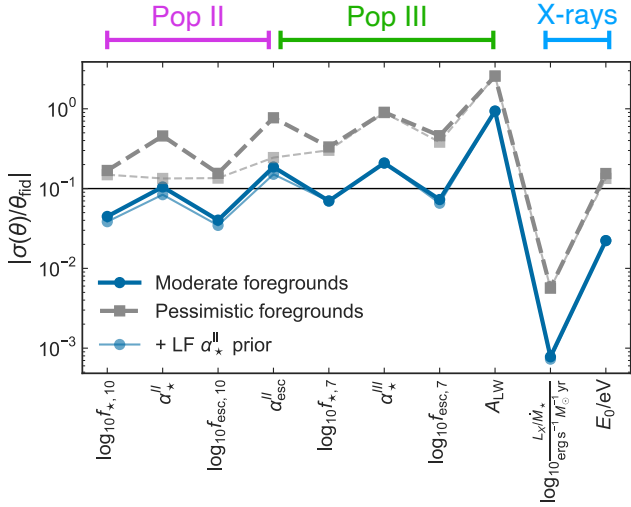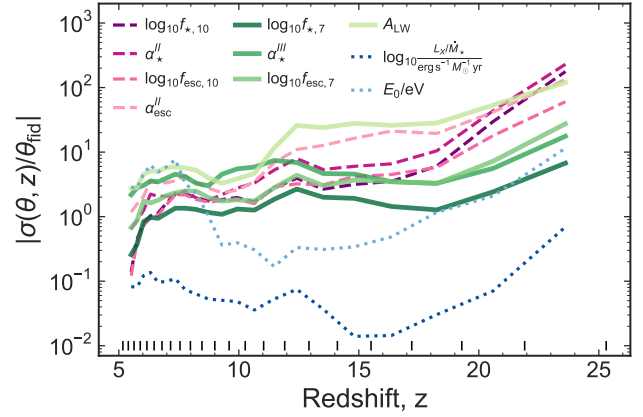
Figure 4. Fractional 1σ error in each parameter as estimated from the Fisher matrix, for the moderate (blue solid line) and pessimistic (grey dashed line) foreground noise models (see Section 2.4). In both cases, the thick lines show the HERA-only constraints and the thin solid lines in the same color and style show the fractional error estimate after including a prior on $\alpha_\star^{\rm II}$ from the UV LF using Hubble data. The $\alpha_\star^{\rm II}$ prior reduces the uncertainty on population II parameters by approximately a factor of three. We mark the 10% fractional uncertainty with a solid horizontal line for ease of comparison.
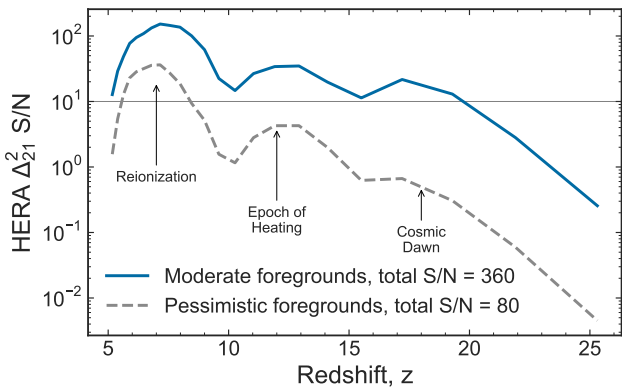
Figure 5. Signal to noise ratio of our fiducial EOS21 21-cm power spectrum as a function of redshift, using the moderate (blue solid) and pessimistic (grey dashed) expected noise from 1 year of HERA data, as described in Section 2.4. Due to the frequency/redshift dependence of the noise, the pessimistic noise level would only enable high S/N measurements during the Epoch of Reionization, thus constraints on population III galaxies would be very limited, as seen in Figure 3.

the lowest errors on X-ray parameters are obtained during the Epoch of Heating ($z \sim 10 - 18$ for the EOS21 model).

The Fisher forecast demonstrates there are strong degeneracies between the normalization and low-mass slope of the stellar-halo mass relation for both population II and III galaxies. There is also a very strong degeneracy between $f_{\star,10}^{\rm II}$ and $f_{\rm esc,10}^{\rm II}$ as noted by Park et al. (2019). $L_X/\dot{M}_\star$ is most strongly degenerate with population II parameters, but not as strongly with population III parameters. This stronger degeneracy is likely because population II galaxies are expected to dominate the star formation rate density during the Epoch of Heating ($z \lesssim 15$, see e.g., Muñoz et al. 2022) and so the relative abundance of these galaxies is thus degenerate with their


Figure 6. Fractional 1σ error in each parameter as estimated from the Fisher matrix in rolling redshift bins corresponding to two adjacent HERA frequency channels, demonstrating the redshift evolution of measurable constraints on each parameter. All lines show the constraints using the moderate foreground noise model (see Section 2.4). Pink dashed lines show population II parameters, green solid lines show population III parameters, and blue dotted lines show X-ray parameters. Vertical black ticks show the redshift bin edges for our mock HERA observations.

X-ray emission, whereas population III galaxies dominate the SFR density at higher redshifts, where X-ray heating has a smaller effect on the 21-cm signal.

In terms of an interplay between population II and III parameters, the strongest degeneracies are between $f_{\rm esc,7}^{\rm III}$ and population II parameters ($\alpha_\star^{\rm II}$, $f_{\star,10}^{\rm II}$, $\alpha_{\rm esc}$, $f_{\rm esc,10}^{\rm II}$), attributed to the uncertainty of the relative role of population III-hosting galaxies in reionization. Population III star formation parameters are also degenerate with $A_{\rm LW}$, which is expected as it sets the strength of Lyman-Werner feedback which suppresses star formation in molecular cooling halos.

We also produce forecasts for HERA observations adding a prior on $\alpha_\star^{\rm II}$. Park et al. (2019) demonstrated that adding other astrophysical observations to the 21-cm signal, in their case the UV luminosity function (LF), provides independent constraints on the model parameters. In their case they found uncertainties in population II star formation: $f_{\star,10}^{\rm II}$ and $\alpha_\star^{\rm II}$, were reduced by a factor of three when adding UV LF data to their likelihood, as the UV LF provides important information to constrain the stellar-to-halo mass relation (Equation 8). Here we use the LF-only constraint on $\alpha_\star^{\rm II}$ inferred by Park et al. (2019) from Hubble Space Telescope observations, $\sigma(\alpha_\star^{\rm II}) = 0.07$, as a prior, by adding $1/\sigma^2(\alpha_\star^{\rm II})$ to the diagonal element corresponding to $\alpha_\star^{\rm II}$ in the Fisher matrix (e.g., Coe 2009). The impact on the total uncertainties in each parameter is shown in Figure 4. Using the moderate foreground noise model, the impact of the $\alpha_\star^{\rm II}$ prior is negligible, as population II parameters can already be measured to similar precision to the prior, but the prior significantly reduces uncertainties on population II parameters in the case of the pessimistic foreground noise.

Our forecasts demonstrate that observations of the 21-cm power spectrum could be able to determine the efficiency of both population II and population III star formation, if the foreground noise follows the 'moderate' prescription, and determine the strength of X-ray production by the first galaxies to sub-percent precision even under 'pessimistic' noise. However, an understanding of the observational noise level is crucial. We note that although the 'pessimistic' foreground model is comparable to the noise model
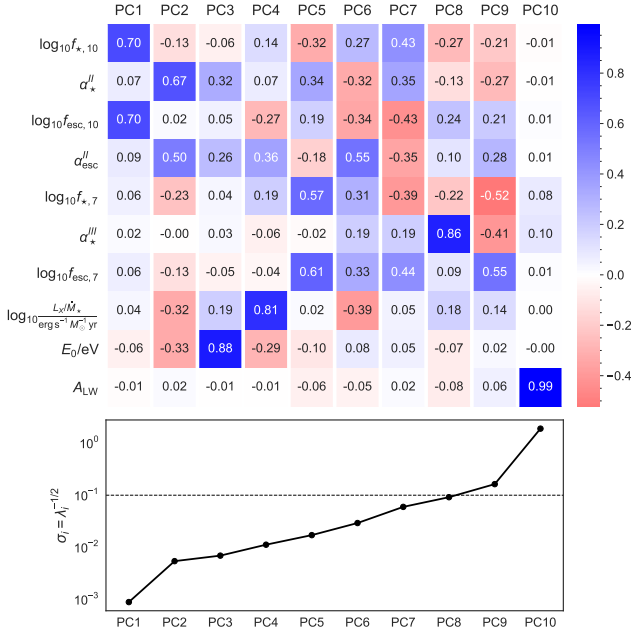
Figure 7. (Upper) Eigenvectors associated with each principal component, in order of decreasing importance/increasing uncertainty. For example, PC1 is $\approx f^{\rm II}_{\rm esc,10} \times f^{\rm II}_{\star,10}$. (Lower) Uncertainty for each principal component.

used by Park et al. (2019) our forecasted uncertainties for the EOS21 model can be higher because the fiducial parameters are different, and the increased number of free parameters adds additional uncertainty due to degeneracies between parameters.

3.3 Principal Components of the 21-cm Power Spectrum

Given the parametric model for galaxy formation in `21cmFAST`, we may ask which combination of parameters can be best constrained by upcoming data. To obtain the most important linear combinations of parameters in the 21-cm signal we carry out a principal component analysis of the Fisher matrix (e.g., Efstathiou 2002; Munshi & Kilbinger 2006).

To ensure all parameters are compared at approximately the same order of magnitude we rescale E_0 from eV to keV. Thus, we rescale the Fisher elements corresponding to E_0 by a factor 10^3.

We can diagonalize our adjusted Fisher matrix $\mathcal{F}'$:

$$\mathcal{F}' = \mathbf{U}\Lambda\mathbf{U}^{\rm T} \tag{18}$$

where Λ is a diagonal matrix with elements λ_i and $\mathbf{U}$ is the matrix of principal component eigenvectors, from which we can obtain the orthogonal linear combinations of the parameters θ:

$$\mathbf{X} = \mathbf{U}^{\rm T}\theta \tag{19}$$

and the variance of the principal component (PC) X_i is $1/\lambda_i$ (e.g. Efstathiou 2002).

Figure 7 shows each PC and its associated uncertainty ($\sigma_i = \lambda_i^{-1/2}$) as a function of the PC number i, in order of decreasing accuracy (lowest variance is the first PC). The first 7 out of 10 principal components can be measured to $< 10\%$ accuracy, and only one, dominated by $A_{\rm LW}$, to $> 100\%$.

These principal components demonstrate which combinations of parameters can be best constrained by HERA-like 21-cm power spectra. We obtain the principal components for the entire redshift

range of HERA ($z \sim 5 - 28$) and note that the lowest variance components detectable in the full redshift range are mainly set by the frequency range with the lowest noise – i.e. the PCs that contribute the most at $z \lesssim 8$ are heavily upweighted due to the much lower noise in that frequency range relative to the Cosmic Dawn window at $z \gtrsim 12$. We explore the PCs in a few redshift windows in Appendix A.

Several of the PCs have intuitive physical interpretations:

(i) PC1 is dominated by $f^{\rm II}_{\star,10}$ and $f^{\rm II}_{\rm esc,10}$. As these two quantities are constrained in log-space, this tells us that the product $f^{\rm II}_{\star,10} \times f^{\rm II}_{\rm esc,10}$ will be well-constrained by HERA. This is unsurprising as this product determines the number of ionizing photons which are available for reionization and these quantities are very degenerate (see Figure 3).

(ii) PC2 is dominated by a combination of population II and X-ray parameters, as these parameters have a strong impact on the signal during the Epoch of Reionization.

(iii) PC3 is dominated by E_0, which will be easily measured to high precision with minimal degeneracies with other parameters as it determines the hardness of the X-ray spectrum that emerges from early galaxies. Values of E_0 below ~ 1keV have a strong impact on the inhomogeneity of X-ray heating (and thus the associated 21-cm power spectrum) since the mean free path is a strong function of the photon energy (e.g., Pacucci et al. 2014; Fialkov & Loeb 2016; Das et al. 2017).

(iv) PC4 is dominated by $L_X/\dot{M}_\star$, which governs the relative timing of the X-ray heating epoch

(v) Higher principal components have combinations of population II and population III parameters, with the combinations of population III parameters having the highest variance. As discussed in Appendix A the principal components are dominated by which parameters can be measured in the redshift windows with lowest noise – thus when looking at the data over the full HERA observing window $z \sim 5 - 28$ population II reionization parameters will be best measured.

(vi) PC10 is dominated by $A_{\rm LW}$. As discussed in Section 3.2 this parameter is least constrained by 21-cm power spectra observations as it has a relatively low impact on the signal (see also Figure 17 by Muñoz et al. 2022).

3.4 Inference using the Fisher matrix

So far we have focused on the ability of the Fisher-matrix formalism to forecast errors with mock data. This can potentially save the running of computationally expensive MCMCs for enlarged parameter spaces (for instance if new parameters such as ETHOS are included, Muñoz et al. 2020; Mason et al. in prep.), and re-use the same Fisher matrix as above. However, the same formalism could be used for parameter inference from real data. We now explain the method for inference and its limitations.

The Fisher formalism assumes that the likelihood is Gaussian within the parameter range under consideration. For that, a requirement is that the observable (Δ^2_{21}) changes linearly with the parameters, i.e. $\Delta^2_{21}(k,z) \approx \sum_i \theta_i f_i(k,z)$, where $f_i(k,z)$ are not necessarily linear. Under that assumption, we can infer shifts $\delta\theta_i$ on parameters of a given data set, when compared to our fiducial (e.g., Muñoz et al. 2016). A major caveat is that this assumption of linearity is likely to break down when the observed parameters are a sufficient distance from the fiducial parameters.

We start by writing the observed 21-cm power spectrum as a linear function of the fiducial model power spectrum and the

parameter shifts:

$$\Delta^2_{21,\text{obs}}(k, z; \{\theta'\}) = \Delta^2_{21,\text{fid}} + \sum_i \frac{\partial \Delta^2_{21,\text{fid}}}{\partial \theta_i}(\theta'_i - \theta_{i,\text{fid}}), \quad (20)$$

for a set of parameters $\{\theta'\}$ which differ from the fiducial parameters, $\{\theta_{\text{fid.}}\}$. This formula will be exact only for small parameter differences ($\theta'_i - \theta \ll 1$), but can allow us to quickly test the deviation between the parameters of an observation and those of our fiducial.

Given the difference $\Delta^2_{21,\text{diff}} = \Delta^2_{21,\text{obs}}(\theta') - \Delta^2_{21,\text{fid.}}$ between the observed 21-cm PS with parameters θ', and that of the fiducial, we define the difference vector as

$$D_j = \sum_{i_k, i_z} \Delta^2_{21,\text{diff}}(k, z) \frac{\partial \Delta^2_{21,\text{fid}}(k, z)}{\partial \theta_j} \frac{1}{\sigma^2_{\Delta^2}(k, z)} \quad (21)$$

akin to Equation 16 (which defined the Fisher matrix), but with one of the derivatives swapped for $\Delta^2_{21,\text{diff}}$.

Then, taking into account the covariance matrix $C = \mathcal{F}^{-1}$ between parameters, we can find the shifts between the fiducial parameters and the observed ones as

$$\delta\theta_i = \theta'_i - \theta_{i,\text{fid}} = \sum_j C_{ij} D_j. \quad (22)$$

We test this formalism by comparing two different choices of mock observations against our baseline (EOS21) fiducial one.

(i) In the first we only slightly modified the parameters, randomly sampling the parameters from within 1σ of our expected errors for the EOS fiducial with moderate noise (i.e. they fall within the blue ellipses in Figure 3). In this case we expect the likelihood to be roughly Gaussian, and the shifts to be consistent with the inputs.

(ii) In the second we aim to infer the OPT EOS model, which has an enhanced population III contribution to star formation and X-ray production (see Muñoz et al. 2022 for details). In the OPT model we change the following parameters: $f^{\text{II}}_{\star,10} = -1.25$, $f^{\text{II}}_{\text{esc},10} = -1.2$, $f^{\text{III}}_{\star,7} = -1.75$, $f^{\text{II}}_{\text{esc},7} = -2.35$ and $E_0 = 200$ eV. These are significantly different from our EOS21 fiducial (see Table 1), so they will test how far from the fiducial the parameters can be before the formalism breaks down.

In both cases we generate the new mock observed 21-cm power spectra, $\Delta^2_{21,\text{obs}}(\theta')$, using identical settings as for our fiducial EOS run, changing only the astrophysical parameters.

A subtlety in this analysis is that, as well as the Poisson noise in the power spectrum – due to measuring the power spectrum from a finite number of k-space modes in the simulation boxes – both our fiducial and "observed" mocks have sample (cosmic) variance due to the scatter between independent realizations of the initial conditions. Both the Poisson noise and cosmic variance depend on the volume of the simulations (see e.g., Iliev et al. 2014; Kaur et al. 2020 for the power spectrum and Muñoz & Cyr-Racine 2021 for the global signal). Thus the total noise on the power spectrum is the Poisson, cosmic variance and instrumental noise terms added in quadrature: $\sigma^2_{\text{tot}}(k, z) = \sigma^2_{\text{poiss.}}(k, z) + \sigma^2_{c.v.}(k, z) + \sigma^2_{\text{instr.}}(k, z)$. As such, using the Fisher matrix as defined above undercounts the error bars, for it assumes the mock has the cosmic variance from a HERA-like cosmic observation, rather than the much smaller volume of our mock simulation.

To circumvent the issue of cosmic variance, we first run the two cases with the same initial conditions as our EOS fiducial. Using simulations with the same initial conditions as the fiducial we find we can recover the 'observed' parameters within 1σ for

case (i), but the inferred parameters are not well-recovered for case (ii), especially for parameters where the shift is greater than the 1σ errors on the fiducial value. We surmise that the inference approach outlined above only works well when the assumption that the power spectrum varies linearly with the parameters holds – i.e. when the likelihood is most like a Gaussian, and the 'true' parameters are within 1σ of the peak likelihood for our specific fiducial.

When varying the initial conditions we find a poor recovery of the parameter shifts for both cases. This is ameliorated when increasing the instrumental error bars (using our 'pessimistic' noise case), as that increases the variance of each k number, which partially accounts for the undercounted cosmic variance. Given that we cannot "match" the initial conditions of a real power spectrum detection in the sky, we conclude that using the Fisher matrix for inference is challenging in our formalism, even if the likelihood is approximately Gaussian. Thus a dedicated variance analysis may be required but is beyond the scope of this paper.

4 DISCUSSION

4.1 Computational efficiency

The key advantage of using Fisher matrices over MCMC is the reduced computational time, due to requiring far fewer individual simulation runs. For the Fisher-matrix calculation we require $2N+1$ simulations, where N is the number of free parameters, whereas for an MCMC with 21CMMC of order $\sim 10^5$ simulations are required to reach convergence (Greig & Mesinger 2017).

For comparison, the MCMC run by Park et al. (2019) required 70,000 individual simulations, including burn-in steps, whereas our Fisher-matrix calculation using the same fiducial parameters required only 17 individual simulations. While the simulation generation can be parallelised, we still expect the Fisher-matrix approach can greatly reduce the computation time compared to an MCMC. The Fisher matrix also has the advantage that additional parameters can be varied by just running a couple of new simulations, rather than having to re-run an MCMC to explore the joint posterior. While an MCMC should be the gold standard for analysis of observations (e.g., Ghara et al. 2020, 2021; Greig et al. 2021b,a; The HERA Collaboration et al. 2022b,a), we have demonstrated that a Fisher-matrix analysis can accurately recover the same parameter degeneracies and uncertainties.

This means a Fisher-matrix approach could be useful for testing and prototyping analyses to forecast parameter uncertainties. Potential applications could be: adding new parameters – for example cosmological and dark-matter parameters, which will be the topic of a future work (Mason et al. in prep); or testing the impact of instrument/observing designs (described more below).

The limitations of the Fisher-matrix approach are that, as described in Section 3.4, it is difficult to use the Fisher-matrix approach for actual inference, thus MCMC or nested sampling will still be required to map full posterior when we have 21-cm observations. But, for forecasting, the Fisher-matrix approach is extremely efficient for estimating parameter uncertainties assuming a fiducial model (e.g. using the maximum *a posteriori* model from current, non-21-cm observations (e.g., Park et al. 2019; Qin et al. 2021b). Additionally, parameter estimates from Fisher matrices are always optimistic, as the Fisher matrix will always produce the lower bound on the possible observed covariance matrix due to the assumption of a multi-variate Gaussian posterior (see Section 2.3), so it is possible that some of the uncertainties will be underestimated.

4.2 Importance of understanding observational noise

As demonstrated in Section 3.2, assumptions about the noise level of the 21-cm observations can have order of magnitude effects on the inferred parameter constraints. Figure 4 shows the fractional error in each parameter under the assumption of the moderate and pessimistic foreground noise models, with and without the inclusion of the UV LF prior on $\alpha_\star^{II}$.

An advantage of the Fisher-matrix approach is that it is very efficient to rerun the forecast analysis under different assumptions of the foreground noise model, whereas using an MCMC this would require re-calculating the likelihood for every simulation. As a good understanding of the noise is crucial to measuring astrophysics and cosmology from 21-cm observations, our approach provides an efficient way to assess the impact of assumptions in instrumental noise and foreground models.

5 CONCLUSIONS

We have created a Fisher-matrix wrapper for the public code `21cm-FAST`, `21cmfish`, enabling rapid parameter forecasts from the cosmic 21-cm signal. Our conclusions are as follows:

- We verify that our Fisher-matrix analysis recovers the same parameter degeneracies and produces comparable uncertainties as the 21cm-only MCMC of Park et al. (2019), requiring only $\sim 0.03\%$ of the individual simulations. This means that, under the assumption of a multi-variate Gaussian posterior, a Fisher matrix can be used to rapidly explore parameter constraints with 21-cm observations (see also e.g., Pober et al. 2014; Liu & Parsons 2016; Shimabukuro et al. 2017; Muñoz et al. 2020; Jones et al. 2021), significantly reducing the requirement for expensive and energy intensive computation when prototyping analyses (and does not require training, c.f. emulator approaches e.g., Kern et al. 2017).

- Using our Fisher-matrix approach, we perform the first joint analysis of population II and population III galaxy parameters as well as those characterizing feedback from radiative backgrounds. We find that under the assumption of a 'moderate' foreground noise floor (using the 'moderate' foregrounds model in `21cmsense`, with a super-horizon buffer $a = 0.1\,h\,\text{Mpc}^{-1}$ and expected HERA system temperature, DeBoer et al. 2017), both population II and population III parameters can be constrained by future HERA data to $\lesssim 10\%$ precision, due to the relative importance of the parameters as a function of redshift and the wide frequency coverage of HERA.

- Adding priors on parameters from independent measurements of the UV luminosity function from Hubble Space Telescope data to 21-cm observations improves constraints on population II parameters by approximately a factor of three in the case of pessimistic foreground noise, but does not significantly improve estimates in the case of moderate foreground noise as those constraints are already comparable to the UV LF information.

- A principal component analysis demonstrates that at least 7 combinations of the parameters could be well-measured by HERA. The first four principal components are dominated by $f_{\text{esc},10}^{II} f_{\star,10}^{II}$ and combinations of population II and X-ray emission properties, as these are the most important features in determining the strength of the 21-cm signal at $z \lesssim 10$ where the signal will be best measured. Our analysis shows the ionizing photon escape fraction and stellar-to-halo mass fraction are highly degenerate using 21-cm observations alone (see also, Park et al. 2019).

- We attempt to use our Fisher-matrix approach to perform an inference of simulated data and find that it is only possible to reliably

recover parameters that were within 1σ of the fiducial simulation and when using the same initial conditions for the fiducial and mock data simulations. Thus, we caution that while the Fisher matrix is useful for forecasting, within our formalism it is not a good choice for inference: Bayesian inference techniques such as MCMC or Nested sampling are more appropriate for mapping out the full posterior for unknown parameters, once detections are available.

Our results show that Fisher-matrix analyses provide a fast and realistic way to estimate parameter uncertainties from future 21-cm observations, under the assumption the posterior is a multi-variate Gaussian and also assuming the noise is separable from the cosmic signal. This greatly reduces the requirement for expensive MCMC runs when prototyping new analyses or adding additional parameters. This framework enables us to more easily explore the effects of foreground noise and the degeneracies between astrophysical and cosmological, e.g. dark-matter, parameters, which will be the topic of a future work.

DATA AVAILABILITY

Our Fisher analysis package for `21cmFAST` is available in a public python package `21cmfish`: https://github.com/charlottenosam/21cmfish. The code and accompanying data to reproduce plots in this paper are included in the package.

SOFTWARE

`21cmFAST` (Mesinger et al. 2011; Park et al. 2019; Qin et al. 2020; Murray et al. 2020; Muñoz et al. 2022), `21cmSense` (Pober et al. 2014), `corner` (Foreman-Mackey 2016), `IPython` (Pérez & Granger 2007), `matplotlib` (Hunter 2007), `NumPy` (Van Der Walt et al. 2011), `SciPy` (Oliphant 2007), `Astropy` (Robitaille et al. 2013).

APPENDIX A: PRINCIPAL COMPONENT ANALYSIS IN REDSHIFT WINDOWS

Our principal component analysis in Section 3.3 demonstrates the combinations of parameters that will be best-constrained by the full HERA dataset. However, due to the strong frequency and therefore redshift dependence of the noise (see Figure 5), where the background is higher at lower frequencies (e.g., DeBoer et al. 2017),

the principal components are dominated by the parameters most important for the Epoch of Reionization – at $z \lesssim 10$ where the noise is lowest.

To provide more insight into the redshift dependence of the principal components we do a PCA in three redshift bins, corresponding to: the Epoch of Reionization ($z \sim 5 - 10$), where the global 21-cm signal peaks and declines to zero); the Epoch of Heating ($z \sim 10-15$), the signal goes from absorption to emission as the IGM is heated (likely by X-ray emission from the first galaxies); and early Cosmic Dawn ($z \sim 15-30$), where the global signal starts to drop due to feedback from the first stars, until it reaches the trough before heating dominates.

We plot the principal components from these three redshift bins in Figure A1. These plots demonstrate which combinations of parameters the observed redshift windows are most sensitive to. For the Reionization redshift window, we see the first principal component is very similar to the one obtain for the whole redshift range: $f_{\star,10}^{\rm II} \times f_{\rm esc,10}^{\rm II}$. The following principal components are dominated by population II and X-ray parameters, as expected during the Reionization epoch. In the Epoch of Heating the primary principal components are dominated by X-ray parameters, as expected, as in the model, X-rays are the dominant driver of IGM heating. For the Cosmic Dawn redshift window, the first principal components are dominated by X-ray parameters and a combination of population III parameters.

REFERENCES

Albrecht A., et al., 2009, arXiv e-prints, p. arXiv:0901.0721
Barkana R., Loeb A., 2001, Phys. Rep., 349, 125
Battaglia N., Trac H., Cen R., Loeb A., 2013, ApJ, 776, 81
Behroozi P. S., Silk J., 2015, ApJ, 799, 32
Behroozi P. S., Marchesini D., Wechsler R. H., Muzzin A., Papovich C., Stefanon M., 2013, ApJ, 777, L10
Bosman S. E. I., Fan X., Jiang L., Reed S., Matsuoka Y., Becker G., Haehnelt M., 2018, MNRAS, 479, 1055
Bouwens R. J., et al., 2015, ApJ, 803, 34
Bowman J. D., Rogers A. E. E., Monsalve R. A., Mozdzen T. J., Mahesh N., 2018, Nature, 555, 67
Choudhury T. R., Paranjape A., 2018, MNRAS, 481, 3821
Coe D., 2009, arXiv e-prints, p. arXiv:0906.4123
Das A., Mesinger A., Pallottini A., Ferrara A., Wise J. H., 2017, MNRAS, 469, 1166
Davies F. B., Becker G. D., Furlanetto S. R., 2018, ApJ, 860, 155
DeBoer D. R., et al., 2017, PASP, 129, 045001
Efstathiou G., 2002, MNRAS, 332, 193
Ewall-Wice A., Hewitt J., Mesinger A., Dillon J. S., Liu A., Pober J., 2016, MNRAS, 458, 2710
Fialkov A., Loeb A., 2016, ApJ, 821, 59
Fialkov A., Barkana R., Visbal E., Tseliakhovich D., Hirata C. M., 2013, MNRAS, 432, 2909
Foreman-Mackey D., 2016, The Journal of Open Source Software, 1, 24
Fragos T., Lehmer B. D., Naoz S., Zezas A., Basu-Zych A., 2013, ApJ, 776, L31
Furlanetto S. R., 2006, MNRAS, 371, 867
Furlanetto S. R., Oh S. P., Briggs F. H., 2006, Phys. Rep., 433, 181
Gehlot B. K., et al., 2019, MNRAS, 488, 4271
Ghara R., Choudhury T. R., Datta K. K., 2015, MNRAS, 447, 1806
Ghara R., et al., 2020, MNRAS, 493, 4728
Ghara R., Giri S. K., Ciardi B., Mellema G., Zaroubi S., 2021, MNRAS, 503, 4551
Greig B., Mesinger A., 2015, MNRAS, 449, 4246
Greig B., Mesinger A., 2017, MNRAS, 472, 2651
Greig B., Mesinger A., 2018, MNRAS, 477, 3217
Greig B., Mesinger A., Haiman Z., Simcoe R. A., 2017, MNRAS, 466, 4239
Greig B., Mesinger A., Bañados E., 2019, MNRAS, 484, 5094
Greig B., Trott C. M., Barry N., Mutch S. J., Pindor B., Webster R. L., Wyithe J. S. B., 2021a, MNRAS, 500, 5322
Greig B., et al., 2021b, MNRAS, 501, 1
Greig B., Ting Y.-S., Kaurov A. A., 2022, MNRAS, 513, 1719
Hills R., Kulkarni G., Meerburg P. D., Puchwein E., 2018, Nature, 564, E32
Hirata C. M., 2006, MNRAS, 367, 259
Hoag A., et al., 2019, ApJ, 878, 12
Hunter J. D., 2007, Comput. Sci. Eng., 9, 99
Hutter A., Dayal P., Yepes G., Gottlöber S., Legrand L., Ucci G., 2021, MNRAS, 503, 3698
Iliev I. T., Mellema G., Ahn K., Shapiro P. R., Mao Y., Pen U. L., 2014, MNRAS, 439, 725
Jones D., Palatnick S., Chen R., Beane A., Lidz A., 2021, ApJ, 913, 7
Kaur H. D., Gillet N., Mesinger A., 2020, MNRAS, 495, 2354
Kern N. S., Liu A., Parsons A. R., Mesinger A., Greig B., 2017, ApJ, 848, 23
Koopmans L., et al., 2015, in Advancing Astrophysics with the Square Kilometre Array (AASKA14). p. 1 (arXiv:1505.07568), doi:10.22323/1.215.0001
Kulkarni M., Visbal E., Bryan G. L., 2021, ApJ, 917, 40
Liu A., Parsons A. R., 2016, MNRAS, 457, 1864
Liu A., Shaw J. R., 2020, PASP, 132, 062001
Liu A., Tegmark M., Bowman J., Hewitt J., Zaldarriaga M., 2009, MNRAS, 398, 401
Liu A., Pritchard J. R., Tegmark M., Loeb A., 2013, Phys. Rev. D, 87, 043002
Loeb A., Zaldarriaga M., 2004, Phys. Rev. Lett., 92, 211301
Machacek M. E., Bryan G. L., Abel T., 2001, ApJ, 548, 509
Mason C. A., Trenti M., Treu T., 2015, ApJ, 813, 21
Mason C. A., Treu T., Dijkstra M., Mesinger A., Trenti M., Pentericci L., de Barros S., Vanzella E., 2018, ApJ, 856, 2
Mason C. A., et al., 2019, MNRAS, 485, 3947
McGreer I. D., Mesinger A., D'Odorico V., 2015, MNRAS, 447, 499
McQuinn M., 2012, MNRAS, 426, 1349
Mertens F. G., et al., 2020, MNRAS, 493, 1662
Mesinger A., 2019, The Cosmic 21-cm Revolution; Charting the first billion years of our universe, doi:10.1088/2514-3433/ab4a73.
Mesinger A., Furlanetto S., 2007, ApJ, 669, 663
Mesinger A., Furlanetto S., Cen R., 2011, MNRAS, 411, 955
Mineo S., Gilfanov M., Sunyaev R., 2011, Astronomische Nachrichten, 332, 349
Mirocha J., 2014, MNRAS, 443, 1211
Morales M. F., Wyithe J. S. B., 2010, ARA&A, 48, 127
Morales A. M., Mason C. A., Bruton S., Gronke M., Haardt F., Scarlata C., 2021, ApJ, 919, 120
Moster B. P., Naab T., White S. D. M., 2013, MNRAS, 428, 3121
Muñoz J. B., 2019a, Phys. Rev. D, 100, 063538
Muñoz J. B., 2019b, Phys. Rev. Lett., 123, 131301
Muñoz J. B., Cyr-Racine F.-Y., 2021, Phys. Rev. D, 103, 023512
Muñoz J. B., Grin D., Dai L., Kamionkowski M., Kovetz E. D., 2016, Phys. Rev. D, 93, 043008
Muñoz J. B., Dvorkin C., Cyr-Racine F.-Y., 2020, Phys. Rev. D, 101, 063526
Muñoz J. B., Qin Y., Mesinger A., Murray S. G., Greig B., Mason C., 2022, MNRAS, 511, 3657
Munshi D., Kilbinger M., 2006, A&A, 452, 63
Murray S. G., Greig B., Mesinger A., Muñoz J. B., Qin Y., Park J., Watkinson C. A., 2020, Journal of Open Source Software, 5, 2582
Mutch S. J., Geil P. M., Poole G. B., Angel P. W., Duffy A. R., Mesinger A., Wyithe J. S. B., 2016, MNRAS, 462, 250
Oh S. P., Haiman Z., 2002, Astrophys. J., 569, 558
Oliphant T. E., 2007, Comput. Sci. Eng., 9, 10
Ouchi M., et al., 2010, ApJ, 723, 869
Pacucci F., Mesinger A., Mineo S., Ferrara A., 2014, MNRAS, 443, 678
Park J., Mesinger A., Greig B., Gillet N., 2019, MNRAS, 484, 933
Pérez F., Granger B. E., 2007, Comput. Sci. Eng., 9, 21
Planck Collaboration et al., 2016, A&A, 596, A108
Planck Collaboration et al., 2020, A&A, 641, A6
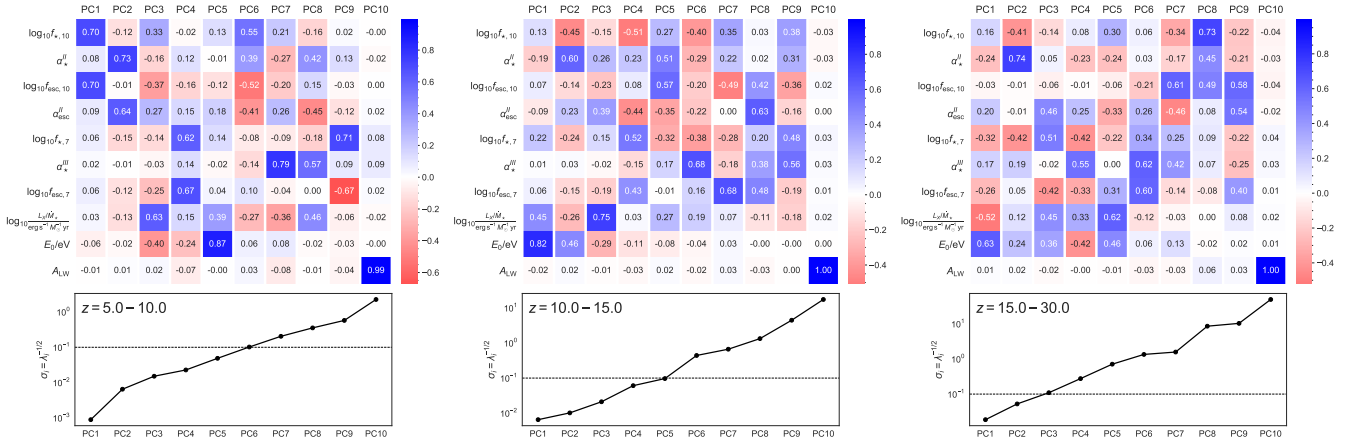
Figure A1. (Left) Eigenvectors associated with each principal component, in order of decreasing importance/increasing uncertainty, in the Epoch of Reionization redshift window $z \sim 5 - 10$. (Center) Eigenvectors associated with each principal component, in order of decreasing importance/increasing uncertainty, in the Epoch of Heating redshift window $z \sim 10 - 15$. (Right) Eigenvectors associated with each principal component, in order of decreasing importance/increasing uncertainty, in the early Cosmic Dawn redshift window $z \sim 15 - 30$.

Pober J. C., et al., 2013, ApJ, 768, L36

Pober J. C., et al., 2014, ApJ, 782, 66

Pritchard J. R., Loeb A., 2012, Reports on Progress in Physics, 75, 086901

Qin Y., Mesinger A., Park J., Greig B., Muñoz J. B., 2020, MNRAS, 495, 123

Qin Y., Mesinger A., Greig B., Park J., 2021a, MNRAS, 501, 4748

Qin Y., Mesinger A., Bosman S. E. I., Viel M., 2021b, MNRAS, 506, 2390

Robitaille T. P., et al., 2013, A&A, 558, A33

Sabti N., Muñoz J. B., Blas D., 2022a, Phys. Rev. D, 105, 043518

Sabti N., Muñoz J. B., Blas D., 2022b, ApJ, 928, L20

Santos M. G., Ferramacho L., Silva M. B., Amblard A., Cooray A., 2010, MNRAS, 406, 2421

Schauer A. T. P., Glover S. C. O., Klessen R. S., Clark P., 2021, MNRAS, 507, 1775

Sheth R. K., Mo H. J., Tormen G., 2001, MNRAS, 323, 1

Shimabukuro H., Yoshiura S., Takahashi K., Yokoyama S., Ichiki K., 2017, MNRAS, 468, 1542

Singh S., et al., 2022, Nature Astronomy, 6, 607

Sitwell M., Mesinger A., Ma Y.-Z., Sigurdson K., 2014, MNRAS, 438, 2664

Skinner D., Wise J. H., 2020, MNRAS, 492, 4386

Sobacchi E., Mesinger A., 2014, MNRAS, 440, 1662

Sobacchi E., Mesinger A., 2015, MNRAS, 453, 1843

Sun G., Furlanetto S. R., 2016, MNRAS, 460, 417

Tacchella S., Bose S., Conroy C., Eisenstein D. J., Johnson B. D., 2018, ApJ, 868, 92

The HERA Collaboration et al., 2022a, arXiv e-prints, p. arXiv:2210.04912

The HERA Collaboration et al., 2022b, ApJ, 924, 51

The HERA Collaboration et al., 2022c, ApJ, 925, 221

Tingay S. J., et al., 2013, Publ. Astron. Soc. Australia, 30, e007

Trott C. M., et al., 2020, MNRAS, 493, 4711

Trotta R., 2008, Contemporary Physics, 49, 71

Tseliakhovich D., Hirata C., 2010, Phys. Rev. D, 82, 083520

Van Der Walt S., Colbert S. C., Varoquaux G., 2011, Comput. Sci. Eng., 13, 22

Visbal E., Loeb A., 2010, J. Cosmology Astropart. Phys., 2010, 016

Visbal E., Haiman Z., Terrazas B., Bryan G. L., Barkana R., 2014, MNRAS, 445, 107

Wang F., et al., 2020, ApJ, 896, 23

Whitler L. R., Mason C. A., Ren K., Dijkstra M., Mesinger A., Pentericci L., Trenti M., Treu T., 2020, MNRAS, 495, 3602

Yung L. Y. A., Somerville R. S., Finkelstein S. L., Popping G., Davé R., Venkatesan A., Behroozi P., Ferguson H. C., 2020, MNRAS, 496, 4574

van Haarlem M. P., et al., 2013, A&A, 556, A2

This paper has been typeset from a TeX/LaTeX file prepared by the author.