Title

Deciphering the Thermal and Ionization State of the Intergalactic Medium over the Past 10 Billion Years

Permalink

https://escholarship.org/uc/item/0nx696np

Author

Hu, Teng

Publication Date

2024

University of California
Santa Barbara

Deciphering the Thermal and Ionization State of the Intergalactic Medium over the Past 10 Billion Years

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Teng Hu

Committee in charge:

Professor Joseph Hennawi, Chair
Professor Peng Oh
Professor Timothy Brandt

September 2024

The Dissertation of Teng Hu is approved.

Professor Peng Oh

Professor Timothy Brandt

Professor Joseph Hennawi, Committee Chair

June 2024

Deciphering the Thermal and Ionization State of the Intergalactic Medium over the Past 10 Billion Years

by

Teng Hu

Dedicated to Guoqiang, my grandpa, who cooks good meals with extraordinary imagination and critical analysis.

Acknowledgements

I would like to express my deepest gratitude to my advisor, Joe. When I began my PhD, I literally had 0 knowledge, and it was Joe who taught me from zero and guided through this project. His dedication, diligence, and curiosity have been a perpetual source of inspiration for me. I am also thankful to my co-advisor, Vikram, who has provided substantial support both academically and personally. I am also grateful to my collaborators and the members of ENIGMA. Working alongside such interesting and brilliant individuals has been a profound pleasure. They have enriched my understanding not only in scientific matters but also in life. Many thanks to my friends. Although I am an introvert, the time I spend with you is deeply cherished and means so much to me. In addition, your charming personalities have deeply affected me, inspiring me constantly along my journey. Most importantly, I must thank my family for being helpful and supportive in my life. This also includes my cats, Universe and Seven, and my grandpa Guoqiang, who always cook good meals with extraordinary imagination and critical analysis. Though he passed away ten years ago in August 2014, his influence remains substantial. Without them, none of my achievements would have been possible. Thank you all for being with me during this journey.

Curriculum Vitæ

Teng Hu

Education

2024	Ph.D. in Physics (Expected), University of California, Santa Barbara.
2020	M.A. in Physics, University of California, Santa Barbara.
2017	B.S. in Astronomy and Mathematics (Double Major), Physics minor, University of Illinois, Urbana-Champaign

Publications

Measuring the Thermal and Ionization State of the Low-z IGM using Likelihood-Free Inference

Hu, Teng; Khaire, Vikram; Hennawi, Joseph F.; Walther, Michael; Hiss, Hector; Alsing, Justin; Oñorbe, Jose; Lukic, Zarija ; Davies, Frederick 2022, MNRAS, 515, 2188; arXiv:2207.07151

The Impact of the WHIM on the IGM Thermal State Determined from the Low-z Lyman-Forest

Hu, Teng; Khaire, Vikram ; Hennawi, Joseph F. ; Onorbe, Jose ; Walther, Michael ; Lukic, Zarija ; Davies, Frederick 2024, MNRAS, 527, 11338; arXiv:2308.14738

Measurements of the Thermal and Ionization State of the Intergalactic Medium during the Cosmic Afternoon

Hu, Teng; Khaire, Vikram ; Hennawi, Joseph F. ; Tripp, Todd M.; Onorbe, Jose ; Walther, Michael ; Lukic, Zarija Submitting to MNRAS; arXiv:2311.17895

Measurements of the IGM Thermal and Ionization State across $0 < z < 0.5$

Hu, Teng et al. 2024 in prep

Can the Low Redshift Lyman Alpha Forest Constrain AGN Feedback Models?

Khaire, Vikram ; **Hu, Teng** ; Hennawi, Joseph F. ; Walther, Michael ; Davies, Frederick Accepted by MNRAS; arXiv:2306.05466

Searching for the Imprints of AGN Feedback on the Lyman Alpha Forest Around Luminous Red Galaxies

Khaire, Vikram ; **Hu, Teng**; Hennawi, Joseph F. ; Burchett, Joseph N. ; Walther, Michael ; Davies, Frederick MNRAS Minor revision; arxiv:2311.08470

Abstract

Deciphering the Thermal and Ionization State of the Intergalactic Medium over the Past 10 Billion Years

by

Teng Hu

One of the great successes of modern cosmology is the percent-level concordance between theory and observations of the intergalactic medium (IGM) at $z \gtrsim 1.7$. Yet, the Lyα forest at $z < 1.7$, which can only be studied via HST UV spectra, has pointed out a puzzling discrepancy, i.e., the Doppler b-parameters of these absorption lines are, on average, $\sim$ 10 km/s wider than those in any existing hydrodynamic simulation. This discrepancy implies that the low-z IGM might be substantially hotter than expected, contradicting one of the fundamental predictions in current cosmology that the IGM should cool down owing to the Hubble expansion after He II reionization ($z < 2.5$). Moreover, the IGM thermal state degenerates with its ionization state characterized by the UV background (UVB) photoionization rate, $\Gamma_{\rm HI}$, which dictates the abundance of the Lyα absorbers, $\mathrm{d}N/\mathrm{d}z$. Such a degeneracy requires any reliable measurement to adopt a careful statistical inference procedure. To overcome these difficulties, in this thesis, a novel machine-learning-based inference framework is employed to jointly measure the thermal and ionization state of the IGM, using the 2D distribution of b-parameter and H I column density and $\mathrm{d}N/\mathrm{d}z$. This method effectively resolves the degeneracies between the thermal and ionization state of the IGM and achieves high precision, even with limited-sized data. I apply this method to 94 archival HST COS and STIS quasar spectra distributed across the seven redshift bins, yielding a comprehensive evolutionary history of the IGM thermal and ionization state at $z < 1.5$. The results suggest that the IGM may be significantly hotter than previously expected at low-z and is potentially isothermal, with IGM temperature at mean density, $T_0 \sim 30,000$K and power-law index of the temperature-density, $\gamma \sim 1.0$ at $z = 0.1$. The inferred thermal history suggests that this unexpected IGM temperature possibly emerges around $z \sim 1$. Additionally, while the $\Gamma_{\rm HI}$ measurements align with the theoretical model at $z \sim 1$, the values measured at $z < 0.5$ are substantially lower than predicted, posing challenges to low-z UV background synthesis models.

Contents

Chapter 1

Introduction

1.1 The Intergalactic Medium and Cosmic Evolution

Since the dawn of civilization, we humans have been enchanted by the mysteries of the origin and evolution of the Universe. Today, modern cosmology allows us to explore the deepest recesses of the cosmos, revealing the earliest observable moments of the Universe. According to current cosmological models, all baryons were highly ionized until "recombination," which took place 378,000 years after the Big Bang (at redshift $z \sim 1100$, see Ryden 2003, Tanabashi et al. 2018). This event occurred as the universe expanded and cooled enough for free electrons to combine with protons, forming neutral hydrogen atoms. This transition made the universe transparent for the first time. The atoms—primarily hydrogen and helium—then transition to their lowest energy state (ground state) by emitting photons, a process known as "photon decoupling." These photons are still observable today as the Cosmic Microwave Background (CMB), which is the oldest direct observation of the Universe(Penzias & Wilson 1965, Planck Collaboration et al. 2014).

Following recombination and decoupling, the universe entered the "dark ages," a period characterized by the absence of electromagnetic radiation sources, lasting about one billion years. During this epoch, as the gas cooled, it collapsed into dark matter halos due to gravitational forces, resulting in the formation of stars and proto-galaxies. These objects began to emit photons into the space between galaxies filled with low-density gas, commonly known as the intergalactic medium (IGM). During this process, photons with sufficient energy may ionize the neutral hydrogen atoms in the IGM. More specifically, photons with an energy $E \geq 13.6$ eV, corresponding to the ionization energy of the hydrogen(HI→HII), re-ionize the neutral hydrogen in the IGM and convert them into protons and electrons. This phase transition, known as "reionization" or "hydrogen reionization," is believed to have completed around $z \sim 6$ (Madau et al. 1998, Fan et al. 2006, Faucher-Giguère et al. 2008, Robertson et al. 2015, McGreer et al. 2015).

Alongside hydrogen reionization, neutral helium atoms in the IGM are also re-ionized by photons with higher energy. The first ionization of helium (from HeI to HeII) requires 24.6 eV, while the double reionization of helium (from HeII to HeIII) demands even more energy, at 54.4 eV. This substantial difference in ionization energies suggests that helium reionization might occur at different epochs in the cosmic history compared to hydrogen reionization. Interestingly, the first ionization of helium actually occurred simultaneously with hydrogen reionization. While the first ionization of helium requires nearly twice the ionization energy of hydrogen, the universe contains an order of magnitude fewer helium atoms at the moment. Therefore, even though there were fewer photons with energy sufficient to ionize helium (i.e., $E \geq 24.6$ eV) compared to those needed for hydrogen ($E \geq 13.6$ eV), the lower abundance of helium means that there are still enough helium-ionizing photons to ionize helium during the era of hydrogen reionization.

Meanwhile, the first-generation stars did not produce enough photons energetic enough for the double ionization of helium, causing this process to pause temporarily until more powerful sources appeared. These sources were the extremely luminous active galactic nuclei (AGN), also known as quasars or QSOs. They are powered by supermassive black holes surrounded by accretion disks, where gas falls into the black hole, releasing vast amounts of energy through jets and electromagnetic radiation. The intense radiation from quasars dramatically alters the gas in the IGM. However, due to their radiation mechanisms, quasars are relatively short-lived, being active for only a few tens of millions of years. Moreover, because of their substantial mass and structural requirements, quasars do not emerge until later in cosmic history compared to hydrogen reionization. The luminosity density of quasars peaks around $z \sim 3$ (Worseck et al. 2011, Khaire 2017, Worseck et al. 2018, Kulkarni et al. 2019), making this epoch the period for the double ionization of helium (Madau & Meiksin 1994, Miralda-Escudé et al. 2000, McQuinn et al. 2009, Dixon & Furlanetto 2009, Syphers & Shull 2014).

These two reionization events significantly heat the IGM while altering its ionization state, as the excess energy from the photons is ultimately converted into kinetic energy. For typical IGM consisting of low-density gas, photoionization heating during these reionization events boosts the gas temperature to the order of $\mathcal{O}(10^4\mathrm{K})$. These heating events, therefore, leave substantial imprints on the thermal state of the IGM throughout cosmic history.

After hydrogen reionization ($z < 6$) the thermal state of the IGM is determined by the balance between heating from photoionization by the extragalactic UV background (UVB) and cooling mechanisms, including adiabatic cooling because of the Hubble expansion, radiative recombination cooling, and inverse Compton scattering where electrons interact with the cosmic microwave background. As a result of these processes, after the epoch of reionization, the IGM subsequently adheres to the power-law temperature-density (T-Δ) relation:

$$T(\Delta) = T_0 \Delta^{\gamma - 1}, \tag{1.1}$$

where $\Delta = \rho/\bar{\rho}$ is the overdensity, T_0 is the temperature at mean density $\bar{\rho}$, and γ is the power-law index (Hui & Gnedin 1997, McQuinn & Upton Sanderbeck 2016). These two parameters $[T_0, \gamma]$ thus characterize the thermal state of the IGM, and enable us to impose constraints on its thermal history at various epochs (Davé & Tripp 2001, Becker et al. 2011, Rorai et al. 2017, Hiss et al. 2018, Gaikwad et al. 2021), which enhance our understanding of the IGM thermal evolution and illustrate the intrinsic heating and cooling mechanisms of the Universe. Fig. 1.1 shows the distribution of T and Δ for the Nyx (Almgren et al. 2013) simulation with standard HeII reionization model at $z = 1.6$. The power-law fit of the IGM T-Δ relationship is shown as white dashed lines, and the best-fit thermal state $[T_0, \gamma]$ are given in the texts. The peak temperature in each $\log\Delta$ bin ($\log T_{\rm peak,i}$,$\log\Delta_i$) used for the power law fit are plotted as black dots with the 1-$\sigma_{T,i}$ error bars are also shown. It can be seen that the tight T-Δ relation is readily apparent. The detailed fitting method is presented in §2.1.1 The gas cells are divided into four phases depending on the temperature and density, namely the Warm-Hot-Intergalactic-Medium (WHIM), Diffuse Lyα, Hot Halo gas, and Condensed [1]. The density-weighted gas phase fractions are shown in the annotation.

Moreover, within the past 10 billion years ($z \lesssim 1.7$), shock heating caused by large-scale velocity flows, and possibly violent feedback from galaxy formation, converts a significant fraction of the cool gas ($T \sim 10^4$ K) in IGM into WHIM with $T > 10^5$K, resulting in a noticeable deviation from the tight power-law IGM T-Δ relationship. We will discuss these WHIM gas and their impacts on the IGM thermal state throughout this work, mainly in §5.

---

[1] Here we follow the definition used in Davé et al. (2010), where the cutoffs are set to be $T = 10^5$K and $\Delta = 120$, more discussion about the different cutoff used in literature can be found in Gaikwad et al. (2017b).

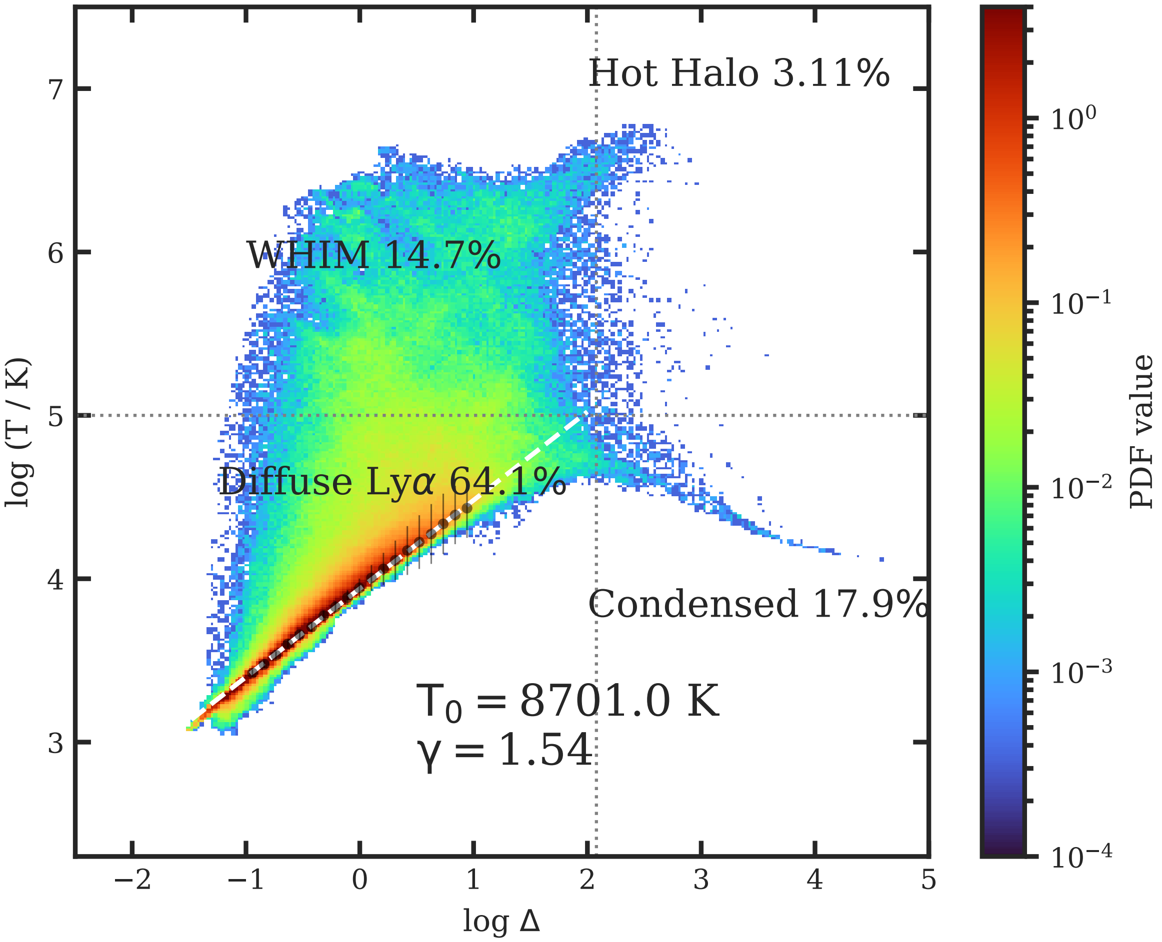

Figure 1.1: The distribution of the IGM gas for Nyx simulation at $z = 1.6$. The power-law fit of relationship is shown as the dashed white line, and the best-fit thermal state $[T_0, \gamma]$ are given in the texts. The peak temperature in each $\log\Delta$ bin ($\log T_{\rm peak,i}$,$\log\Delta_i$) used for the power law fit are plotted as black dots with the 1-$\sigma_{T,i}$ error bars are also shown. The density-weighted gas phase fractions are shown in the annotation.

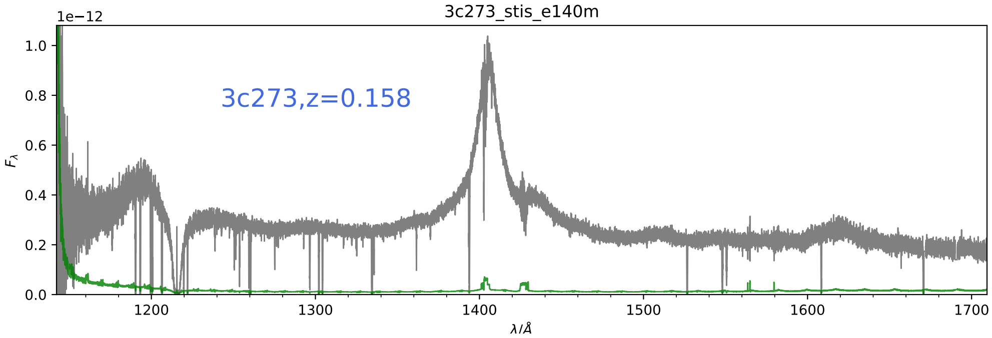

Figure 1.2: An example of Lyα forest quasar 3C 273 at z=0.158 observed by the HST STIS E140M grating.

1.2 The Lyman-α Forest and Line Decomposition

The diffuse nature of the IGM makes its corresponding emission lines difficult to observe, necessitating reliance on absorption features to probe its properties. Even though The IGM is predominantly ionized after ionization, there remains residual neutral hydrogen in the diffuse IGM (with a neutral fraction $n_{\rm HI}$ of the order of $\mathcal{O}(10^{-5})$). Owing to the high oscillator strength of the Lyα transition, the Lyα absorption ($\lambda_{\rm Ly\alpha} = 1215.67$ Å) caused by the residual neutral hydrogen is pronounced despite the low neutral fraction. As a result, these intervening residual neutral hydrogen across the IGM together give rise to the Lyα forest, a series of absorption lines in the spectra of distant galaxies or quasars arising from the Lyα transition of the neutral hydrogen atom in the IGM(Gunn & Peterson 1965, Lynds 1971). These absorption lines are sensitive to the temperature, density, and velocity of the residual neutral hydrogen in the IGM, making the Lyα forest a premier tool for probing the IGM and its thermal state. Fig. 1.2 shows an example of Lyα forest quasar 3C273 at z=0.158 observed by the Hubble Space Telescope (HST) space telescope imaging spectrograph (STIS) E140M grating. The figure clearly illustrates how neutral hydrogen atoms in the IGM create multiple absorption lines as the photons emitted by the quasar traverse space towards us.

Various statistical properties of the Lyα forest are used to measure the IGM thermal state, including the power spectrum (Theuns et al. 2000, Zaldarriaga et al. 2001, McDonald et al. 2001, Walther et al. 2017, Khaire et al. 2019, Gaikwad et al. 2021), the flux probability density function (PDF) (Bolton et al. 2008, Viel et al. 2009, Lee et al. 2015), the transmission curvature (Becker et al. 2011, Boera et al. 2014), the wavelet decomposition of the forest (Theuns & Zaroubi 2000, Theuns et al. 2002, Lidz et al. 2010, Garzilli et al. 2012, Wolfson et al. 2021), the quasar pair phase angle distribution (Rorai et al. 2013; 2017), and the decomposition of the Lyα forest (Schaye et al. 1999, Ricotti et al. 2000, McDonald et al. 2001, Hiss et al. 2018). However, These measurements are typically performed using Lyα forest spectra from ground-based telescopes at $z > 1.6$, where the Lyα transition lies above the atmospheric cutoff ($\lambda \sim 3300$Å), which explains why there are currently very few measurements of the IGM thermal state at redshift below such limit (i.e. $z < 1.6$), which is, however, an essential epoch for galaxy formation. Currently, the only measurements of the IGM thermal state for $z < 1.5$ come from Ricotti et al. (2000, see. Fig. 1.4) at $z = 0.1$, utilizing a dataset with very limited size (34 Lyα absorption lines). Due to the limited size of this dataset, the associated error margins on measured temperatures are substantial, with $\sigma_{T_0} \gtrsim$ 10,000K, the large error margin implies that our understanding of the IGM thermal state at low-z remains imprecise.

As an important technique used in this work, the decomposition of the Lyα forest into the Voigt profiles detailed below. The intrinsic absorption line shape of the Lyα transition, dictated by the scattering cross-section, approximates a Lorentzian profile. Concurrently, the thermal motions within the absorbing clouds are described by a Maxwell-Boltzmann velocity distribution, which exhibits a Gaussian shape. Thus, the observed Lyα lines can be effectively modeled by a Voigt profile, a convolution of a Lorentz distribution and a Gaussian distribution:

$$V(x) = \int G(y;\sigma)L(x-y;\gamma)dy \tag{1.2}$$

where

$$L(x;\gamma) = \frac{\gamma}{\pi(\gamma^2) + x^2}, \tag{1.3}$$

is the Lorentz profile, with γ being the scale factor, and the

$$G(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp(\frac{-x^2}{2\sigma^2}), \tag{1.4}$$

is the Gaussian profile, where is the standard deviation.

For a transition line caused by intrinsic quantum effect, its line profile is:

$$\phi_\nu = \frac{4\gamma_{ul}}{16\pi^2(\nu-\nu_{lu})^2 + {\gamma_{ul}}^2}, \tag{1.5}$$

where the ν_{lu} is the transition frequency, and the γ_{ul} is a scale factor determined by the lifetime of the excited state. As for Lyα transition, the intrinsic width caused by the intrinsic profile is

$$(\Delta v)_{\mathrm{Ly}\alpha,\mathrm{FWHM}} = c\frac{(\Delta\nu)_{\mathrm{FWHM}}}{\nu_{ul}} = \frac{c\gamma_{ul}}{2\pi\nu_{ul}} \sim 0.0121 \text{ km/s}. \tag{1.6}$$

For thermal broadening, the distribution of velocities of the gas particles, $p(v)$, follows the Maxwell distribution, which is Gaussian,

$$p(v) = \frac{1}{\sigma_v\sqrt{2\pi}} \exp(\frac{-(v-v_0)^2}{b^2}), \tag{1.7}$$

where v_0 is the mean velocity of the gas particles, $b = \sqrt{2}\sigma_v$ is the *Doppler parameter* (or *Doppler width*), and $\sigma_v = \sqrt{k_{\mathrm{B}}T/m}$ is the standard deviation derived from the pure thermal motion of the particles. Here, m represents the particle mass, T is the gas temperature, and k_{B} is the Boltzmann constant. For typical Lyα transition caused by hydrogen atoms with $T \sim 10{,}000$ K, the Doppler width, $b \sim 13$ km/s.

Combing eqns. 1.5, 1.7, and 1.2, we can express the Voigt profile as:

$$\phi_{\mathrm{Voigt}}(\nu) = \int p(v)\phi_{\mathrm{intr}}(\nu)dv = \frac{1}{\sqrt{2\pi}\sigma_v}\int \frac{4\gamma_{ul}}{16\pi^2(\nu - (1 - v/c)\nu_{lu})^2 + {\gamma_{ul}}^2}\exp(\frac{-v^2}{2{\sigma_v}^2})dv. \tag{1.8}$$

Given the thermal broadening is much larger than the intrinsic broadening for typical Lyα lines, i.e., $b \gg (\Delta v)_{\mathrm{Ly}\alpha,\mathrm{FWHM}}$, we can see $\gamma_{ul} \ll \frac{\sigma_v \nu_{ul}}{c}$. The Lorentz profile can thus be approximated by a δ function, i.e.,

$$\frac{4\gamma_{ul}}{16\pi^2(\nu - (1 - v/c)\nu_{lu})^2 + {\gamma_{ul}}^2} \approx \delta(v - c(1 - \nu/\nu_{ul})) \tag{1.9}$$

Therefore, the Voigt profile becomes:

$$\phi_{\mathrm{Voigt}}(\nu) = \frac{1}{\sqrt{2\pi}\sigma_\nu}\exp(\frac{-(\nu - \nu_{ul})^2}{2{\sigma_\nu}^2}), \tag{1.10}$$

where $\sigma_\nu = \sigma_v \nu_{ul}/c$. This approximation is generally valid for typical low-redshift Lyα forest observations ($z < 3$), which are optically thin. However, it is not appropriate for high column-density absorbers, as treating the intrinsic profile as a delta function fails to account for the significant wings of the Lorentzian profile. These high column-density absorbers are less common in the Lyα forest and are, therefore, beyond the scope of this thesis.

To model the shape of an absorption profile observed in the Lyα forest spectra, we begin by calculating the optical depth of the Lyα transmission, $\tau_{\mathrm{Ly}\alpha}$, which can be calculated by:

$$\tau_{\mathrm{Ly}\alpha} = \sigma_{\mathrm{Ly}\alpha} N_{\mathrm{HI}} = \frac{\pi e^2}{m_e c} f_{\mathrm{Ly}\alpha} \phi_{\mathrm{Ly}\alpha} N_{\mathrm{HI}} \tag{1.11}$$

where $\sigma_{\mathrm{Ly}\alpha}$ is the cross-section for the Lyα transition, N_{HI} is the neutral hydrogen column-density, e is the electron charge, m_e is the electron mass, and $f_{\mathrm{Ly}\alpha}$ is the oscillator strength of the Lyα transition. Given the Voigt profile, Eq. 1.10, we thus have

$$\tau_{\mathrm{Ly}\alpha} = \tau_{0,\mathrm{Ly}\alpha} \exp(-v^2/b^2), \tag{1.12}$$

where the velocity v corresponds to the frequency shift $v = c(\nu_{\mathrm{Ly}\alpha} - \nu)/\nu_{\mathrm{Ly}\alpha}$, and the τ_0 is the optical depth at line center,

$$\tau_{0,\mathrm{Ly}\alpha} = \sqrt{\pi} \frac{e^2}{m_e c} \frac{f_{\mathrm{Ly}\alpha} \lambda_{\mathrm{Ly}\alpha} N_{\mathrm{HI}}}{b} \simeq 0.758 (N_{\mathrm{HI}}/10^{13}\mathrm{cm}^{-2})(10\mathrm{km\ s}^{-1}/b). \tag{1.13}$$

Given the optical depth τ, the observed transmission is then

$$T = F/C = \exp(-\tau), \tag{1.14}$$

where F is the observed flux and C is the spectrum continuum.

Given the above, an absorption profile in a Lyα spectrum can be parameterized by the neutral hydrogen column density N_{HI}, the thermal broadening Doppler parameter b, and its position in redshift space, z. The statistics of the Lyα forest line parameters $\{b, N_{\mathrm{HI}}\}$, therefore, can be used to constrain the thermal state of the IGM. In this thesis, the joint 2D b-N_{HI} distribution is employed to constrain the IGM thermal state. The method is elaborated further in §2.

1.3 The b-parameter Discrepancy at low-z

At high redshifts ($z > 2$), our theoretical model of the IGM, particularly concerning its thermal and ionization state, aligns remarkably well with high-resolution Lyα forest observations (e.g. Bolton & Haehnelt 2007, Becker et al. 2011, Becker & Bolton 2013a, Hiss et al. 2018, Walther et al. 2019a, Gaikwad et al. 2021). This concordance has established confidence in these IGM models, leading to further research utilizing the Lyα forest, e.g., probing cosmology (e.g. Busca et al. 2013), measuring neutrino masses (e.g. McDonald 2006, Yèche et al. 2017, Garny et al. 2021), and testing alternate models of dark matter (e.g. Viel et al. 2013, Armengaud et al. 2017, Palanque-Delabrouille et al. 2020, Iršič et al. 2024). However, in the recent decade, low-z ($z < 0.5$) observations of the IGM facilitated by the Cosmic Origin Spectrograph (COS) on board the Hubble Space Telescope (HST) have introduced discrepancies that question our current understanding of the IGM, and signal an intriguing issue with otherwise stellar concordance between theory and observations.

This particularly pressing issue in the low-z IGM is the discrepancy between the observed and simulated distributions of the Dopper width, b, of the Lyα forest(Gaikwad et al. 2017b, Viel et al. 2017, Nasir et al. 2017, Bolton et al. 2022b). Such a discrepancy is illustrated in Fig. 1.3, in which the b-parameters from three simulations Nyx, Illustris(Genel et al. 2014), and IllustrisTNG(Weinberger et al. 2017) are compared with those in the low-z Lyα forest quasar spectra Danforth et al. (2016). All these b parameters are obtained from Lyα forest decomposition using automated fitting program `VPFIT`, which will be further explained in §2.1.4. The b-parameter is observed to peak at $\simeq 32$ km/s, which is significantly higher than predicted by simulations ($\simeq 21$ km/s). This discrepancy holds regardless of whether the simulations account for galaxy formation with various degrees of feedback (e.g Illustris and IllustrisTNG) or without any galaxy forma-

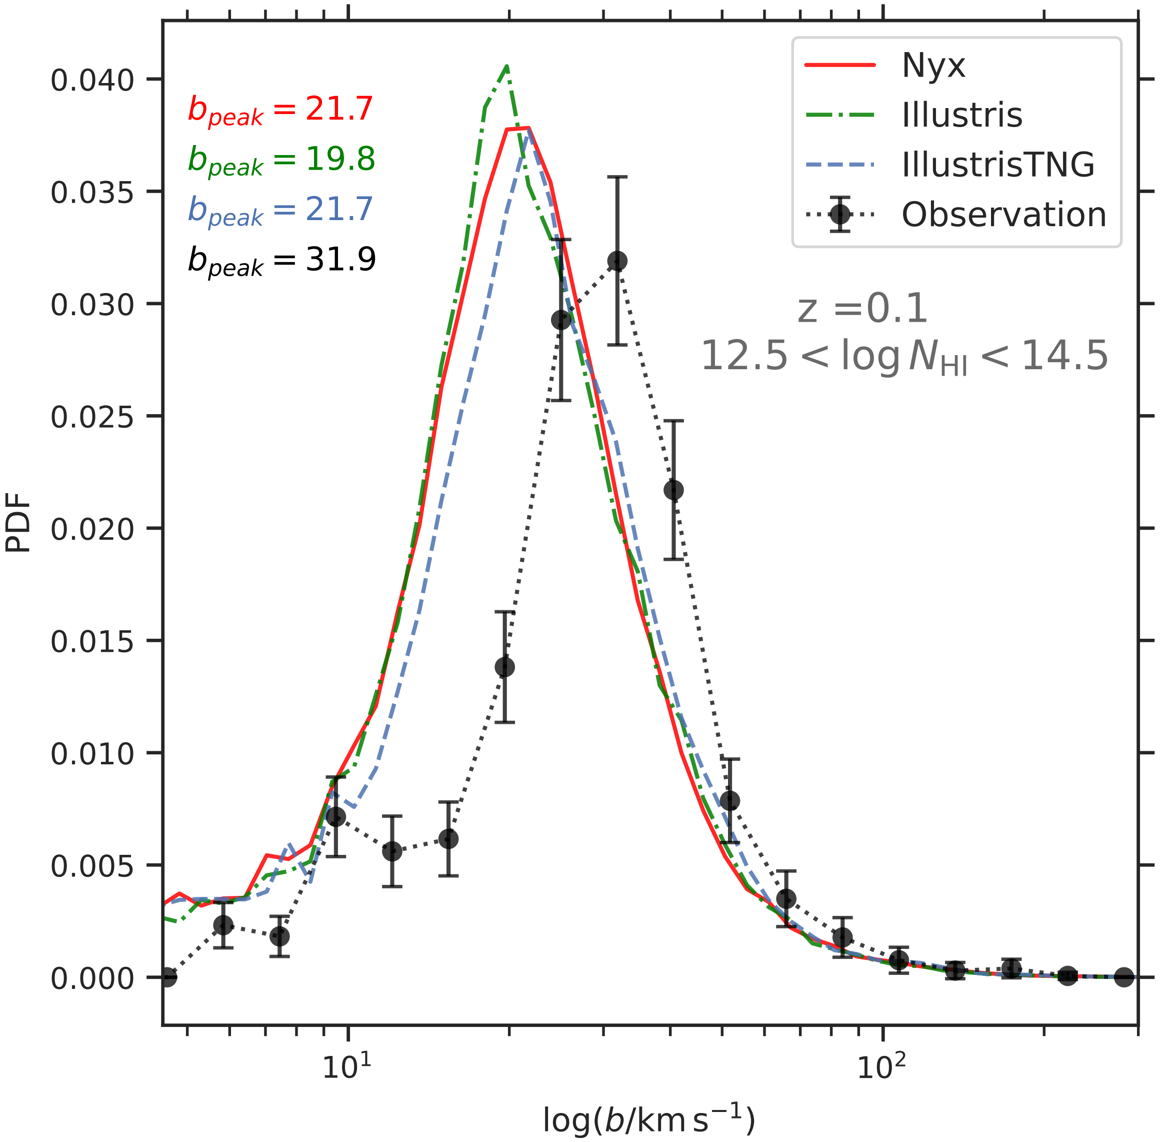

Figure 1.3: Comparisons of distributions of b parameters between the observation and three simulations Nyx, Illustris, and IllustrisTNG. The observations are obtained by fitting a sample of COS spectra (Danforth et al. 2016) across $0.06 < z < 0.16$. The b distributions from three simulations are obtained by VP-fitting mock spectra forward-modelled based on the same COS spectra as the observation.

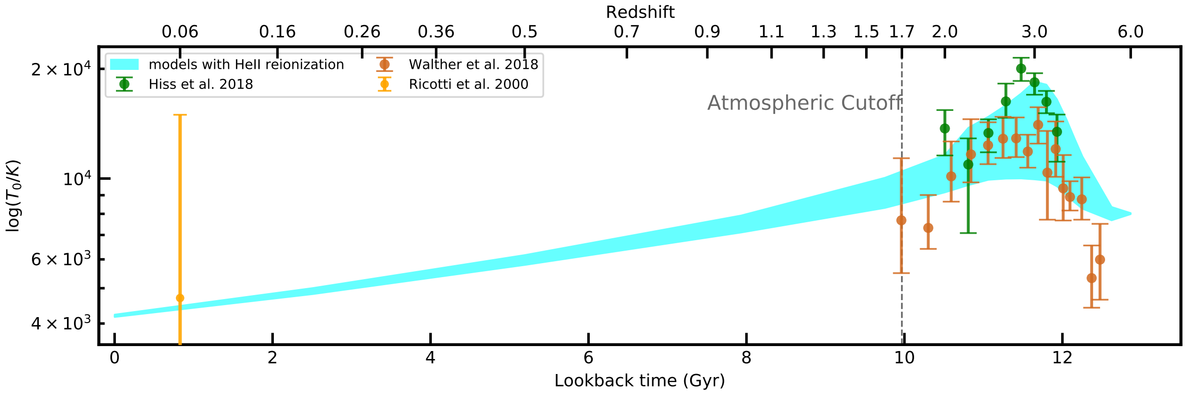

Figure 1.4: The IGM T_0 evolution history based on theory and previous observations. The blue-shaded region represents the range spanned by T_0 and γ from hydrodynamical simulations of a large family of different He II reionization models. Measurements from other studies at are displayed in different colours.

tion (Nyx in Fig. 1.3). The larger line widths in observations as compared to simulations suggest either a significant source of additional turbulence in the IGM of the order of ~ 10 km/s, which is not captured by current simulations (e.g Nasir et al. 2017, Bolton et al. 2022b) or, alternatively, the low-z IGM is substantially hotter than expected (e.g Viel et al. 2017, Nasir et al. 2017), with $T_0 > 10,000$K at $z = 0.1$, at least twice the temperature predicted by theoretical models.

In fact, one of the fundamental predictions of the theory of the IGM is that, a few hundred million years after the completion of He II reionization (i.e., at $z \sim 2.5$) the IGM should cool down primarily because of the adiabatic cooling resulting from Hubble expansion as illustrated in Fig. 1.4. This figure illustrates the evolution of IGM temperature at the mean density, T_0, where the cyan-coloured swath, generated from hydrodynamical simulations with a diverse set of He II reionization models, shows a remarkable convergence towards a consistent T_0 at $z < 1$. This convergence illustrates that the IGM quickly loses memory of the thermal impact of He II reionization, and asymptotes towards a temperature $T_0 \simeq 5000$ K. However, as discussed in §1.2, the lack of observational data over a span of 10 billion years in cosmic time results in a void in the IGM thermal history, leaving the anticipated cooling of the IGM at low-z unconfirmed.

1.4 The UVB and 'Photon Under-Production Crisis'

In addition to the discrepancy in b-parameters, the low-z IGM presents another puzzle: the nature of the UVB, characterized by the H I photoionization rate, $\Gamma_{\rm HI}$, which directly affects the abundance of Lyα absorbers in the low-z IGM as well as crucial for studying the circum-galactic medium (e.g Lehner et al. 2013, Hussain et al. 2017, Chen et al. 2017, Wotta et al. 2019, Acharya & Khaire 2022). A notable deviation between the $\Gamma_{\rm HI}$ deduced from the Lyα forest at $z \sim 0.1$ and the forecasts from previous UVB synthesis models (e.g Haardt & Madau 2012, Faucher-Giguère et al. 2009) lead Kollmeier et al. (2014) to introduce the problem of a "photon under-production crisis", which suggest that the measured $\Gamma_{\rm HI}$ exceeds the value predicted by UVB synthesis models, which is shown in the Fig. 1.5. However, such a result has however not been confirmed by other studies (Shull et al. 2015, Gaikwad et al. 2017a, Fumagalli et al. 2017, Khaire et al. 2019) and recent UVB models (Khaire & Srianand 2015; 2019b, Puchwein et al. 2019, Faucher-Giguère 2020).

The recent UVB models agree to the extent that the low-z $\Gamma_{\rm HI}$ measurements favour UVB dominated by H I ionizing photons from quasars alone and the fraction of ionizing photons from galaxies at $z < 2$ is negligibly small (Khaire & Srianand 2019b, Puchwein et al. 2019, Faucher-Giguère 2020). However, at higher redshifts, $z > 3$, a substantial increase in the ionizing escape fraction from galaxies from less than one percent to 15-20 percent is needed (Khaire et al. 2016) even in the presence of a high fraction of low-luminosity quasars claimed to be present at high-z (Khaire 2017, Finkelstein et al. 2019). This transition of escape fraction hinges only on the $\Gamma_{\rm HI}$ measurements at $z > 2$ and $z < 0.5$ whereas there are no measurement of $\Gamma_{\rm HI}$ at $0.5 < z < 1.8$, with a substantial

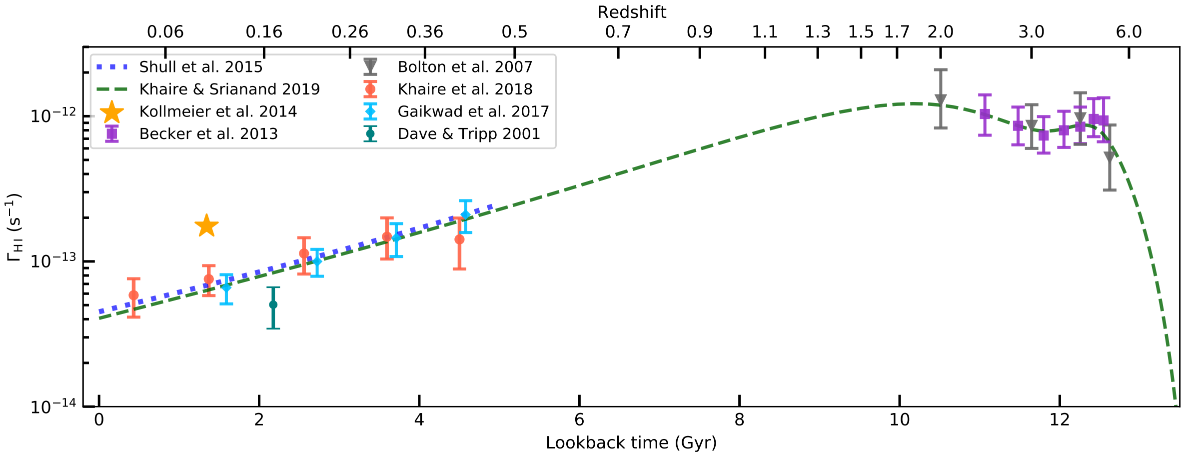

Figure 1.5: The $\Gamma_{\rm HI}$ evolution history based on theory and previous observations. The higher-than-expected value measured by Kollmeier et al. (2014) is shown as the orange star.

void of almost five billion years of cosmic time. A part of this lack of measurement, besides the limited data from HST at these redshifts, is caused by the potential degeneracy between the IGM thermal and ionization state. To overcome this problem, for $z < 0.5$ $\Gamma_{\rm HI}$ measurements, previous studies (Gaikwad et al. 2017a, Khaire et al. 2019) leveraged either post-processing simulations to generate the thermal histories (Gaikwad et al. 2018) or a huge grid of Nyx simulations (Walther et al. 2017) performed with different thermal histories of the IGM. It is important to recognize, a full description of the Lyα forest depends on three parameters T_0, γ, and $\Gamma_{\rm HI}$. Degeneracies among these variables require that any reliable data-model comparison must adopt a careful statistical inference procedure.

1.5 Outline of this Thesis

In this thesis, I investigate the thermal and ionization state of the IGM across seven redshift bins across $0 < z < 1.5$, covering approximately 10 billion years of cosmic history. To address the challenges outlined earlier, I utilize a novel inference method based on neural networks to jointly measure the thermal and ionization state of the IGM based on the 2D b-N_{HI} distribution. I apply this method to 94 archival HST COS and STIS quasar spectra distributed across the seven redshift bins, yielding a comprehensive evolutionary history of the IGM thermal and ionization state. This analysis suggests that the IGM may be significantly hotter at low-z than previously expected, with this discrepancy possibly emerging around $z \sim 1$. Additionally, I have carefully verified that this discrepancy is not attributable to feedback from AGN activities or galaxy formation processes. The detailed statistical methods, data handling, simulations, and results of these analyses are extensively discussed across three separate chapters in this thesis.

In §2, the inference methodology is extensively detailed and rigorously evaluated. This section also presents the corresponding hydrodynamical simulations and relevant post-processing techniques, such as generating mock spectra, forward-modeling based on observational data, and adjustments to the ultraviolet background (UVB).

In §3, I carry out the first-ever measurements of the IGM thermal and ionization state at three redshift bins at $z = 1, 1.2$ and 1.4, based on 12 archival HST STIS quasar spectra. The results suggest that the IGM might be hotter than expected at $z = 1$, while the thermal and ionization state of the IGM are coherent with the theoretical model at $z = 1.2$ and 1.4. However, these conclusions are tentative due to the considerable uncertainties.

In §4, I conduct measurements of the IGM's thermal and ionization state at redshift bins of $z = 0.1, 0.2, 0.3$, and 0.4, based on 82 archival HST COS quasar spectra compiled by Danforth et al. (2016). The results consistently indicate that the IGM is substantially hotter than predicted by theoretical models at low-z. Alternative explanations other than the IGM being too hot are carefully investigated, such as turbulence or resolution effects.

In §5, I explore the impact of the WHIM on measurements of the low-z IGM thermal state based on the b-N_{HI} distribution of the Lyα forest, and check whether the discrepancy in b parameter can be caused by the feedback. I conclude that the b-N_{HI} distribution is not sensitive to the WHIM under realistic conditions, and while the Lyα forest hardly probes feedback, it is significantly influenced by the photoionization rate used in the simulation.

Throughout this thesis, we write log in place of $\log_{10}$. Cosmology parameters used in this study ($\Omega_m = 0.319181, \Omega_b h^2 = 0.022312, h = 0.670386, n_s = 0.96, \sigma_8 = 0.8288$) are taken from Planck Collaboration et al. (2014) .

1.6 Permissions and Attributions

The content of §2 and Appendix A is the result of a collaboration with Vikram Khaire, Joseph F. Hennawi, Michael Walther, Hector Hiss, Justin Alsing, Jose Onorbe, Zarija Lukic and Frederick Davies, and has previously appeared in Hu et al. (2022).

The content of §3 and Appendix B is produced in collaboration with Vikram Khaire, Joseph F. Hennawi, Todd M. Tripp, Jose Onorbe, Michael Walther, and Zarija Lukic. It has previously been published in Hu et al. (2023b).

The content of §4 is done in collaboration with Vikram Khaire, Joseph F. Hennawi, Todd M. Tripp, Jose Onorbe, Michael Walther, and Zarija Lukic. This study is currently being prepared for publication.

The content of §5 and Appendix C is created in collaboration with Vikram Khaire, Joseph F. Hennawi, Jose Onorbe, Michael Walther, Zarija Lukic, and Frederick Davies. It has previously been published in Hu et al. (2023a).

Chapter 2

Measuring the Thermal and Ionization State of the Low-z IGM using Likelihood-Free Inference

In this Chapter, we follow the method for measuring the IGM thermal state based on Voigt profile decomposition of the Lyα forest (Schaye et al. 1999, Ricotti et al. 2000, McDonald et al. 2001). In this approach, a transmission spectrum is treated as a superposition of multiple discrete Voigt profiles, with each line described by three parameters: redshift $z_{\rm abs}$, Doppler broadening b, and neutral hydrogen column density $N_{\rm HI}$. By studying the statistical properties of these parameters, i.e. the b-$N_{\rm HI}$ distribution, one can recover the thermal information encoded in the absorption profiles. The majority of past applications of this method constrained the IGM thermal state by fitting the low-b-$N_{\rm HI}$ cutoff of the b-$N_{\rm HI}$ distribution (Schaye et al. 1999; 2000, Ricotti et al. 2000, McDonald et al. 2001, Rudie et al. 2012a, Bolton et al. 2014, Boera et al. 2014, Garzilli et al. 2015; 2018, Rorai et al. 2018, Hiss et al. 2018). The motivation for this approach is that the Lyα lines are broadened by both thermal motion and non-thermal broadening resulting from combinations of Hubble flow, peculiar velocities and turbulence. By isolating the narrow lines in the Lyα forest that constitutes the lower-cutoff in b-$N_{\rm HI}$ distributions, of which the line-of-sight component of non-thermal broadening is expected to be zero, the broadening should be purely thermal, thus allowing one to constrain the IGM thermal state. However, this method has three crucial drawbacks. First, the IGM thermal state actually impacts all the lines besides just the narrowest lines. Therefore, by restricting attention to data in the distribution outskirts, this approach throws away information and reduces the sensitivity to the IGM thermal state significantly(Rorai et al. 2018, Hiss et al. 2019). Second, in practice, determining the location of the cutoff is vulnerable to systematic effects, such as contamination from the narrow metal lines (Rorai et al. 2018, Hiss et al. 2018). Lastly, the results from this approach critically depend on the choice of low-b cutoff fitting techniques, where different techniques might result in inconsistent T_0 and γ measurements (Rorai et al. 2018, Hiss et al. 2018).

To overcome these limitations, Hiss et al. (2019) developed a new approach to measure the IGM thermal state from the full b-$N_{\rm HI}$ distribution based on density estimation and Bayesian analysis. I further advance the b-$N_{\rm HI}$ distribution emulation by employing a novel density estimation technique based on machine learning, namely Density-Estimation Likelihood-Free Inference (DELFI) (see Papamakarios & Murray 2016, Alsing et al. 2018, Papamakarios et al. 2018, Lueckmann et al. 2018, Alsing et al. 2019). In addition, I augment the likelihood function to take into account the absorber number density $\mathrm{d}N/\mathrm{d}z$, making the improved method far more sensitive to the photoionization rate of hydrogen $\Gamma_{\rm HI}$ sourced by the UV background.

In this chapter, I introduce our new method, demonstrate its robustness, and perform an analysis using realistic mock datasets to illustrate the sensitivity to IGM parameters. The inference is built on a suite of cosmological hydrodynamic simulations with different thermal parameters at redshift $z \sim 0.1$. While this method can be applied to the Lyα

forest at any redshift where the opacity is low enough to make it amenable to Voigt profile decomposition (e.g. $z \lesssim 3.4$, see Hiss et al. 2018), we choose to focus on $z \sim 0.1$ because we want to quantify the sensitivity of archival *Hubble Space Telescope* spectra, so as to perform the first measurements of the IGM thermal state at $z < 1.6$ in future work. Such a measurement would directly test the prediction that the IGM cools down at low-z, which has been challenged by recent observations. To this end, we run a set of cosmological hydrodynamic simulations with different thermal parameters at redshift $z \sim 0.1$, from which we create mock datasets with the same properties as the Danforth et al. (2016) low redshift Lyα forest dataset observed with the *Cosmic Origins Spectrograph* (COS, Green et al. 2012) on the HST. I demonstrate that the method applied to such a dataset can reliably and accurately determine the thermal state of the IGM.

This chapter is structured as follows. In §2.1 I introduce our hydrodynamic simulations, parameter grid, and data processing procedures, which include generating Lyα forest from simulation, forward-modeling and our method to fit Voigt profiles (`VPFIT`). In §2.2 we present our inference algorithm, including likelihood, emulators, inference results, and a set of inference tests. Finally, we discuss these results and summarize the highlights of this study in §2.3.

2.1 Simulations

A set of Nyx cosmological hydrodynamic simulations (see Lukić et al. 2015, Almgren et al. 2013) is used to model the low-redshift IGM. Nyx is a massively-parallel, cosmological simulation code primarily developed to simulate the IGM. [1] In Nyx simulations, the evolution of dark matter is traced by treating dark matter as self gravitating La-

---

[1] Nyx simulation is able to run with Adaptive Mesh Refinement (AMR). However, the AMR feature is not used in this work, since this work focus on the Lyα forest, which distribute nearly the entire simulation domain rather than isolated concentrations of matter where AMR is more effective.

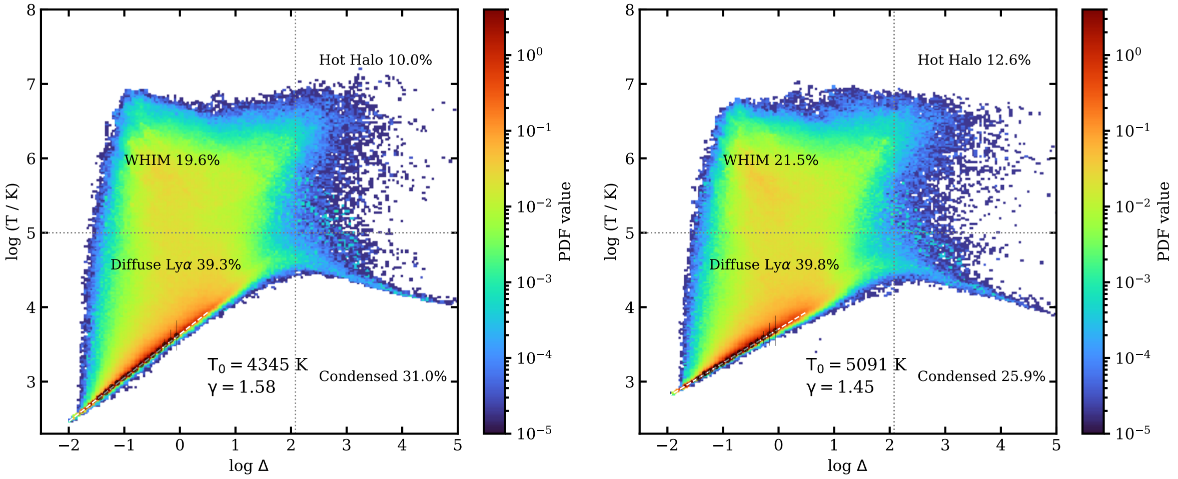

Figure 2.1: Temperature-density (T-Δ) distribution for the IGM gas in two different models from Nyx simulation. White dashed lines is the power-law fit to the T-Δ relation, and legends show the best fit values of T_0 and γ. Dotted T $= 10^5$ K lines divide the phase diagram into hot and cold region, while only cold gas is used for the fitting. Density peaks ($\log T_{\rm peak,i}$, $\log \Delta_i$) for each bin are plotted as black dot, and 1-$\sigma_{T,i}$ error bars are shown as black bars. The left-hand panel shows a model with $T_0 = 4345$ K and $\gamma = 1.58$, and the right-hand panel shows a model with $T_0 = 5091$ K and $\gamma = 1.45$. The density weighted gas phase fractions are shown in the annotation.

grangian particles, while baryons are modeled as as ideal gas on a uniform Cartesian grid following an Eulerian approach. The Eulerian gas dynamics equations are solved following a second-order piece-wise parabolic method (PPM), which captures shock waves accurately.

Nyx includes the main physical processes relevant for modeling the Lyα forest. First of all, gas in the Nyx is assumed to have a primordial composition with a hydrogen mass fraction of 0.76, and helium mass fraction of 0.24 and zero metallicity. The recombination, collisional ionization, dielectric recombination, and cooling are implemented based on prescriptions given in Lukić et al. (2015). Nyx keeps track of the net loss of thermal energy resulting from atomic collisional processes and takes into account the inverse Compton cooling off the microwave background. Ionizing radiation is modeled by a spatially homogeneous but time-varying ultraviolet background radiation field (from Haardt

& Madau 2012) that changes with redshift, while assuming all cells in the simulation are optically thin. We later make the UV background a free parameter for generating Lyα forest in post-processing (See §2.1.2). Since Nyx simulations are developed mainly to study the IGM, no feedback or galaxy formation processes are included, significantly reducing the computational requirement allowing us to run a large ensemble of simulations varying the thermal parameters (see 2.2.3).

Each Nyx simulation used in this study is initialized at $z = 159$ and evolves down to $z = 0.03$ in a $L_{\rm box} = 20$ cMpc/h simulation domain, using $N_{\rm cell} = 1024^3$ Eulerian cells and 1024^3 dark matter particles. The box size is chosen as the best compromise between computational cost and the need to be converged at least to $< 10\%$ on small scales (large k). More discussion about resolutions and box sizes can be found in Lukić et al. (2015). We also performed box size convergence tests at low redshift as explained in appendix A.4.

2.1.1 Thermal Parameters and Simulation Grid

To model the IGM with different thermal state, we use part of the publicly available Thermal History and Evolution in Reionization Models of Absorption Lines (Thermal History and Evolution in Reionization Models of Absorption Lines (THERMAL))[2] suite of Nyx simulations (see Hiss et al. 2018, Walther et al. 2019a). We make use of in total 48 models with different thermal histories, and for each model, we generate a simulation snapshot at $z = 0.1$, from which we measure the thermal state [$\log T_0$,γ]. The thermal grid is illustrated in the left panel of Fig.2.2, which shows that $\log(T_0/\mathrm{K})$ spans from 3.2 to 3.95, and γ ranges from 0.86 to 2.41. Here different thermal histories are achieved by artificially changing the photoheating rates (ϵ) following the method presented in Becker

---

[2] For details of THERMAL suite, see http://thermal.joseonorbe.com

et al. (2011). In this method, ϵ is treated as a function of overdensity, i.e.

$$\epsilon = \epsilon_{\mathrm{HM12}} A \Delta^{B}, \tag{2.1}$$

where ϵ_{HM12} represents the photoheating rate per ion tabulated in Haardt & Madau (2012), and A and B are parameters that are varied to obtain different thermal histories. It is noteworthy that the thermal state tends to converge towards low redshifts due to the cooling and other physical processes in the evolution, and it is therefore difficult to generate models with a uniform grid of T_0 and γ (for more details, see Walther et al. 2019a). Moreover, it is especially challenging to generate models with low $T_0 (< 10^{3.5}$ K) and high $\gamma (> 1.9)$ at low-z, because when one reduces the photoheating rates to obtain lower T_0, the cooling rate from Hubble expansion dominates, and γ asymptotically approaches values near 1.6 (see McQuinn & Upton Sanderbeck 2016). As a result, the T_0-γ grid has an irregular shape, and there are no models in the high γ low T_0 regions. In addition, such an irregular T_0-γ grid is also a result of the original grid of the THERMAL suite, which is driven by the high-z thermal state analysis in Walther et al. (2019a) that obtains relatively high temperatures.

To measure the thermal state for each of the 48 models, we fit temperature-density (T-Δ) relation (see Eq. 1.1) to the temperatures and densities in the simulation domain. While fitting the T-Δ relationship, we noticed broader distributions of the IGM temperatures in low redshift ($z < 0.5$) compared to high redshift ($z > 3$). Examples of low-z IGM temperature-density distributions are illustrated in Fig. 2.1, where we show 2D histograms of T-Δ of gas in each cell for two of our simulations on the THERMAL grid at $z = 0.1$. The gas cells are divided into four phases depending on the temperature and density, and the density-weighted gas phase fractions are shown in the legends of the figure, where the diffuse Lyα phase representing the densities and temperatures probed

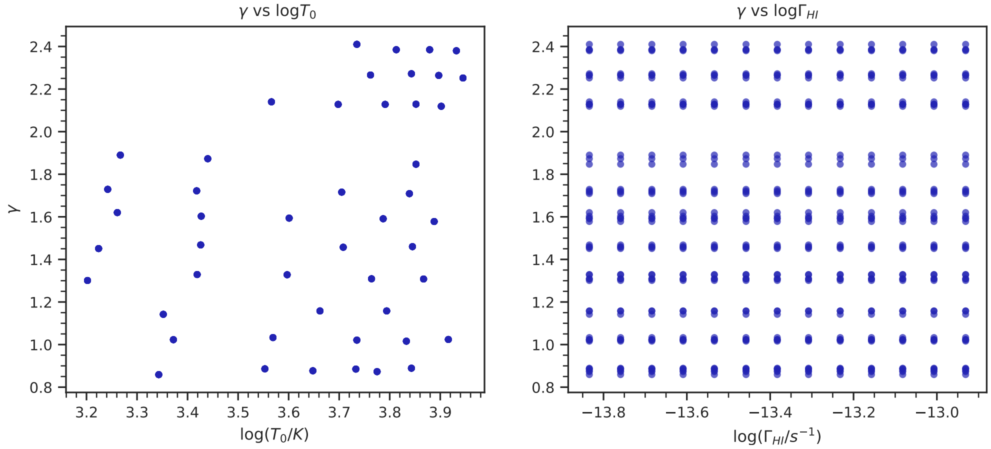

Figure 2.2: Thermal grid (blue circles) from snapshots of hydrodynamic simulations of the THERMAL suite at $z = 0.1$. The left-hand panel is the γ - T_0 grid, whose shape is determined by the parameters of thermal grid at and the evolution of the thermal state of the IGM. The right-hand panel is γ - Γ_{HI} grid, showing the thirteen Γ_{HI} values for each point on the 2D γ - T_0 grid.

by the Lyα forest occupies about 40% of the total gas mass, while this percentage can be up to about 80% or higher at high-z. Therefore at high-z most of the gas lies on or around the T-Δ power-law relation. Whereas the high-temperature low-density WHIM phase is negligible at high-z, it appears significantly at low-z, resulting in puffy-looking gas distribution around the T-Δ power-law (see Fig. 2.1), which makes T-Δ power-law fitting non-trivial at low-z.

We address this issue by implementing an improved fitting procedure following Villasenor et al. (2021). First, we extract the temperature T and the overdensity Δ for each cell of a simulation and then select gas with $-1.5 < \log\Delta < 0$ and $T < 10^5K$ to avoid regions significantly deviant from the expected power-law T-Δ relationship. Afterward, we divide the selected region into 15 equal-width bins in $\log\Delta$, where the overdensity $\log\Delta_i$ for each bin i is given by the median value of overdensity in the bin. Here we define the bin temperature $\log T_i$ to be the maximum of the marginal temperature distribution $P(\log T\,|\,\log\Delta_i)$ and its effective 1-$\sigma_{T,i}$ interval to be 1/2 of the temperature range containing the 68% (16% $\sim$ 84%) highest probability density. The temperature-density relationship Eq. (1.1) is then fitted using a least squares linear fit on these ($\log \Delta_i, \log T_\mathrm{i}$) pairs weighted by $1/{\sigma_\mathrm{T,i}}^2$. Examples of temperature-density relationships for two models in our thermal grid are illustrated in Fig.2.1. Power-law fits of the T-Δ relationship of our simulations are shown as white dashed lines while their values are given in the legends texts. The peak temperature in each $\log \Delta$ bin ($\log T_\mathrm{peak,i}$,$\log \Delta_i$) are plotted as black dots and 1-$\sigma_{T,i}$ error bars are also shown. Left panel shows a model with $T_0 = 4353$ K and $\gamma = 1.58$, while right panel shows another model with $T_0 = 5091$ K and $\gamma = 1.45$. Finally, as will be discussed later in §2.1.2, we let the HI photoionization rate Γ_HI be a free parameter when generating Lyα forest skewers from our simulations. As such, we add an additional parameter $\log \Gamma_\mathrm{HI}$ to our thermal grid, extending it to [$\log T_0$, γ, $\log \Gamma_\mathrm{HI}$]. The value of Γ_HI we used in this study spans from $\log(\Gamma_\mathrm{HI}/\mathrm{s}^{-1})$ = -13.834 to -12.932 in logarithmic steps of 0.075 dex, which gives 13 values in total (see right-hand panel of Fig.2.2). In total, the 3D thermal grid consists of $48 \times 13 = 624$ models.

2.1.2 Skewers

We generate simulated Lyα spectra by calculating the Lyman-α optical depth (τ) array along the line-of-sight, which hereafter will be referred as skewers for simplicity. For each model on the thermal grid, a set of 60,000 skewers are constructed parallel to the x, y, z axes of the simulation box (20,000 skewers in each direction). For each cell on these skewers, we extract properties needed for optical depth calculation, including temperature T, overdensity Δ, and the velocity along the line-of-sight v_z. The hydrogen neutral fraction x_HI, which is also needed to generate the synthetic Lyα forest skewers, is calculated by assuming ionization equilibrium while considering both collisional ionization and photoionization. Here the collisional ionization rate is computed based on the

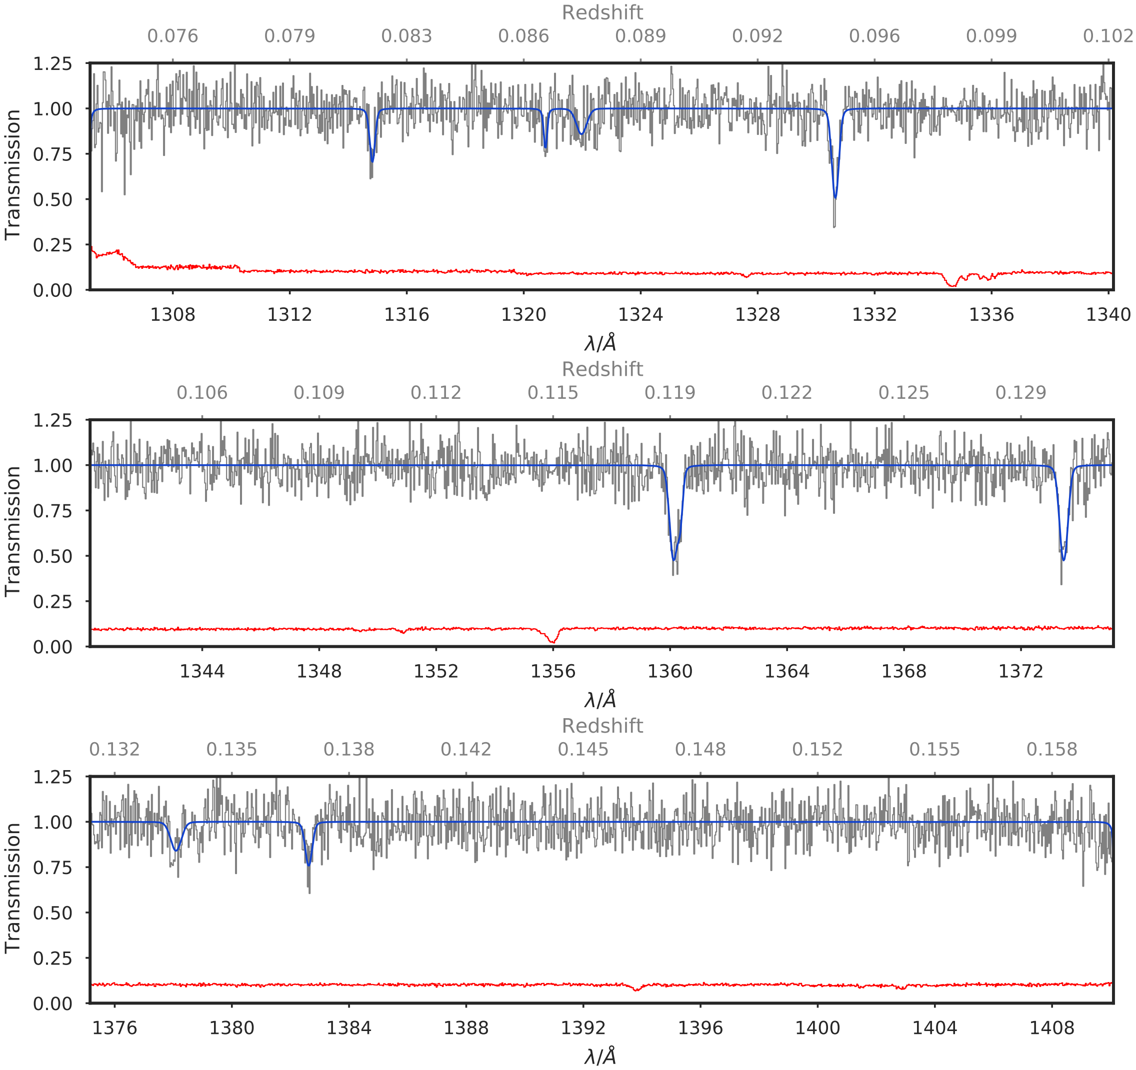

Figure 2.3: Illustration of one of our forward modeled spectra from Nyx simulation. The simulated raw spectrum is shown in gray, while a model spectrum based on `VPFIT` line fitting (described in § 2.1.4) is shown in blue and the noise vector is plotted in red. This particular spectrum is forward-modeled in order to model the instrumental effect and noise properties of one of the HST COS spectra in Danforth et al. (2016) low redshift dataset.

gas temperature T. Whereas the photoionization rate Γ_{HI} is set to be a free parameter in the post-processing of the simulation. Since Nyx does not model radiative transfer, we approximate the self-shielding of the UV background for optically thick gas following the method given by Rahmati et al. (2013), which amounts to attenuating Γ_{HI} for cells containing dense gas.

Given x_{HI}, T, Δ, v_z, and Γ_{HI}, we then calculate the optical depth τ in redshift space by summing contributions from all cells in real space along the line-of-sight following the full Voigt profile approximation by Tepper-García (2006). Then $F = e^{-\tau}$ gives us continuum normalized flux of Lyα forest along with skewers. Lastly, we redo the procedure described above for each Γ_{HI} value to recalculate the skewers. More specifically , we do not re-scale the τ to obtain skewers for a different Γ_{HI}, which is the standard procedure at higher redshifts. This is because, whereas the high-z IGM is predominantly photoionized, there is significantly more shock-heated WHIM gas at low-z, rendering the contribution from collisional ionization important as shown by Khaire et al. (2019) for the case of Lyα flux power spectrum. Although, it may not be essential for studying Lyα forest absorption lines, to be more precise we recalculate skewers for each value of Γ_{HI}.

2.1.3 Forward Modeling of Noise and Resolution

As discussed in §1, we are interested in understanding the constraints on the IGM achievable with realistic data. To this end, we generate mock datasets with properties consistent with the Danforth et al. (2016) low redshift Lyα forest dataset, which comprises 82 unique quasar spectra with S/N > 5 observed with the Cosmic Origins Spectrograph (COS) on the HST. To avoid proximity regions and contamination from lower Ly-β, we use rest-frame wavelength range $1050 - 1180$ Å to identify Lyα forest for each of these spectra. As a result, we select 34 of Danforth et al. (2016) quasar spectra covering the redshift range $0.06 < z < 0.16$ of our interest for the study, comprising a total redshift pathlength of $\Delta_{\rm ob} = 2.136$, which corresponds to our observational dataset for forward modelling. We choose this redsfhit bin to be the same as the redshift bin used for power spectrum calculation by Khaire et al. (2019) at $z = 0.1$ so that we can compare our future analysis with the results obtained with power spectrum measurements.

The COS has a nominal resolution $R \sim 20000$, which corresponds to roughly 15 km/s, and a non-Gaussian line spread function (LSF) exhibiting significantly broad Lorentzian wings, which could alter the shape of absorption lines on velocity scales larger than the resolution quoted above. For low-z IGM with temperatures at mean density $T_0 \sim 5000$ K, the b-values for pure thermal broadening (i.e. the narrowest lines in the Lyα forest) are about $10 \sim 20$ km/s, which means that the corresponding absorption features can not be fully resolved by COS. Thus, it is crucial to treat the instrumental effect carefully, including the peculiar shape of COS LSF. Therefore, we forward model noise and resolution to make our simulation results statistically comparable with the observation data.

In practice, we make use of tabulated COS LSF and noise vectors from Danforth et al. (2016) data. For any individual quasar spectrum from the observation dataset, we first stitch randomly selected simulated skewers without repetition to cover the same wavelength (in the rest frame $1050 - 1180$ Å) of that quasar and then rebin the skewers onto the pixels of the observed spectra. Then we convolve the simulated spectra with the HST COS line spread function (LSF) while taking into account the grating and life-time positions used for that specific data spectrum. Here the COS LSF is obtained from `linetools`[3] and is tabulated for up to 160 pixels in each direction. We interpolate the LSF onto the wavelengths of the mock spectrum (segment) to obtain a wavelength dependent LSF. Each output pixel is then modeled as a convolution between the input stitched skewers and the interpolated LSF for the corresponding wavelength. Afterward,

---

[3] For more information, visit https://linetools.readthedocs.io

the newly generated spectrum is interpolated to the wavelength of the selected COS spectra. The noise vector of the quasar spectrum is propagated to our simulated spectrum pixel-by-pixel by sampling from a Gaussian with $\sigma = \psi_i$, with ψ_i being the data noise vector value at the i^{th} pixel. In the end, a fixed floor is added to the error vector for all simulated spectra to avoid an artificial effect in post-processing, which will be discussed later in §2.1.4.

For each model, we generated 2000 forward-modeled spectra, corresponding to a total pathlength $\Delta z_{\mathrm{tot}} \sim 125$, from the 60,000 raw skewers[4], and fit voigt profiles to each line in the spectra to obtain the $\{b, N_{\mathrm{HI}}\}$ pairs for our dataset (as described in section § 2.1.4). For the purpose of illustration, an example of a forward-modeled spectrum is shown in Fig.2.3 where the simulated spectrum is shown in gray, the model spectrum based on `VPFIT` line fitting (see § 2.1.4) is in blue, and the noise vector in red.

2.1.4 VPFIT

To perform the analysis based on the b-N_{HI} distributions, we have to fit the Lyα lines in our simulated spectra to obtain a set of $\{b, N_{\mathrm{HI}}\}$ pairs for each model. To this end, we run a line-fitting program on our forward-modeled mock spectra to obtain a set of b-N_{HI} pairs for each simulation model in our thermal grid. In this work, we use the line-fitting program `VPFIT`, which fits a collection of Voigt profiles convolved with the instrument LSF to spectroscopic data (Carswell & Webb 2014)[5]. Here, we employ a fully automated `VPFIT` wrapper adapted from Hiss et al. (2018), which is built on the `VPFIT` version 10.2. The wrapper routine controls `VPFIT` with the help of the `VPFIT` front-end/back-end programs `RDGEN` and `AUTOVPIN` and fit our simulated spectra automatically.

`VPFIT` identifies lines automatically and fits each line with three parameters: the

---

[4] For each Nyx model, 2000 spectra needs about 20,000 raw skewers, i.e, we randomly pick 20,000 skewers from 60,000.

[5] VPFIT: `http://www.ast.cam.ac.uk/~rfc/vpfit.html`

absorption redshift z_{abs} of the line, its Doppler parameter b, and column density N_{HI}. `VPFIT` obtains these parameters for a collection of lines by minimizing the χ^2 between the data and the model spectrum generated from all the fitted lines. While fitting, `VPFIT` restrict b and N_{HI} to $1 \leq b(\mathrm{km/s}) \leq 300$ and $11.5 \leq \log(N_{\mathrm{HI}}/\mathrm{cm}^{-2}) \leq 18$, respectively. Our `VPFIT` wrapper allows us to fit spectra with a custom LSF[6]. Since we are working at $0.06 \leq z \leq 0.16$, the Ly α forest lies completely in the wavelength range covered exclusively by the COS G130M grating having a central wavelength 1300 Å. We fit our forward-modeled spectra with the same G130M LSF. Furthermore, the effective resolution of the grating also depends on the COS lifetime position during the observations, and they are also taken into account while running `VPFIT` as well as in forward modelling. An example of model spectrum generated by combining lines fitted using `VPFIT` is shown in Fig.4.2 as blue lines.

Moreover, we notice the presence of a significant number of absorption lines with very low Doppler-b parameters and low column densities N_{HI} after fitting mock as well as real data with high signal-to-noise ratios (SNR). These weak narrow absorption lines, however, are not seen in our simulated and forward-modeled spectra. Visual inspection of these lines indicates that they are spurious and introduced by `VPFIT` while attempting to fit artifacts due to flat-fielding, continuum placement, or errors in the data reduction. These lines are only introduced in spectra of the highest quality, where the extremely high signal-to-noise ratio (SNR) leads to over-fitting by `VPFIT`. To avoid this problem, a fixed floor of value 0.02 is added in quadrature to the error vector of the continuum normalized flux for all simulated spectra without adding additional noise to the normalized flux. With such a noise 'floor', these weak features are essentially removed from the `VPFIT` output.

---

[6] Although our `VPFIT` wrapper allows us to implement an LSF in `VPFIT`, only a single LSF can be used at once, i.e. the wavelength dependence can not be taken into accout. As such, for the input into `VPFIT` we use the LSF at the lifetime of the data and evaluated it at the central wavelength of the spectrum that we are trying to fit. Such treatment is applied to both observed (mock) spectra and stimulated spectra so as to make sure our statistics are not biased.

We find this floor value 0.02 via trial and error. In practice, this additional noise floor mainly removes lines with $\log(N_{\mathrm{HI}}/\mathrm{cm}^{-2}) < 12.5$ from our dataset, which is outside our limits used in likelihood calculations (which will be discussed in §2.2.2) and therefore not used in this study.

Furthermore, we follow the convention and apply another filter for both b and N_{HI} in this study, using only b-N_{HI} pairs in region $12.5 \leq \log(N_{\mathrm{HI}}/\mathrm{cm}^{-2}) \leq 14.5$ and $0.5 \leq \log(b/\mathrm{km\ s}^{-1}) \leq 2.5$ in our analysis (Schaye et al. 2000, Rudie et al. 2012a, Hiss et al. 2018). Such an limitation is chosen to include the b-N_{HI} distributions for all of our Nyx models while guaranteeing that the absorbers are not strongly saturated, which maximizes the sensitivity to IGM thermal state and minimizes the impact of poorly understood strong absorbers arising from the circumgalactic medium of galaxies.

2.2 Inference Algorithm

Hiss et al. (2019) introduced a Bayesian method to estimate the IGM thermal parameters from the joint b-N_{HI} distribution. In this paper, we adopt a similar approach while employing a new method for b-N_{HI} distribution emulation, namely Density-Estimation Likelihood-Free Inference (DELFI). In addition, we also include the absorber number density along the line-of-sight $\mathrm{d}N/\mathrm{d}z$ in our analysis, i.e. the number of absorption lines (in some range of b and N_{HI}) per unit path-length along the line-of-sight, which helps us to better constrain the UV background photoionization rate Γ_{HI}. The reason behind this is that the b-N_{HI} distribution is less sensitive to Γ_{HI} compared with thermal parameters T_0 and γ (see Fig.2.5 and §2.2.3), whereas the number density of absorbers (see Fig.2.4) depends strongly on Γ_{HI}. It is analogous to the fact that the mean flux of the Lyα forest is sensitive to Γ_{HI}. In this work, we emulate the $\mathrm{d}N/\mathrm{d}z$ using a Gaussian process emulator based on our simulations and employ it as a normalization factor in our likelihood function. More discussion about this modification is presented in §2.2.2 and Appendix A.1.

This section is organized as follows, we first introduce our new b-N_{HI} distribution emulator and then discuss the modifications to the likelihood function. Afterward, we investigate the relationship between thermal parameters and b-N_{HI} distribution in §2.2.3. Finally, we present our inference results in §2.2.4 and apply a series of inference tests to evaluate the statistical validity of our method in §2.2.5.

2.2.1 Emulating the b-N_{HI} distribution with DELFI

In this work, we build our b-N_{HI} distribution emulator following the density-estimation likelihood-free inference (DELFI) method (Papamakarios & Murray 2016, Alsing et al. 2018, Papamakarios et al. 2018, Lueckmann et al. 2018, Alsing et al. 2019), which turns inference into a density estimation task by learning the sampling distribution of the data as a function of the parameters. Compared with the previously used Kernel Density Estimation (KDE) method in Hiss et al. (2019), this method provides a flexible framework for conditional density estimation and does not implicitly apply a smoothing kernel to the training data. It hence is able to deliver higher-fidelity conditional density estimators given the same training data.

We make use of `pydelfi`[7] – the publicly available `python` implementation of DELFI based on neural density estimation (NDE)s and active learning (Alsing et al. 2019). `pydelfi` makes use of NDEs to learn the sampling conditional probability distribution $P(\mathbf{d} \mid \boldsymbol{\theta})$ of the data summaries $\mathbf{d}$, as a function of parameters $\boldsymbol{\theta}$, from a training set of simulated data summary-parameter pairs $\{\mathbf{d}, \boldsymbol{\theta}\}$. In this work, the parameters $\boldsymbol{\theta}$ are $\log T_0$, γ and $\log \Gamma_{\mathrm{HI}}$, and the data summaries $\mathbf{d}$ are $\log N_{\mathrm{HI}}$ and $\log b$[8], and the b-N_{HI} dis-

---

[7] See https://github.com/justinalsing/pydelfi

[8] `pydelfi` also has the option to apply different data compression methods (e.g., Alsing & Wandelt

tribution is considered as a conditional probability distribution $P(b\ ,\ N_{\rm HI}\ |\ T_0, \gamma, \Gamma_{\rm HI})$ learned from our simulations. More specifically, the b-$N_{\rm HI}$ distribution is modeled as a Masked Autoregressive Flow (MAF; Papamakarios et al. 2017) neural density estimator, which is constructed as a stack of five Masked Autoencoders for Density Estimation, (MADE; Germain et al. 2015), each with two hidden layers with 50 units each and tanh activation functions. The NDEs are trained by stochastic gradient descent. For more technical details about MAF and MADE neural network architectures see Germain et al. (2015), Papamakarios et al. (2017) and Alsing et al. (2019). To prevent over-fitting, the NDEs are weighted by their relative cross-validation losses and are trained with early-stopping (see Alsing et al. 2019 for details). For convenience, in this paper we will refer to the b-$N_{\rm HI}$ distribution emulator discussed above as the DELFI emulator.

As mentioned above, the DELFI emulator is trained on the data summary-parameter pairs $\{[\log T_0, \gamma, \log \Gamma_{\rm HI}], [b, \log N_{\rm HI}]\}$. For each model, we fit (VPFIT) 2000 simulated spectra, corresponding to a total pathlength $\Delta z_{\rm tot} \sim 123$, to get $\{b, N_{\rm HI}\}$ pairs for the model, and label these $\{b, N_{\rm HI}\}$ pairs with their simulation parameters $[\log T_0, \gamma, \log \Gamma_{\rm HI}]$. Our training set therefore consists of all these labeled $\{b, N_{\rm HI}\}$ pairs for all models on the thermal grid. Here we quantify the size of data by its total pathlength rather than number of lines[9], because the latter depends on the $\mathrm{d}N/\mathrm{d}z$ that varies among different models.

2.2.2 Likelihood function

Hiss et al. (2019) used only the shape of b-$N_{\rm HI}$ distribution to constrain IGM thermal

---

2018) and active learning methods to optimize the data and parameter space sampling. Here we do not exploit these features since we have pre-chosen our summary statistics and simulation grid (the b-$N_{\rm HI}$ distribution) at a fixed grid of thermal parameters.

[9] It means that the learned b-$N_{\rm HI}$ distribution has a resolution that depends on the $\mathrm{d}N/\mathrm{d}z$ of the model. We could instead set the number of $\{b, N_{\rm HI}\}$ pairs to be fixed while using different total panthlength for each model. However such a change does not affect the results of our inference method.

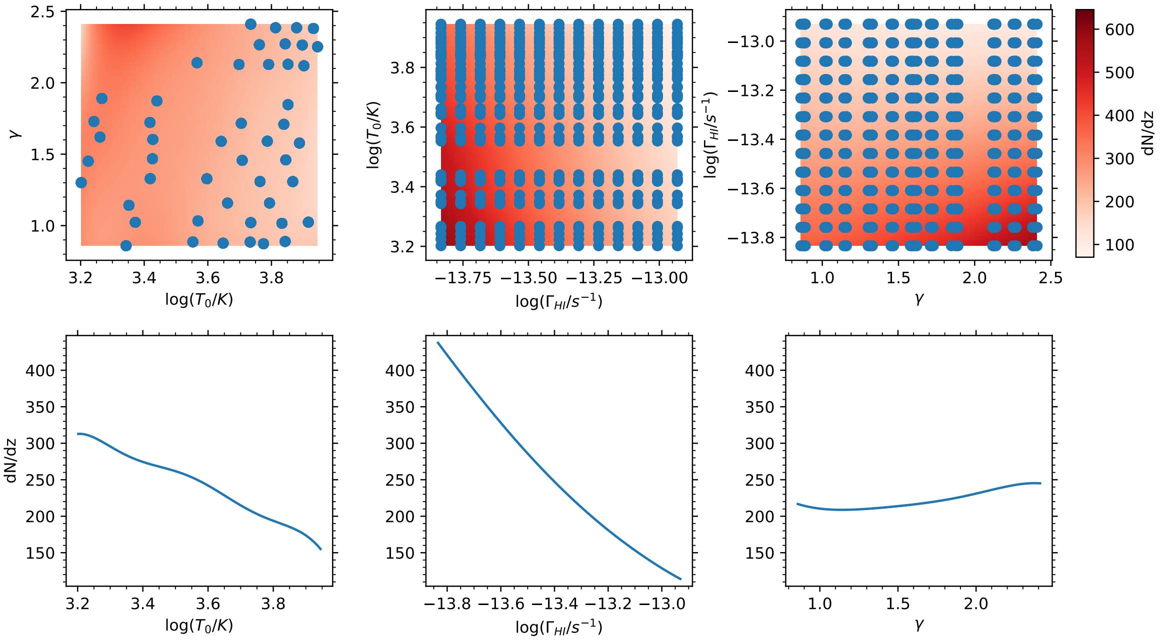

Figure 2.4: An example of emulation of the absorber density dN/dz generated by the Gaussian emulator sliced at the median value of the posterior from the MCMC process, where $\log(T_0/\mathrm{K})$ = 3.69, γ = 1.55, $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$ = -13.30. Top panels are the 2D dN/dz distributions where Nyx models are shown in blue circle. The top left panel is the dN/dz on γ-$\log T_0$ plane at $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$= -13.30. The top middle is $\log T_0$-$\log \Gamma_{\mathrm{HI}}$ plane at $\gamma = 1.55$. The top right is $\log \Gamma_{\mathrm{HI}}$-$\gamma$ plane at $\log(T_0/\mathrm{K}) = 3.69$. Bottom panels are marginalized 1D dN/dz distributions at the thermal parameters mentioned above. From left to right: dN/dz vs $\log T_0$, dN/dz vs $\log \Gamma_{\mathrm{HI}}$, and d$N$/d$z$ vs γ.

parameters, but ignored the normalization, which can be thought of as the total number of absorption lines in the dataset or equivalently as the line density dN/dz. Here we generalize the likelihood formalism introduced in Hiss et al. (2019) to include the information contained in the absorber density $\mathrm{d}N/\mathrm{d}z$ (see also Hiss 2019). Our goal is to find the likelihood of observing a set of absorption lines $\{b_i, N_{\mathrm{HI},i}\}$ given a model with a set of thermal parameters $[\log {T_0}', \gamma', \log {\Gamma_{\mathrm{HI}}}']$. We first assume that the probability density function (PDF)s are normalized such that

$$\iint P(b, N_{\mathrm{HI}})\, \mathrm{d}N_{\mathrm{HI}}\, \mathrm{d}b = 1, \tag{2.2}$$

where $P(b, N_{\mathrm{HI}})$ is the conditional probability distribution function $P(b\ , N_{\mathrm{HI}} \mid {T_0}', \gamma', {\Gamma_{\mathrm{HI}}}')$, for simplicity we write it as $P(b, N_{\mathrm{HI}})$ in the rest of this subsection. We imagine dividing the b-N_{HI} into a set of infinitesimally fine grid cells, such that the occupation number of each grid cell is either one or zero. Knowing that our set of observational/mock dataset $\{b_i, N_{\mathrm{HI},i}\}$ is comprised of n lines, and assuming that there are N_{g} grid cells in total, the likelihood for a model with thermal parameters $[\log {T_0}', \gamma', \log {\Gamma_{\mathrm{HI}}}']$ can thus be written as the following product of Poisson probabilities[10]

$$\begin{aligned} \mathcal{L} =& P(\mathrm{data}|\mathrm{model}) \qquad (2.3)\\ =& \left(\prod_{i=1}^{n} \mu_i\, e^{-\mu_i}\right) \left(\prod_{j\neq i}^{N_{\mathrm{g}}} e^{-\mu_j}\right), \end{aligned}$$

where the first product is over the occupied cells, and the second product is over the empty cells. Here the μ_i is the Poisson rate of occupying a cell in the b-N_{HI} plane with

---

[10] In assuming the probability distribution for each grid cell is Poisson, we are implicitly assuming each b-N_{HI} pair is an uncorrelated draw from the b-N_{HI} distribution. This assumption, also made by Hiss et al. (2019), amounts to ignoring the spatial correlations between absorption lines. Hiss et al. (2019) showed that this is a very good approximation and yields unbiased inference as we will also demonstrate in § 2.2.5)

area $\Delta \mathrm{N}_{\mathrm{HI}i} \times \Delta b_i$, i.e.

$$\mu_i = \left(\frac{\mathrm{d}N}{\mathrm{d}z}\right)_{\mathrm{model}} P(b_i, N_{\mathrm{HI},i})\, \Delta N_{\mathrm{HI}}\, \Delta b\, \Delta z_{\mathrm{data}}, \tag{2.4}$$

where $P(b_i, N_{\mathrm{HI}i})$ is the probability distribution function evaluated at the point $(b_i, N_{\mathrm{HI},i})$ using the DELFI b-N_{HI} distribution emulator described in § 5.2.1, and Δz_{data} is the total redshift path covered by the data spectra from which we obtain our data set $\{b, N_{\mathrm{HI}}\}$, whereas $(\mathrm{d}N/\mathrm{d}z)_{\mathrm{model}}$ is the absorber density of the model which will be further discussed later in this subsection.

Afterwards, it is easy to show that Eq. (2.3) implies

$$\ln \mathcal{L} = \sum_{i=1}^{n} \ln(\mu_i) - \sum_{k=1}^{N_g} \mu_k. \tag{2.5}$$

Above, the second sum over k is simply an integral of Eq. (2.4) over the b-N_{HI} plane, while the integral of $P(b, N_{\mathrm{HI}})\, \mathrm{d}N_{\mathrm{HI}}\, \mathrm{d}b$ over the plane is unity according to Eq. (2.2). As a result, we can write our likelihood function as

$$\ln \mathcal{L} = \sum_{i=1}^{n} \ln(\mu_i) - \left(\frac{\mathrm{d}N}{\mathrm{d}z}\right)_{\mathrm{model}} \Delta z_{\mathrm{data}}. \tag{2.6}$$

Since Hiss et al. (2019) did not consider the absorber density, the likelihood in their analysis is simply given by $\ln \mathcal{L}_{\mathrm{Hiss}} = \sum_{i=1}^{n} \ln P(b_i, N_{\mathrm{HI},i})$. In comparison, our likelihood function Eq. (2.6) can be written as

$$\ln \mathcal{L} = \sum_{i=1}^{n} \ln P(b_i, N_{\mathrm{HI},i}) + n \ln \xi - \xi, \tag{2.7}$$

where $\xi = (\mathrm{d}N/\mathrm{d}z)_{\mathrm{model}} \Delta z_{\mathrm{data}}$. We can see that the first term remains the same, and our modification (the implementation of absorber density $\mathrm{d}N/\mathrm{d}z$) can be considered as a correction term based on the absorber density of the model, the number of lines observed, and the pathlength of the data set $\Delta z_{\rm data}$.

As a result of our modification, the likelihood of observing a line with certain line parameter $(b, N_{\rm HI})$ now depends not only on the b-$N_{\rm HI}$ distributions of models but also on absorber densities of the models. Consequently, to evaluate the likelihood $\mathcal{L}$ on the parameter space, we need the ability to evaluate $({\rm d}N/{\rm d}z)_{\rm model}$ at an arbitrary location on the parameter space. To this end, a Gaussian process emulator (based on `George`, see Ambikasaran et al. 2016) is employed to emulate $({\rm d}N/{\rm d}z)_{\rm model}$ by interpolating the ${\rm d}N/{\rm d}z$ of models from Nyx simulations based on their $N_{\rm model}/\Delta z_{\rm model}$, where $\Delta z_{\rm model}$ is the total pathlength of simulated spectra that are fed into `VPFIT`, and $N_{\rm model}$ is the total number of lines identified by `VPFIT` from these spectra. The Gaussian process emulator is constructed with smoothing lengths of 40% of our thermal grid length[11] in $\log T_0$ and $\log \Gamma_{\rm HI}$ and a smoothing length of 80% of thermal grid length in γ. The longer smoothing length in γ is set to prevent the emulator from over-fitting the noise, considering that γ has less effect on the absorber density ${\rm d}N/{\rm d}z$ compared with T_0 and $\Gamma_{\rm HI}$ (see Fig.2.4), which makes small fluctuations induced by noise more significant.

The results of our ${\rm d}N/{\rm d}z$ emulation are shown in Fig.2.4, where both $\log T_0$ and $\log \Gamma_{\rm HI}$ (left and middle column) have negative correlations with absorber density ${\rm d}N/{\rm d}z$. This dependence can be explained qualitatively by the fluctuating Gunn-Peterson approximation (FGPA, see Weinberg et al. 1997)

$$\tau_{\rm Ly\alpha} \propto n_{\rm HI} \propto x_{\rm HI} n_{\rm H} \propto \frac{n_{\rm H}^2 T^{-0.7}}{\Gamma_{\rm HI}}, \tag{2.8}$$

where the $\tau_{\rm Ly\alpha}$ denotes the Lyα optical depth and the $n_{\rm H}$ is the hydrogen number density. This equation implies that both higher temperatures and higher photoionization rates

---

[11] The smoothing length is input as an initial guess, which is then refined later in the routine. In addition, all dimensions in the thermal grid are rescaled to unity in the Gaussian process emulator.

reduce the Lyα optical depth of gas absorbers in the IGM, leading to lower absorber density. The wiggles shown in dN/dz vs T_0 plot (bottom left panel of Fig.2.4) are effects of poor interpolation due to lack of models at $\gamma \sim 1.5$ (see top left panel). Moreover, we notice a weak correlation between γ and dN/dz (see the bottom right panel of Fig.2.4). However, such γ dependence is relatively weak compared with T_0 and $\Gamma_{\rm HI}$ dependencies, and is likely caused by artifacts due to the emulation. As shown in the top left-hand panel, we do not have models in low T_0 high γ region, the absorber density dN/dz could thus be over-extrapolated in these regions, further biasing the γ dependence on the whole parameter space. We performed some tests and found that the weak correlations in γ - dN/dz vanishes if we do not include the high γ simulations. Therefore, in conclusion, the marginalized γ - dN/dz correlation shown in Fig. 2.4 is an artifact introduced by our Gaussian emulator, however, it is too weak to affect our inference results.

2.2.3 Parameter study

A new feature of the DELFI b-$N_{\rm HI}$ distribution emulator is its ability to emulate b-$N_{\rm HI}$ distributions continuously on the parameter space. With such a feature, we are now able to illustrate the parameter dependence of the b-$N_{\rm HI}$ distribution and investigate the physics behind these dependence. Fig.2.5 shows emulated b-$N_{\rm HI}$ distributions with different values of thermal parameters [$\log T_0$, γ , $\log \Gamma_{\rm HI}$]. The top panel shows b-$N_{\rm HI}$ distributions with increasing T_0, where $\log(T_0/{\rm K})$ = 3.25 (left), 3.60 (middle) and 3.95 (right) respectively, while γ=1.55 and $\log(\Gamma_{\rm HI}/{\rm s}^{-1})$=-13.36 for all three plots. Increasing T_0 results in the upward shifting of the b-$N_{\rm HI}$ distributions, which can be explained by the thermal component of the b parameter and the T-Δ relationship Eq. (1.1), i.e.

$$b_T \propto (2kT/m)^{1/2} \propto (T_0 \Delta^{\gamma-1})^{1/2}, \tag{2.9}$$

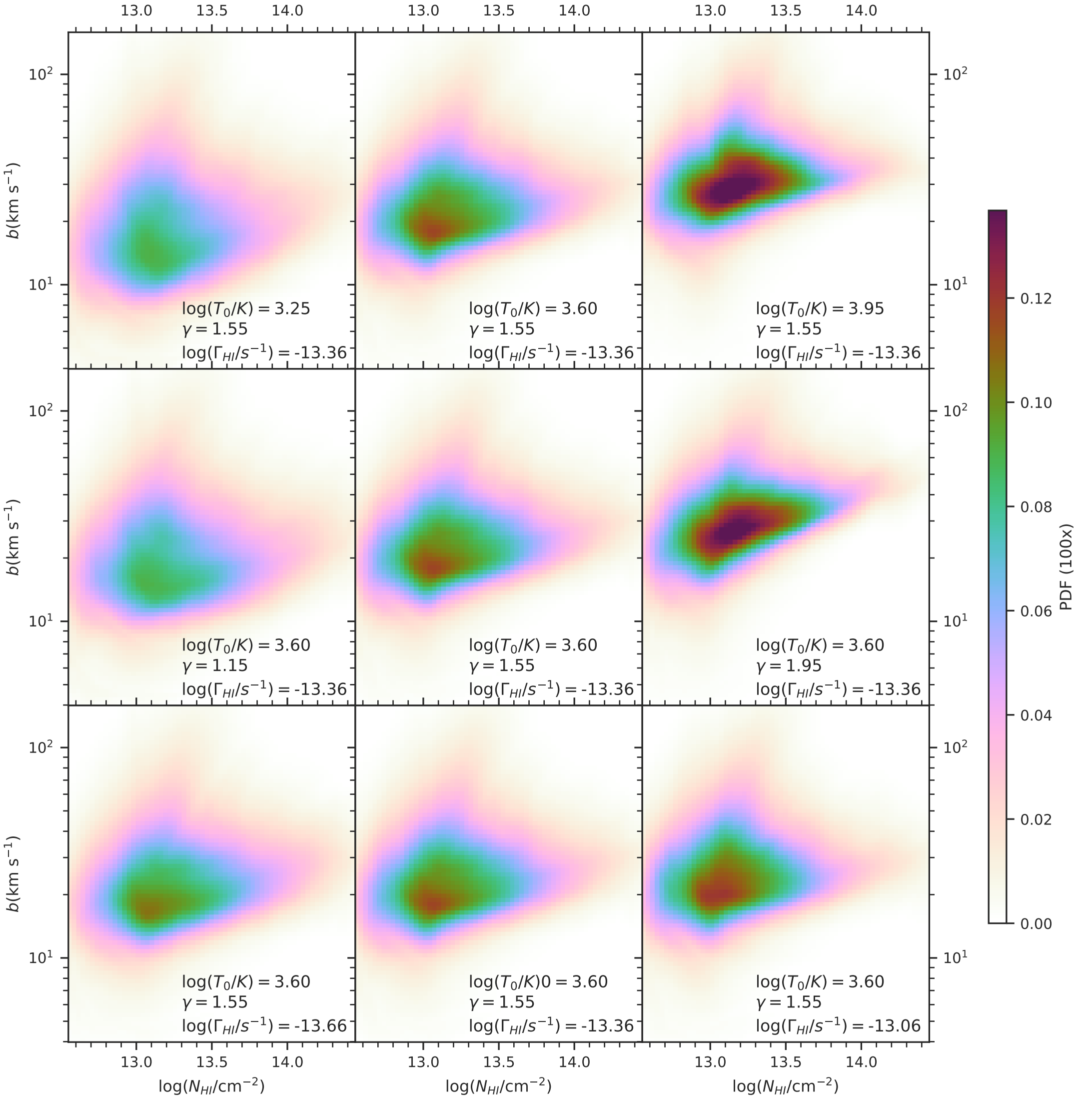

Figure 2.5: Comparisons of b-N_{HI} distributions modeled by DELFI emulator with different thermal parameters. Top panel shows changes in the b-N_{HI} distribution with increasing $\log T_0$, where $\log(T_0/\mathrm{K})$ = 3.25 (left), 3.60 (middle) and 3.95 (right) respectively, while γ=1.55 and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$=-13.36 for all three plots. The middle panel shows changes of the b-N_{HI} distribution where γ =1.15 (left), 1.55 (middle) and 1.95 (right) respectively, while $\log(T_0/\mathrm{K})$=3.60 and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$=-13.36 are fixed. The bottom panel shows b-N_{HI} distributions with decreasing UV background. $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$= -13.66 (left), -13.36 (middle) and -13.06 (right), while $\log T_0$ and γ remain unchanged. All pdfs here are normalized to unity. For illustration purposes, values of pdf are multiplied by 100 in the color bar.

where higher T_0 results in higher IGM temperature, leading to larger b parameters. In addition, we notice that as the T_0 goes up, the b-N_{HI} distribution becomes more concentrated, i.e. the distribution becomes tighter, and the pdf values increases. Such behavior might be explained as follows. There are two components contributing to b parameter, namely thermal motion and non-thermal broadening. The thermal component is associated with the IGM temperature and thus follows a distribution determined by T_0. On the other hand, as a result of the small-scale motion of the gas, the non-thermal component is independent of the temperature and has a large dispersion, leading to broader distribution. At low temperatures, where the thermal contribution is weak, the b parameter is dominated by the non-thermal component, resulting in broad distribution. As the temperature goes up, the thermal component dominates over non-thermal broadening, and the b parameter thus concentrates on a central value of b determined by the IGM temperature.

The middle panel of Fig.2.5 shows the b-N_{HI} distribution with increasing γ, where γ =1.15 (left), 1.55 (middle), and 1.95 (right), respectively, while $\log(T_0/\mathrm{K})$ =3.65 and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$=-13.36 are fixed. These plots indicate that there are degeneracies between γ and T_0, where an increasing γ also shifts b-N_{HI} distributions upwards, which can be understood from Eq. (2.9) and the fact that at low-z, the Lyα lines originate predominantly from gas with $\Delta_{\mathrm{abs}} > 1$ ($\Delta_{\mathrm{abs}} \sim 10$, see Gaikwad et al. 2017b), which results in higher temperatures at densities of absorbers for models with larger γ. The concentration effect is also seen in the middle panel, which can be explained in the same way as the upward shifting of the b-N_{HI} distribution due to increasing γ. It can also be seen from the middle panel that the γ is correlated with the slope of the low-b cutoff of the b-N_{HI} distribution, which is consistent with the analytical fit of the low-b cutoff, where the slope can be approximated by $\Delta \log b/\Delta \log N = (\gamma - 1)/3$ (see Rudie et al. 2012a).

The bottom panel of Fig.2.5 shows b-N_{HI} distributions with increasing photoionization rate Γ_{HI}, where $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$= -13.66 (left), -13.36 (middle) and -13.06 (right), while $\log T_0$ and γ remain unchanged. We observe that increasing Γ_{HI} results in a similar but much weaker effect compared with increasing T_0, i.e. the b-N_{HI} distribution slightly shifts upward and becomes more concentrated with increasing Γ_{HI}. Such effects are because the photoionization rate Γ_{HI} alters the Lyα optical depth of the IGM. Since the Lyα forest typically probes regions with optical depth $\tau_{\mathrm{Ly}\alpha} \sim 1$, given higher Γ_{HI}, it probes regions with higher temperatures and densities, which can be derived from Eq.(2.8), causing effects similar to increasing T_0. However, such effects are relatively weak, making the b-N_{HI} distribution less sensitive to the photoionization rate Γ_{HI}.

All these aforementioned parameter dependences (except Γ_{HI}, which is not considered in previous works) of the b-N_{HI} distribution are consistent with previous works that measure the IGM thermal state based on the full b-N_{HI} distribution[12] (Hiss et al. 2019) and low-b cutoff (Schaye et al. 1999, Rudie et al. 2012a, Bolton et al. 2014, Rorai et al. 2018, Hiss et al. 2018), indicating that our DELFI emulator successfully reproduce the parameter dependences of the b-N_{HI} distribution. Furthermore, it also implies that our understanding of the b-N_{HI} distribution agrees with the physics prediction.

2.2.4 Mock Inference Results

Sets of mock spectra are created from our Nyx simulations to test the performance of our inference algorithm under realistic conditions. These mock spectra sets are generated from a set of simulated spectra following a forward-modeling approach designed to match the pathlength, resolution, and noise properties of the Danforth et al. (2016) low-redshift quasar spectra in one-to-one correspondence as described in §2.1.3. Consequently, each mock spectra set consists of 34 forward-modeled spectra, which has exactly the same noise

---

[12] In Hiss et al. (2019), at $z \sim 2$, the Lyα lines originate predominantly from gas with $\Delta < 0$, causing different effects when changing γ. However, the physics explanations behind the effect are coherent.

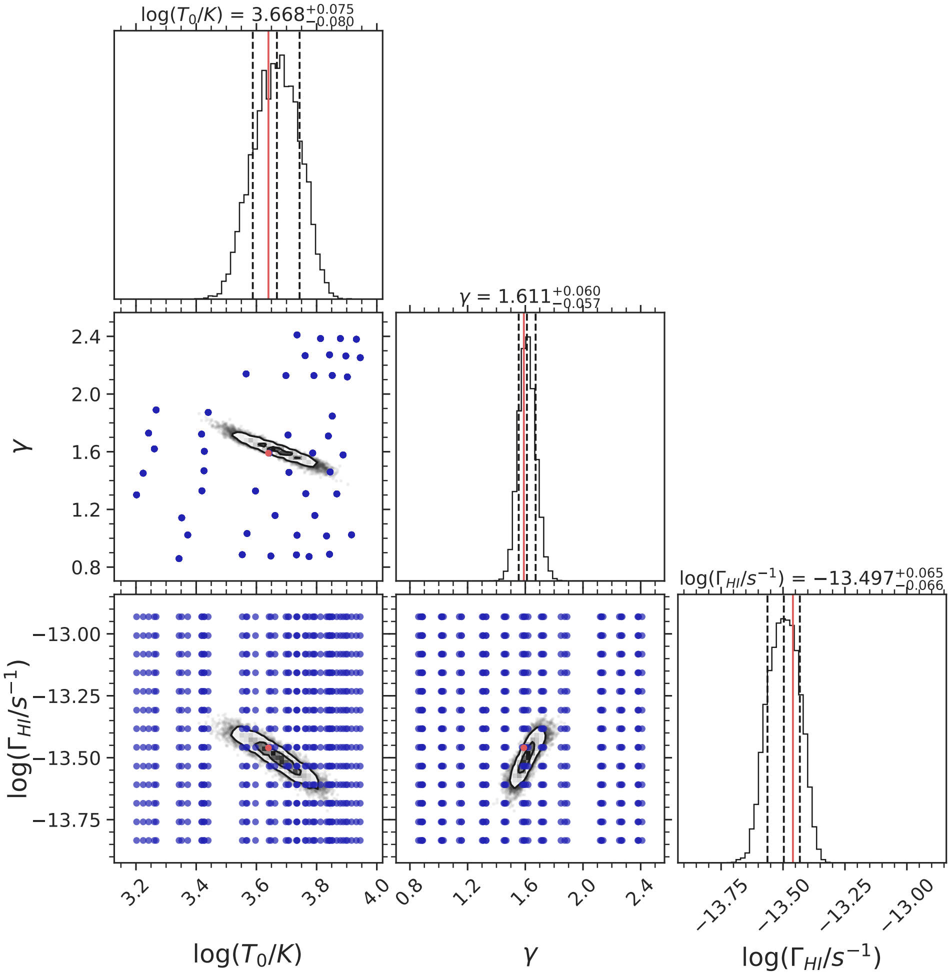

Figure 2.6: MCMC posterior for one of the models from Nyx simulation (absorbers shown in Fig.2.7) using the likelihood function Eq. 2.6. Projections of the thermal grid used for generating models are shown as blue dots, while the true model is shown as red dot. Inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The parameters of true model are indicated by red lines in the marginal distributions, while the dashed black lines indicates the 16, 50, and 84 percentile values of the posterior. The true parameters are: $\log(T_0/\mathrm{K})$ = 3.643, γ = 1.591 and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$ = -13.458.

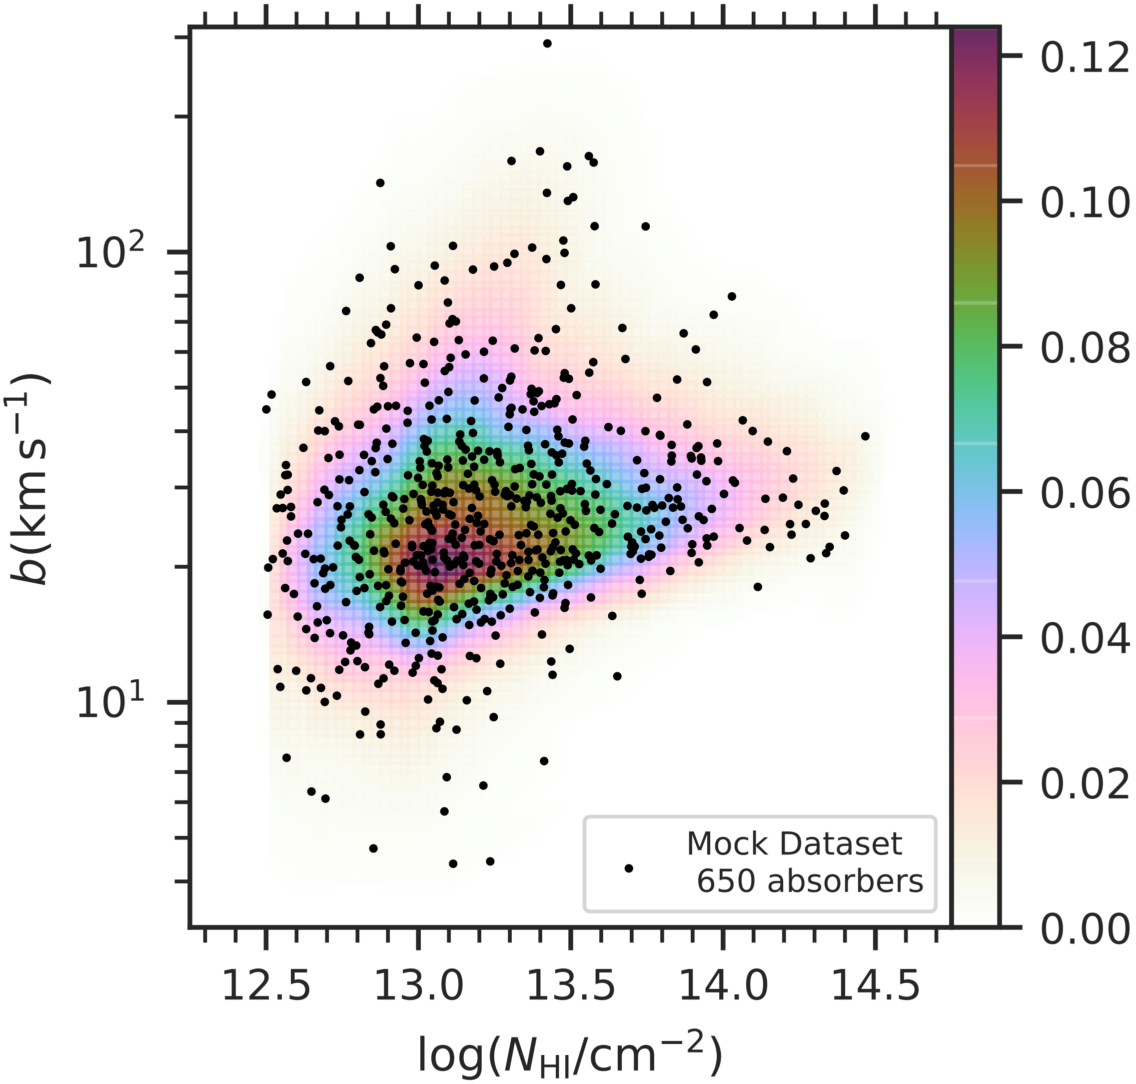

Figure 2.7: The color map is the full b-N_{HI} distribution recovered from the Nyx mock dataset, which is emulated by our DELFI emulator based on the best-fit parameters (median values of the marginalized MCMC posterior), where $\log(T_0/\mathrm{K}) = 3.668$, $\gamma = 1.611$ and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$ =-13.498. Black dots are the mock datasets we used in the inference. For illustration purposes, values of pdf are multiplied by 100 in the color bar.

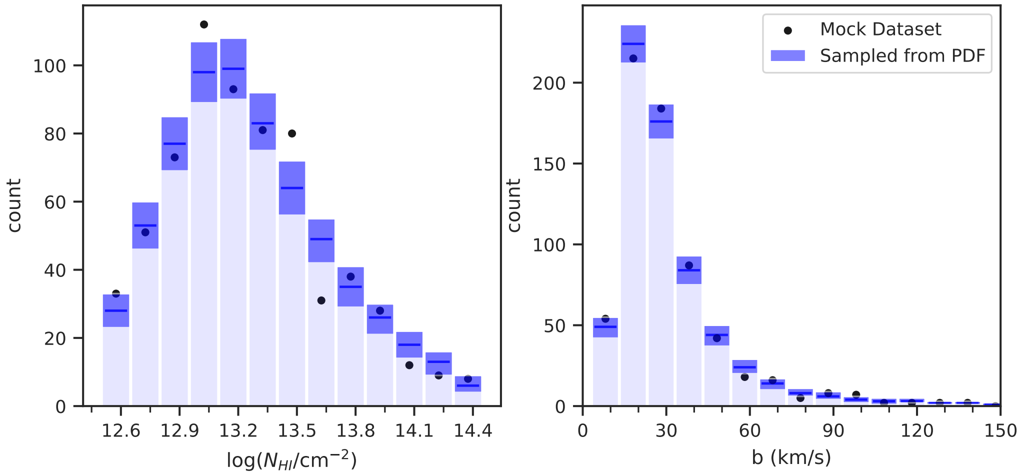

Figure 2.8: Marginalized 1D distributions of N_{HI} (*left-hand panel*) and b (*right-hand panel*) for the mock dataset (black dots) and the sampling from emulated b-N_{HI} distribution (blue bars) at $\log(T_0/\mathrm{K}) = 3.699$, $\gamma = 1.549$ and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$ =-13.506. Blue bars show the average of 5000 sampling from the emulated b-N_{HI} distribution using MCMC, while the (dark) blue shaded regions represent the 1-σ fluctuation (16%-84% percentile among 5000 samples).

vectors, instrumental effects, and total pathlength (Δz_{data}=2.136) as the real observed dataset, which ensures that the accuracy of our analysis is realistic and achievable when the method is applied to real data. A set of $\{b, N_{\mathrm{HI}}\}$ pairs, obtained by fitting these spectra using `VPFIT` (see §2.1.4), is then used as the 'data' in the likelihood function (see Eq.2.3) to infer the posterior distribution for IGM thermal parameters for this mock dataset.

In this work, we perform inference via Markov chain Monte Carlo (MCMC) sampling using the `python` package `emcee` (Foreman-Mackey et al. 2013), which implements the affine-invariant sampling technique (Goodman & Weare 2010) to sample the posterior probability distribution. Here the posterior is calculated based on the likelihood in Eq. (2.5), which takes into account the absorber density dN/dz as described in §2.2.2, while assuming uniform (flat) priors for $\log T_0$, γ and $\log \Gamma_{\mathrm{HI}}$, where the boundaries are chosen to be the range of each respective parameter in 1D. MCMC posteriors obtained from the aforementioned mock datasets ($\{b, N_{\rm HI}\}$ pairs) are shown in Fig.2.6. We obtain $\log(T_0/{\rm K}) = 3.668^{+0.075}_{-0.080}$, $\gamma = 1.611^{+0.060}_{-0.055}$ and $\log(\Gamma_{\rm HI}/{\rm s}^{-1}) = -13.497^{+0.065}_{-0.066}$ from the marginalized distributions, whereas the true parameters are: $\log(T_0/{\rm K}) = 3.643$, $\gamma = 1.591$ and $\log(\Gamma_{\rm HI}/{\rm s}^{-1}) = -13.458$ (red dot and red vertical lines). We recover the input parameters in very high precision with errors $\Delta\log(T_0/{\rm K}) = +0.025$dex, $\Delta\gamma = +1.3\%$, and $\Delta\log(\Gamma_{\rm HI}/{\rm s}^{-1}) = -0.039$dex, while true parameters (red dot/solid lines) are all in the 1-σ interval (inner black contours/ black dashed lines) of the posterior. Here the degeneracy between T_0 and γ can be quantitatively understood by the T-Δ relationship Eq. (1.1) and the typical overdensity of absorbers $\Delta_{\rm abs} \sim 10$. More specifically, both higher T_0 and γ result in higher temperature of the absorbers, shifting the b-$N_{\rm HI}$ distribution upward (see Fig.2.5 and relevant discussion in §2.2.3). The degeneracy between T_0 and $\Gamma_{\rm HI}$ is mainly a result of the degeneracy in the absorber density ${\rm d}N/{\rm d}z$ with respect to the two parameters (see Fig.2.4 and Fig.A.1 as comparison), which is explained in §2.2.2. It is noteworthy that our inference algorithm provides preeminent accuracy for all three parameters even under a very realistic condition, where the resolution of spectra is rather low (with lines not fully solved), and the number of data is limited (with a total pathlength $\Delta_z = 2.136$). Such a high sensitivity and precision makes our inference method a powerful tool in the study of the low-z IGM and Lyα forest.

Fig.2.7 shows the full b-$N_{\rm HI}$ distribution recovered from the mock dataset, which is emulated by our DELFI emulator based on the best-fit parameters (median values of the marginalized MCMC posterior). It appears that the PDF (color map) successfully represents the density distribution of the data points. Furthermore, marginalized 1D distributions of b and $N_{\rm HI}$ are given in Fig.2.8 for both the mock dataset (black dots) and random samples from the emulated b-$N_{\rm HI}$ distribution (blue bars). It can be seen that our emulator successfully reproduces the 1D marginalized b and $N_{\rm HI}$ distribution, though there is a fluctuations in $N_{\rm HI}$ for the mock dataset at around $\log(N_{\rm HI}/{\rm cm}^{-2}) \sim 13.5$. We figured out that such fluctuation is caused by the random error during the generation of the mock dataset, which can be reduced by increasing the size of the mock datasets. However, to test the performance of our inference method under realistic conditions, we fix the size of the mock datasets and bear with such fluctuation in this work.

2.2.5 Inference test

As discussed above, the likelihood function used in our inference algorithm involves several approximations and emulation/interpolation procedures. Most importantly, our inference ignores correlations between the lines (see the discussion in Hiss et al. 2019), and we emulate the b-$N_{\rm HI}$ distribution and the $\mathrm{d}N/\mathrm{d}z$ with our DELFI and Gaussian emulators respectively, while both emulations involve interpolations. These procedures might induce additional uncertainties that are counted in our error budget[13], we hence want to make sure our inference results are valid under these assumptions, and our interpolation procedures work correctly. Therefore, we perform a series of inference tests to evaluate the robustness of the entire inference method. An inference test is to carry out a set of realizations of the inference algorithm based on the mock dataset and inspect the results to reveal if the inference method returns valid posterior probability distributions, i.e. whether the 'true model' is included in a set of probability contours following the ratio indicated by the posterior.

The inference test is done as follows. First of all, we adopt the same prior as described in 2.2.4, and construct a regular uniform grid in the parameter space spanning the range set by our prior. For each realization, we pick a model (set of parameters) on the above grid, which we refer to as the 'true model'. and we refer its thermal parameters as 'true

---

[13] The uncertainty of the b-$N_{\rm HI}$ distributions emulated by DELFI is also ignored in our analysis. Such uncertainty is caused by the randomness in the training process, and has not been included in the results. But since our inference method (and the toy model) does well in the inference test, such randomness should be smaller than stochastic error shown in our analysis, and should not dominate our error budget.

Table 2.1: Table of results of the inference test

models	Total	68(%)	95 (%)
random models	480	290 ($60.42 \pm 2.29\%$)	439 ($94.67 \pm 1.25\%$)
single model	200	134 ($67.00 \pm 3.50\%$)	190 ($95.00 \pm 1.50\%$)

parameters' $\boldsymbol{\theta}_{\text{true}}$, We then create a corresponding mock dataset following the prescription described in §2.2.4. Given the mock dataset, since our priors are flat, we can determine the corresponding posterior probability distribution by evaluating the likelihood function $\mathcal{L} = P(\text{data}|\text{model})$ on the whole parameter space. We then normalize the posterior function to unity and determine 3D posterior probability contours based on the posterior (likelihood) distribution. Knowing that the likelihood function is continuous on the whole domain, the 3D volume integral can hence be substituted by a 1D integral over the sorted likelihood function. Here we define the probability contours C_P and the likelihood thresholds $\mathcal{L}_P$ in the following way,

$$\iiint_{C_P} \mathcal{L} \mathrm{d}V = \int_{\mathcal{L}_P}^{\infty} \mathcal{L} \mathrm{d}\mathcal{L} = P, \tag{2.10}$$

such that a probability contour C_P is simply where $\mathcal{L} = \mathcal{L}_{\mathrm{P}}$, and any 'model' with parameter $\boldsymbol{\theta}$ being inside a contour C_P thus becomes equivalent to $\mathcal{L}(\boldsymbol{\theta}) > \mathcal{L}_{\mathrm{P}}$. We further define the effective 1σ (68%) and 2σ (95%) intervals as the volume between contour pairs $(C_{0.16}, C_{0.84})$ and $(C_{0.025}, C_{0.975})$ respectively. Finally, we judge the performance of our inference method based on how often the parameters of the 'true model' $\boldsymbol{\theta}_{\text{true}}$ falls in these 1(2)-σ interval contour pairs compared to the expectation based on the corresponding probabilities, i.e. if our posterior distribution is perfect, the true model should land within the 1σ (2σ) contours 68% (95%) of the time. An example of the distribution of the likelihood function is shown in Fig.A.2, and more details about the calculation of the likelihood distribution is presented in Appendix A.2.

In practice, we perform an inference test on a set of random models on the thermal grid to test the overall performance of our inference algorithm. We pick 12 models and execute 40 realizations per model. The result shows that the true values are within the 1-σ (68%) interval for 60.42 ± 2.29 % (290/480) of the time, and in the 2σ(95%) interval for 94.67 ± 1.25 % (439/480) of the time, while the upper and lower limits are given by the $\pm 1\sigma_{\rm bi}$ error for corresponding binomial distributions. In addition, we carry out a cross-validation test to ensure our emulators are not affected by over-fitting problem. Here we select a single model near the center of the parameter space ($\log(T_0/{\rm K}) = 3.643$, $\gamma = 1.591$, and $\log(\Gamma_{\rm HI}/{\rm s}^{-1})$ = -13.458.), and exclude the model[14] from the training dataset. We train our emulators (both b-$N_{\rm HI}$ distribution and $\mathrm{d}N/\mathrm{d}z$) based on the new dataset, and run 200 realizations of our inference method. We observe that the true values are inside the 1σ (68%) interval for 67.00 ± 3.50% (134/200) of the time, and inside the 2σ (95%) interval for 95.0 ± 1.50% (190/200) of the time. Results are presented in Table 2.1. The overall performance indicates that our algorithm passes the inference[15].

In the end, to further demonstrate and elaborate on the effectiveness of our inference algorithm, we created a toy model, which involves entire inference pipeline, (in Appendix A.3) to test the whole inference algorithm under more controlled conditions, where the toy b-$N_{\rm HI}$ distribution is analytical, and the parameter dependence is known. Here the toy b-$N_{\rm HI}$ distribution consists of a multivariate Gaussian distribution parameterized by three mock parameters following the parameter dependence discussed in § 2.2.3. Moreover, these mock parameters also control the line density $\mathrm{d}N/\mathrm{d}z$ of the model based

---

[14] In practice we exclude all models with the same T_0 and γ $\log(T_0/{\rm K})$ =3.643 and γ = 1.591), since we mostly want to test the performance of the b-$N_{\rm HI}$ distribution emulator on the T_0-γ plane.

[15] Our inference method performs better when the model is close to the center of the grid. This might be because our emulators, both DELFI and Gaussian process emulator, perform better at the center of the grid where the interpolation is more accurate. Besides, our thermal grid has an irregular shape on the T_0-γ plane, and might thus make the interpolation even harder or distorted when there are no or only a few models around. Such a problem might be addressed by adding more simulation models, extending the thermal grid to make sure the region we are interested in always lies at the center of the grid.

on the dN/dz map generated by the Gaussian emulator from our Nyx simulation models (see Appendix A.3 for more details). As a result of this toy model and also the inference test, we conclude that our inference algorithm is sound.

2.3 Summary

In this hapter, we have presented and evaluated our new method of measuring the thermal state $[T_0, \gamma]$ and the photoionization rate $\Gamma_{\rm HI}$ of the low redshift IGM using its b-$N_{\rm HI}$ distribution and absorber density dN/dz. We made use of a novel machine learning technique DELFI to build a b-$N_{\rm HI}$ distribution emulator and used a Gaussian process emulator to simulate the absorber density dN/dz. We trained both emulators on a dataset generated from a set of Nyx simulations on a large parameter grid. To test the performance of our inference algorithm under realistic conditions, we applied forward modeling techniques to model the noise and instrumental effects based on the HST COS quasar spectra from Danforth et al. (2016). We showed using extensive tests that our inference method is proficient and reliable. Here we conclude by discussing the performance and summarizing the essential elements of our new algorithm.

- We used mock datasets to simulate the measurement of the thermal state $[T_0, \gamma]$ of the low redshift IGM from the full joint b-$N_{\rm HI}$ distribution, for the first time taking the absorber density dN/dz into account. The latter enables us to constrain the photoionization rate $\Gamma_{\rm HI}$, since only the shape of the b-$N_{\rm HI}$ distribution is insensitive to this parameter (see Fig.2.5). We also confirm that the dN/dz term we introduced is consistent with our inference based on the b-$N_{\rm HI}$ distribution alone, and improves the performance of our inference method (see Appendix A.1).

- Our new inference method successfully recovers thermal parameters of models from the Nyx simulation with small uncertainties (in our example, $\sigma_{\log T_0} \sim 0.08$ dex, $\sigma_\gamma \sim 0.06$, and $\sigma_{\log \Gamma_{\rm HI}} \sim 0.07$ dex), using a relatively small dataset with $\Delta z = 2.316$. Furthermore, these results are obtained under realistic conditions as we forward-model the observational effects and noise from the Danforth et al. (2016) low-z COS quasar spectra while setting the size of our mock datasets to be the same as the observational dataset (i.e. having the same total pathlength $\Delta z_{\rm ob}$). Considering all these factors, the accuracy and sensitivity we attained in this study should be achievable when our inference method is applied to real observational data, making it a powerful tool for studying the Lyα forest.

- Our algorithm passes the inference test (see §2.2.5), indicating that our approximation and emulation/interpolation are reliable. We also demonstrate the robustness of our inference method by testing the entire inference pipeline, including emulation and interpolation procedures on a toy model under better-controlled conditions(see Appendix A.3).

- The b-$N_{\rm HI}$ distribution (DELFI) emulator successfully emulates both the 2D b-$N_{\rm HI}$ distributions and 1D marginalized distributions of b and $N_{\rm HI}$. We find that the 2D b-$N_{\rm HI}$ distribution shifts upward (towards higher b values) with increasing T_0 and γ, while larger γ also tilts up the low-b cut off. We explain these effects qualitatively in section § 2.2.3 and show that they are consistent with previous work.

Chapter 3

Measurements of the IGM Thermal and Ionization State during the Cosmic Afternoon

3.1 Observational Data

In this chapter, we employ the aforementioned method to measure both the thermal and ionization state of the IGM using quasar spectra obtained from STIS on board HST. We opt for STIS due to its superior resolution compared with COS and available archival data. We utilize 12 HST STIS quasar spectra covering $0.9 < z < 1.5$, which are selected from the STIS archive based on their redshift coverage, SNRs, and the availability of metal identification. For the identification of metal lines, we import the metal identification from the COS Absorption Survey of Baryon Harbors (CASBaH) project (Tripp 2014, Prochaska et al. 2019, Burchett et al. 2019, Haislmaier et al. 2021) for five of our spectra, and make use of the metal identification from Milutinović et al. (2007) for the remaining seven spectra. We fit these spectra to obtain our $\{b, N_{\rm HI}\}$ sample using `VPFIT` (see § 2.1.4)

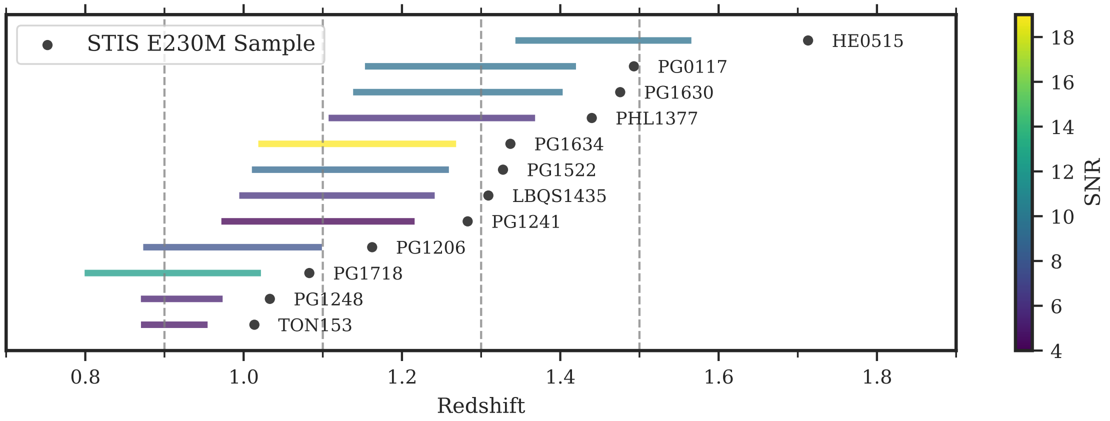

Figure 3.1: The HST STIS E230 spectra used in this study. The quasar are shown as black dots, and the Lyα spectra, with proximity zones removed, are shown as line segments with their colour indicating the mean SNR (per pixel). The three redshift bins used in this study are shown by the vertical dashed lines.

and apply the Hu22 method to measure the thermal and ionization state of the IGM in three redshift bins centering on $z = 1$, 1.2, and 1.4.

This chapter is structured as follows. We introduce our observational data in § 3.1 together with the data processing procedure, including continuum fitting, Voigt profiles fitting, and metal masking. In §3.2 we describe our hydrodynamic simulations, parameter grid, and mock data processing procedures, including generating Lyα forest from simulation, creating mock sightlines, and forward-modeling. In §3.4 we discuss our results. Lastly, we summarize the chapter in §3.5.

To measure the thermal state of the IGM around z $\sim$ 1, we make use of the quasar spectra observed with the HST STIS (Woodgate et al. 1998) using the E230M echelle mode, which provide spectroscopic coverage from $\sim$ 1600 Å to 3100 Å. We select such echelle mode for two reasons. First, as discussed in § 1, its high spectral resolution is beneficial for our analysis, with R $\sim$ 30,000, corresponding to $\sim$ 10 km/s (Kimble et al. 1998, Medallon & Welty 2023), and its LSF is close to Gaussian and has a weak dependence on the wavelength, which makes both the Voigt profile fitting (see § 2.1.4) and the generation of forward models easier (see § 2.1.3). Secondly, the echelle modes have higher wavelength coverage compared with first-order grating modes, enabling us to measure the $\{b, N_{\rm HI}\}$ of the Lyα absorption lines across a wider redshift range with constant instrumental effects such as LSF, which makes our analysis across different redshift bins more robust. We search the archival HST STIS E230M data observed in the 0.2" × 0.2" slit, and retrieve 12 spectra with average SNR $\gtrsim 5$. The details of the observation, from which our quasar samples are obtained, are summarized in Table 3.1, and Fig. 3.1 depicts the redshift coverage of the spectra used in this study. The quasars are shown as black dots, and the spectra are shown as line segments with their colour indicating the SNR. The redshift bins considered for the measurements are shown by the vertical dashed lines in Fig. 3.1.

To reduce and combine the STIS spectra, we used the procedure of Tripp et al. (2001) with `CALSTIS v3.4.2`. In brief, starting with the `CALSTIS x1d` files, for each quasar we combined all exposures, including the coaddition of overlapping regions of adjacent echelle orders, all with appropriate weighting and using the STIS flags to mask out bad pixels (see Tripp et al. 2001, for details). We then fit the continuum of these spectra using the interactive continuum fitting program imported from `linetools`[1]. Since we focus on the Lyα forest in this study, we make use of only the Lyα regions, excluding Lyβ and higher Lyman series absorption lines at $\lambda < 1050$ Å, while also masking the quasar proximity zones at $\lambda > 1180$ Å (see Fig.3.1). As a result, we only use the spectral segment with rest frame wavelength $1050 < \lambda_{\rm rest} < 1180$ Å. The quasar sightlines are chopped and padded by white noise based on the noise vector of the spectrum before passing into the VP-fitting program to avoid any complications arising from the edges of the spectra, and the padded regions are later masked in post-processing. This treatment of the edges is also applied to the mock forward models to ensure our analysis is consistent.

---

[1] For more information, visit https://linetools.readthedocs.io

Table 3.1: Summary of HST STIS sightlines used in the study

ID	z_{qso}	Wavelength (Å)	Obs. date	Exp. time (ksec)	Ave. SNR per pix. (full)	(Lyα)
TON153	1.014	2275 - 3110	2001 Jan.	5.3	5.0	4.8
			2002 Jun.	8.2		
PG1248+401	1.033	2275 - 3110	2002 Jul.	25.2	5.9	5.0
			2001 Oct.	28.8		
PG1718+481	1.083	1841 - 2673	1999 Nov.	14.1	7.9	9.8
PG1206+459[a]	1.162	2273 - 3110	2001 Jan.	17.3	7.3	6.4
LBQS1435-0134[a]	1.309	1985 - 2781	2015 Jun.	20.9	10.6	5.5
PG1241+176	1.283	2275 - 3110	2002 Jun.	19.2	4.7	4.4
PG1522+101[a]	1.328	1985 - 2781	2015 Mar.	7.7	9.5	7.1
			2015 May.	13.2		
PG1634+706	1.337	1858 - 2673	1999 May.	14.5	12.9	18.7
		2275 - 3110	1999 Jun.	14.5		
		1858 - 2673	1999 Jun.	26.4		
PHL1377[a]	1.440	2275 - 3110	2002 Jan.	14.0	7.2	5.3
			2002 Feb.	28.0		
PG1630+377[a]	1.476	2275 - 3110	2001 Feb.	5.3	10.6	7.5
			2001 Oct.	28.8		
PG0117+213	1.493	2275 - 3110	2000 Dec.	42.0	7.2	7.5
HE0515-4414	1.713	2275 - 3110	2000 Jan.	31.5	7.9	7.6

[a] The quasar sightlines on which we use the metal identification from the COS Absorption Survey of Baryon Harbors (CASBaH).

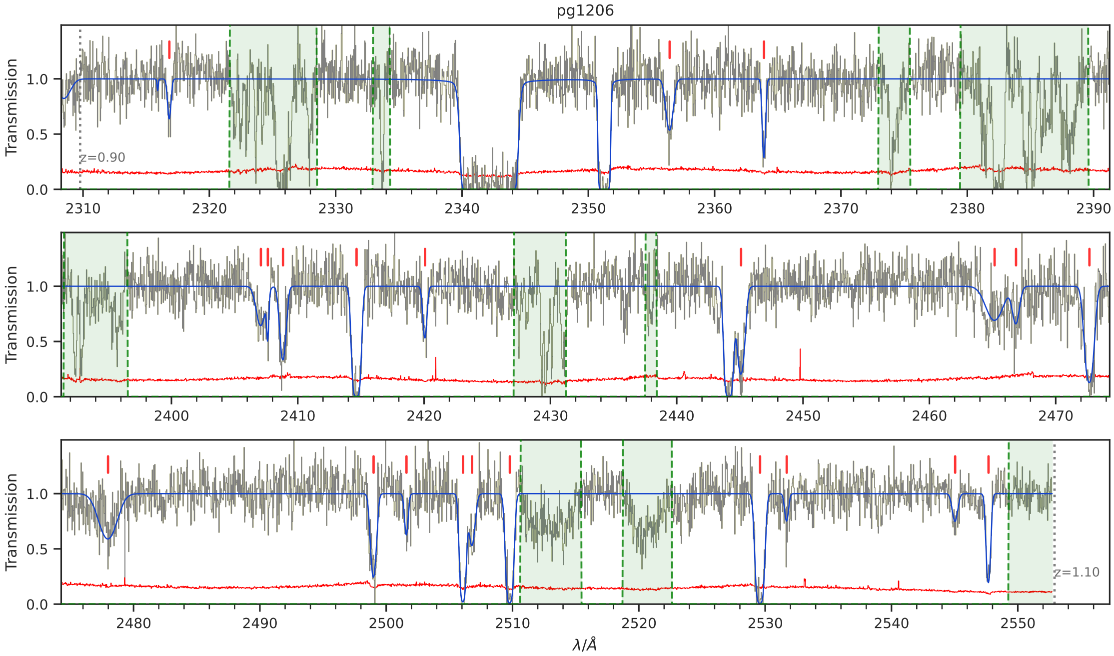

Figure 3.2: Illustration of the processed STIS spectrum of PG1206+459. The original spectrum is shown in gray, while a model spectrum based on VP-fitting is shown in blue. The noise vector is shown in red, and the masked regions are shown as green shaded regions. The Lyα lines used for our $\{b, N_{\rm HI}\}$ dataset are labelled by red vertical lines.

3.1.1 Voigt-Profile Fitting

In this work, we use the line-fitting program presented in §2.1.4. We implement the STIS E230M LSF to fit our observational data. Notice that the VP-fitting procedure is applied to the whole spectral segment, fitting both the Lyα lines and metal lines, including both intervening metal lines and those from interstellar medium of Milky Way (MW); for simplicity, hereafter we refer to these collectively as metal lines. The removal of these metal lines is later discussed in § 3.1.2.

Our `VPFIT` wrapper is designed to fit spectra using a custom LSF. However, it is important to note that it accommodates only a single LSF, without accounting for any wavelength dependency. To address this, we extract the STIS E230M LSF from `linetools` and interpolate it to match the central wavelength of the spectrum we aim to fit. As previously detailed in § 3.1, the STIS 230M exhibits a Gaussian-like LSF, which shows minimal variation across different wavelengths. Consequently, our approach of employing a singular LSF in the VP-fitting process does not introduce significant errors. To ensure consistency and avoid statistical biases, we apply the same fitting methodology to both our observational data and forward-modelled mock.

One of our STIS spectra, PG1206 is shown as an example of the VP-fitting procedure in Fig. 3.2. The original spectrum is shown in grey, and the model based on VP-fitting is shown in blue. The noise vector of the original spectrum is shown in red, and the masked regions due to metal line detection are shown as green shaded regions. The Lyα lines used for our $\{b, N_{\rm HI}\}$ dataset (after all filters) are labelled by red vertical lines.

3.1.2 Metal Identification

As previously mentioned, our VP-fitting procedure fits all absorption lines including Lyα lines and metal lines. For our analysis based on the $\{b, N_{\rm HI}\}$ of the Lyα forest, it is critical to filter out these metal lines. To this end, we make use of archival metal identification data presented in Milutinović et al. (2007) for seven of our quasar sightlines and use metal identification from the CASBaH survey (Tripp 2014, Prochaska et al. 2019, Burchett et al. 2019, Haislmaier et al. 2021) for the rest five spectra (see notes of the Table. 3.1). For each spectrum, we create a mask to cover the vicinity of each metal line based on the aforementioned metal identification. These masked regions are initially aligned with the central wavelength of the metal lines reported in the literature, while their initial widths are set to be $\Delta v = 30$ km/s in velocity space. Such a value is chosen based on the resolution of STIS E230M, which corresponds to $\sim$ 10 km/s. We then apply the masks to our VP-fit results to filter out potential metal lines. To do so, we first locate the absorption line region characterized by $F_{\rm line,fit} \leq 0.99$, where the $F_{\rm line,fit}$

stands for the normalized flux given by the VP-fit model (the blue line in Fig. 3.2). If any absorption line region overlaps with the initial mask, we increase the width of the mask to cover the detected line, while the increment is given by the full width at half maximum (FWHM) of the detected line, approximated by FWHM $= b/0.6$, where the b is given by `VPFIT`. Lastly, we adjust the masks manually to fill the small gaps (with $\Delta v = 30$ km/s) between the masked regions and make sure all absorption lines close to (the original) metal masks reported by our VP-fitting procedure are masked. The aforementioned masking procedure is needed based on the fact that our VP-fitting procedure does not match the line identified in the literature exactly, due to the different spectra[2] used for metal identification and different post-processing procedures, including coaddition, continuum fitting, and data smoothing used in our data. The aforementioned masking procedure makes sure that all potential metal contamination is removed. Afterwards, we manually masked a few gap regions in our quasar spectra resulting in the failure of the `VPFIT` caused by Damped Lyα absorption systems (DLAs). These masks are generated in post-processing, which means that we first apply `VPFIT` to the spectra assuming all lines are HI Lyα and remove the absorption lines that fall within the masked regions, same as done for finding overlapped lines with metal masks. In the end, we subtract the metal mask from our total pathlength and obtain $\Delta z = 2.097$. Our full sample of quasar segments and their corresponding masks are presented in Appendix B.1.

With our imposed cuts on the $\{b, N_{\mathrm{HI}}\}$, we find that 40 out of 341 lines are masked for our whole sample, and that leaves us with a $\{b, N_{\mathrm{HI}}\}$ dataset consisting of 301 Lyα absorption lines. We divide the 301 Lyα absorbers into three redshift bins: $0.9 < z < 1.1$, $1.1 < z < 1.3$ and $1.3 < z < 1.5$ centered at $z = 1$, 1.2 and 1.4, respectively, according to their central wavelength as determined by `VPFIT`. This provides us with the number of Lyα lines to be 102, 160 and 39 and redshift path of 0.762, 0.972 and 0.363 in the

---

[2] HST COS spectra are used in CASBaH project to identify the metal lines.

Table 3.2: Summary of the of the observational dataset

z bins	Δz	Number	$b_{\rm m}/{\rm km\ s^{-1}}$	$\log(N_{\rm HI,m}/{\rm cm^{-2}})$
$0.9 \leq z \leq 1.1$	0.762	102	31.74	13.48
$1.1 < z \leq 1.3$	0.972	160	28.83	13.37
$1.3 < z \leq 1.5$	0.363	39	29.69	13.48

The numbers of identified Lyα lines in each redshift, the total pathlength Δz, and the median value $b_{\rm m}$ and $\log N_{\rm HI,m}$.

bins centred at $z = 1$, 1.2 and 1.4, respectively. In Table 3.2 we summarize our $\{b, N_{\rm HI}\}$ dataset for each redshift bin, with redshift pathlength, number of final Lyα lines as well as median values for the b and $N_{\rm HI}$ in each bin.

3.2 Simulations

We utilize a set of Nyx cosmological hydrodynamic simulations from THERMAL suite, both described in§2.1. To model the Lyα forest at $z \sim 1$. We rescale the temperature for a set of models to model the Lyα forest in extremely hot IGM, which is favoured by our inference result, as described in §3.3. In addition, we implement the STIS LSF in the forward-modelling procedures, and use the different UVB photoionization rate $\Gamma_{\rm HI}$ in the post-processing due to the different redshift compared with §2. For other relevant post-processing procedures, such as generating mock spectra, and forward-modeling we follow the prescription discussed in §2.1, if no special treatment is needed.

3.2.1 T_0-rescaling models

As will be discussed later in § 3.3, our data favour models with high T_0 at $z = 1.0$ and $z = 1.4$, which is hard to generate based on the aforementioned procedure. This is because, as suggested by Eq. 2.1, our method alters the IGM thermal history of the simulation model by varying the heat released by the H I photoionization. However, the results of such a heating procedure fade away in low z, where the IGM is dominated by the adiabatic cooling caused by Hubble expansion (McQuinn 2016). As a result, the T_0 of the IGM at $z < 1.5$ becomes insensitive to the heat input in our method for models with high T_0. To this end, we rescale the IGM temperature to model the IGM with high temperature. For $z = 1.0$, we select six simulation snapshots with $3.75 \leq \log T_0 \leq 3.9$, which has T_0 close to the Nyx model 00 with $A = 1, B = 0$ (see Eq. 2.1) at $z = 1.0$, and multiply their temperature T (at each simulation cell) by 2.5 and 3 respectively to generate 12 new models. The other properties of the simulation remain unchanged, and since we rescaled the temperature of all simulation cells uniformly the whole Δ-T distribution of the simulation model still follows the power law Δ-T relationship Eq. 1.1 with the T_0 rescaled. The [T_0, γ] of original models and models with rescaled T_0 are illustrated in Fig. 3.6, where THERMAL Nyx models are plotted as blue dots, and the models with rescaled temperature are shown as orange dots. the model rescaled to $2.5\times T_0$ and $3.0\times T_0$ are shown in orange and red respectively. Such temperature rescaling procedures are also applied to $z = 1.4$ models, where our preliminary results also favour hot models, and the corresponding models are shown in Fig. 3.4.

3.2.2 Varying the UVB Γ_{HI}

Since we want to measure the ionization state of the IGM, we let the HI photoionization rate Γ_{HI} be a free parameter when generating Lyα forest skewers from our simulations. As such, we add an additional parameter $\log \Gamma_{\mathrm{HI}}$ to our thermal grid, extending it to [$\log T_0$, γ, $\log \Gamma_{\mathrm{HI}}$]. Such procedure is done in the post-processing of the simulation, at the time when the simulated slightlines are generated (see § 2.1.2). The value of Γ_{HI}

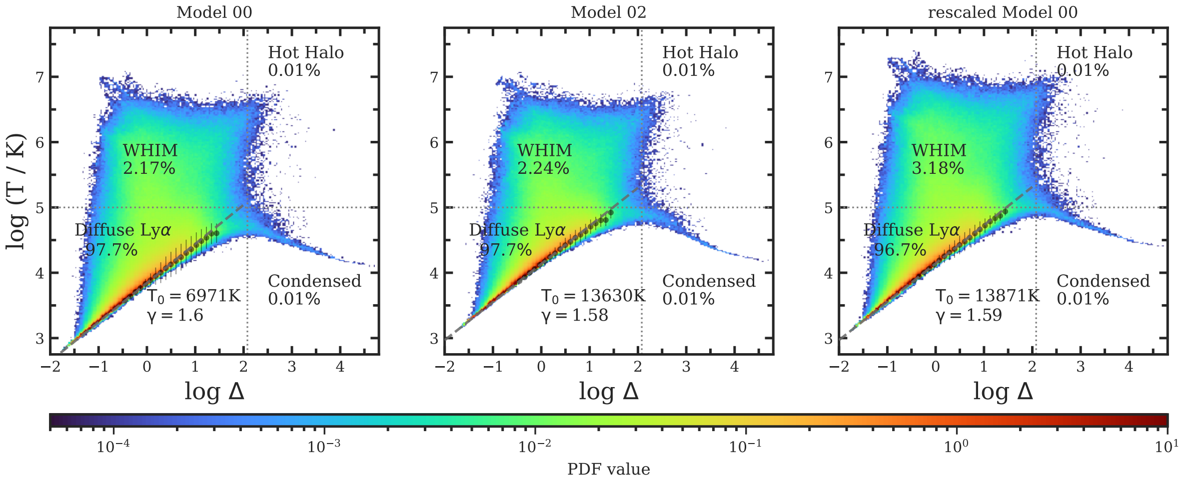

Figure 3.3: Volume weighted T-Δ distribution for three simulations models at $z = 1.0$. The left panel is the Nyx model 00 with $T_0 = 6971$ K, $\gamma = 1.60$, and the middle panel is the Nyx model 02 with $T_0 = 13630$ K, $\gamma = 1.58$. The right panel shows the model generated by multiplying the temperature in model 00 by two, resulting in a $T_0 = 13871$ K and γ =1.59 according to our $\Delta - T$ fitting procedure. The best-fit power-law relationship is shown as grey dashed lines. The $\log T$ for each bin are plotted as black dots, with the 1-σ_T error bars shown as black bars. The volume-weighted gas phase fractions are shown in the annotations. The fraction of diffuse Lyα gas and the values of T_0 and γ in the rescaled model (the right panel) agree within a percent level to the actual model shown in the middle panel.

we used in this study spans from $\log(\Gamma_{\rm HI}/{\rm s}^{-1})$ = -11.2 to -12.8 in logarithmic steps of 0.2 dex, which gives 9 values in total. These values are fixed for all redshift bins.

3.2.3 Forward Modeling of Noise and Resolution

In this section, we aim to measure the thermal and ionization state of the IGM at $z \sim 1$. To this end, we generate mock datasets with properties consistent with our STIS E230M quasar spectra, which comprise 12 unique quasar spectra. To this end, we use the forward-modeling procedure as described in §2.1.3, but with altered LSF and noise based on the observational data.

For low-z IGM with temperatures at mean density $T_0 \sim 5000$ K, the b-values for pure thermal broadening (i.e. the narrowest lines in the Lyα forest) are $b \sim 9$ km/s, corresponding to a full width at half maximum (FWHM)$\sim b/0.6 \sim 15$ km/s. Such absorption features can not be fully resolved by STIS, which has a resolution of roughly 10 km/s. Thus, it is crucial to treat the instrumental effect carefully. Therefore, we forward model noise and resolution to make our simulation results statistically comparable with the observation data. In practice, we make use of tabulated STIS E230M LSF obtained from `linetools` and noise vectors from our quasar sample. For any individual quasar spectrum from the observation dataset, we first stitch randomly selected simulated skewers without repetition to cover the same wavelength of the quasar and then rebin the skewers onto the wavelength grid of the observed spectra. Then we convolve the simulated spectra with the HST STIS LSF while taking into account the grating and slits used for that specific data spectrum. The STIS LSF is tabulated for up to 160 pixels in each direction. We interpolate the LSF onto the wavelengths of the mock spectrum (segment) to obtain a wavelength-dependent LSF. Each output pixel is then modelled as a convolution between the input stitched skewers and the interpolated LSF for the corresponding wavelength. Afterwards, the newly generated spectrum is interpolated to the wavelength of the selected STIS spectra. The noise vector of the quasar spectrum is propagated to our simulated spectrum pixel-by-pixel by sampling from a Gaussian with $\sigma = \psi_i$, with ψ_i being the data noise vector value at the i^{th} pixel. In the end, a fixed floor of 0.02 in quadrature is added to the error vector for all simulated spectra to avoid artificial effects in post-processing, as discussed in §2.1.4.

For each model, including both Nyx model from the THERMAL suite and those generated by rescaling the temperature, we generated 1000 mock spectra, from the 15,000 raw skewers[3]. The total pathlength of the dataset for each model is roughly $\Delta z_{\mathrm{tot}} \sim 100$. We then fit Voigt profiles to each line in the spectra to obtain the $\{b, N_{\mathrm{HI}}\}$ dataset used for the training of the b-N_{HI} distribution emulator, which will be discussed in § 5.2.1. For the purpose of illustration, an example of a forward-modelled spectrum is shown in Fig. 4.2 where the simulated spectrum is shown in grey, the model spectrum based on `VPFIT` line fitting is in blue, and the noise vector in red.

3.3 Inference Results

We applied the inference method described in §2.2 to our dataset at three redshift bins to measure the IGM thermal and ionization state at $z = 1.4$, 1.2, and 1.0. The difference is that, in this work, we train the neural networks based on $\{b, N_{\mathrm{HI}}\}$ dataset forward-modelled based on the STIS spectra described in §3.1. The neural networks are trained on mock dataset at $z = 1.4$, 1.2, and 1.0 individually.

The resulting MCMC posteriors are presented in Fig. 3.4, Fig. 3.5 and Fig. 3.6 respectively. Projections of the thermal grid used for generating models are shown as blue dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval.

---

[3] Generating 1000 spectra requires about 10,000 raw skewers, which are randomly selected from the total 15,000 skewers for each model.

Table 3.3: Summary of the inference results

z bins	$\log(T_0/\mathrm{K})$	γ	$\log\ (\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$
$1.3 < z \leq 1.5$	$4.119^{+0.152}_{-0.253}$	$1.341^{+0.208}_{-0.258}$	$-11.789^{+0.181}_{-0.147}$
$1.1 < z \leq 1.3$	$3.791^{+0.106}_{-0.107}$	$1.704^{+0.092}_{-0.094}$	$-11.984^{+0.089}_{-0.088}$
$0.9 \leq z \leq 1.1$	$4.132^{+0.115}_{-0.103}$	$1.357^{+0.102}_{-0.151}$	$-12.320^{+0.103}_{-0.115}$

The inference results i.e., median values of the marginalized 1D posteriors for each parameter, for all three redshift bins. The errors are given by the 1-σ error (16-84%) of the marginalized 1D posteriors.

The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior. For $z = 1.0$ and 1.4, our preliminary results indicate that the observational data favour models with high temperature, and the MCMC posterior is truncated at the boundary of the parameter space. As described in § 3.2.1, these models with high temperatures are hard to model due to the heating mechanism used in the Nyx simulation. We thus manually rescale the temperature of some of the Nyx models and extend the parameter grid for our inference procedure. With these rescaled models, we are able to measure the thermal and ionization state of the IGM. The parameter grids that contains the rescaled models are shown in Fig. 3.4 and Fig. 3.6 for z=1.4 and z=1.0 respectively. The Nyx models used for temperature rescaling are shown as green dots, and the models with 2.5 and 3.0 times T_0 are shown as orange and red dots respectively.

We summarize the inference results (median values of the marginalized 1D posteriors for each parameter) in Table. 3.3. The $\{b, N_{\mathrm{HI}}\}$ data and the corresponding b-N_{HI} distributions emulated by our DELFI emulator are shown in Fig. 4.8, and the likelihood contours corresponding to 80, 60, 40, and 20 cumulative percentiles are plotted as grey dashed lines. These plots show good agreement between the observational data and the emulated b-N_{HI} distributions. We notice the inference result at $z = 1.4$ has huge uncertainty due to the lack of observational data. However, the precision is still satisfactory, given the fact that our $\{b, N_{\rm HI}\}$ sample at this redshift bin contains only 39 data points. Such a size is comparable with the one used in Ricotti et al. (2000), whereas the error bar is much smaller (see Fig. 3.11), which mainly because of our novel method using full b-$N_{\rm HI}$ distribution (see Hiss et al. 2019, for the relevant discussion).

Based on the marginalized 2D posteriors, we observe that our results across all redshift bins exhibit the anticipated degeneracies between parameters. Specifically, T_0 is degenerate with both γ and $\Gamma_{\rm HI}$, as indicated in Hu22. To further assess the goodness of our inference results, we plot the marginalized 1D b and $N_{\rm HI}$ distributions of our sample in Fig. 3.8, Fig. 3.9, and Fig. 3.10 for each redshift bin, and compare them with 5000 mock datasets with the same size, sampled from the b-$N_{\rm HI}$ distributions emulated based on the median values of the MCMC posteriors. The blue bars indicate the mean value of the number of lines that fall in each bin for the 5000 datasets, whereas the blue shaded regions represent the 1-σ uncertainty calculated from the 5000 datasets. From the results, it is evident that our inference method adeptly recovers both the 2D and marginalized 1D distributions of $\{b, N_{\rm HI}\}$, even though the limited data size, particularly at $z = 1.4$, leads to noticeable fluctuations, which are underscored by the substantial 1-σ error bar in the marginalized 1D distributions in both b and $N_{\rm HI}$ distributions.

As illustrated in Fig. 3.8, Fig. 3.9, and Fig. 3.10, our 1D b-parameter distributions emulated for best fit [T_0, γ, $\Gamma_{\rm HI}$] are in good match with the observations, highlighting the robustness of our inference and suggesting that there is no severe discrepancy in b distribution as opposed to the what is seen at $z < 0.5$ (Gaikwad et al. 2017b, Viel et al. 2017). Note that this $z < 0.5$ b-parameter discrepancy arises from studies based on COS low-z Lyα spectra (Danforth et al. 2016), however, in reality, the spectral resolution and LSF of COS may not be very good for accurate b-parameter measurements, especially for small b values. In contrast, old studies on higher-resolution STIS spectra, although with high uncertainty, found observed b-parameter in good agreement with predictions from cosmological simulations (see Fig. 3 in Davé & Tripp 2001). This consistency implies that the b-parameter discrepancy found in the literature may be an artifact of the limited spectral resolution provided by COS, which will be further investigated in our future work. It also suggest that it might be beneficial to study the Lyα forest with the higher-resolution spectra obtained with STIS.

3.3.1 Evolution of the thermal state of the IGM

In Fig. 3.11 we summarize the T_0, γ evolution across three redshift bins, and compare them with archival from previous studies at higher z (Ricotti et al. 2000, Hiss et al. 2018, Walther et al. 2019b, Gaikwad et al. 2021). Our results and their 1-σ uncertainties are shown as filled red data points and error bars. As a benchmark for current theoretical models, we plot the IGM thermal history spanned by all potential Helium reionization models (Oñorbe et al. 2017b;a) as the cyan-shaded region.

To further assess how well do our low-z results agree with previous results, in Fig. 3.12, we fit a power law relationship between T_0 and z (blue dashed line), i.e., $\log T_0(z) = c_1 z + c_2$, where c_1, c_2 are fitting coefficients obtained from a least squares linear fit based on all previous T_0 measurements in between $1.5 \leq z \leq 3.0$ (i.e., not including our measurements). Such a power law fit is a reasonable approximation in between $1.0 \leq z \leq 3.0$ (see the prediction of low-z T_0 in Upton Sanderbeck et al. 2016, McQuinn & Upton Sanderbeck 2016). The power law relationship (blue dashed line in Fig. 3.12) suggests that our measurements at $z = 1.2$ and 1.4 are consistent with previous results. However, a noticeable discrepancy in T_0 emerges at $z = 1.0$, where our measurement of T_0 13500 K is significantly higher than best-fit power-law relationship predicted by previous measurements. Such a discrepancy suggests that the IGM may be far hotter

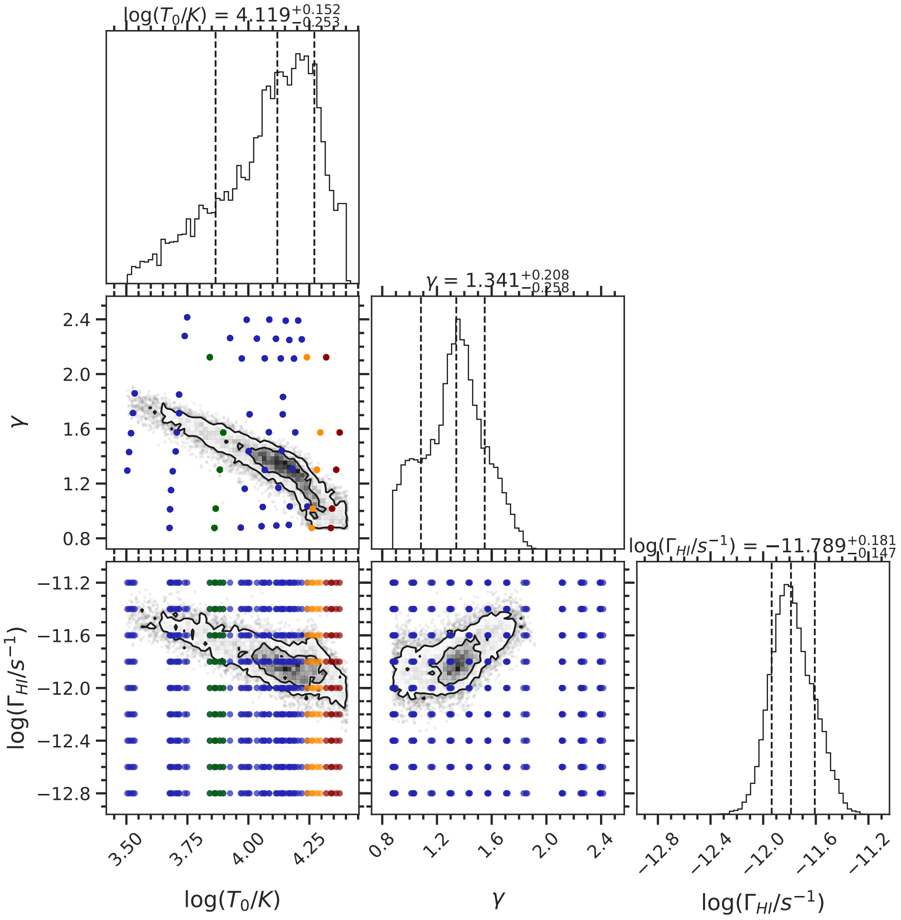

Figure 3.4: The MCMC posterior obtained by our inference method using our $\{b, N_{\rm HI}\}$ dataset at $z = 1.4$. Projections of the thermal grid used for generating models are shown as blue dots. The Nyx models used for temperature rescaling are shown as green dots, and the models with 2.5 and 3.0 times T_0 are shown as orange and red dots respectively. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

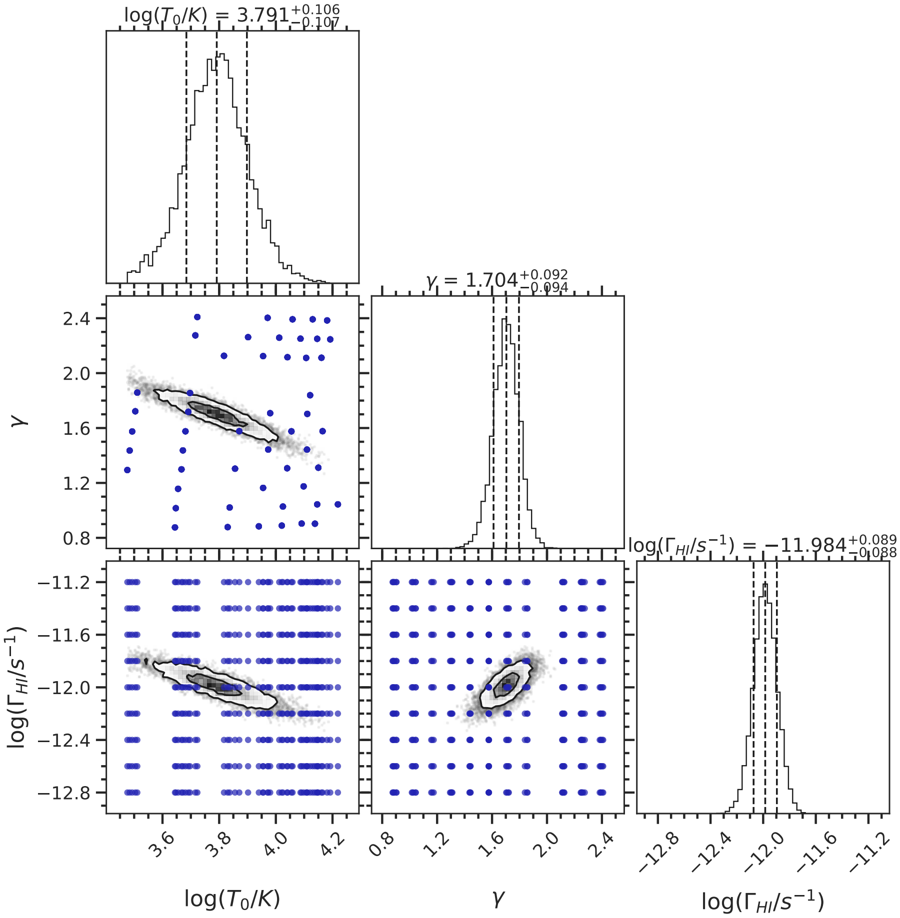

Figure 3.5: The MCMC posterior obtained by our inference method using our $\{b, N_{\rm HI}\}$ dataset at $z = 1.2$. Projections of the thermal grid used for generating models are shown as blue dots, while the true model is shown as red dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. Red lines in the marginal distributions indicate the parameters of true models, while the dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

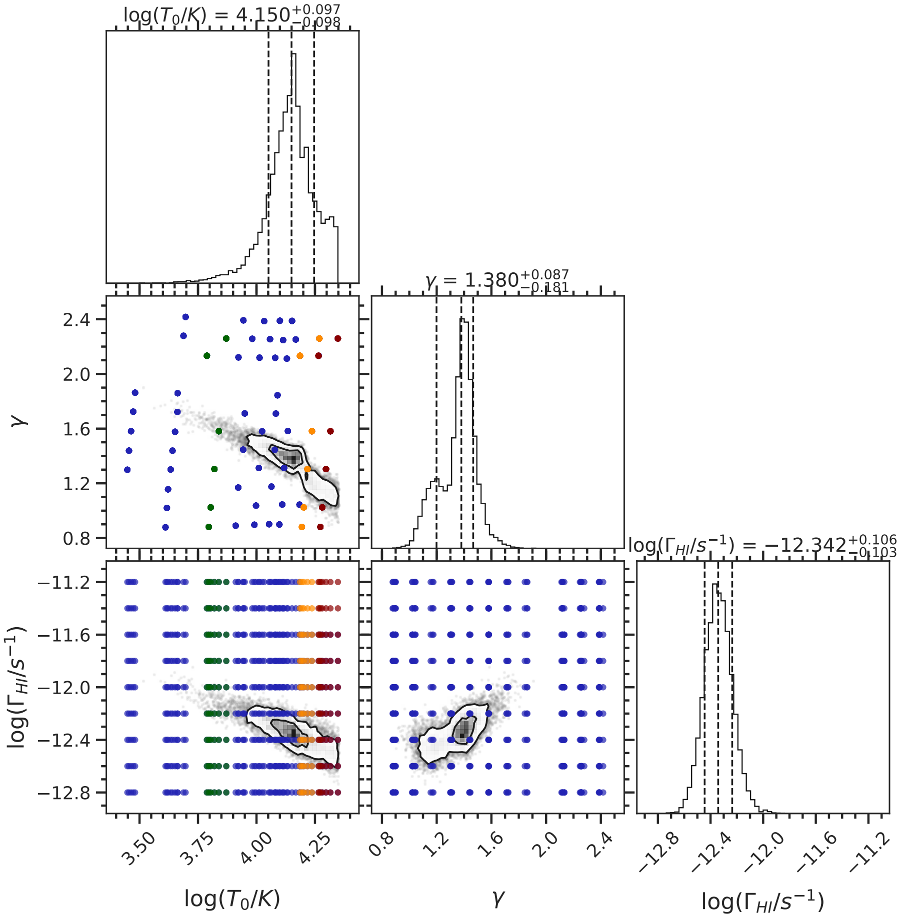

Figure 3.6: The MCMC posterior obtained by our inference method using our $\{b, N_{\mathrm{HI}}\}$ dataset at $z = 1.0$. Projections of the thermal grid used for generating models are shown as blue dots. The Nyx models used for temperature rescaling are shown as green dots, and the models with 2.5 and 3.0 times T_0 are shown as orange and red dots respectively. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

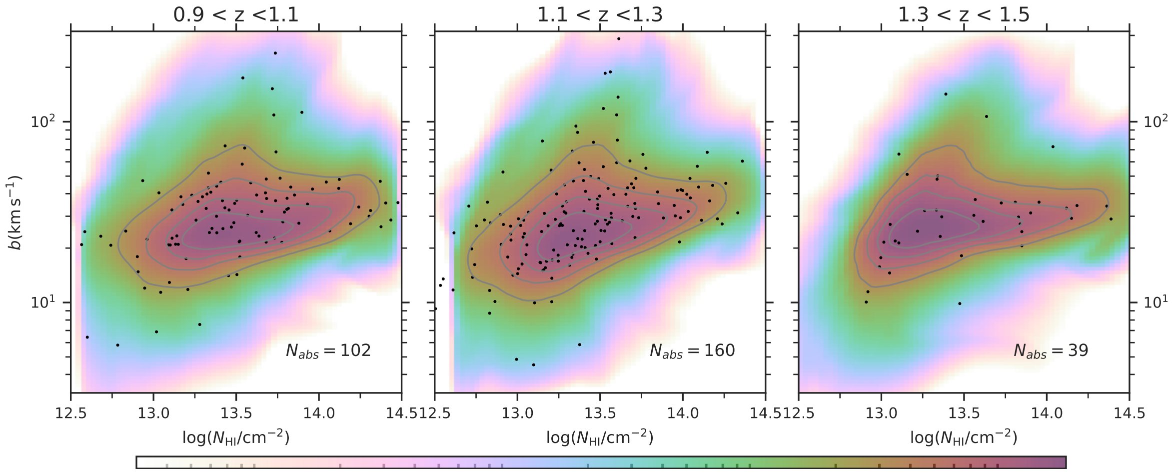

Figure 3.7: Joint b-$N_{\rm HI}$ distributions emulated by our DELFI emulator based on the median values of the marginalized MCMC posterior at $z = 1.0$, 1.2, and 1.4. Black dots are the $\{b, N_{\rm HI}\}$ data. The likelihood contours corresponding to 80,60,40, and 20 cumulative percentiles CDF are plotted as gray solid lines. For illustration purposes, the values of the pdf are multiplied by 100 in the colour bar.

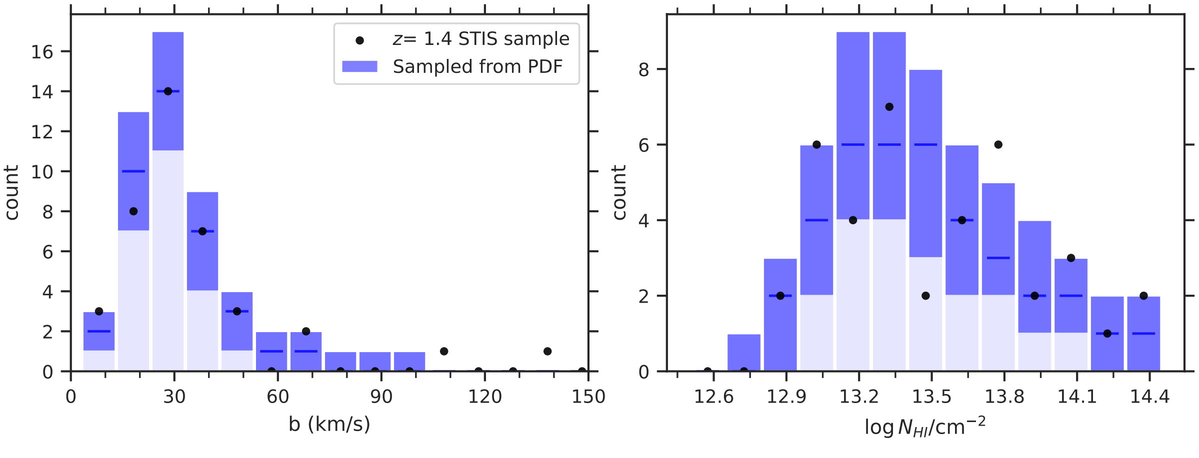

Figure 3.8: The marginalized 1D b and $N_{\rm HI}$ distributions of our are compared with 5000 mock datasets with the same size, sampled from the b-$N_{\rm HI}$ distributions emulated based on the median values of the MCMC posteriors. The black dots represent our $\{b, N_{\rm HI}\}$ data at z =1.4. The blue bars indicate the mean value of the number of lines that fall in each bin for the 5000 datasets, whereas the blue shaded regions represent the 1-σ uncertainty calculated from the 5000 datasets.

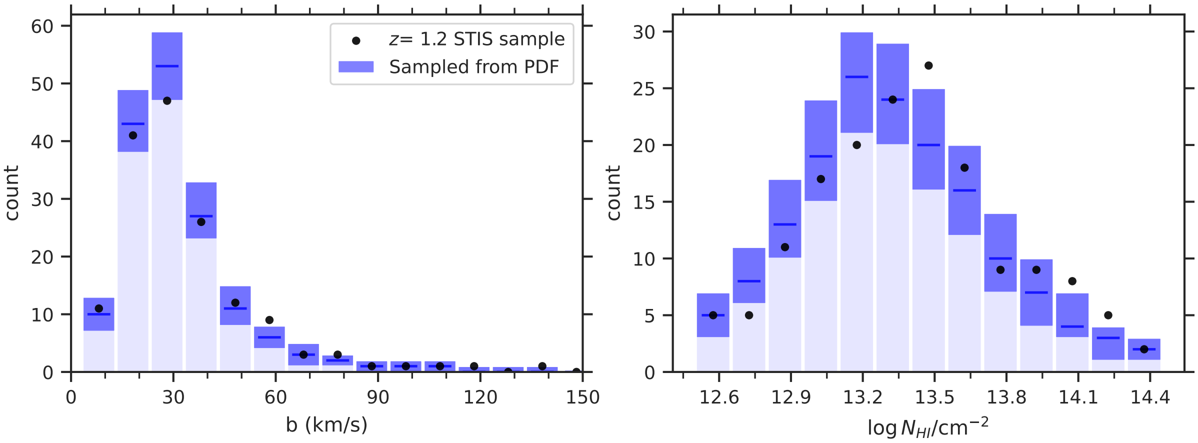

Figure 3.9: The marginalized 1D b and $N_{\rm HI}$ distributions of our are compared with 5000 mock datasets with the same size, sampled from the b-$N_{\rm HI}$ distributions emulated based on the median values of the MCMC posteriors. The black dots represent our $\{b, N_{\rm HI}\}$ data at z =1.2. The blue bars indicate the mean value of the number of lines that fall in each bin for the 5000 datasets, whereas the blue shaded regions represent the 1-σ uncertainty calculated from the 5000 datasets.

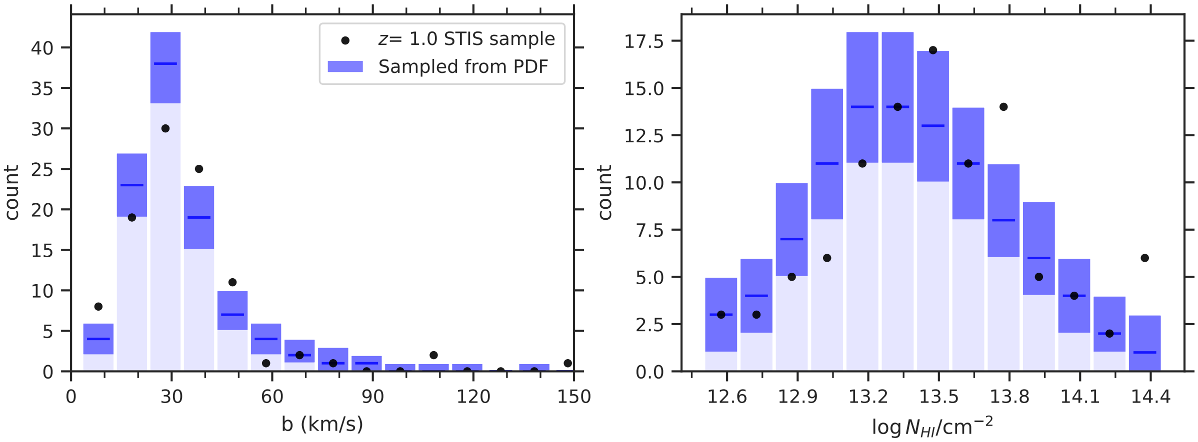

Figure 3.10: The marginalized 1D b and $N_{\rm HI}$ distributions of our are compared with 5000 mock datasets with the same size, sampled from the b-$N_{\rm HI}$ distributions emulated based on the median values of the MCMC posteriors. The black dots represent our $\{b, N_{\rm HI}\}$ data at z =1.0. The blue bars indicate the mean value of the number of lines that fall in each bin for the 5000 datasets, whereas the blue shaded regions represent the 1-σ uncertainty calculated from the 5000 datasets.

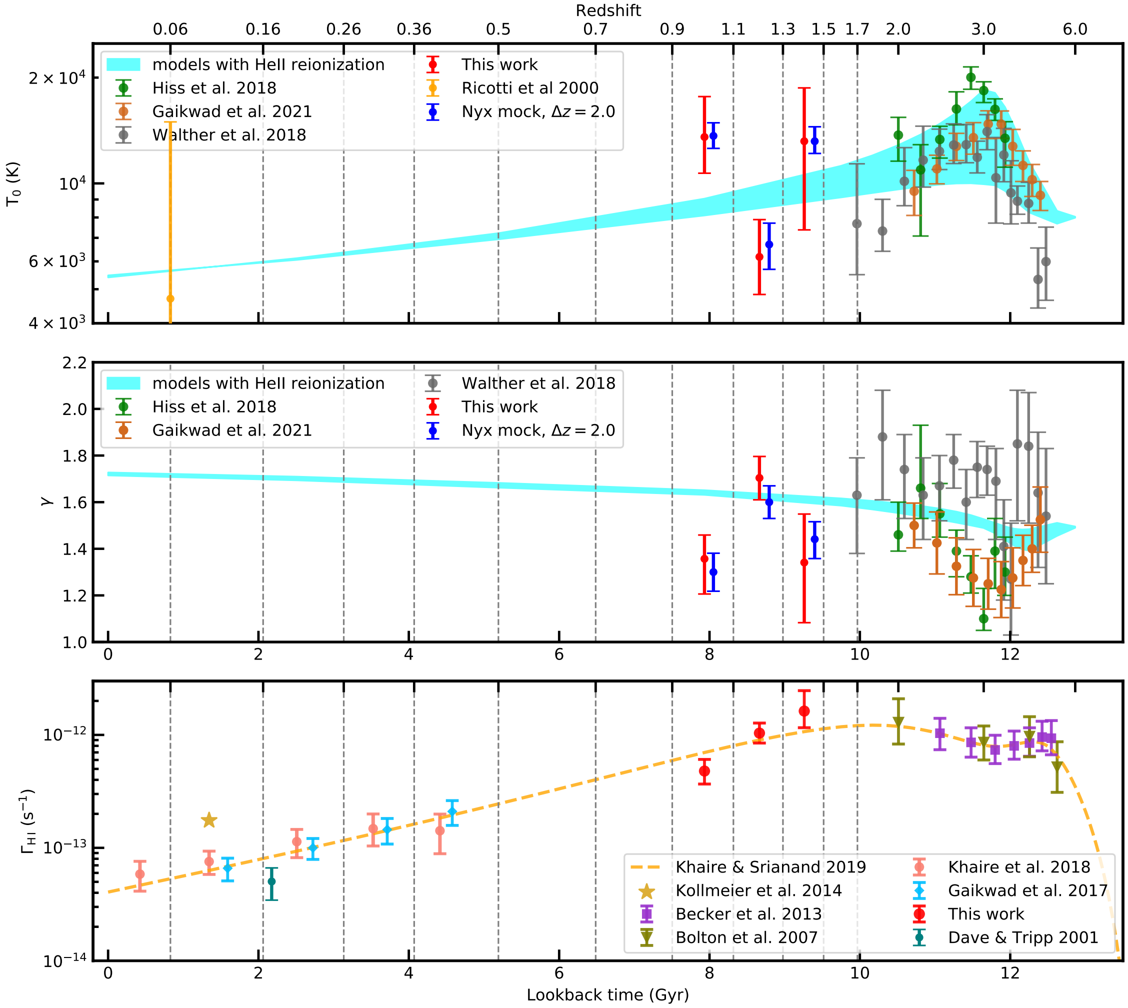

Figure 3.11: Evolution history of T_0 (top) and $\log\Gamma_{\rm HI}$ (bottom) based on our inference results using the STIS data. Our results are shown as red dots, while measurements from other studies are displayed in different colours. The error bars stand for the 1-σ error. The blue-shaded region in the top panel represents the range spanned by T_0 from hydrodynamical simulations of a large family of different HeII reionization models. The mock measurements based on Nyx simulation are shown in blue.

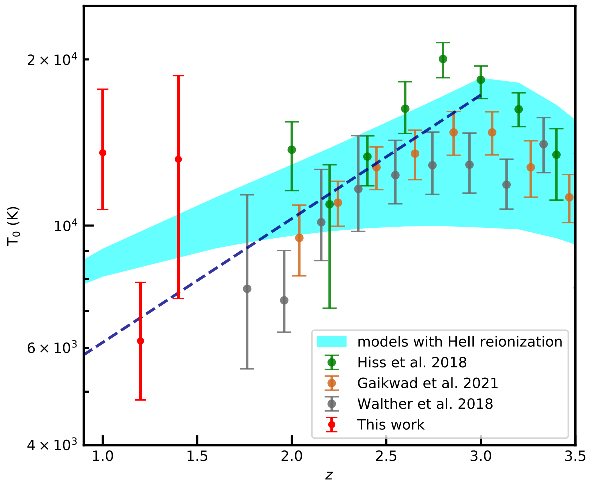

Figure 3.12: The evolution of T_0 in the IGM across $0.9 < z < 3.5$, with results from previous studies shown in comparison. The power law fit of $\log(T_0/\mathrm{K})$ obtained by fitting all previous results in between $0.9 < z < 3.0$ are plotted as dark cyan dashed line.

than expected at $z \sim 1.0$, implying the existence of extra heating sources that are not included in our current IGM model, which becomes crucial at $z \sim 1.0$. Summarizing the T_0 measurements across all three redshift bins, two potential thermal histories for the IGM emerge: (1) The IGM might undergo a cooling phase around $z \sim 1.2$ before heating up to 13500 K at $z \sim 1.0$, which is not unfeasible given the significantly large time span of $\sim$ 700 Myr between these two redshifts. (2) Alternatively, the IGM could consistently maintain a high temperature since $z \sim 1.5$. However due to the substantial error bars in T_0 in all three redshift bins, no definitive conclusion can be made until further investigation with larger datasets.

To further investigate the possibe change of the IGM thermal state from $z = 1.2$ to $z = 1.0$, in Fig. 3.13, we over-plot the likelihood contours of the b-N_{HI} distribution at $z = 1.2$ on top of the $\{b, N_{\mathrm{HI}}\}$ dataset and the corresponding b-N_{HI} distribution at $z = 1.0$. It can be seen that the $\{b, N_{\mathrm{HI}}\}$ dataset and the corresponding b-N_{HI} distribution at $z = 1.0$ lies above the likelihood contours of the b-N_{HI} distribution at $z = 1.2$, suggesting

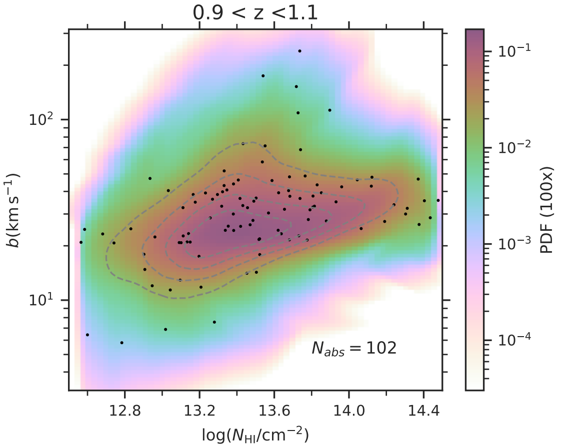

Figure 3.13: The likelihood contours of the b-$N_{\rm HI}$ distribution at $z = 1.2$ on top of the $\{b, N_{\rm HI}\}$ dataset and the corresponding b-$N_{\rm HI}$ distribution at $z = 1.0$. The likelihood contours corresponding to 80,60,40, and 20 cumulative percentiles are plotted as gray dashed lines.

that our observational data indeed favour a rapid change in the IGM thermal state between $1.0 < z < 1.2$. More discussion on this unexpected high T_0 is present in § 3.4.1.

As for γ, our results for $z = 1.4$ and 1.2 align with this trend as outlined in McQuinn (2016), in which the value of γ tends to decrease towards lower redshifts. However, the result at $z = 1.0$ indicates a reduced γ. The cause of this discrepancy remains unclear, but it is worth noting that such a trend of T_0 and γ, i.e., high T_0, low γ, is consistent with the T_0-γ degeneracy shown in the inference posterior (see the 2D marginalized posterior contours in T_0-γ plane in Fig. 3.6). As a result, it is likely that the inference results at $z = 1.0$, which yields high T_0 and low γ are caused by inference uncertainty and degeneracy. On the other hand, it is also possible that the IGM starts to heat up at $z \sim 1.0$, leading to both increasing T_0 and decreasing γ. In this case, the inconsistencies observed in both γ and T_0 have a common root cause.

To illustrate the evolution of the IGM thermal and ionization state, we over-plot the three MCMC posterior on top of each other in Fig. 3.14, where the 2D marginalized posterior for $z = 1.4$ are shown in green, the one for $z = 1.2$ is plotted in blue, and the

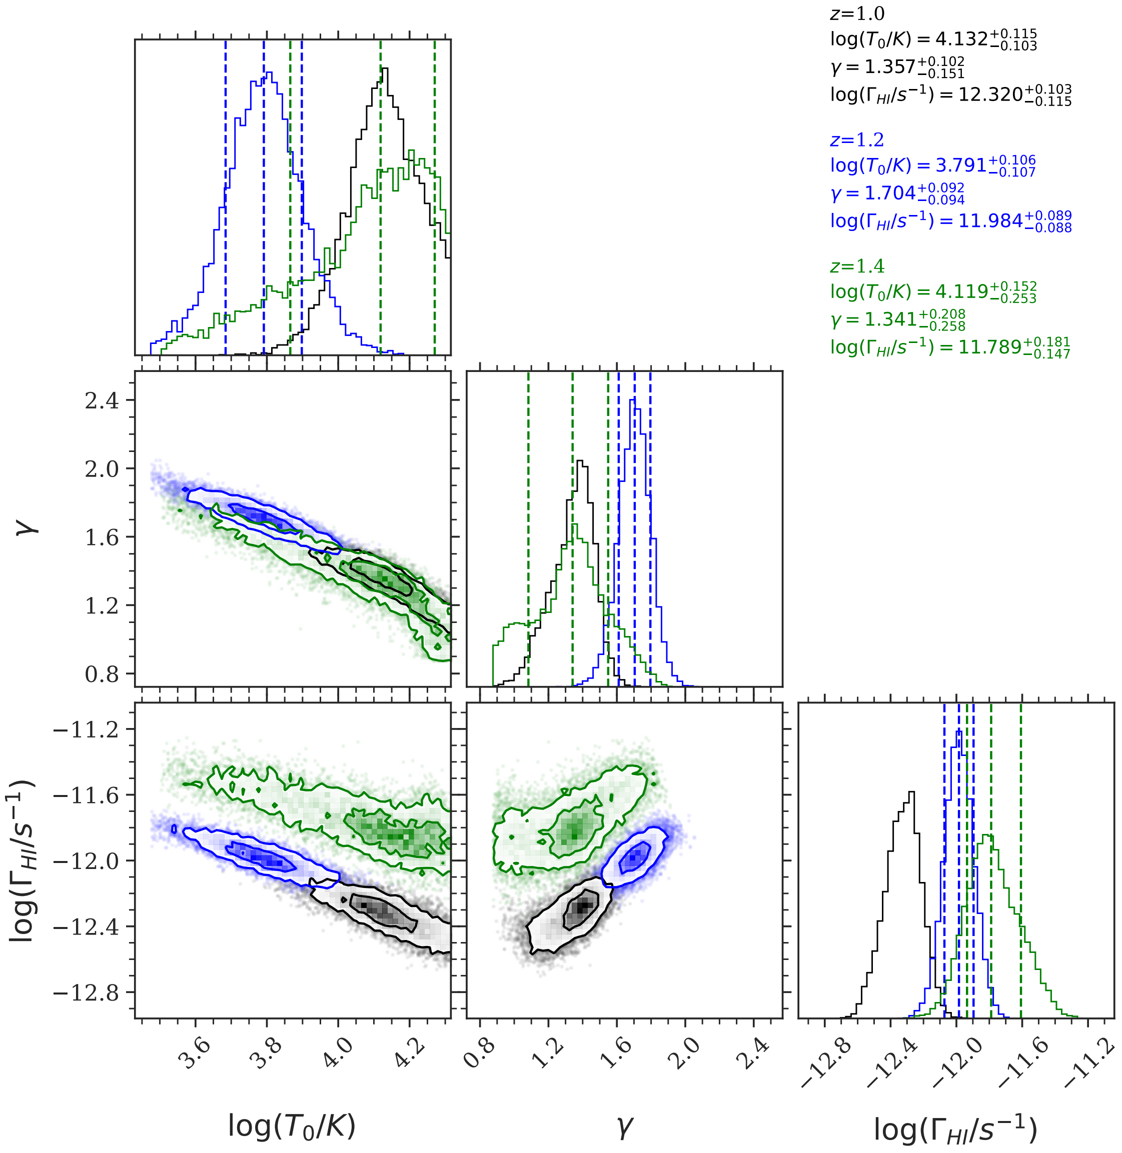

Figure 3.14: The MCMC posteriors obtained by our inference method for all three redshift bins. The $z = 1.0$ posterior is shown in black, the $z = 1.2$ posterior is shown in blue, and the $z = 1.4$ posterior is shown in green.

one for $z = 1.0$ is shown in black. From the T_0-γ plane, we observe a clear turnover for both T_0 and γ at $z = 1.0$, suggesting a reverse evolution trend at $z \sim 1.0$. Such synchronization between the evolution of T_0 and γ is important for us to understand the origin of the discrepancy, and relevant discussion is presented in § 3.4.1.

3.3.2 Evolution of the H I photoionization rate and UVB

Our measurements fill in the $\Gamma_{\rm HI}$ evolution history between $0.0 < z < 1.7$. In the bottom panel of Fig. 3.11 we show our $\Gamma_{\rm HI}$ measurements across our three redshift bins, compared with previous studies (Davé & Tripp 2001, Bolton 2007, Becker & Bolton 2013b, Kollmeier et al. 2014, Gaikwad et al. 2017b, Khaire et al. 2019). Our inference results indicate that the $\Gamma_{\rm HI}$ is in good agreement with the UVB model presented in Khaire & Srianand (2019b) in all three redshift bins.

It is worth noting that for $z < 3$, the UVB model of Khaire & Srianand (2019b) is dominated by photons emitted by quasars alone i.e., the escape fraction on ionizing photons from galaxies is negligible at $z < 3$. Our $\Gamma_{\rm HI}$ measurements support the same conclusion that galaxies are not the main source of ionizing photons at $z < 3$. The same conclusion can be drawn from the new UVB models of Puchwein et al. (2019) and Faucher-Giguère (2020) because their $\Gamma_{\rm HI}$ values align very well with the UVB model of (Khaire & Srianand 2019b) at $z < 2$. This is mainly because all three UVB models use updated quasar luminosity functions at $z < 3$ (as presented in Croom et al. 2009, Ross et al. 2013, Palanque-Delabrouille et al. 2013) after Khaire & Srianand (2015) pointed out that previous UVB models (Haardt & Madau 2012, Faucher-Giguère et al. 2009) used old quasar luminosity functions that predict factor two smaller ionizing emissivity. The consistency of our new $\Gamma_{\rm HI}$ measurements in the previously unexplored redshift range with recent UVB models attests to the robustness of these UVB synthesis models, especially in the aspect of hydrogen ionizing part of the UVB.

3.4 Discussion

3.4.1 The discrepancy in T_0

In this section, we delve into the observed discrepancy in IGM thermal state at $z \sim 1.0$. First of all, we notice a coherence between the high T_0 measured at $z \sim 1$ and the high b-values observed at $z \sim 0.1$, based on the COS Lyα forest dataset (Danforth et al. 2016), where the observed b-parameter significantly surpass the predicted value based on various simulations (Gaikwad et al. 2017b, Viel et al. 2017, Nasir et al. 2017, Bolton et al. 2022b;a). Quantitatively, Viel et al. (2017) compares the marginalized b distribution with various simulations, showing that the b distribution at $z \sim 0.1$ can be best recovered by the hydrodynamic simulations (`P-GADGET-3`, see Springel et al. 2005) with $T_0 \gtrsim 10000$ K, while the theoretical model dictates that the $T_0 \sim 5000$ at $z = 0.1$. The similarity of required IGM temperature at both $z = 0.1$ and 1.0 suggests that the discrepancy at $z \sim 0.1$ may be related the one at $z \sim 1$, indicating a persistent trend from $z \sim 1.0$ to 0.1. Additionally, it also suggests that the discrepancy observed at $z = 0.1$ may not be attributable to the limited resolution of the COS.

The simplest explanation for these discrepancies is the thermal broadening caused by a higher-than-expected IGM temperature, which requires the existence of extra heating sources. If this is true, our understanding of IGM physics will be changed drastically, highlighting a severe need to investigate processes that are possibly responsible for it, such as dark matter annihilation (Araya & Padilla 2014, Bolton et al. 2022a), gamma-ray sources (Puchwein et al. 2012), or feedback from galaxy formation, whose effects are not fully understood in low-z (see Springel et al. 2005, Croton et al. 2006a, Sijacki et al.

2007a, Hopkins et al. 2008a, Tillman et al. 2023b;a, Hu et al. 2023a).

Another possible explanation instead of extra heating is the presence of unexpected non-thermal broadening mechanisms affecting the b-parameter of the Lyα forest, such as micro-turbulence motion in the IGM induced by jet or feedback (Gaikwad et al. 2017b, Viel et al. 2017, Nasir et al. 2017, Bolton et al. 2022b). However, these non-thermal broadening models fail to account for the unexpected trend in γ observed in our results, where the γ are lower than expected at $z = 1.0$. To further investigate this, we plan to apply our inference method to the COS Lyα forest dataset at $z \leq 0.5$, which should help to break the degeneracy between T_0 and γ, thereby providing deeper insight into the b discrepancy observed at $z \sim 0.1$.

3.4.2 Forecast based on mock observations

In this section, we make realistic forecasts for our future measurements with more abundant observational data. Given the amount of the newfound bright objects expected in upcoming surveys including *Gaia* DR3 (Gaia Collaboration et al. 2016; 2023). With a realistic amount of the observation from HST STIS, i.e., $\sim$ 50 orbits, we expect the path length coverage for each redshift bin to be significantly extended. Here we assess the constraining power based on total pathlength $\Delta z = 2$ for each redshift bin, corresponding to three times the current data size or roughly 15 spectra for each redshift bin, while assuming the characteristic SNRs of the data do not change. We pick forward-modelled mock spectra from our mock dataset at each redshift bin and generate mock observational data with total pathlength Δz=2. The Nyx model used here is the one with the thermal state that is closest to the inference results presented in § 3.3. The inference results obtained from these mock observations and their 1-σ error bars are shown in Fig. 3.11 as blue dots. It can be seen that with $\Delta z = 2$, the 1-σ errors for T_0 become roughly $\sim$ 1500

K, and the 1-σ errors for γ become roughly ~ 0.08. These results will help us to confirm whether the IGM cools down as predicted.

3.4.3 The effect of potential contamination

In spite of the careful masking procedure, our $\{b, N_{\rm HI}\}$ dataset still encounters potential contaminants, including blended lines and unidentified metal lines, especially for the metal masks obtained from Milutinović et al. (2007), since their metal identification might not be complete. Here we briefly discuss the potential effects of these contaminants. It is well known that ionic metal line transitions mainly contribute to narrow absorption lines with $b \leq 10$ km/s (Schaye et al. 1999, Rudie et al. 2012b, Hiss et al. 2018). As a result, the metal line contaminants tend to bias our inference toward lower T_0. To this end, these contaminants shall not affect the main and most important result of this paper, i.e., the IGM seems to be hotter than expected at low-z, especially at $z = 1.0$. For these blended lines, in this paper, we adopt a more conservative metal masking, where we manually filter out all suspicious lines close to the masked regions (see the masks in Appendix B.1). As for a more detailed quantitative analysis, we plan to identify all Lyα lines using the Lyβ (or higher transitions) forest (see e.g. Rudie et al. 2012b). We plan to do this in future by combining our data set with other archival and upcoming data form HST.

3.4.4 The effect of SNRs of the spectra

We notice that a few quasar sightlines in our sample have relatively low SNRs (see Table. 3.1), and it is unknown whether our results are biased by these spectra. Hence, in this section, we test the effect of these low-SNR sightlines on our inference results. To do that, we exclude three quasar spectra from our sample which have relatively lower SNR

(≤ 5), while the remaining spectra all have SNR >5. We exclude TON153, PG1248+401 and PG1241+176 from the observational data and obtain a new $\{b, N_{\rm HI}\}$ dataset, which provides 25 fewer Lyα lines compared with the old one and reduces the total pathlength Δz by 0.24. We generate new mock datasets based on the nine spectra with SNR > 7, and train our emulators based on the new dataset. The outcomes indicate that even after excluding low SNR spectra from our data (and correspondingly in our mock data), the results remain consistent across each redshift bin. Such a result is important for our future work, suggesting that it is possible to make use of relatively low SNR data to obtain higher total pathlength and analyse the evolution of the thermal and ionization state on finer redshift bins, such that we could pinpoint the onset of the discrepancy in T_0 (or b-parameter) between the observation and simulation more precisely.

3.5 Summary

In this chapter, we make use of 12 archival STIS E230M quasar spectra, from which we obtain the b-$N_{\rm HI}$ distribution distribution and line density $\mathrm{d}N/\mathrm{d}z$ over the redshift range $0.9 < z < 1.5$ in three redshift bins. We then measure the thermal and ionization state of the IGM following a machine-learning-based inference method presented in Hu22 for this redshift range for the first time. Below we summarize our resutls:

- We Voigt-profile fit the Lyα in all 12 quasar spectra using a fully automated `VPFIT` wrapper and obtain $\{b, N_{\rm HI}\}$ for 341 lines. We use the metal identifications from the CASBaH project and combine them with the metal identification from Milutinović et al. (2007) to generate our metal masks, filtering out 40 contaminants besides Lyα absorption lines, and obtain a final sample of 301 Lyα lines across a total pathlength of Δz =2.097.

- We employ the Hu22 inference method, which simultaneously measures $[T_0, \gamma, \Gamma_{\rm HI}]$ from the b-$N_{\rm HI}$ distribution and $\mathrm{d}N/\mathrm{d}z$, with the help of neural density estimators and Gaussian process emulators trained on a suite of 51 Nyx simulations each having a different IGM thermal history. It enables us to measure the IGM thermal and ionization state with high precision even with limited data.

- We obtain $[\log(T_0/K), \gamma] = [4.119^{+0.152}_{-0.253}, 1.341^{+0.208}_{-0.258}]$ at $z = 1.4$ and $[\log(T_0/K), \gamma] = [3.791^{+0.106}_{-0.107}, 1.704^{+0.092}_{-0.094}]$ at $z = 1.2$. These two measurements agree with the theoretical model (Fig. 3.11 and 3.12), suggesting that the thermal state of the IGM evolves as expected from $z = 1.4$ to $z = 1.2$.

- Nevertheless, our results yield $[\log(T_0/K), \gamma] = [4.132^{+0.115}_{-0.103}, 1.357^{+0.102}_{-0.151}]$ at $z = 1$, suggesting an unexpectedly high IGM temperature and low γ, which is against the trend predicted by the current theoretical models of the IGM. Such high T_0 potentially suggests the existence of extra heating or unexpected non-thermal broadening at $z \sim 1.0$.

- Based on our measurements, it is possible that the IGM experiences a cooling phase until $z \sim 1.2$ from $z \sim 3$, and then it gets heated up to 13500 K at $z = 1$ in approximately 700 Myr. Alternatively, the IGM temperature might have remained consistently high since $z \sim 2$. However, due to significant uncertainties in T_0 for all three redshift bins, a definitive conclusion cannot be reached without further investigation.

- The inference results of γ suggest that it also goes through unexpected evolution at $z \sim 1$. However, while it is likely that such a trend is caused by extra heating that causes the discrepancy in T_0, it is also possible that it is due to inference degeneracy between T_0 and γ.

- We compare our findings with previous work, which reports unanticipated high b-parameters compared with various simulations based on observational data at $z \sim 0.1$. These high b values, if caused by thermal broadening, correspond to an IGM temperature with $T_0 \sim 10000$ K. This convergence towards a higher IGM temperature aligns with our findings and suggests that the discrepancy in b-parameter observed at $z \sim 0.1$ (Gaikwad et al. 2017b, Viel et al. 2017) could be related to the one we have identified in this study. It further implies that the observed discrepancy may emerge around $z \sim 1.0$ and persist down to $z \sim 0$.

- We successfully measure the $\Gamma_{\rm HI}$ at three redshif bins, reporting $\log(\Gamma_{\rm HI}/s^{-1}) = -11.789^{+0.181}_{-0.147}$, $-11.984^{+0.089}_{-0.088}$, and $-12.320^{+0.103}_{-0.115}$ at $z = 1.4, 1.2$ and 1.0 respectively. These measurements agree well with the theoretical values based on recent UVB synthesis models (Khaire & Srianand 2019b, Puchwein et al. 2019, Faucher-Giguère 2020), reinstating the fact that low-z UV background (at $z < 3$) is dominated by radiation from quasars alone.

- By excluding three spectra with relatively low SNRs from our quasar sample, we confirm that our results are not sensitive to the SNR of the dataset, suggesting that it is feasible to conduct our analysis on larger quasar samples with lower SNR to make finer measurements of the IGM thermal and ionization state, so as to pinpoint the onset of the discrepancy in the IGM thermal state between the observation and simulation more precisely.

- We perform mock measurements using realistic datasets based on Nyx simulation to forecast the constraining power for our future work. The results demonstrate that with redshift pathlength $\Delta z \sim 2.0$ for each redshift bin, three times the current data size, we will be able to constrain the T_0 within ± 1500 K. This precision will help us to constrain the thermal history of the IGM in $0.9 < z < 1.5$, and confirm whether the IGM cools down as expected at $z \sim 1.0$.

Chapter 4

Measurements of the IGM Thermal and Ionization State IGM at $z < 0.5$

In this Chapter, we employ the inference method presented in §2.2 to precisely measure both the thermal and ionization state of the IGM in four redshift bins from $z = 0.06$ to 0.48, using 82 archival HST COS quasar spectra reduced by Danforth et al. (2016, hereafter D16), which also provide the corresponding metal identification. To keep consistency in our analysis, instead of using the $\{b, N_{\rm HI}\}$ dataset from D16, we fit these spectra to obtain our $\{b, N_{\rm HI}\}$ sample using our own line fitting program, which is also applied to our forward-modelled mock spectra (see § 2.1.4). Our measurements suggest that the IGM at $z < 0.5$ are way much hotter than expected and close to isothermal, with $T_0 \sim 30000$K, and $\gamma \sim 1.1$ at z=0.1. Interestingly, owing to the degeneracy between IGM thermal parameters and $\Gamma_{\rm HI}$, if the IGM temperatures are higher than expected, the resulting lower $\Gamma_{\rm HI}$ values inferred could pose a challenge to existing UV background synthesis models (Khaire & Srianand 2019a, Puchwein et al. 2019, Faucher-Giguère 2020).

As an alternative explanation, small-scale turbulence could have similar effects on the b-$N_{\rm HI}$ distribution of the Lyα forest. Therefore, we model the IGM small-scale turbulence by a Gaussian velocity component along the line-of-sight, $v_{\rm tur}$, and add it to our simulation in post-processing. We assume a standard IGM evolution, i.e, with T_0 approaching 5000K and γ approaching 1.6 at $z = 0$. We generate a simulation parameter grid consisting of different $v_{\rm tur}$ and $\Gamma_{\rm HI}$ values. we apply the aforementioned inference method on this parameters grid and obtain that,

This paper is organized as follows. In Section 4.1, we introduce our observational data and the data processing procedures, which include continuum fitting, Voigt profile fitting, and metal masking. Section 4.2 describes our hydrodynamic simulations, the parameter grid, and mock data processing procedures, such as generating the Lyman-alpha forest from simulations, creating mock sightlines, and forward-modelling. Our results are discussed in Section 4.4, and we summarize the key findings of this study in Section 4.5.

4.1 Observational Data

The dataset we analyze is the publicly available compilation[1] of high signal-to-noise ratio HST/COS spectra published by Danforth et al. (2016). Consisting of 82 quasar spectra observed between 2009 and 2013 with the G130M ($900 \sim 1450$ Å) and G160M ($1360 \sim 1775$ Å) gratings, this dataset represents the largest publicly available low redshift ultraviolet (UV) survey of the Lyα forest to date. The nominal resolution of COS is $R \sim 12,000 - 20,000$ depending on the wavelength and grating, which corresponds to roughly $15 \sim 25$ km s^{-1}, and has a non-Gaussian LSF[2]. Individual spectra were co-added, taking into account all exposures and gratings, and then continuum-fitted by Danforth et al. (2016). We then use our automated VP(Voigt-Profile)-fitting program,

---

[1] `http://archive.stsci.edu/prepds/igm/`

[2] For reference see `https://hst-docs.stsci.edu/display/COSIHB/3.3+The+COS+Line-Spread+Function`

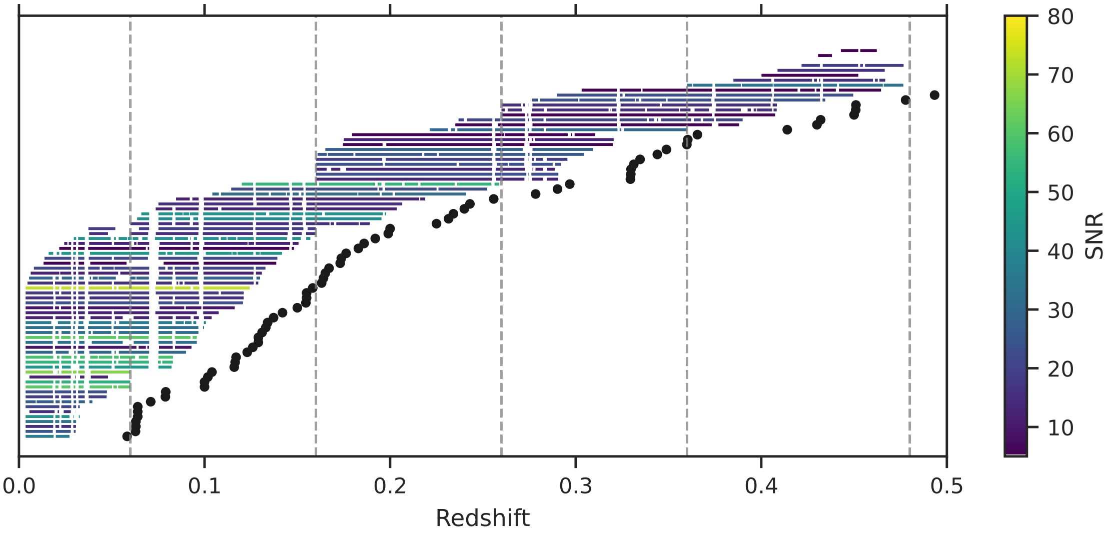

Figure 4.1: The HST COS spectra used in this study. The quasar are shown as black dots, and the Lyα spectra, with proximity zones removed, are shown as line segments with their colour indicating the mean SNR (per pixel), and the gaps represent the masked regions. The four redshift bins used in this study are shown by the vertical dashed lines.

which is described in § 2.1.4, to identify and fit all absorption and emission lines in the aforementioned spectra set.

To ensure that our $\{b, N_{\rm HI}\}$ dataset contains only the Lyα forest line, we mask all intervening metal absorbers identified by Danforth et al. (2016), which include both intervening metal lines and the lines arising from the MW absorbers at $z = 0$. We mask all emission lines, whereas the mask is adjusted by eye to include the full emission profiles, and gaps in the wavelength coverage. The procedures for generating masks are detailed in §4.1.2.

In this study, we focus exclusively on the Lyα forest. Therefore, we utilize only the Lyα regions, omitting Lyβ and higher Lyman series absorption lines at wavelengths shorter than 1050 Å, and masking the quasar proximity zones at wavelengths greater than 1180 Å (see Fig.4.1). This limits our analysis to the spectral segment with rest frame wavelengths between 1050 Å and 1180 Å. To mitigate edge effects at the spectral edges, quasar sightlines are segmented and padded with white noise based on the noise vector of the spectrum before being fed into the VP-fitting program. These padded regions are subsequently masked in post-processing. We apply similar edge treatments to the mock forward models to maintain consistency in our analysis.

We show the redshift path that is covered by the data segments (after all masking) in Fig. 4.1. The lines are coloured based on the sightline SNR, and the gaps in the spectra correspond to masked regions. It is noticeable that some gaps appear at the same wavelength for different sightlines. These lines are caused by metal absorbers in the MW, which always lies at $z \sim 0$.

4.1.1 Voigt-Profile Fitting

In this work, we use the line-fitting program presented in §2.1.4. Such a VP-fitting procedure is applied to the whole spectral segment, fitting both the Lyα lines and metal lines, including both intervening metal lines and those from interstellar medium of MW; for simplicity, hereafter we refer to these collectively as metal lines. The removal of these metal lines is later discussed in § 4.1.2.

Our `VPFIT` wrapper is designed to incorporate with custom LSF, which is important for COS due to its non-Gaussian LSF. In addition, the effective resolution of the HST COS gratings also depends on its lifetime position during the observations, and they are also taken into account in our VP-fitting program[3], as well as in our forward modelling procedures (see §2.1.3). Such VP-fitting program is applied to both observed and stimulated spectra so as to make sure our statistics are not biased.

The D16 Lyα forest quasar spectra consists of sightlines from both COS G130M and G160M grating, which have wavelength coverage $900 \sim 1450$ Å and $1360 \sim 1775$ Å

---

[3] Although our `VPFIT` wrapper is compatible with LSF in `VPFIT`, only a single LSF without wavelength dependence can be used at once. As such, for the input into `VPFIT`, we use the LSF at the lifetime of the data and evaluate it at the central wavelength of the spectrum that we are trying to fit.

respectively. These spectra are co-added, making the actual LSF/resolution complicated at the overlapping wavelength. After inspecting the flux and noise of individual exposure covering the overlapping region, we decide to VP-fit the spectra segment at $0.06 < z < 0.16$ using G130M LSF only and fit the spectra in all other bins using the G160M LSF only. In practice, such an arrangement only affects the redshift bin centering at $z = 0.2$, and it prevents us from chopping the spectra in the middle of the redshift bin, which causes more edge effects, inducing potential errors. This arrangement is also applied to our forward modelling procedures (see §2.1.3).

During our VP-fitting procedure, we observed weak artificial lines in the HST COS spectra that were absent from our forward-modelled mock spectra. A visual assessment suggests that these minor features are artefacts arising from issues like flat fielding, continuum placement, or data reduction, especially in spectra with high SNR. Therefore, we introduced a fixed 'floor' of 0.02 in quadrature to the normalized flux noise vector across all spectra, effectively adding robustness without increasing the noise in the normalized flux. This adjustment was derived through trial and error, according to the presence of absorption lines with low b and N_{HI} identified by `VPFIT` in the highest SNR spectra. These faint, narrow lines, absent in our simulated and forward-modelled sightlines, primarily affect lines with $\log N_{\mathrm{HI}}/\mathrm{cm}^{-2} < 12.5$ in our dataset, which are excluded from inference. For consistency, we applied the same noise floor to the simulated datasets in data processing (see § 2.1.3).

One of our COS spectra, phl1811, and the corresponding VP-fitting profile with metal masking are shown in Fig, 4.2. The continuum normalized spectrum is plotted in gray, and the fit to the unmasked spectrum is shown in blue and consists of the identified Lyα absorbers. The parts of the fitted spectrum shown in orange illustrates the segments that were masked based on the metal identifications reported in Danforth et al. (2016). The Lyα lines used for our $\{b, N_{\mathrm{HI}}\}$ dataset (after all filters) are labelled by red vertical lines.

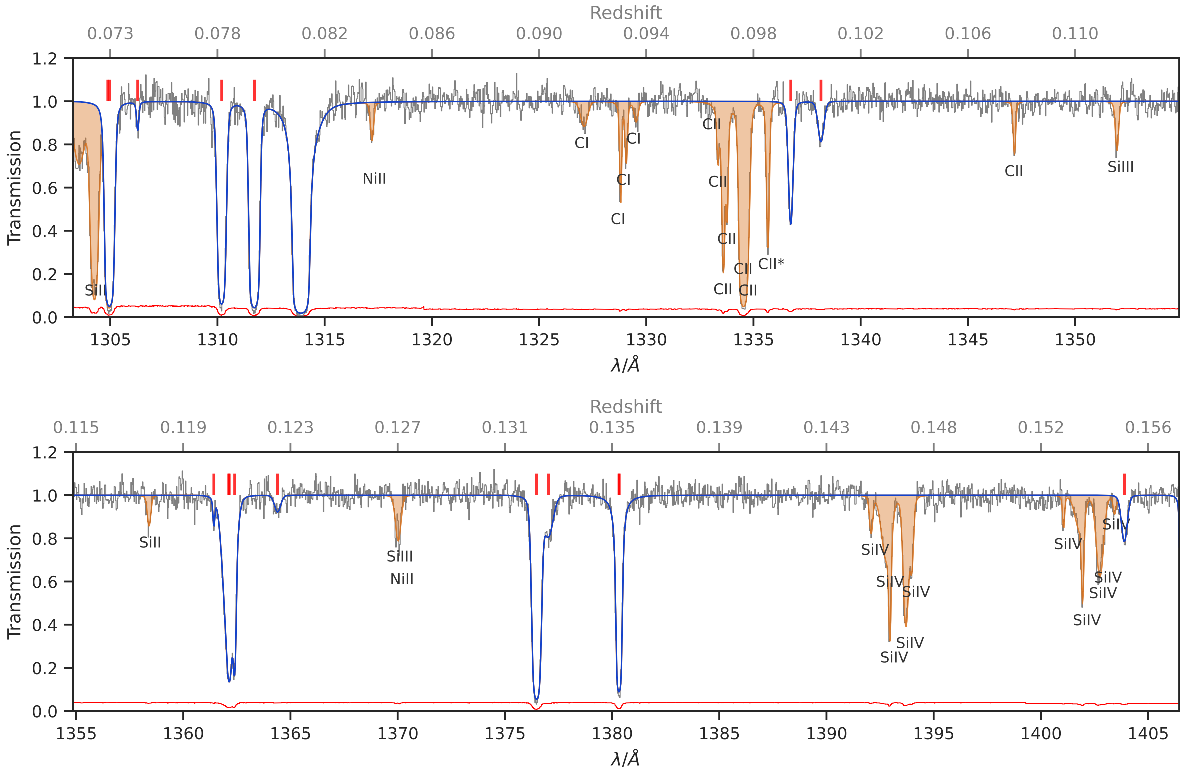

Figure 4.2: A segment of HST COS quasar spectra phl1811, with flux shown in grey and noise plotted in red. Metal line identification and masks are indicated by orange shaded regions, and the fit models of Lyα forest are shown in blue, while locations of the corresponding Lyα lines are indicated by red vertical lines. The fitting procedure is done by the automated program `VPFIT`, with the corresponding COS LSF taken into account.

4.1.2 Metal Identification

As previously mentioned, our VP-fitting procedure fits all absorption lines, including Lyα and metal lines. For our analysis, which focuses on the $\{b, N_{\rm HI}\}$ of the Lyα forest, it is crucial to exclude these metal lines. We utilize archival metal identification data presented in D16. For each spectrum, we create a mask around each identified metal line, which initially centers on the reported wavelengths of the metal lines and has a default width of Δv = 50 km/s, which is chosen based on the resolution of COS, i.e., 15 $\sim$ 25 km/s. We then refine these masks according to our VP-fit results, which is needed since our VP-fitting results do not all match the D16 metal IDs precisely, due to the different VP-fitting procedures used in this work compared to D16. We start by locating absorption regions where the fitted normalized flux, $F_{\rm line,fit}$, is less than or equal to 0.99 (as depicted by the blue line in Fig. 4.2). If an absorption line region overlaps with a mask, we extend the mask's width to fully cover the line, using the full width at half maximum (FWHM) of the detected line, calculated as $2\sqrt{\ln 2}b$, where b is determined by `VPFIT`. We then manually adjust the masks to bridge small gaps ($\lesssim$ 50 km/s) between masked regions to ensure full coverage of all lines close to the identified metal lines to ensure that all potential contamination is eliminated. Additionally, we manually mask regions affected by Damped Lyα absorption systems (DLAs), which disrupt `VPFIT` results. These post-processing masks are created after initially applying `VPFIT` under the assumption that all absorption lines are Lyα, and later, any absorption lines within these regions are excluded from our $\{b, N_{\rm HI}\}$ dataset. The metal masks are subtracted from our total pathlength, resulting in a net pathlength of Δz =4.42 for all four redshift bins. The complete set of quasar segments and their corresponding masks are detailed in Appendix B.1.

With our imposed limit on the $\{b, N_{\rm HI}\}$, we find that 84 out of 741 lines are masked

Table 4.1: Summary of the of the observational dataset

z bins	Δz	Number	$b_{\rm m}/{\rm km\ s}^{-1}$	$\log(N_{\rm HI,m}/{\rm cm}^{-2})$
$0.06 \le z \le 0.16$	1.79	270	34.27	13.22
$0.16 < z \le 0.26$	1.30	201	36.05	13.14
$0.26 < z \le 0.36$	0.78	102	32.44	13.30
$0.36 < z \le 0.48$	0.56	84	32.29	13.31

The numbers of identified Lyα lines in each redshift, the total pathlength Δz, and the median value $b_{\rm m}$ and $\log N_{\rm HI,m}$.

for our whole sample, and that leaves us with a $\{b, N_{\rm HI}\}$ dataset consisting of 657 Lyα absorption lines. We divide the 657 Lyα absorbers into four redshift bins: $0.06 < z < 0.16$, $0.16 < z < 0.26$, $0.26 < z < 3.6$, and $0.36 < z < 0.48$ respectively, according to their central wavelength as given by `VPFIT`. Such ranges for these bins are selected to ensure their pathlength-weighted redshift centers at $z = 0.1, 0.2, 0.3$, and 0.4, matching the central redshift of our simulation snapshot. This provides us with the number of Lyα lines to be 270, 201,102 and 84 and redshift pathlength $\Delta z = 1.79$, 1.30, 0.77 and 0.56 in the bins centred at $z = 0.1$, 0.2 0.3 and 0.4, respectively. In Table 4.1 we summarize our $\{b, N_{\rm HI}\}$ dataset for each redshift bin, with redshift pathlength, number of final Lyα lines as well as median values for the b and $N_{\rm HI}$ in each bin.

4.2 Simulations

We utilize a set of Nyx cosmological hydrodynamic simulations from THERMAL suite, both described in§2.1. To model the Lyα forest at $z < 0.5$. We rescale the temperature for a set of models to model the Lyα forest in extremely hot IGM, which is favoured by our inference result, which will be discussed in below. For other relevant post-processing procedures, such as generating mock spectra, forward-modeling, and UVB variation, we follow the prescription discussed in §2.1.

4.2.1 T_0-rescaling models

As will be discussed later in § 4.3, our data favour models with high T_0 at $z = 0.1, 0.2, 0.3$ and 0.4. These extremely hot models are challenging to generate at low-z solely by varying the H I photoheating rate, since the IGM is dominated by the adiabatic cooling caused by Hubble expansion at this epoch, and the heat injection caused by H I photoionization fades away quickly (McQuinn 2016), As a result, the T_0 of the IGM at low-z is insensitive to the H I photoheating. To address this, we rescale the IGM temperature to model the IGM with high temperature. For each redshift bin, we divide the models into γ bins with $\Delta\gamma = 0.1$ and select the simulation with highest T_0 in each γ bins, and multiply their temperature T (at each simulation cell) by $[\sqrt{2}, 2, 2\sqrt{2}, 4, 4\sqrt{2}, 8]$ respectively to generate 66 new models. The other properties of the simulation remain unchanged, and since we rescaled the temperature of all simulation cells uniformly the whole Δ-T distribution of the simulation model still follows the power law Δ-T relationship Eq. 1.1 with the T_0 rescaled. The rescaling procedure are demonstrated in Fig. 4.3, and the $[T_0, \gamma]$ of original models and models with rescaled T_0 are illustrated in Fig. 4.4, where the original models are shown as blue dots, and the model with rescaled T_0 are shown in orange. Such temperature rescaling procedures are also applied to all four redshift bins.

4.3 Infernce Results

We applied the inference method described in §2.2 to our dataset at four redshift bins to measure the IGM thermal and ionization state at $z = 0.1$, 0.2, 0.3 and 0.4. In this

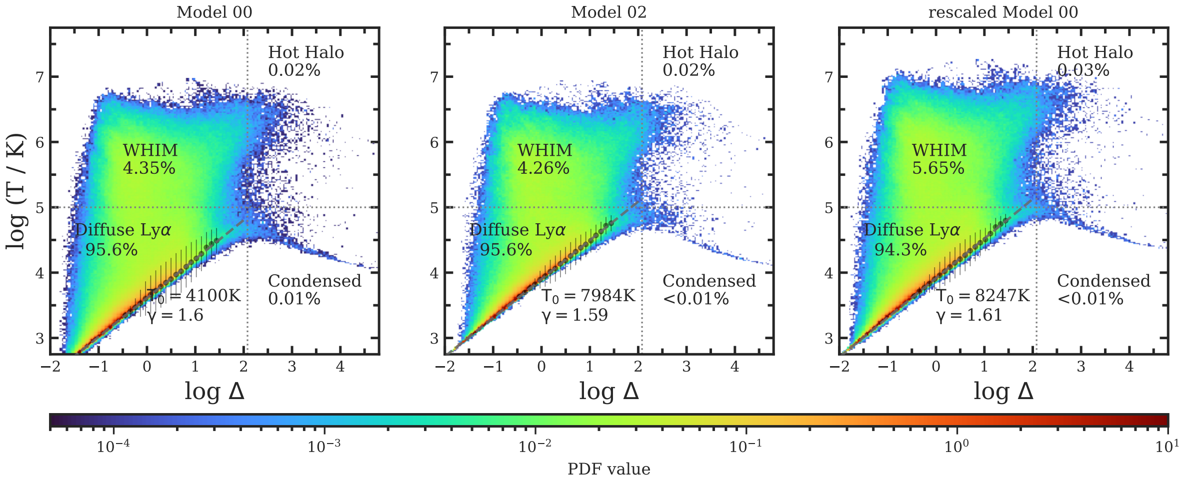

Figure 4.3: Volume-weighted T-Δ distribution for four simulation models at $z = 0.1$. The left panel presents the Nyx model 00 with $T_0 = 4100$ K and $\gamma = 1.60$, while the middle panel features the Nyx model 02 with $T_0 = 7984$ K and $\gamma = 1.59$. The right panel displays a model derived by doubling the temperature in model 00, resulting in $T_0 = 8247$ K and $\gamma = 1.61$, as determined by our $\Delta - T$ fitting procedure. The best-fit power-law relationships are depicted as grey dashed lines. The $\log T$ for each bin is represented by black dots, accompanied by 1-σ_T error bars in black. Volume-weighted gas phase fractions are annotated.

work, the neural networks are trained based on $\{b, N_{\rm HI}\}$ dataset forward-modelled based on the COS dataset described in §3.1. The neural networks are trained on mock datasets at $z = 0.1$, 0.2, 0.3 and 0.4 separately.

The MCMC posteriors are presented in Fig. 4.4, Fig. 4.5, Fig. 4.6 and Fig. 4.7 respectively. In these figures, projections of the Nyx models from the THERMAL suite are plot as blue dots, while models with rescaled temperatures are depicted as orange dots (detailed in § 4.2.1). The inner and outer black contours indicate the 1-sigma and 2-sigma projected 2D intervals, respectively. The dashed black lines mark the 16th, 50th, and 84th percentile values of the marginalized 1D posteriors. From the marginalized 2D posteriors, it is evident that our results across all redshift bins show the expected degeneracies between parameters. Specifically, T_0 exhibits degeneracy with both γ and $\Gamma_{\rm HI}$ (see Hu et al. 2022, for relevant discussion).

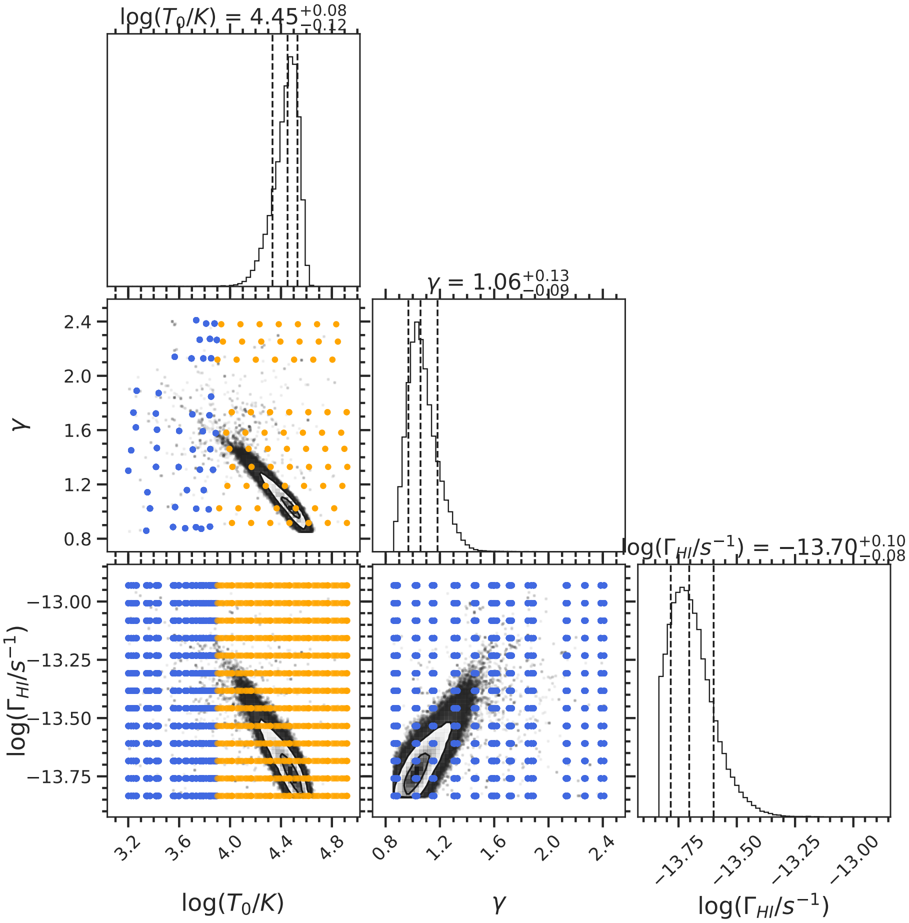

Figure 4.4: The MCMC posterior obtained by our inference method using our $\{b, N_{\rm HI}\}$ dataset at $z = 0.1$. Projections of the thermal grid used for generating models are shown as blue dots. The THERMAL Nyx models are plotted as blue dots, and the models with rescaled temperature are shown as orange dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

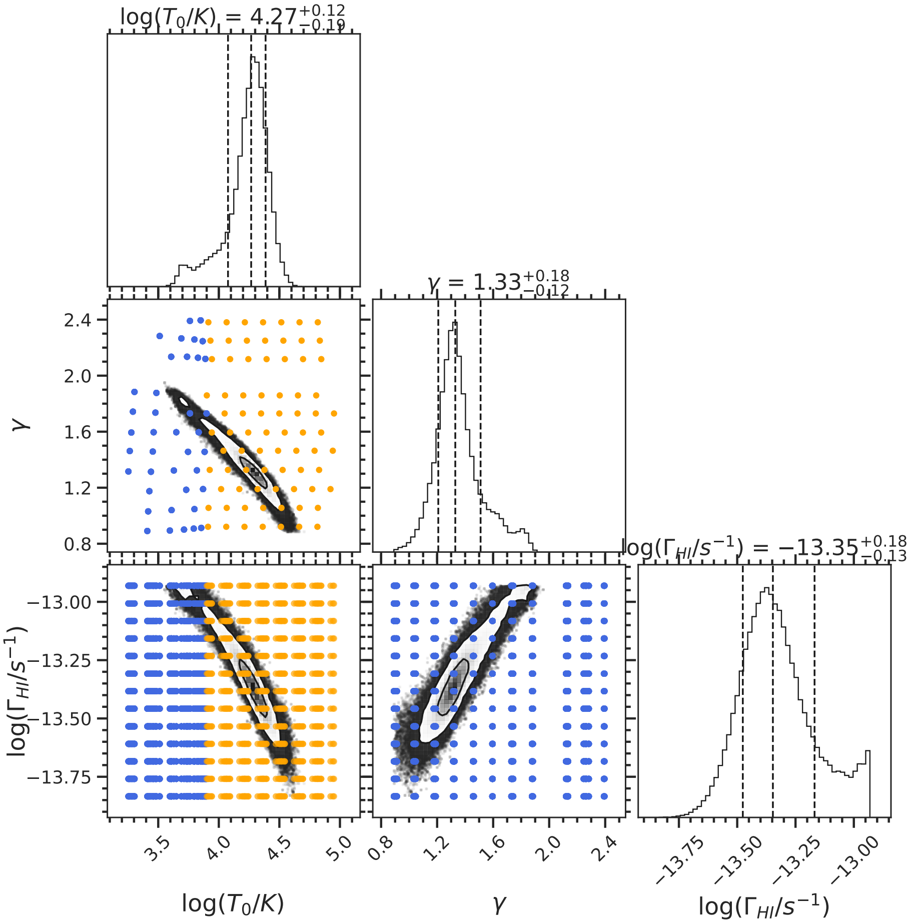

Figure 4.5: The MCMC posterior obtained by our inference method using our $\{b, N_{\rm HI}\}$ dataset at $z = 0.2$. Projections of the thermal grid used for generating models are shown as blue dots. The THERMAL Nyx models are plotted as blue dots, and the models with rescaled temperature are shown as orange dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

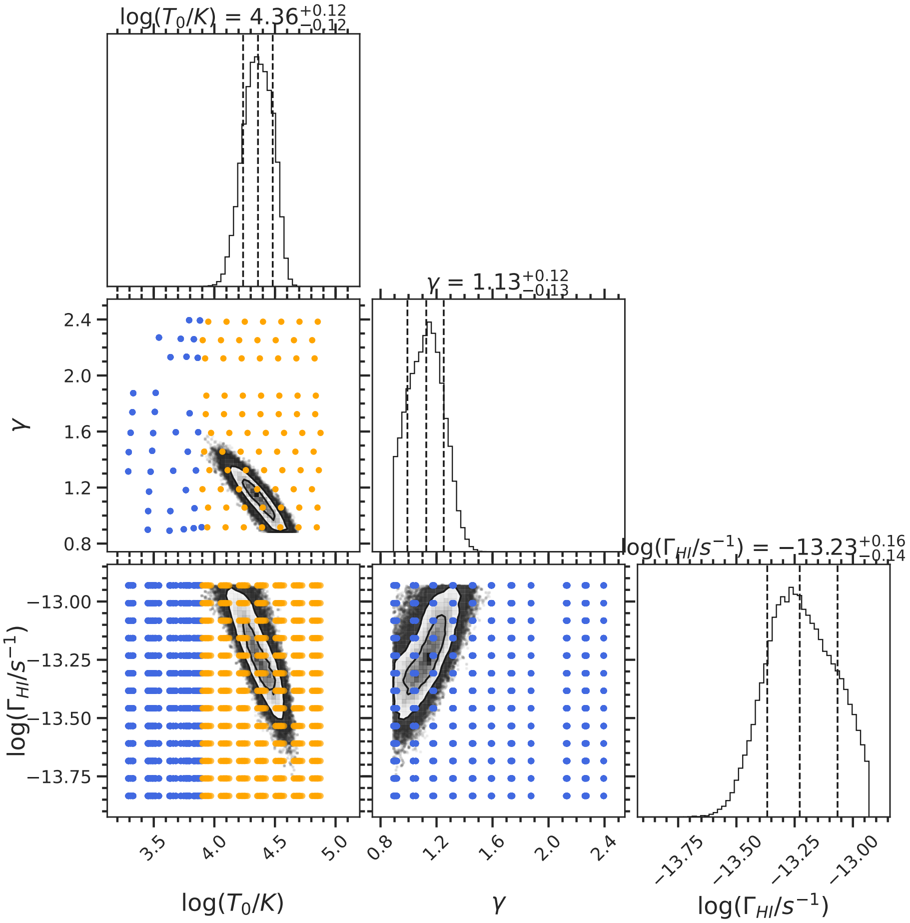

Figure 4.6: The MCMC posterior obtained by our inference method using our $\{b, N_{\rm HI}\}$ dataset at $z = 0.3$. Projections of the thermal grid used for generating models are shown as blue dots. The THERMAL Nyx models are plotted as blue dots, and the models with rescaled temperature are shown as orange dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

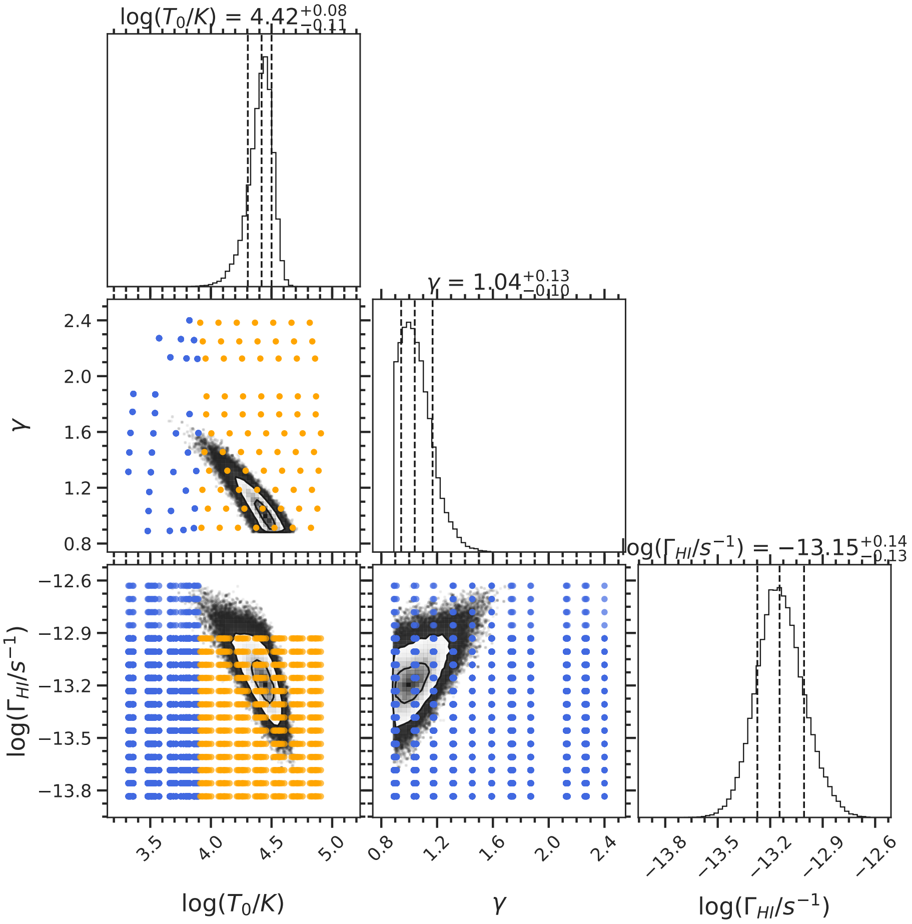

Figure 4.7: The MCMC posterior obtained by our inference method using our $\{b, N_{\rm HI}\}$ dataset at $z = 0.4$. Projections of the thermal grid used for generating models are shown as blue dots. The THERMAL Nyx models are plotted as blue dots, and the models with rescaled temperature are shown as orange dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

Table 4.2: Summary of the inference results

z bins	$\log(T_0/\mathrm{K})$	γ	$\log\,(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$
$0.06 < z \leq 0.16$	$4.45^{+0.08}_{-0.12}$	$1.06^{+0.13}_{-0.09}$	$-13.70^{+0.10}_{-0.08}$
$0.16 < z \leq 0.26$	$4.27^{+0.12}_{-0.19}$	$1.33^{+0.18}_{-0.12}$	$-13.35^{+0.18}_{-0.13}$
$0.26 \leq z \leq 0.36$	$4.36^{+0.12}_{-0.12}$	$1.13^{+0.12}_{-0.13}$	$-13.23^{+0.16}_{-0.14}$
$0.36 \leq z \leq 0.48$	$4.42^{+0.08}_{-0.11}$	$1.04^{+0.13}_{-0.10}$	$-13.15^{+0.14}_{-0.13}$

The inference results i.e., median values of the marginalized 1D posteriors for each parameter, for all four redshift bins. The errors are given by the 1-σ error (16-84%) of the marginalized 1D posteriors.

We tabulate our inference results (median values of the marginalized 1D posteriors for each parameter) in Table. 4.2. Our inference results show that the temperature of the IGM is much higher than expected and is nearly isothermal, with T_0 approaching $30,000$K, and γ approaching 1.0 at $z = 0.1$. In addition, the Γ_{HI} values we inferred are lower than the theoretical model (Khaire & Srianand 2019b). Such a result is a manifestation of the aforementioned b-parameter distribution discrepancy, but now expressed as a quantitative measurement, which fully accounts for the parameter degeneracies with γ and Γ_{HI}. The results will be further discussed in the following section.

The $\{b, N_{\mathrm{HI}}\}$ data obtained from the COS spectra and the corresponding b-N_{HI} distributions emulated by our DELFI emulator are shown in Fig. 4.8, and the likelihood contours corresponding to 80, 60, 40, and 20 cumulative percentiles are plotted as grey dashed lines. These plots show good agreement between the observational data and the b-N_{HI} distributions emulated by DELFI, while the precision is satisfactory, given the size of our $\{b, N_{\mathrm{HI}}\}$ sample is fairly limited especially for the bin at $z = 0.4$.

To further evaluate the reliability of our inference results, we plot the marginalized 1D distributions of b and N_{HI} for our sample at the $z = 0.1$ bin in Fig. 4.9. We compare

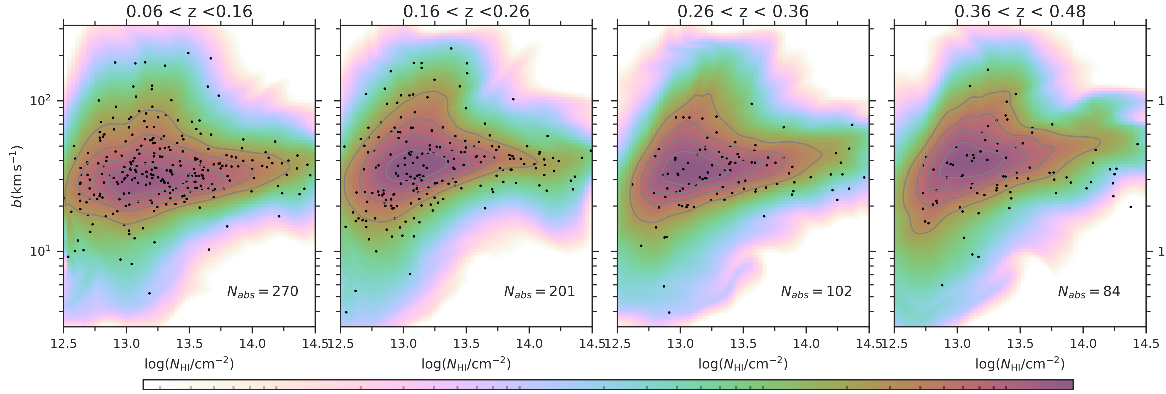

Figure 4.8: Joint b-$N_{\rm HI}$ distributions emulated by our DELFI emulator based on the median values of the marginalized MCMC posterior at $z = 0.1$, 0.2, 0.3 and 0.4. Black dots are the $\{b, N_{\rm HI}\}$ data. The likelihood contours corresponding to 80,60,40, and 20 cumulative percentiles CDF are plotted as gray solid lines. For illustration purposes, the values of the pdf are multiplied by 100 in the colour bar.

the 1D marginalized b distributions to these from 5000 mock datasets of the same size, each sampled from the b-$N_{\rm HI}$ distributions emulated at the median values of the MCMC posteriors. The blue bars represent the average number of lines per bin across the 5000 datasets, while the blue shaded areas denote the 1-σ uncertainty derived from these datasets. The results clearly demonstrate that our inference method effectively recaptures the marginalized 1D distributions of $\{b, N_{\rm HI}\}$. Due to the limited data size, there are fluctuations in the results, which are reflected by the 1-σ error bars in the marginalized 1D distributions for both b and $N_{\rm HI}$.

4.3.1 Evolution of the thermal state of the IGM

We summarize the evolution of T_0, γ and $\Gamma_{\rm HI}$ across four redshift bins in Fig. 4.10, in which we also present results from previous studies at higher redshifts (Hu et al. 2023b, Hiss et al. 2018, Walther et al. 2019b, Gaikwad et al. 2021). Our measurements, along with their 1-σ uncertainties, are shown as filled red data points with error bars. As a reference for current theoretical cosmology models, we plot the IGM thermal histories

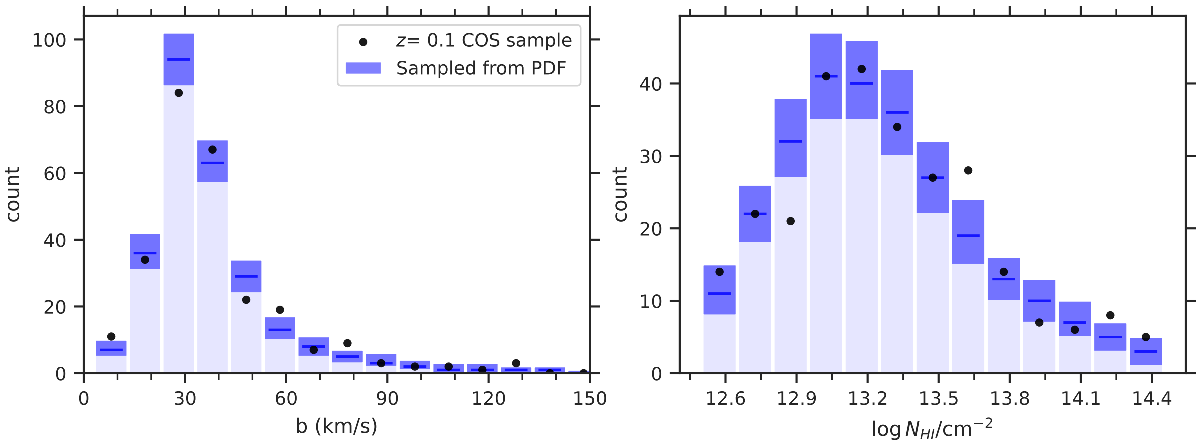

Figure 4.9: The marginalized 1D b and N_{HI} distributions of the data sample at $z = 0.1$ compared with 5000 mock datasets with the same size, sampled from the b-N_{HI} distributions emulated based on the median values of the MCMC posteriors. The black dots represent our $\{b, N_{\mathrm{HI}}\}$ data, and the blue bars indicate the mean value of the number of lines that fall in each bin for the 5000 datasets, whereas the blue shaded regions represent the 1-σ uncertainty calculated from the 5000 datasets.

permitted by different Helium reionization models (Oñorbe et al. 2017b;a) as cyan-shaded region.

Our measurements indicate a significant discrepancy in T_0 at $z < 0.5$, where our observed $T_0 \sim 30,000$ K at $z = 0.1$ is substantially higher than the values predicted by cosmological simulations. We also notice that these higher-than-expected IGM temperatures, which exhibit an increasing trend towards lower redshifts, align with Hu et al. (2023b, blue data points in Fig. 4.10), which suggests an IGM $T_0 \sim$ 13,500 K at $z = 1.0$. If the IGM is indeed much hotter than expected, such high IGM temperature requires the existence of additional heating sources not accounted for in current IGM models, particularly relevant around $z \sim 1.0$. Further discussion on this unexpected high temperature can be found in § 3.4.1.

In addition, our measurements indicate that γ is significantly lower than expected at low-z, with $\gamma \sim 1.0$ at $z = 0.1$, suggesting that the low-z IGM may be nearly isothermal, which might put important constraints on the aforementioned heating mechanism that

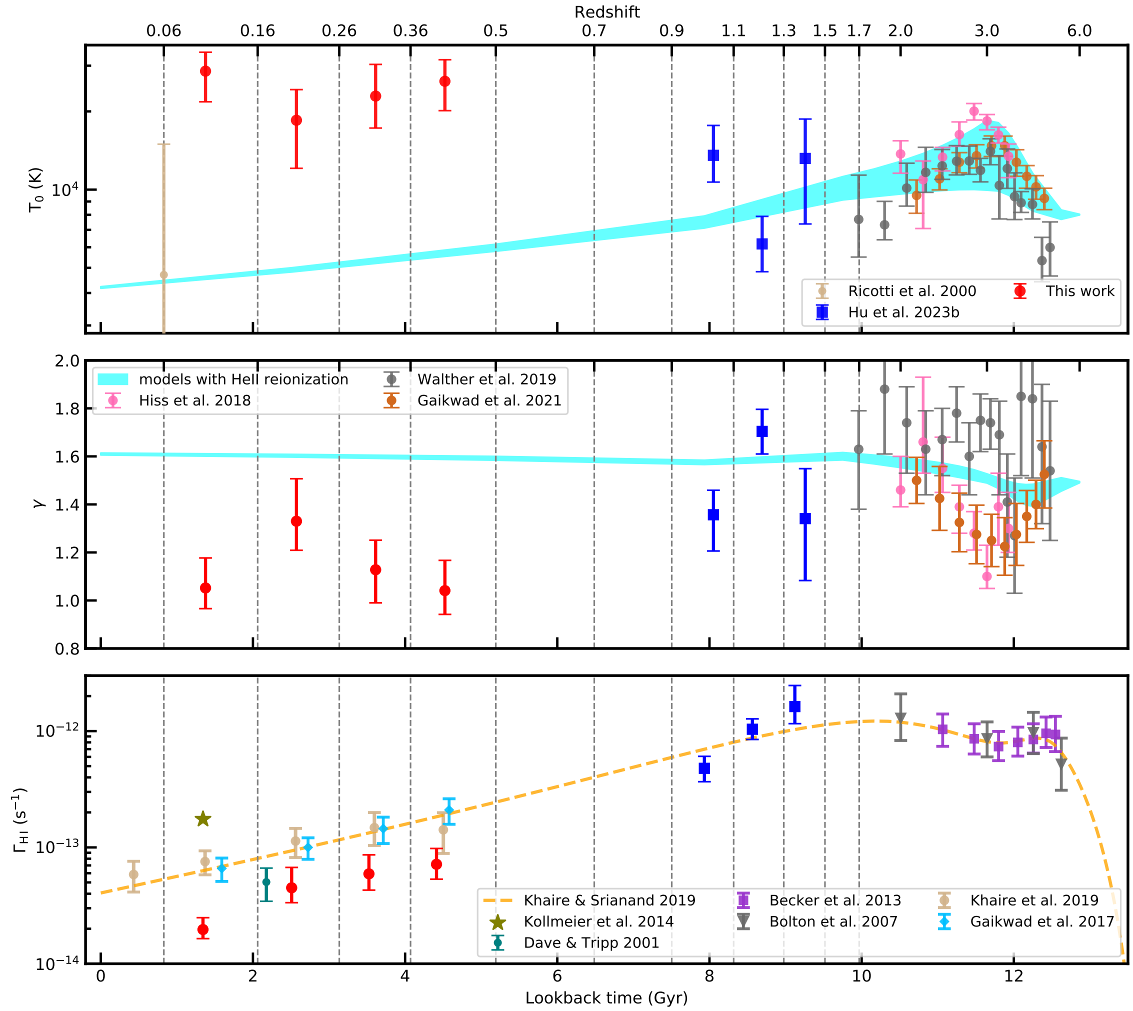

Figure 4.10: Evolution history of T_0 (top), γ(middle) and $\log \Gamma_{\mathrm{HI}}$ (bottom) based on our inference results obtained from the COS data. Our results are shown as red dots, while measurements from other studies are displayed in different symbols and colors. The error bars stand for the 1-σ error. The blue-shaded region in the top panel represents the range spanned by T_0 from hydrodynamical simulations of a large family of different HeII reionization models. The mock measurements based on Nyx simulation are shown in blue.

caused the observed extremely high IGM temperature. However, due to the uncertainties in our measurements and the known degeneracy between T_0 and γ (Hu et al. 2022), it remains uncertain whether the IGM is truly isothermal at low redshifts.

4.3.2 Evolution of the H I photoionization rate and UVB

Our measurements also provide insights into the $\Gamma_{\rm HI}$ evolution at $z < 0.5$. In the bottom panel of Fig. 4.10, we display our $\Gamma_{\rm HI}$ measurements across four redshift bins, compared with measurements from previous studies (Davé & Tripp 2001, Bolton 2007, Becker & Bolton 2013b, Kollmeier et al. 2014, Gaikwad et al. 2017b, Khaire et al. 2019, Hu et al. 2023b). We report $\Gamma_{\rm HI}$=$-13.70^{+0.10}_{-0.08}$, $-13.35^{+0.18}_{-0.13}$, $-13.23^{+0.16}_{-0.14}$, and $-13.15^{+0.14}_{-0.13}$ at $z = 0.1, 0.2$, 0.3 and 0.4 respectively, which are noticeably lower than the predictions of the UVB model presented in Khaire & Srianand (2019b), which align well with other low-z measurements derived using the Lyα power spectrum (Gaikwad et al. 2017a, Khaire et al. 2019) using the D16 low-z Lyα forest spectra. It is important to note, however, that the studies by Gaikwad et al. (2017a) and Khaire et al. (2019) do not fully take into account the degeneracy between the ionization and thermal state of the IGM. In their analyses, $\Gamma_{\rm HI}$ is measured using cosmological simulations with a standard thermal history, specifically with a $T_0 \sim 5000$ K and $\gamma \sim 1.6$ at $z = 0.1$. Nevertheless, both higher IGM temperatures and $\Gamma_{\rm HI}$ levels suppress the formation of Lyα absorbers, leading to increased flux and altering the power spectrum. If the IGM is indeed hotter than expected, the corresponding $\Gamma_{\rm HI}$ value required to match the observed power spectrum would be lower. Therefore, it is crucial to break the degeneracy between the ionization and thermal state of the IGM.

In addition, our $\Gamma_{\rm HI}$ measurements agree well with Davé & Tripp (2001) measurement using HST STIS data, which, however, suggests that the b parameter of the low-z

Lyα forest align well with the simulation. It is possible that the different constraints on the b parameter are caused by the resolution effects, which differ for the COS and STIS, where the latter has a significantly higher resolution and a better LSF shape. We will leave the relevant discussion to §4.4.3.

4.4 Discussion

4.4.1 The discrepancy in T_0 and γ

Many previous studies based on the low-z Lyα forest have pointed out that the observed b-parameter significantly surpasses the predicted value based on various simulations (Gaikwad et al. 2017b, Viel et al. 2017, Nasir et al. 2017, Bolton et al. 2022b;a). Quantitatively, Viel et al. (2017) compares the marginalized b distribution with various simulations, showing that the b distribution at $z \sim 0.1$ can be best recovered by the hydrodynamic simulations (`P-GADGET-3`, see Springel et al. 2005) with $T_0 \gtrsim 10000$ K, while the theoretical model dictates that the $T_0 \sim 5000$ at $z = 0.1$.

Our results favour an even hotter and isothermal IGM with $T_0 \sim 30,000$K and $\gamma \sim 1.0$ at $z = 0.1$, and our results are mainly consistent across four different redshift bins. Such results are obtained by taking the full b-$N_{\rm HI}$ distribution into account, rather than only matching the b parameter. To investigate this issue, in Fig. 4.11 we plot the b-$N_{\rm HI}$ distribution recovered from our inference results and compare it with a model with parameters $[\log T_0,\ \gamma,\ \log(\Gamma_{\rm HI}/s^{-1})] = [4.0,1.55,-13.3]$, which represent model favoured by previous study based solely on the b distribution. It can be seen that, the model with higher T_0 and lower γ fit the observational data better at high $N_{\rm HI}$ end.

Furthermore, as mentioned in §4.3, the unexpected thermal state of the IGM at low redshifts may be attributed to an unknown heating mechanism. It is also plausible that

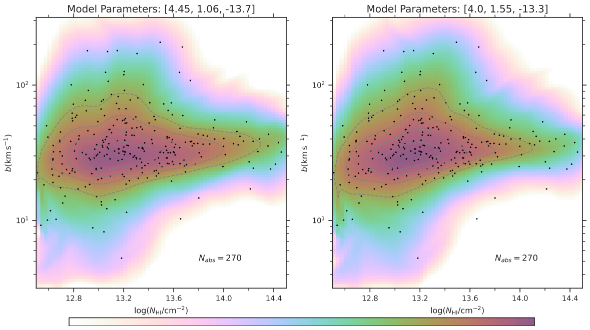

Figure 4.11: Joint b-N_{HI} distributions emulated by our DELFI emulator based on the median values of the marginalized MCMC posterior at $z = 0.1$ vs the one recovered from $[\log T_0,\ \gamma,\ \log(\Gamma_{\mathrm{HI}}/s^{-1})] = [4.0,1.55,-13.3]$, which represent model favoured by previous study based solely on the b distribution. Black dots are the $\{b, N_{\mathrm{HI}}\}$ data. The likelihood contours corresponding to 80,60,40, and 20 cumulative percentiles CDF are plotted as gray solid lines.

the discrepancies observed at $z \sim 0.1$ and $z \sim 1$ originate from the same heating source. Thus, we hypothesize that this heating mechanism becomes significant around $z \sim 1$ and persists down to $z = 0$. If this hypothesis holds true, it would dramatically alter our understanding of IGM physics, underscoring an urgent need to explore potential causes, such as dark matter annihilation (Araya & Padilla 2014, Bolton et al. 2022a), gamma-ray sources (Puchwein et al. 2012), or feedback mechanisms from galaxy formation, which remain poorly understood at low redshifts (see Springel et al. 2005, Croton et al. 2006a, Sijacki et al. 2007a, Hopkins et al. 2008a, Tillman et al. 2023b;a, Hu et al. 2023a).

Another potential heating mechanism is the dust heating Inoue & Kamaya (2010), Bolton et al. (2022c), which also also reduce the γ of the IGM since dust heating scales as $u_{\mathrm{dust}} \propto \Delta^{1/3-(1-\gamma)/6}$. As suggested by Ménard et al. (2010), the dust might be more abundant in the IGM than we previously thought. In addition, Chen & Oh (2024) shows that the dust can survive the galactic wind and thus possibly be transported to the IGM.

4.4.2 The effect of the Turbulence

An alternative explanation of the observed higher-than-expected b-parameter is the existence of small-scale turbulence in the low-z IGM, which increases the width of the observed Lyα lines (Nasir et al. 2017, Viel et al. 2017, Gaikwad et al. 2017b, Bolton et al. 2022a). In this section, we assess such a hypothesis quantitatively by applying our inference method to the COS Lyα forest dataset at $z \leq 0.5$, with fiducial thermal history and flexible small scale IGM turbulence.

In practice, we model the small turbulence in the IGM by adding a Gaussian component $N(0, \sigma_v)$ to the peculiar velocity along line-of-sight, where σ_v is the standard deviation in km/s. Such a velocity component is added to each simulation grid with ΔL= 0.024 Mpc/h. To quantitatively constrain the turbulence, we post-process the simulation and generate skewers with σ_v = 3,6,9,12,15,18,21,24,27 km/s. Since we consider the turbulence as an alternative explanation to the high-than-expected IGM temperature, here, we use the standard Nyx model with $[T_0,\gamma] \sim$ [4000K, 1.6] at $z = 0$. For each σ_v value, we generate forward-modelled mock spectra with 11 different UVB photoionization rate, $\Gamma_{\rm HI}$, following the prescription given in §2.1. We then apply our inference framework on the σ_v-$\Gamma_{\rm HI}$ parameter gird following the procedure discussed in §5.2. To monitor the evolution of such turbulence, we conduct the inference at all four redshift bins individually, and obtain that σ_v = 14,17,12,11 km/s for $z = 0.1, 0.2, 0.3$ and 0.4; while the corresponding $\log(\Gamma_{\rm HI}/s^{-1}) = -13.1, -12.9, -12.8$, and -12.7. Such a σ_v is coherent with the one derived in Bolton et al. (2022c), which yield $\sigma_v \sim$ 15 km/s at $z = 0.1$. The inference results are shown in Fig. 4.12, and the evolution history of $v_{\rm tur}$ and $\log \Gamma_{\rm HI}$ are shown in Fig. 4.13. Interestingly, it can be seen that the $v_{\rm tur}$ required to match the observation increases toward low-z, suggesting that the discrepancy between the observation and simulation in b-parameter must be caused by continuous sources that increase toward low-z.

In addition, we notice that with the standard thermal model and altered small-scale velocity, our inference method suggests higher $\Gamma_{\rm HI}$ values. This is mainly because, while the σ_v has no noticeable impact on the $\mathrm{d}N/\mathrm{d}z$, both T_0 and $\Gamma_{\rm HI}$ have similar correlation on the $\mathrm{d}N/\mathrm{d}z$, i.e., both higher T_0 and $\Gamma_{\rm HI}$ suppress the formation of the HI absorbers in the IGM, causing degeneracy in the inference results.

4.4.3 The effect of the Resolution

Furthermore, we assess whether the observed discrepancy in the b-parameter could be caused by the overestimation or underestimation of the HST COS resolution. Firstly, we aim to determine the required resolution of HST COS to observe the aforementioned discrepancy if the incorrect COS LSF is the sole factor, i.e., if the IGM is neither hotter than expected nor affected by additional turbulence.

Similar to the analysis presented in §4.4.2, we perform our inference method on a 2D parameter grid consisting of resolution and UV background $\Gamma_{\rm HI}$. To do so, we assume that the 'True' COS LSF is unknown and forward-model our mock spectra using 10 different Gaussian LSF with resolution (FWHM) $\sim$ 10-100 km/s. Such procedure is applied all standard Nyx models (model T000, with $T_0 \sim 4000$ and $\gamma \sim 1.6$ at z=0.1) with 11 different $\Gamma_{\rm HI}$ discussed in §2.1. To make these forward-modelled mocks comparable to our observation, our VP-fitting program still uses the tabulated COS LSF for both data and mocks, as presented in §2.1.4. In Fig. 4.14, we plot the distributions of the b parameters for both observed and mock data, which are forward-modeled based on various resolutions across the redshift range $0.06 < z < 0.16$. Both the observational and mock datasets are fitted using our VP-fitting program, based on the reported resolution of HST COS. We observe that the peaks of the b distributions shift to larger values as

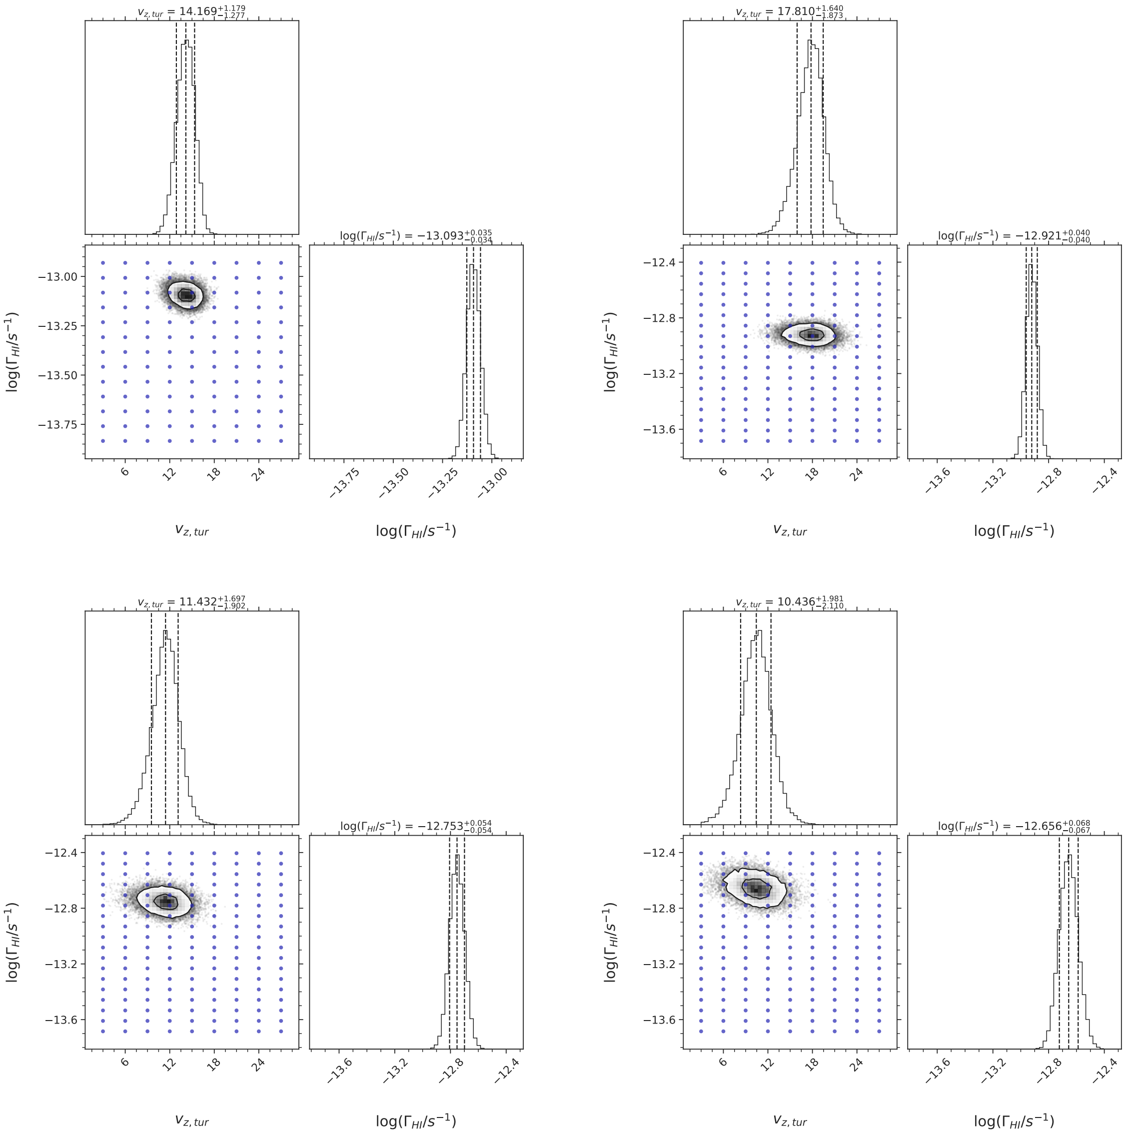

Figure 4.12: Posteriors obtained by applying our inference method on the $v_{\rm tur}$-$\Gamma_{\rm HI}$ grid at $z = 0.1, 0.2, 0.3$ and 0.4. Projections of the parameter grid used for generating models are shown as blue dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

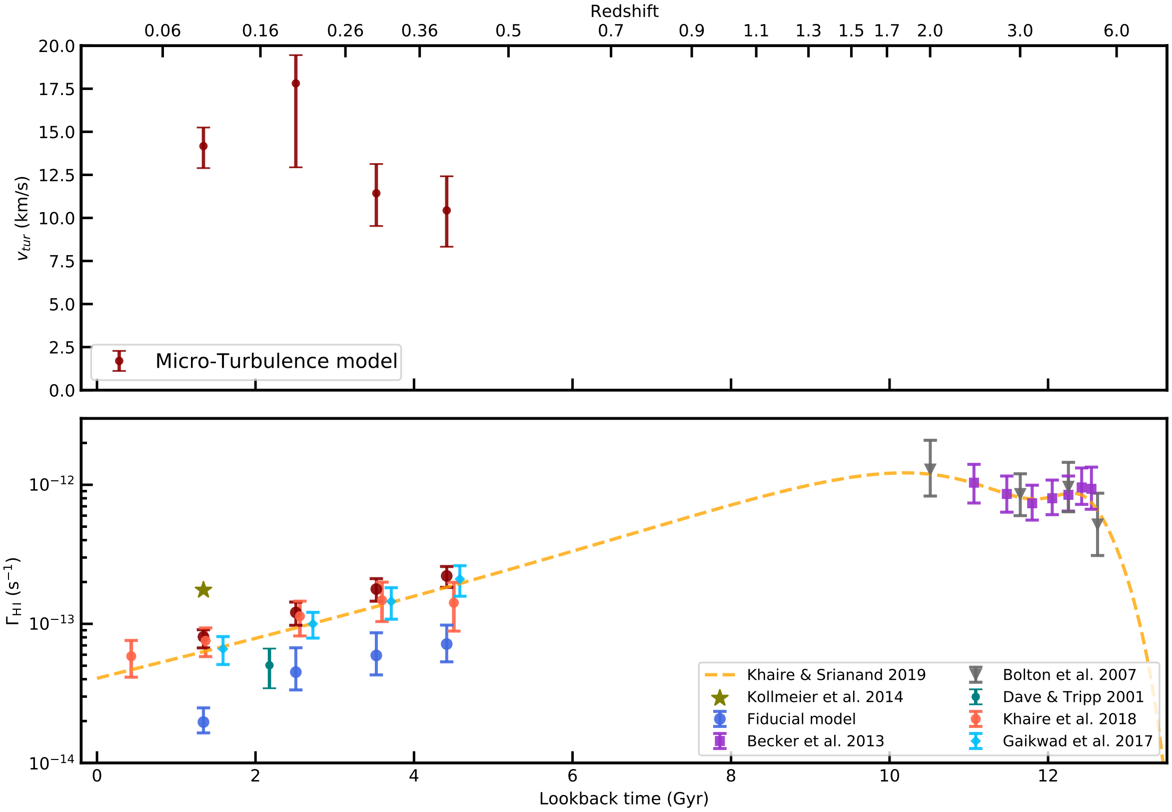

Figure 4.13: Evolution history of $v_{\rm tur}$ (top) and $\log \Gamma_{\rm HI}$ (bottom) based on standard thermal model and altered small-scale velocity. The results are shown as dark red dots, while measurements from other studies are displayed in different colours. The error bars stand for the 1-σ error.

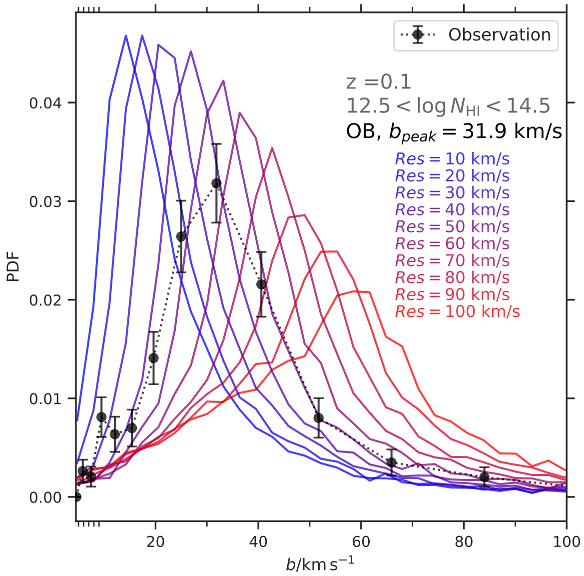

Figure 4.14: Comparisons of distributions of b parameters of the observation and mock data forwarded-modelled based on various resolutions over a redshift range $0.06 < z < 0.16$. Both the observation and the mock data are fitted using our VP-fitting program based on the reported HST COS resolution.

the resolution increases. Notably, the peak of the observed b distribution aligns with the mock data forward-modeled with a resolution of approximately $40 \sim 50$ km/s.

We then run our inference method to find the resolution required to obtain the observed b-$N_{\rm HI}$ distribution, and the inference result is shown in Fig. 4.15. Given the observed b distribution peaking at ~ 30 km/s, our inference method suggests that the required resolution (FWHM) is about 40km/s, which corresponds to a resolution $R \sim$ 7,000, while the reported R for HST COS is 12,000 to 16,000.

In Ghavamian et al. (2009), the reported COS LSF is carefully examined using the COS spectra of the O9 Ib supergiant star Sk 155 in the Small Magellanic Cloud (V=12.4), where the COS spectra are compared with those observed with STIS E140H spec (R∼114,000), which are then convolved with the modeled COS LSF. The close alignment between these spectra confirms the high accuracy of the reported COS LSF.

Here, we follow the aforementioned method and compare the COS spectra with the

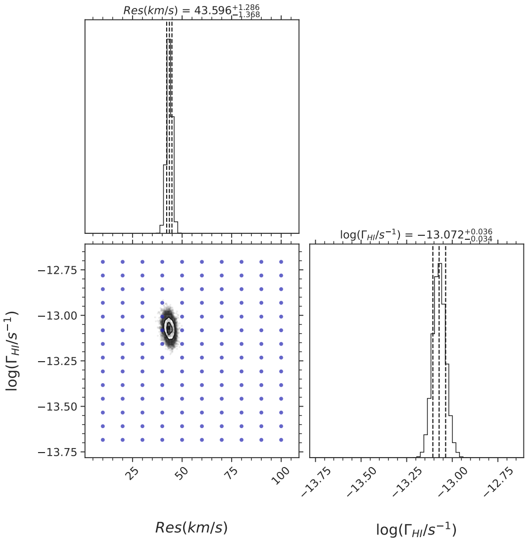

Figure 4.15: Posteriors obtained by applying our inference method on the Resolution-$\Gamma_{\rm HI}$ grid at $z = 0.1$. Projections of the parameter grid used for generating models are shown as blue dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

higher-resolution STIS E140M spectra, which covers 1144∼1729 Å and has a reported resolution of approximately 45,000—roughly three times that of the COS. We use the STIS E140M as intrinsic spectra and convolved them with the HST COS LSF as tabulated in `linetools`. This allows us to compare the reported COS LSF with the actual COS spectra to determine if the reported HST COS resolution is accurate.

To this end, we utilize STIS E140M spectra of the object PHL1811, PG1216, 3c273, and H1821 (Lacki & Charlton 2010, Tripp et al. 2008), whose COS spectra are also examined in this work, and perform a detailed comparison. Visual inspection suggests that convolving these very high-resolution STIS E140H spectra with the COS LSF models leads to an excellent match to the observed COS FUV spectra. Fig. 4.16 shows a segment of the COS G130M spectrum of PHL1811 (red) compared with the corresponding STIS E140M spectrum convolved COS G130M LSF (blue). The green shaded region indicates

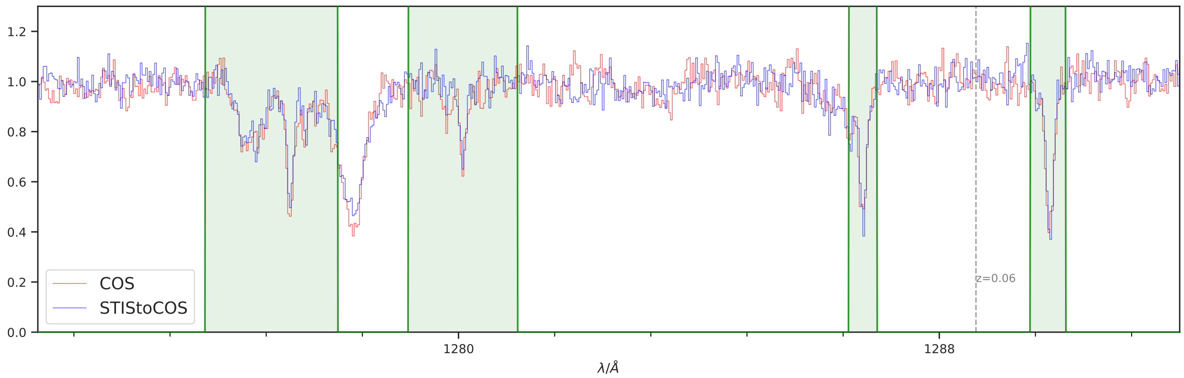

Figure 4.16: A segment of the COS G130M spectrum of PHL1811 (red) compared with the corresponding STIS E140M spectrum convolved COS G130M LSF (blue). The green shaded region indicates the masked region due to the existence of metal absorption lines.

the masked region due to the existence of metal lines, and it can be seen that for these narrow metal lines, the high-resolution STIS spectrum convolved with COS LSF matches well with the observed COS spectrum, suggesting that the reported COS LSF model is accurate.

In addition, we also compare the b-parameters fitted from both COS and STIS spectra and check if there are systematic bias, which indicates the over- or underestimation of the COS LSF. In Fig. 4.17, we plot the b parameters of absorption lines observed in both the HST COS spectrum and the HST STIS spectrum of the object PHL1811, while the STIS spectrum is convolved with the COS LSF. The red dots represent the b parameters fitted from the COS spectrum observed with the G130M grating, while the blue dots come from the STIS E140M spectrum after convolution with the COS G130M LSF. Both spectra are fitted based on the reported HST COS G130M LSF. The comparison reveals that the b-parameters from the STIS spectrum, once convolved with the COS LSF, are not systematically lower than those derived from the direct COS spectrum. This suggests that the reported COS LSF is reliable, confirming its effectiveness in replicating higher-resolution spectral features. We therefore conclude that the observed discrepancy in the

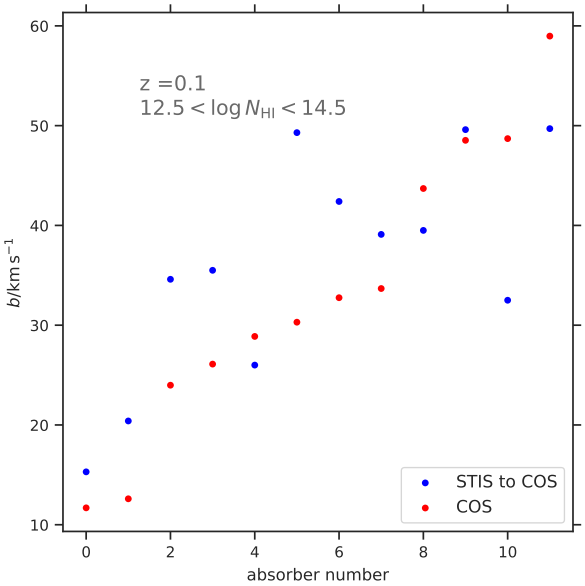

Figure 4.17: The b parameters of the lines observed in the spectrum of PHL1811, where the red dots are fitted from the HST COS spectrum observed with G130M grating, and the blue dots are fitted from the HST STIS E140M spectrum convolved with COS G130M LSF. Both spectra are fitted using our VP-fitting program based on the reported HST COS G130M LSF.

b parameter can not be caused by inaccurate resolution of the HST COS.

4.5 Summary

In this paper, we make use of 82 archival HST G130/G160 quasar spectra, from which we obtain the b-$N_{\rm HI}$ distribution distribution and line density $\mathrm{d}N/\mathrm{d}z$ over the redshift range $0.06 < z < 0.48$ in four redshift bins. We then measure the thermal and ionization state of the IGM following a machine-learning-based inference method presented in Hu et al. (2022) for this redshift range for the first time.

We summarize our results in below:

- We Voigt-profile fit the Lyα in all 82 quasar spectra using a fully automated `VPFIT` wrapper and obtain $\{b, N_{\rm HI}\}$ for 657 lines. We use the metal identifications from the D16 to generate our metal masks, filtering out 84 contaminants besides Lyα absorption lines, and obtain a final sample of 657 Lyα lines across a total pathlength of Δz =4.43.

- We employ the Hu22 inference method, which simultaneously measures $[T_0, \gamma, \Gamma_{\rm HI}]$ from the b-$N_{\rm HI}$ distribution and $\mathrm{d}N/\mathrm{d}z$, with the help of neural density estimators and Gaussian process emulators trained on a suite of 51 Nyx simulations each having a different IGM thermal history. It enables us to measure the IGM thermal and ionization state with high precision even with limited data.
- We obtain $[\log(T_0/K), \gamma]$ = [$4.45^{+0.08}_{-0.12}$,$1.06^{+0.13}_{-0.09}$], [$4.27^{+0.12}_{-0.19}$,$1.33^{+0.18}_{-0.12}$], [$4.36^{+0.12}_{-0.12}$,$1.13^{+0.12}_{-0.13}$] and [$4.42^{+0.08}_{-0.11}$,$1.04^{+0.13}_{-0.10}$] at $z = 0.1, 0.2, 0.3$ and 0.4 respectively. These measurements suggest that the IGM might be much hotter than expected and close to isothermal.
- We compare our findings with previous work, which reports unanticipated high b-parameters compared with various simulations based on observational data at $z \sim 0.1$, and these high b values, if caused by thermal broadening, correspond to an IGM temperature with T_0 ∼10,000 K. However, these research analysed solely the b distribution and ignored the degeneracy between γ and T_0. By comparing the corresponding b-$N_{\rm HI}$ distribution emulated by DELFI, we conclude that our models with $T_0 \sim$ 30,000 K and $\gamma \sim 1$ at $z = 0.1$ fit the observation better for absorber with larger $N_{\rm HI}$.
- Our results might be attributed to an unknown heating mechanism. It is also plausible that the discrepancies observed at $z \sim 0.1$ and $z \sim 1$ originate from the same heating source. Thus, we hypothesize that this heating mechanism becomes significant around $z \sim 1$ and persists down to $z = 0$. If this hypothesis holds true,

it would dramatically alter our understanding of IGM physics, underscoring an urgent need to explore potential causes, such as dark matter annihilation (Araya & Padilla 2014, Bolton et al. 2022a), gamma-ray sources (Puchwein et al. 2012), or feedback mechanisms from galaxy formation, which remain poorly understood at low redshifts (Springel et al. 2005, Croton et al. 2006a, Sijacki et al. 2007a, Hopkins et al. 2008a, Tillman et al. 2023b;a, Hu et al. 2023a).

- We successfully measure the $\Gamma_{\rm HI}$ at four redshif bins, reporting $\log(\Gamma_{\rm HI}/s^{-1}) = -13.70^{+0.10}_{-0.08}$, $-13.35^{+0.18}_{-0.13}$, $-13.23^{+0.16}_{-0.14}$, and $-13.15^{+0.14}_{-0.13}$ at $z = 0.1, 0.2$, 0.3 and 0.4 respectively. These measurements are noticeably lower than the predictions of the UVB model presented in Khaire & Srianand (2019b), and the measurements of Gaikwad et al. (2017b), Khaire et al. (2019) using the Lyα power spectrum based on D16, but agree with the measurements made by Davé & Tripp (2001) based on the STIS data. However, it is worth mentioning that all previous measurements do not take the potential degeneracy between the IGM thermal and ionization state into account.

- An alternative explanation of the observed higher-than-expected b parameter is the existence of small-scale turbulence in the low-z IGM, which increases the width of the observed Lyα lines. To this end, we perform our inference method on a $v_{\rm tur}$-$\Gamma_{\rm HI}$ grid to conclude that if the observed discrepancy is indeed caused by turbulence in small-scale, it would need velocity dispersion with $\sigma_v \sim 14$, 18, 11, and 10 km/s at $z = 0.1$, 0.2, 0.3 and 0.4 respectively. Furthermore, the increase in σ_v towards low-z implies that the discrepancy between observed and simulations in b-parameters, whether caused by turbulence, must be driven by continuous sources that intensify towards low-z.

- In addition, we evaluate whether the observed effect can be caused by overestimation of the COS resolution. We find that it requires an effective resolution of $\sim$ 45 km/s (R $\sim$7,000) to cause the observed b-N_{HI} distribution. Furthermore, we perform a detailed comparison between the HST COS spectrum and the spectrum observed with HST STIS E140M, which has a much higher resolution, for four objects. The comparison suggests that the reported COS resolution (LSF) is reliable, and the observed discrepancy in the b-parameter could not be caused by the resolution effects solely.

Chapter 5

The Impact of the WHIM on the IGM Thermal State Determined from the Low-z Lyman-α Forest

In this paper, we adopt the Hu22 inference method to investigate the impact of the WHIM on measurements of the IGM thermal state, $[T_0, \gamma]$, based on the b-$N_{\rm HI}$ distribution of the Lyα forest. Firstly, we assess the effectiveness of $[T_0, \gamma]$ as IGM parameters at low-z by comparing its performance as neural network training labels against the photoheating labels $[A$,$B]$ (see § C.1). These latter labels are photoheating rate rescaling factors used to generate the Nyx simulation suite with various thermal histories (see e.g Becker et al. 2011). Since our emulators are trained on these Nyx simulations generated by varying $[A, B]$, the inference method is naturally inclined to retrieve these photoheating labels. On the other hand, if shock heating at low-z causes the T-Δ distribution of the Lyα absorbers to deviate from the power-law relationship, the effectiveness of $[T_0, \gamma]$ as labels could be compromised. Thus, our comparison between these two sets of labels provides insight into the robustness of $[T_0, \gamma]$ as IGM parameters at low-z, in the presence of substantial shock heating.

Afterwards, we explore the potential effects of different feedback mechanisms, which are associated with WHIM, on measurements of the IGM thermal state, $[T_0, \gamma]$. In terms of our inference methodology, the question becomes: what would happen if we used a simulation grid without feedback to interpret a Universe that includes feedback? Would this lead to unbiased $[T_0, \gamma]$? To answer these questions, we apply the Hu22 inference methodology to mock data drawn from the Illustris (Genel et al. 2014) and IllustrisTNG (Weinberger et al. 2017) simulations at $z = 0.1$. These two simulations incorporate galaxy formation models and feedback mechanisms that are not included in the Nyx simulation, which heat up the IGM substantially at low-z, and transform the cool diffuse Lyα gas into WHIM more effectively compared with Nyx simulation (see Fig. 5.2). We examine the inference results based on these two simulations and explore whether feedback biases the measurement of the thermal state $[T_0, \gamma]$.

To further investigate this problem, we explore the specific impacts of shock heating and other astrophysical processes, such as AGN feedback and UV background photoionization, on the physical properties of the Lyα forest at $z = 0.1$. Within the three aforementioned simulations, we identify simulated Lyα absorbers in the simulations and establish a direct correlation between the physical properties of these absorbers (including temperature T, overdensity Δ, and line-of-sight velocity $v_{\rm los}$) and the observed Lyα line parameters (b, $N_{\rm HI}$) derived from the absorption lines detected in corresponding mock spectra. We then examine the distributions of Δ and T of these simulated Lyα absorbers across the three aforementioned simulations to study the detailed effects of the feedback and UV background photoionization rate, $\Gamma_{\rm HI}$, on the Lyα forest.

This chapter is organized as follows: In section § 5.1, we outline the simulations and associated processes applied to generate synthetic Lyα forest. It includes post-processing, forward-modeling, and Voigt profile fitting. The inference framework and

Table 5.1: Parameters of cosmology and $T-\Delta$ relation (at $z=0.1$)

Parameters	Nyx	IllustrisTNG	Illustris
Ω_m	0.3192	0.3089	0.2726
Ω_Λ	0.6808	0.6911	0.7274
Ω_b	0.0496	0.0486	0.0456
h	0.670	0.677	0.704
σ_8	0.8288	0.8159	0.809
n_s	0.96	0.97	0.963
T_0	4093 K	4241 K	4292 K
γ	1.588	1.593	1.577

results for all three simulations are then presented in in § 5.2. Section § 5.3 is dedicated to the investigation of the physical characteristics of low-redshift Lyα forest absorbers in all three simulations. Finally, in § 5.4, we present a summary and discussion of our findings. For the sake of brevity, we use log as a shorthand to denote $\log_{10}$ throughout the paper.

5.1 Simulations

In this paper, we utilize the inference framework described in Hu22, which employs the b-N_{HI} distribution emulator built on neural networks trained on a set of Nyx simulations. We also use galaxy formation simulations IllustrisTNG and Illustris to investigate the low-z Lyα forest under different feedback mechanisms. Since this work focuses on the low redshift Lyα forest, we use $z=0.1$ simulation snapshots for all three simulations. In this section, we first provide a description of the simulations and the implemented physical models, followed by the (mock) data processing procedures employed in our study. This includes the generation of simulated line-of-sight (LOS) of Lyα forest (hereafter referred to as skewers for simplicity), forward modelling, and the Voigt profile fitting of Lyα lines. The cosmological parameters and thermal state of the three simulations are summarized in Table 5.1.

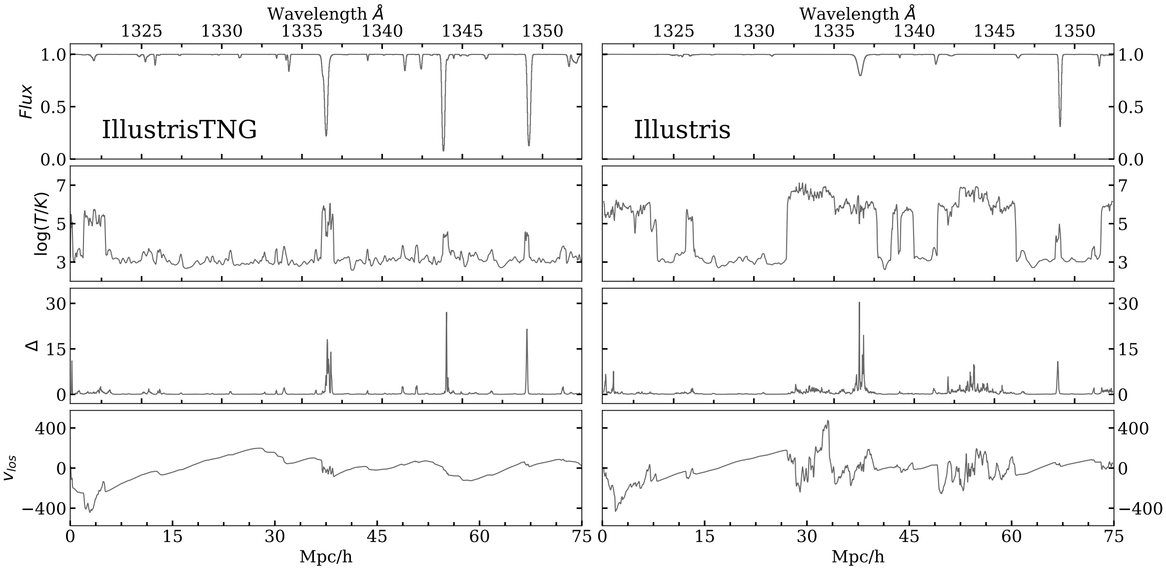

Figure 5.1: Examples of simulation skewers for IllustrisTNG (left) and Illustris (right) simulations, probing the structure generated by the same initial condition, while the two simulations are post-processed to share the same UV backgrounds photoionization rate, $\Gamma_{\rm HI}$. The flux is plotted in black on the top panel, while the temperature T, overdensity Δ, and velocity along LOS $v_{\rm los}$ are shown in black in the second, third, and bottom panels consecutively.

5.1.1 Nyx

Nyx is an adaptive mesh, massively parallel, cosmological simulation code primarily developed to simulate the IGM (Almgren et al. 2013, Lukić et al. 2015).

5.1.2 IllustrisTNG and Illustris

To evaluate the effectiveness of the IGM thermal state [T_0, γ] as the IGM parameters and test the efficacy of our inference framework on the realistic IGM, which can be affected by astrophysical processes that are not included in Nyx simulation such as galaxy formation and AGN feedback, we employ Illustris (Genel et al. 2014) and IllustrisTNG (Weinberger et al. 2017, Nelson et al. 2019) simulations, and use them as mock observational data in our inference method.

The IllustrisTNG and Illustris are cosmological hydrodynamic simulations powered by the AREPO code (Springel 2010). This code employs a moving mesh approach to solve hydrodynamics through the Euler equations, and it computes gravitational forces on a quasi-Lagrangian moving Voronoi mesh via the tree-PM method. Both simulations incorporate a wide range of astrophysical processes for galaxy formation, such as star formation, stellar and AGN feedback, galactic winds, and chemical enrichment (Springel et al. 2018, Naiman et al. 2018, Marinacci et al. 2018). They utilize the UV background detailed in Faucher-Giguère et al. (2009) for photoionization heating and cooling. Other processes for modelling the Lyα forest, like collisional ionization and inverse Compton cooling from the cosmic microwave background, are also taken into account.

The primary distinction between IllustrisTNG and Illustris lies in their AGN feedback mechanisms, especially regarding AGN feedback. Both simulations implement AGN feedback in two modes based on the gas accretion rate onto the central supermassive black hole: the ‘quasar-mode’ at high accretion rates (Springel 2005, Hopkins et al. 2008b, Debuhr et al. 2011) and the ‘radio-mode’ at low rates (Croton et al. 2006b, Bower et al. 2006, Sijacki et al. 2007b). While both use continuous thermal feedback in ‘quasar-mode’, their ‘radio-mode’ implementations differ. Illustris employs a bubble model for radio-mode feedback, accumulating substantial feedback energy for explosive release, often ejecting excessive hot gas (Genel et al. 2014). Conversely, IllustrisTNG models this feedback as a kinetic wind, injecting momentum into neighbouring regions from the central black hole. This approach better replicates astrophysical properties like star formation rates and galaxy colour distributions (Nelson et al. 2018, Pillepich et al. 2018a;b).

Both the IllustrisTNG and Illustris simulations we used in this study have box sizes of 75 cMpc/h and 1820^3 baryon and dark matter particles. Since AREPO is a moving mesh code, we convert the Voronoi mesh outputs to 1820^3 cartesian grids by dumping the smoothed quantities such as temperature[1], density, and velocities on grids to generate Lyα forest skewers. A Gaussian kernel with a size equal to 2.5 times the radius of each Voronoi cell is applied for the smoothing, assuming each Voronoi cell is spherical. We then generate skewers for IllustrisTNG and Illustris simulations following the approach discussed in § 2.1. In Fig. 5.1, we plot two simulation skewers for IllustrisTNG and Illustris respectively, while the two simulations are post-processed to share the same UV backgrounds photoionization rate, $\Gamma_{\rm HI}$ (see § 5.1.3 for more discussion). The flux ($e^{-\tau}$) is plotted in the top panel, and the temperature, over-density, and line-of-sight velocity profiles are shown in the second, third, and bottom panels consecutively. It is worth mentioning that the two skewers probe the structure generated by the same initial condition, suggesting that the differences in T, Δ, and $v_{\rm los}$ are primarily caused by different feedback strengths. Specifically, the Illustris exhibits higher temperatures due to its stronger feedback, which results in weaker absorption features given the same UV backgrounds. More discussion on the differences between Lyα forest in IllustrisTNG and Illustris simulations can be found in Khaire et al. (2023a) and Khaire et al. (2023b).

Such a fitting procedure is applied to all simulations used in this study, including all Nyx models and IllustrisTNG and Illustris simulations. The best-fit power law relationship based on $[T_0, \gamma]$ and the T-Δ distributions are illustrated in Fig. 5.2. The figure shows that although the three simulations yield very different overall T-Δ distributions, their thermal state T_0 and γ are however similar.

Furthermore, as described in Hu22, we vary the UV background photoionization rate, $\Gamma_{\rm HI}$, of the Nyx simulations in post-processing when the simulation skewers are generated, extending the parameter grid to $[\log T_0,\ \gamma,\ \log \Gamma_{\rm HI}]$. The value of $\Gamma_{\rm HI}$ we used in this study

---

[1] As presented in Appendix A of Martizzi et al. (2019), there exists a bug that affects the IGM temperature of the IllustrisTNG simulation. However, its effect on the Lyα forest is minimal because the bug predominantly impacts the gas with the lowest density. Consequently, we continue to use the uncorrected temperature for the IllustrisTNG simulation.

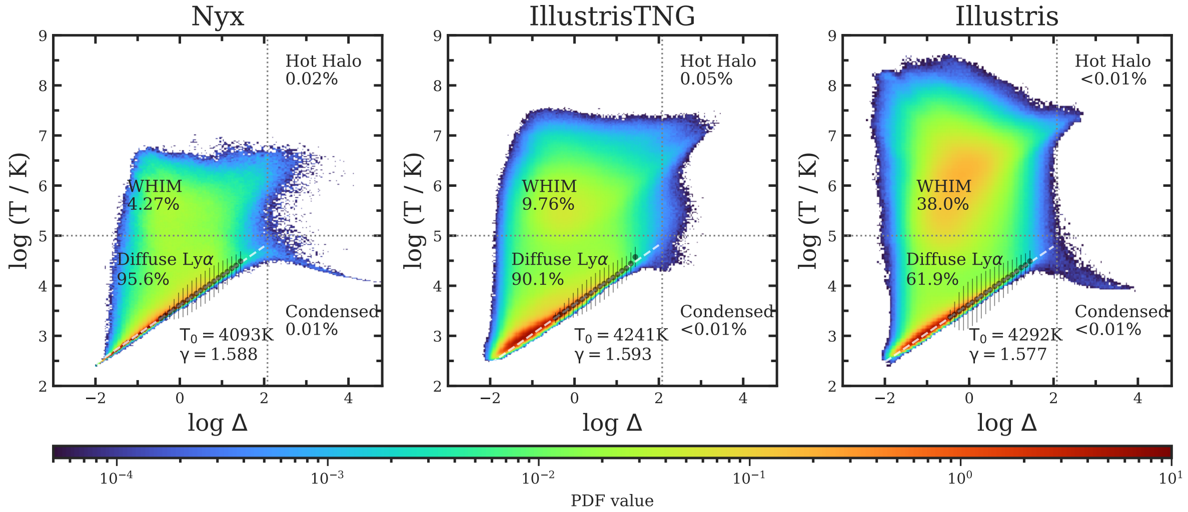

Figure 5.2: Volume weighted T-Δ distribution for all three simulations at $z = 0.1$. The $\log T$ for each bin are plotted as black dots, and the 1-σ_T error bars are shown as black bars. The best-fit power-law relationship is shown as white dashed lines. The Nyx (left) model is the default model which has $\log(T_0/\mathrm{K}) = 3.612$, $\gamma = 1.588$; and IllustrisTNG (middle) yields $\log(T_0/\mathrm{K}) = 3.627$, $\gamma = 1.593$; whereas Illustris (right) has $\log(T_0/\mathrm{K}) = 3.633$, $\gamma = 1.577$. The gas phase fractions are shown in the annotation.

spans from $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$ = -13.834 to -12.932 in logarithmic steps of 0.075 dex, which gives 13 values in total (see the right panel of Fig. 2.2). The range of Γ_{HI} used here covers more than twice the range obtained by UV background models of Khaire & Srianand (2019a) at $z = 0.1$, achieved by varying the spectral energy distribution of quasars. Note that the range also coves more than 2σ uncertainty in the Γ_{HI} measurements (Gaikwad et al. 2017b, Khaire et al. 2019). In total, the 3D thermal grid consists of $48 \times 13 = 624$ Nyx models.

As mentioned earlier, T_0 and γ characterize the IGM thermal state at $z \gtrsim 2$, where the IGM is dominated by the power law T-Δ relationship. However, their efficacy as parameters for the IGM thermal state remains uncertain at $z \lesssim 1$, where a significant fraction of the gas deviates from the power-law T-Δ relationship due to shock heating and feedback. In this paper, we evaluate the effectiveness of the thermal state $[T_0,\gamma]$

as IGM parameterization at low-z using the inference framework presented in Hu22, and we make use of the photoheating parameters [A,B] as an alternative set of labels as a comparison. These labels are particularly relevant since all Nyx models used in the training procedure of our neural network, which is the major component of our inference method, are generated by varying [A,B]. This suggests that our inference framework should be capable of recovering the values of [A,B] efficiently. Therefore, $[A, B]$ are particularly useful in the evaluation of the $[T_0, \gamma]$. More information about the photoheating labels [A,B] is presented in Appendix. C.1.

5.1.3 Photoioniztion rate $\Gamma_{\rm HI}$ and dN/dz

It is noteworthy that the three simulations used in this study by default have different UV background photoionization rates $\Gamma_{\rm HI}$ (for Nyx, here we are referring to the default model with $\log(T_0/{\rm K}) = 3.612$ and $\gamma = 1.588$.) This is because the photoionization rate $\Gamma_{\rm HI}$ are tuned in post-processing across all three simulations to ensure they exhibit the same absorber density dN/dz as the one we measured from D16 dataset at $z = 0.1$. Specifically, we apply the aforementioned VP-fitting procedure to D16 spectra (segments) with $0.06 < z < 0.16$, and obtain dN/dz=167.3 for absorbers within the limits $12.5 \leq \log(N_{\rm HI}/{\rm cm}^{-2}) \leq 14.5$ and $0.5 \leq \log(b/{\rm km\ s}^{-1}) \leq 2.5$. Such matching of d$N$/d$z$ is analogous to the matching of the mean flux of simulations at high-z for optically thin absorbers (Lukić et al. 2015). To match this dN/dz, we tune the photoionization rate, following the prescription described in § 2.1, and set $\log(\Gamma_{\rm HI}, /{\rm s}^{-1})$ = -13.093, -13.021, -13.414 for Nyx, IllustrisTNG and Illustris respectively (see Fig. 5.3). Such mismatch in $\Gamma_{\rm HI}$ is caused by the degeneracy between the photoionization rate and different feedback recipes used in the simulations. Since both the UV background and feedback suppress the formation of Lyα absorbers (Khaire et al. 2023a). More specifically, the feedback heats up the IGM, converting a significant amount of the diffuse Lyα gas into WHIM, which reduces the Lyα transmission caused by the neutral hydrogen H I in the cool diffuse Lyα gas. To this end, simulations with stronger feedback exhibit lower dN/dz under the same $\Gamma_{\rm HI}$.

We measure the dN/dz for the three simulations, including all Nyx simulation models and IllustrisTNG and Illustris, each based on its respective set of 1000 forward-modelled mock spectra. The relationships between UV background photoionization rate and dN/dz for all three models are shown in Fig. 5.3, where the dN/dz for Nyx is plotted in blue, IllustrisTNG in green, and Illustris in red, while the dN/dz for the D16 data at $z = 0.1$ is shown as the horizontal dash-dotted grey line. Fig. 5.3 demonstrates that while Illustis has the strongest feedback, which causes more gas to be collisionally ionized, reducing the Lyα absorption, it requires the lowest $\Gamma_{\rm HI}$ to match the dN/dz to the observed value, and IllustrisTNG, with mild feedback, has higher for the same UV background. Interestingly, whereas Nyx employs no feedback mechanism, it requires slightly lower $\Gamma_{\rm HI}$ compared with IllustrisTNG, which implements feedback, to achieve the same dN/dz. Such a trend, which is opposite to the correlation between dN/dz and feedback strength (as seen between Illustris and IllustrisTNG), is caused by the small difference in the Δ, T distributions in Nyx and IllustrisTNG. More specifically, IllustrisTNG exhibits a slightly higher mass fraction of the diffuse Lyα gas in the particular Δ, T range that is probed by the Lyα forest[2]. It is possible that while the mild feedback in the IllustrisTNG simulation results in a slightly higher WHIM fraction compared with Nyx, it also produces more gas with $T \sim 10^{4.5}$K and $\Delta \sim 10$, which is the Δ, T range probed by the Lyα forest at $z = 0.1$ (see § 5.3). However, the detailed astrophysical mechanism

---

[2] This is different from the mass-weighted or volume-weighted diffuse Lyα fraction, $f_{\rm Ly\alpha}$, which is defined to include all gas with $T \leq 10^5$K and $\Delta \leq 120$ following Davé et al. (2010). Instead, the gas probed by the observed HST COS Lyα Forest has a narrow range of Δ and T which also depends on the $\Gamma_{\rm HI}$ value used in the simulation (See § 5.3 for more details).

leading to this specific Δ-T distribution in IllustrisTNG at $z = 0.1$ is still unclear, and we leave it to our future work.

If not otherwise specified, the three simulations used in this study, including Illustris, IllustrisTNG, and Nyx default model, are tuned to have the same Lyα line densities, with $\mathrm{d}N/\mathrm{d}z = 167.3$, which is the same value we measured from the D16 dataset.

5.2 Inference Method

5.2.1 Emulating the $\{b, N_{\mathrm{HI}}\}$ Distribution

In this work, we make use of the inference framework following Hu22, which measures the thermal state and the photoionization rate Γ_{HI} of the low redshift IGM using its b-N_{HI} distribution and absorber line density $\mathrm{d}N/\mathrm{d}z$.

To perform our analysis under realistic conditions, all tests performed in this paper are based on mock datasets consisting of 34 forward-modelled spectra, each corresponding to one of the 34 D16 quasar spectra, which gives these datasets the same pathlength as the observation dataset with $\Delta z_{\mathrm{ob}} = 2.136$. Each of the mock datasets is constructed by randomly selecting 34 spectra from the 1000 forward-modelled spectra, while making sure that each of the 34 D16 quasar spectra is represented exactly once, thereby maintaining the integrity and representativeness of our mock datasets.

5.2.2 Inference test

An inference test is an effective method to evaluate the robustness of a given inference algorithm, which usually consists of approximations and emulation/interpolation procedures that might induce additional uncertainties, altering the error budget. In practice, an inference test can be conducted by performing a set of realizations of the inference

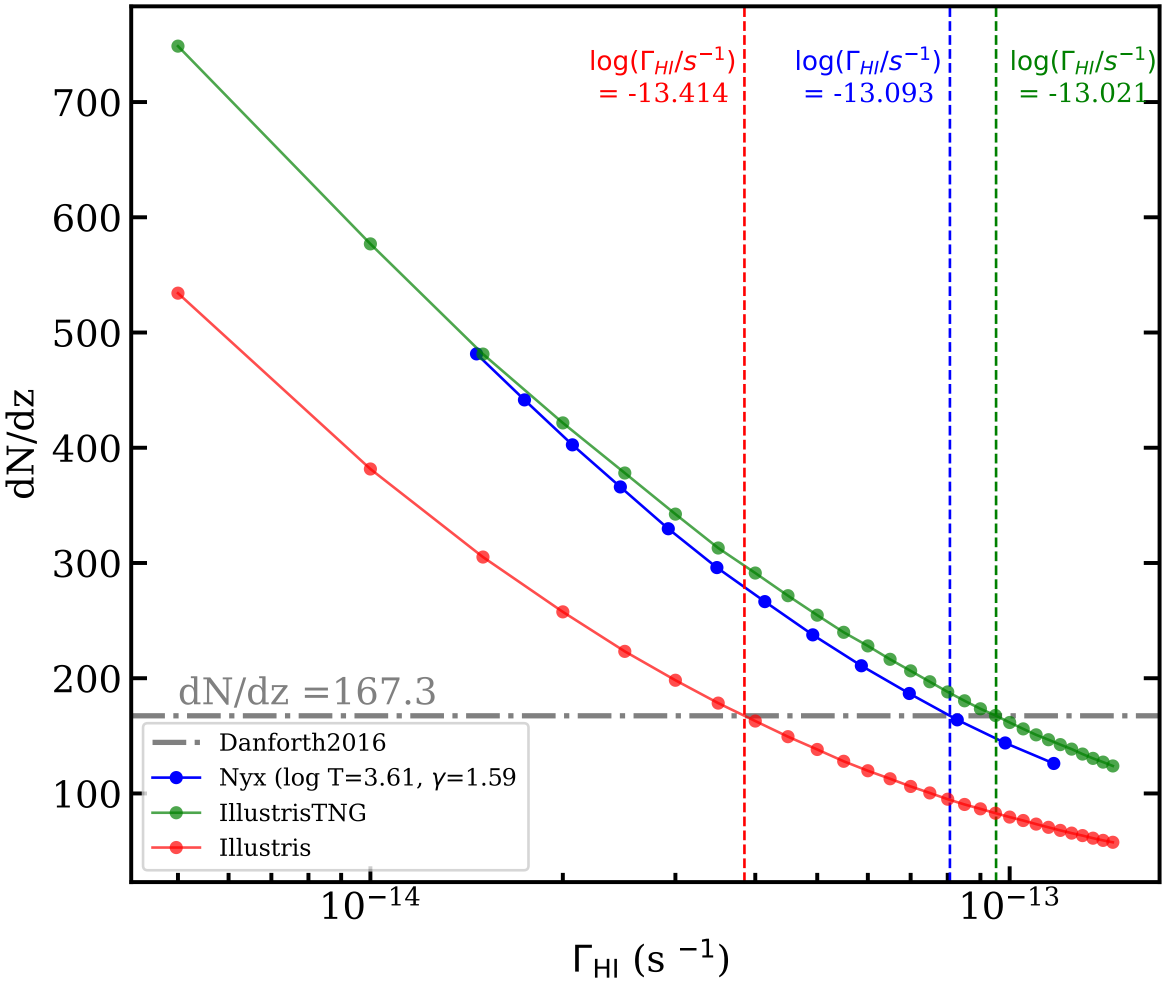

Figure 5.3: dN/dz vs $\Gamma_{\rm HI}$ for all three simulations at z=0.1. Nyx (default model) is shown in blue, IllustrisTNG in green, and Illustris in red, while the observed dN/dz calculated from D16 dataset for the corresponding redshift are shown in the horizontal grey dash-dotted line. The $\Gamma_{\rm HI}$ values used for each simulation to match the observed dN/dz are indicated by vertical dashed lines with the corresponding colour.

method using mock datasets and evaluating the robustness of the resulting posterior probability distributions, which can be quantified by the coverage probability $P_{\rm cov}$ (Prangle et al. 2014, Ziegel & Gneiting 2014, Morrison & Simon 2018, Sellentin & Starck 2019), the proportion of the time that the true parameters used to generate a mock dataset are contained within the posterior contour corresponding to a certain probability level $P_{\rm inf}$. Such calculations can be performed for many different probability levels, resulting in a series of coverage probabilities. For perfect inference, this coverage probability $P_{\rm cov}$ is always equal to the probability level of the chosen posterior contour $P_{\rm inf}$ (shown as the black dashed line in Fig. 5.4).

In this study, we make use of the inference test described in Wolfson et al. (2022), which calculates the coverage probability based on the MCMC posteriors. Compared with the one used in Hu22, this inference test algorithm is more precise and automatically returns full coverage probabilities from 0 to 1 rather than coverage probabilities at only a few specific probability levels.

To evaluate the effectiveness of $[T_0, \gamma, \log \Gamma_{\rm HI}]$ as IGM parameters for inference at low-z, where the IGM T-Δ distribution is no longer characterized by the power-law relationship, we perform inference tests based on different sets of labels. We compare the result of the inference test based on labels $[T_0, \gamma, \log \Gamma_{\rm HI}]$ with the one based on the photoheating labels $[A, B, \log \Gamma_{\rm HI}]$. As discussed in § 2.1, the comparison between these two sets of labels sheds light on the efficacy of $[T_0, \gamma]$ as IGM parameters at low-z, where the pervasive shock heating causes significant dispersion in the IGM T-Δ distribution.

For each set of labels, we ran 300 realizations of our inference method, each based on a model randomly chosen from the grid. We exclude models that are close to the boundaries to mitigate boundary effects caused by the hard cutoff of the inference prior, which leads to the truncation of the posteriors. For $[T_0, \gamma, \log \Gamma_{\rm HI}]$ gird, we specify $3.3 < \log(T_0/{\rm K}) < 3.9$, $1.0 < \gamma < 2.3$, $-13.75 < \log(\Gamma_{\rm HI}/{\rm s}^{-1}) < -13.0$. We then

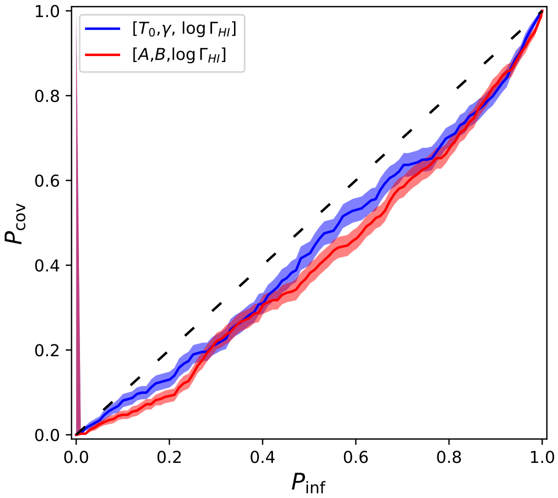

Figure 5.4: Coverage probability $P_{\rm cov}$ for inference tests based on different labels. The x-axis stands for the inferred probability $P_{\rm inf}$, and the y-axis shows the coverage probability $P_{\rm cov}$ for the true parameters to fall in the contour corresponding to $P_{\rm inf}$. Blue: Inference test based on the thermal state $[T_0,\gamma, \log\Gamma_{\rm HI}]$, Red: Inference test based on the photoheating label $[A,B,\log\Gamma_{\rm HI}]$. The shaded regions indicate the 1-σ error for $P_{\rm cov}$.

calculate the full coverage probabilities based on all 300 MCMC posteriors.

The results of the inference tests are shown in Fig. 5.4, where the x-axis stands for inferred probability $P_{\rm inf}$, and the y-axis shows the coverage probability, $P_{\rm cov}$. The shaded regions indicate the 1-σ error for $P_{\rm cov}$, which is calculated based on the binomial distribution. The $y = x$ black-dash line represents a perfect inference test. It can be seen that for Nyx simulations, our inference method is mildly over-confident, and the thermal state $[T_0, \gamma]$ (blue) performs slightly better than the photoheating labels $[A,B]$ (red), i.e., $P_{\rm cov}/P_{\rm inf}$ is closer to unity. Quantitatively, for inference based on the thermal state $[T_0, \gamma]$, the 68% contour contains the true parameters $61.2 \pm 2.8\%$ of the time, and the 95% contour contains the true parameters $90.4 \pm 1.6\%$ of the time. The results show that the $[T_0, \gamma]$ are still robust inference labels for the IGM at low-z, although the shock heating alters the T-Δ distribution. This further suggests that shock heating alone does not significantly change our understanding in determining the thermal state of the

IGM using the Lyα forest. Lastly, while the general efficacy of the inference framework remains robust, we attribute its imperfections to two primary sources: deficiencies within the neural network used in our inference algorithm, and the boundary effects caused by the truncation of the posteriors when hitting the boundary.

5.2.3 Inference results for IllustrisTNG and Illustris

In this section, we employ the IllustrisTNG and Illustris simulations as mock observational data to explore the impacts of feedback, mainly AGN feedback, on the IGM thermal state $[T_0, \gamma]$. More specifically, we evaluate the robustness of our inference method, built on the Nyx thermal grid without galaxy formation and feedback, when applied to observational data derived from a (mock) Universe with substantial feedback associated with galaxy formation and AGN activities. The investigation is broken down into two separate inquiries. First, it explores the extent to which feedback associated with galaxy formation and AGN activities impacts the Lyα forest. Second, it investigates how, given the presence of these effects, the feedback influences the inferred parameters $[T_0, \gamma]$.

Following the forward-modelling prescription described in § 2.1.3, we generate mock datasets with Δz=2.136, the pathlength of D16 dataset at $z = 0.1$, for both simulations (see Fig. 5.6), and run our inference method on each dataset. As discussed in § 5.1.3, a degeneracy exists between the strength of the AGN feedback implemented in the simulations and the UV background photoionization rate $\Gamma_{\rm HI}$, both of which suppress the abundance of absorbers, hence reducing the $\mathrm{d}N/\mathrm{d}z$ (see Khaire et al. 2023a, for more details). Given that our inference method primarily derives the photoionization rate $\Gamma_{\rm HI}$ based on the $\mathrm{d}N/\mathrm{d}z$, the resulting $\Gamma_{\rm HI}$ always aligns with the value that generates the equivalent $\mathrm{d}N/\mathrm{d}z$ in the Nyx simulation (see § 5.1.3). Since here we use IllustrisTNG and Illustris simulations with their $\mathrm{d}N/\mathrm{d}z$ matched to the D16 low-z dataset, the inferred

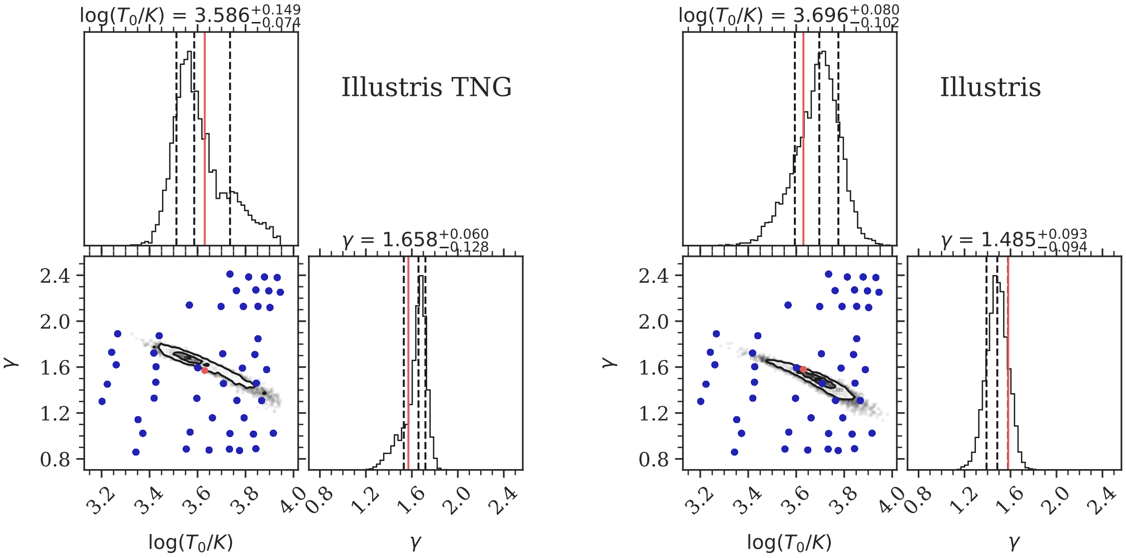

Figure 5.5: Posteriors obtained by using IllustrisTNG (left) and Illustris (right) as mock observational data in our inference method. Projections of the thermal grid used for generating models are shown as blue dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The true parameters for the simulations, obtained by fitting the T-Δ distributions of the simulations, are indicated by the red dot (lines) in the (marginal) distributions, while the dashed black lines indicate the 16, 50, and 84 percentile values of the marginalized 1D posterior.

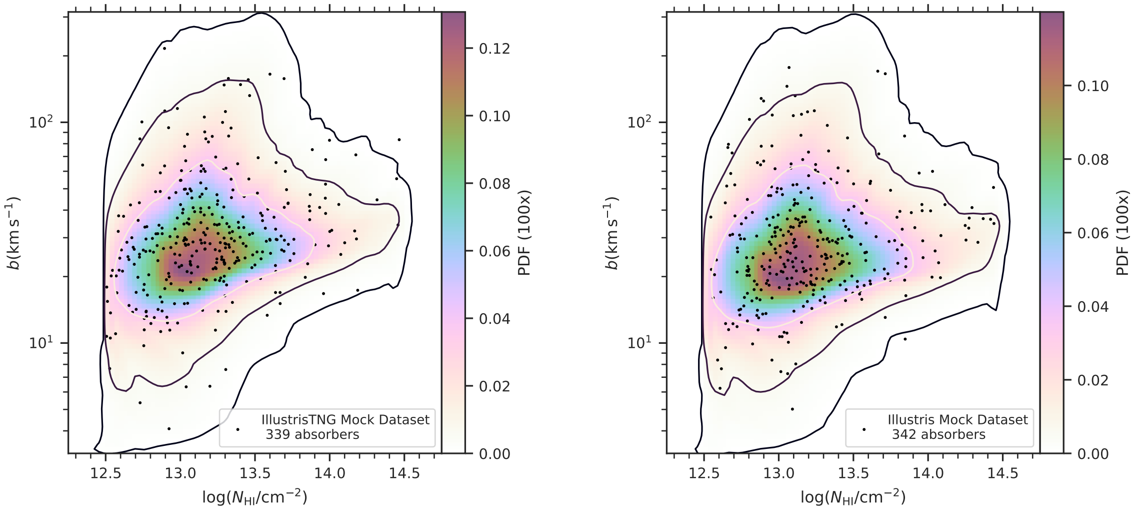

Figure 5.6: Joint b-$N_{\rm HI}$ distributions recovered from the inference results for IllustrisTNG (left) and Illustris (right) simulations, emulated by our DELFI emulator based on the median values of the marginalized MCMC posterior. Black dots are the mock datasets we used in the inference. The contours correspond to cumulative probabilities of 68%, 95% and 99.7%. For illustration purposes, the values of the pdf are multiplied by 100 in the colour bar.

$\Gamma_{\rm HI}$ always disagrees with the true values used to generate the IllustrisTNG and Illustris simulations. To this end, we conduct our inference test in 2D without considering the accuracy with which we recover the photoionization rate $\Gamma_{\rm HI}$. Posterior distributions for the thermal parameters obtained from our inference applied to Illustris and IllustrisTNG are shown in Fig. 5.5, where we have marginalized over $\Gamma_{\rm HI}$. For these two mock datasets, we infer that $[\log(T_0/{\rm K}),\ \gamma] = [3.586^{+0.149}_{-0.074}, 1.658^{+0.060}_{-0.128}]$ for IllustrisTNG ([3.627,1.593]), and $[\log(T_0/{\rm K}),\ \gamma] = [3.696^{+0.080}_{-0.102}, 1.485^{+0.093}_{-0.094}]$ for Illustris ([3.633,1.577]), while the true parameters for the two simulations, $[T_0, \gamma]_{\rm fit}$, are given in parentheses respectively.

It can be seen that the true parameters $[T_0, \gamma]_{\rm fit}$, obtained by fitting the T-Δ distributions of the simulations, are within 1-σ errors (1D marginalized) for both simulations, and the 1-σ errors for both the IllustrisTNG and Illustris simulations are slightly larger than those for Nyx simulations, which is caused by the intrinsic difference between Nyx, IllustrisTNG and Illustris simulations, where the latter two are based on completely different hydrodynamic codes. In Fig. 5.6, we present both the mock datasets used for inference and the b-$N_{\rm HI}$ distributions emulated based on the inference results. The plots highlight strong agreement between the emulated b-$N_{\rm HI}$ distributions and the respective mock dataset for each simulation.

Nevertheless, it is worth mentioning that the inferred thermal state for IllustrisTNG and Illustris presented above are based on realistic conditions, with total pathlength $\Delta z = 2.136$. Such a small Δz makes the inference result vulnerable to randomness induced by the selection of mock datasets. To address this issue, here we conduct our inference on IllustrisTNG and Illustris simulations, using datasets with much larger pathlength, specifically with $\Delta z = 42.47$, which is 20 times the size of the observational dataset. The inference results yield $[\log(T_0/{\rm K}),\ \gamma] = [3.605^{+0.031}_{-0.027}, 1.657^{+0.022}_{-0.024}]$ for IllustrisTNG ([3.627,1.59]), and $[\log(T_0/{\rm K}),\ \gamma] = [3.680^{+0.019}_{-0.020}, 1.483^{+0.021}_{-0.021}]$ for Illustris ([3.633,1.58]), while the true parameters for the two simulations, $[T_0, \gamma]_{\rm inf}$, are given in

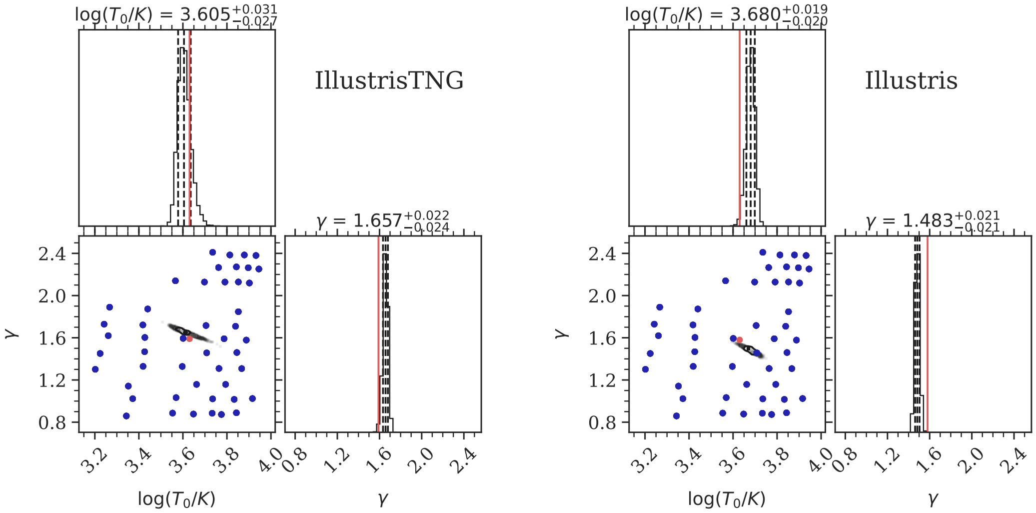

Figure 5.7: Corner plots for IllustrisTNG (left) and Illustris (right), based on the larger mock dataset, with $\Delta z = 42.72$, corresponding 20 times the observational dataset. Projections of the thermal grid used for generating models are shown as blue dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The true parameters for the simulations, obtained by fitting the T-Δ distributions of the simulations, are indicated by the red dot (lines) in the (marginal) distributions, while the dashed black lines indicate the 16, 50, and 84 percentile values of the posterior.

parentheses. The resulting corner plots are presented in Fig. 5.7. These results are used as our inferred thermal state $[T_0, \gamma]_{\rm inf}$ for IllustrisTNG and Illustris simulations in the following part of this study. It is noticeable that the inferred T_0 for Illustris is higher than the true value with an error $\Delta \log(T_0/\mathrm{K}) = 0.047$ dex, while the γ is below the true value, with $\Delta\gamma = -0.094$. For IllustrisTNG, the offsets between the $[T_0, \gamma]_{\rm inf}$ and $[T_0, \gamma]_{\rm fit}$ are smaller, with $\Delta \log(T_0/\mathrm{K}) = -0.022$ dex, $\Delta\gamma = 0.064$. We notice that these offsets are smaller than the typical inference precision obtained based on realistic datasets, as shown in Fig. 5.5, which report the marginalized 1D 1-σ error in $\log T_0$, $\sigma_{\log T_0}$, ~ 0.1 dex and the marginalized 1D 1-σ error in γ, σ_γ, ~ 0.1. For both simulations we observe the offsets $\Delta \log T_0 \lesssim 0.5\sigma_{\log T_0}$, and $\Delta\gamma \lesssim \sigma_\gamma$.

To check the robustness of these results, we use the IllustrisTNG and Illustris simulations as mock observational data and perform inference tests using two different sets of 'true parameters': the $[T_0, \gamma]_{\rm fit}$ obtained from our power law fits the $\Delta - T$ distribution of the simulations (see Fig. 5.2), and the $[T_0, \gamma]_{\rm inf}$ given by our inference method when applied to an extremely large mock dataset, as described above. Given that the inferred $\Gamma_{\rm HI}$ for both IllustrisTNG and Illustris simulations consistently deviates from the actual values, owing to the previously mentioned degeneracy between the photoheating rate and feedback strength, any inference tests incorporating the $\Gamma_{\rm HI}$ from these two simulations will surely fail. To this end, we focus on the inference results on the T_0-γ plane and conduct marginalized inference tests by marginalizing the posteriors over the $\Gamma_{\rm HI}$, in which 2D marginalized contours levels are modelled by Gaussian mixture models. For each simulation, we run 100 realizations on each set of 'true parameters', and run inference tests on the obtained posteriors. The results are shown in Fig. 5.8, indicating that our inference method is over-confident for both sets of 'true parameters'. While the inference method is not able to recover the thermal state $[T_0, \gamma]_{\rm fit}$, the thermal state $[T_0, \gamma]_{\rm inf}$ significantly improves the outcome of the inference test. These results suggest that our inference method is able to robustly recover the $[T_0, \gamma]$ with small biases, for simulations that include feedback mechanisms.

The inference tests imply that there exist offsets for the inferred parameters $[T_0, \gamma]_{\rm inf}$ for IllustrisTNG and Illustris, where $\Delta \log(T_0/{\rm K}) = -0.022$ dex, $\Delta\gamma = 0.064$ for IllustrisTNG and $\Delta \log(T_0/{\rm K}) = 0.047$ dex, $\Delta\gamma = -0.094$ for Illustris. However, these offset are insignificant, with $\Delta \log T_0 \lesssim 0.5\sigma_{\log T_0}$, and $\Delta\gamma \lesssim \sigma_\gamma$. However, it is unclear whether the observed differences between $[T_0, \gamma]_{\rm inf}$ and $[T_0, \gamma]_{\rm fit}$ can be attributable to the intrinsic difference between the Nyx, IllustrisTNG, and Illustris simulations, or if they arise from potential degeneracy between the IGM thermal state $[T_0, \gamma]$ and the feedback mechanism implemented in the simulation. Nevertheless, the latter hypothesis seems to contrast with the results based on the various statistics of the low-z Lyα forest presented in Khaire et al. (2023a), which suggests that the impacts from different feedback models

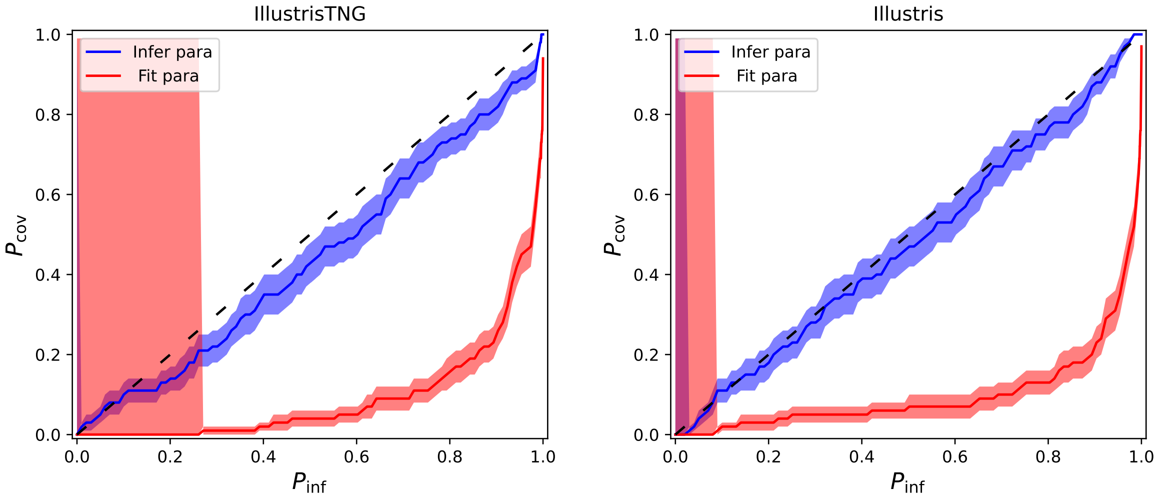

Figure 5.8: Marginalized coverage probability $P_{\rm inf}$ for inference tests using IllustrisTNG (left) and Illustris (right) simulations as mock observational data. The x-axis stands for the inferred probability $P_{\rm inf}$, and the y-axis shows the probability for the parameters of the true model to fall in the contour corresponding to $P_{\rm cov}$. The shaded regions indicate the $1-\sigma$ error for $P_{\rm cov}$. Inference tests with the true parameters set by $[T_0, \gamma]_{\rm inf}$ are shown in blue, while inference tests with the true parameters set by $[T_0, \gamma]_{\rm fit}$ are shown in red.

are not distinguishable via the Lyα forest under realistic scenarios, i.e., forward-modelled using the D16 COS dataset. The only exception is the case of the Lyα flux power spectrum at small scales, where minor deviations are observed in both simulations(see Khaire et al. 2023b, for the dN/dz around massive halos). To further explore this problem, we examine the physical properties of low -z Lyα absorbers in the following section.

5.3 Low-z Lyα Forests and Simulated Absorbers

5.3.1 Identifying the simulated Lyα absorbers

To understand whether the low-z Lyα forest effectively probes the WHIM, we attempt to identify the simulated Lyα absorbers, i.e., the $n_{\rm HI}$ peaks in the simulation skewers, that give rise to the Lyα lines detected in the mock spectra. This approach allows us to

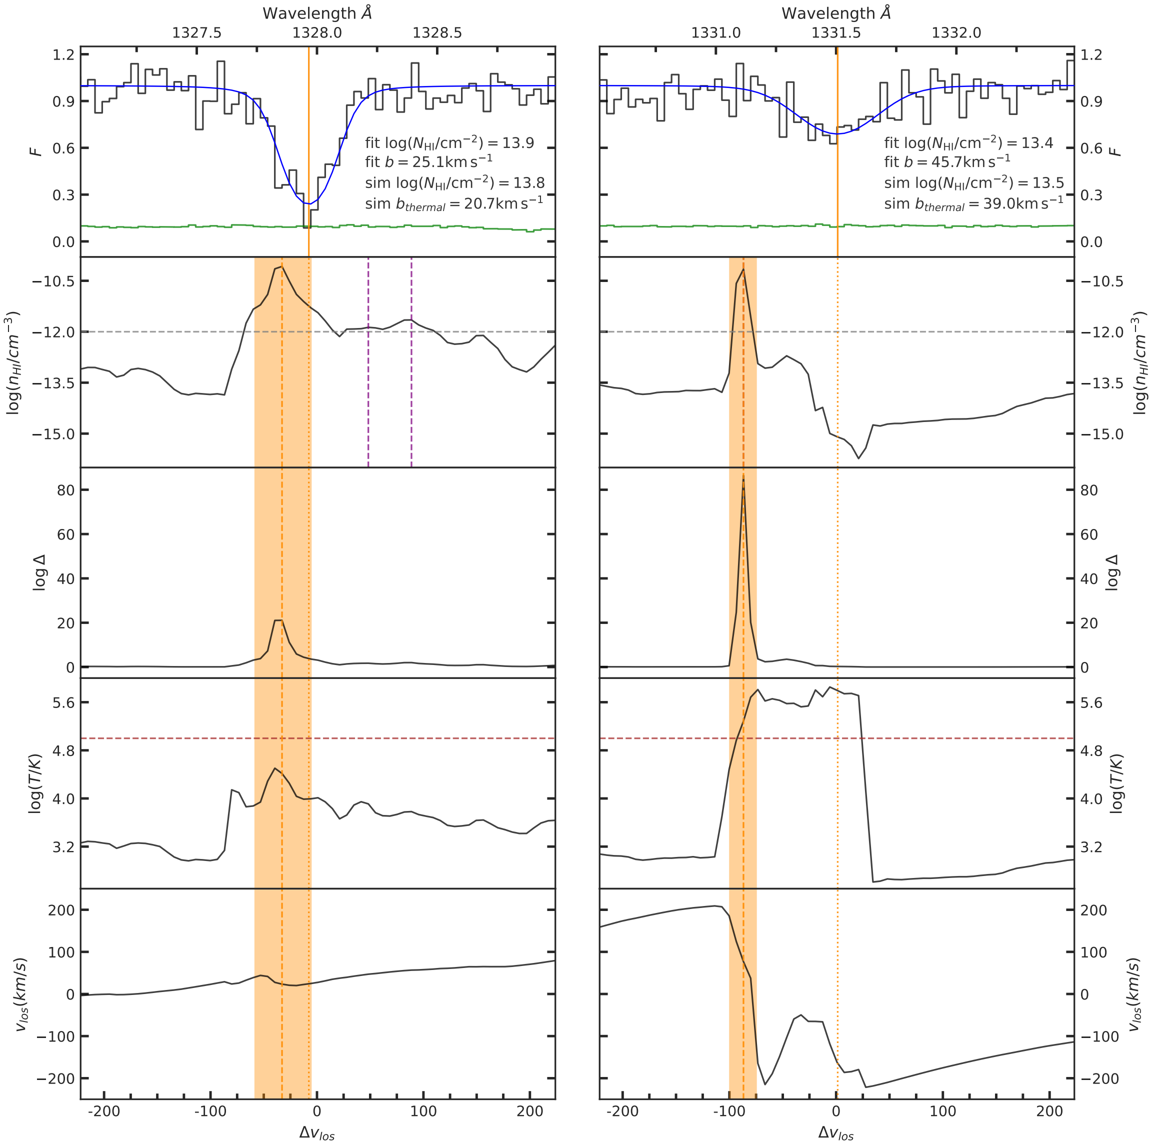

Figure 5.9: Illustration of a segment of one of the forward-modelled mock spectra (top panel) with the absorption lines detected by VPFIT and the corresponding skewer. The top panel depicts the flux (black), noise vector (green), and the model fitted by `VPFIT` (blue). The central wavelength of Lyα lines identified by `VPFIT` are indicated by orange vertical lines, and the corresponding simulated absorbers are indicated by orange dashed lines in the second panel (and all other panels below). The log $N_{\rm HI,fit}$, $b_{\rm fit}$ reported by `VPFIT` are given in the annotation, together with the log $N_{\rm HI,sim}$, $b_{\rm thermal}$ calculated based on the simulation. The second panel depicts the neutral hydrogen density $n_{\rm HI}$, while the shaded regions represent the identified Lyα absorbers along LOS, which are used as the integral ranges while computing the $N_{\rm HI,sim}$. The orange vertical dashed lines show the $n_{\rm HI}$ peaks of the confirmed simulated Lyα absorbers, while the purple vertical dashed lines show the potential simulated Lyα absorbers that do not cause detectable Lyα lines.

Figure 5.9: The grey horizontal dashed line represents the minimal H I peak density, $n_{\mathrm{HI,min}} = 10^{-12}\ \mathrm{cm}^{-3}$. The third, fourth and fifth panels show the overdensity Δ, temperature T and LOS velocity v_{los}. The brown horizontal dashed line in the fourth panel stands for $T = 10^5$ K. *left*: A Lyα absorbers in the diffuse Lyα phase. *right*: A Lyα absorbers in the WHIM phase.

directly examine the physical properties (T, Δ, and n_{HI}) of these simulated Lyα absorbers and draw a direct correspondence between them and the line parameters ($\{b, N_{\mathrm{HI}}\}$) of their corresponding Lyα lines detected in the mock spectra. In this section, we chose to focus on the simulated Lyα absorbers in the Nyx simulation at $z = 0.1$ (with default thermal history, i.e., $\log(T_0/K) = 3.612$, and $\gamma = 1.588$ at $z = 0.1$). For clarification, within the context of this study, the terms 'simulated Lyα absorbers' or simply 'simulated absorbers' are used to denote the n_{HI} peaks that give rise to the Lyα absorption lines in the mock spectra detected by `VPFIT`.

Our approach for identifying simulated Lyα absorbers works as follows. Firstly, we include the physical properties, including temperature T, over-density Δ, velocity along LOS v_{los}, and the neutral fraction x_{HI} in our skewers and stitch them in the forward-modelling procedure (see § 2.1.3). We interpolate the stitched skewer onto the forward-modelled wavelength grid, and calculated the neutral hydrogen density n_{HI} for each simulation cell, based on the neutral fraction x_{HI}, over-density Δ, and the mean hydrogen density $\bar{n}_H$. Subsequently, we scan the stitched skewers (in real space) for n_{HI} peaks, and classify these with $n_{\mathrm{HI}} > 10^{-12}\ \mathrm{cm}^{-3}$ as potential simulated Lyα absorbers. The minimal peak H I density is derived from both the minimal H I column density for the detected lines $N_{\mathrm{HI,min}} = 10^{12.5}\mathrm{cm}^{-2}$ (see § 2.1.3) and the maximal length for simulated absorbers $l_{\mathrm{abs,\ max}} = 0.5$ Mpc/h, which is consistent with previous studies that attempt to characterize the structures giving rise to the Lyα forest at $z = 0.1$ (Bolton et al. 2022b, Tillman et al. 2023b). Given these two parameters, the requisite minimum H I peak density for simulated absorbers to yield observable Lyα absorption lines is computed as

$n_{\mathrm{HI,min}} = \frac{N_{\mathrm{HI,min}}}{l_{\mathrm{abs,\ max}}} = \frac{10^{12.5}\mathrm{cm}^{-2}}{0.5\mathrm{Mpc/h}} \sim 10^{-12}$ cm^{-3}, which effectively filters out n_{HI} peaks that give rise to Lyα absorption lines below our sensitivity. We then determine the physical size for each potential simulated absorber along the LOS, l_{abs}, using a threshold at which n_{HI} drops below 1% of its peak value, while restricting the maximal size to be $l_{\mathrm{abs,\ max}} =$ 0.5 Mpc/h. We calculate the H I column densities of the simulated Lyα absorbers, $N_{\mathrm{HI,\ sim}}$, by integrating the n_{HI} over the ranges set by aforementioned threshold. We observed that the resulting N_{HI} is not particularly sensitive to l_{abs}, because the n_{HI} peak is so narrow that the majority of the neutral hydrogen comes from the peak region (see Fig. 5.9).

After identifying the potential Lyα absorbers, we extract their LOS velocity from the simulation cells, and compute the central wavelength of the expected absorption lines in redshift space, accounting for the redshift caused by its LOS velocity. For each anticipated absorption line originating from an n_{HI} peak, we check whether its central wavelength lies within ± 50km/s of the central wavelength of any Lyα lines detected in the mock spectrum. If so, we confirm the identification of a simulated Lyα absorber, and take the T and Δ at the n_{HI} peak as its values, which is valid since the n_{HI} peak is so narrow that the majority of the N_{HI} comes from the region close to the peak. While theoretically, the Lyα lines are expected to be caused by multiple n_{HI} peaks in real space (Garzilli et al. 2015), we discover that at $z = 0.1$, each Lyα line detected in the mock spectra with $12.5 \leq \log(N_{\mathrm{HI}}/\mathrm{cm}^{-2}) \leq 14.5$ predominantly originates from one single n_{HI} peak in the simulation. It is not a resolution effect, since the simulation (Nyx) has a grid length ~ 0.02 Mpc/h, while the simulated absorbers have sizes ~ 0.5 Mpc/h. As such, we only consider the n_{HI} peak with the highest n_{HI} value if multiple n_{HI} peaks contribute to the same detected absorption line.

Fig. 5.9 showcases examples of the simulated Lyα absorbers, alongside their corresponding absorption lines in the mock spectra and the related simulation skewers. The top panel depicts the flux (black), noise vector (green), and the model fitted by `VPFIT` (blue). The central wavelength of Lyα lines identified by `VPFIT` are indicated by orange vertical lines, and the corresponding simulated absorbers are indicated by orange dashed lines in the second panel (and all other panels below). The $\log N_{\mathrm{HI,fit}}$, b_{fit} reported by `VPFIT` are given in the annotation, together with the $\log N_{\mathrm{HI,sim}}$, b_{thermal} calculated based on the simulation, whereas the $b_{\mathrm{thermal}} = (2kT/m_{\mathrm{HI}})^{1/2}$ is the thermal component of the b-parameters computed based on the T of the simulated Lyα absorbers (see eq. 5.1). The second panel depicts the neutral hydrogen density n_{HI}, while the shaded regions represent the identified Lyα absorbers along LOS, which are used as the integral ranges while computing the $N_{\mathrm{HI,sim}}$. The orange vertical dashed lines show the n_{HI} peaks of the confirmed simulated Lyα absorbers, while the purple vertical dashed lines show the potential simulated Lyα absorbers that do not cause detectable Lyα lines. The grey horizontal dashed line represents the minimal H I peak density, $n_{\mathrm{HI,min}} = 10^{-12}$ cm^{-3}. The third, fourth and fifth panels show the overdensity Δ, temperature $\log T$ and LOS velocity v_{los} (black solid lines). The brown horizontal dashed line in the fourth panel stands for $T = 10^5$ K, which divides the cool diffuse Lyα gas and the WHIM. The left panel shows a simulated Lyα absorber in the diffuse Lyα phase, while the right left panel shows a simulated Lyα absorber arising from the WHIM phase.

We perform the identification procedure for all 1000 mock spectra, discovering 34011 potential simulated Lyα absorbers,i.e., n_{HI} peaks, among which 10510 are identified as simulated Lyα absorbers and matched to their respective absorption lines identified by VPFIT. The discrepancy between potential and confirmed Lyα absorbers is due to the inclusion of minor n_{HI} peaks, that are too weak to cause any detectable Lyα line, which is indicated by purple vertical lines in the left panel of Fig. 5.9. Lastly, approximately 2% of the lines detected by VPFIT could not be matched to any simulated Lyα absorber. These anomalies could potentially result from false identification of the `VPFIT` induced by noise. Nonetheless, given the rarity of these cases, omitting them should not influence our statistical results or conclusions.

To validate our identification method, we compare the observed line parameters, reported by `VPFIT`, with the values calculated from the simulation. In Fig. 5.10, we showcase the $N_{\mathrm{HI,fit}}$ (left) and b_{fit} (right) for all Lyα lines fitted by VPFIT, compared with the $N_{\mathrm{HI,sim}}$ and b_{thermal} respectively, both calculated from the corresponding simulated Lyα absorbers identified in the Nyx simulation. The left panel indicates a strong correlation between the fitted $N_{\mathrm{HI,fit}}$ and the $N_{\mathrm{HI,sim}}$ calculated from the simulation, implying that the n_{HI} peaks identified by our method are indeed the simulated Lyα absorbers responsible for the Lyα lines detected in the mock spectra. The right panel demonstrates that the bulk of b_{fit} lies above the dashed line representing $b_{\mathrm{fit}} = b_{\mathrm{thermal}}$. This result aligns with the nature of the b-parameter, as given by

$$b = \sqrt{b_{\mathrm{thermal}}^2 + b_{\mathrm{notherm}}^2}, \tag{5.1}$$

where the b_{notherm} is the non-thermal component of the b-parameter resulting from combinations of Hubble flow, peculiar velocities and turbulence in the IGM. Eq. 5.1 demonstrates that the b_{thermal} is the lower limit of the b-parameter, which corresponds to the lower right cutoff of the b-N_{HI} distribution (see the colour maps in Fig. 2.7 and Fig. 5.6 as examples. More discussions on this topic can be found in Schaye et al. 1999, Rudie et al. 2012a, Bolton et al. 2014, and Hu22). Furthermore, the right panel of Fig. 5.10 gives a rough correlation between the b_{fit} and b_{thermal} and provides an approximate estimation of the strength of the non-thermal broadening of the Lyα lines at $z = 0.1$. It suggests that for the Nyx simulation, the non-thermal contribution to the b-parameter can be modelled by a 'turbulent' motion in the IGM with $b_{\mathrm{notherm}} \sim 20$ km/s (indicated by the black dash-dot line in Fig. 5.10).

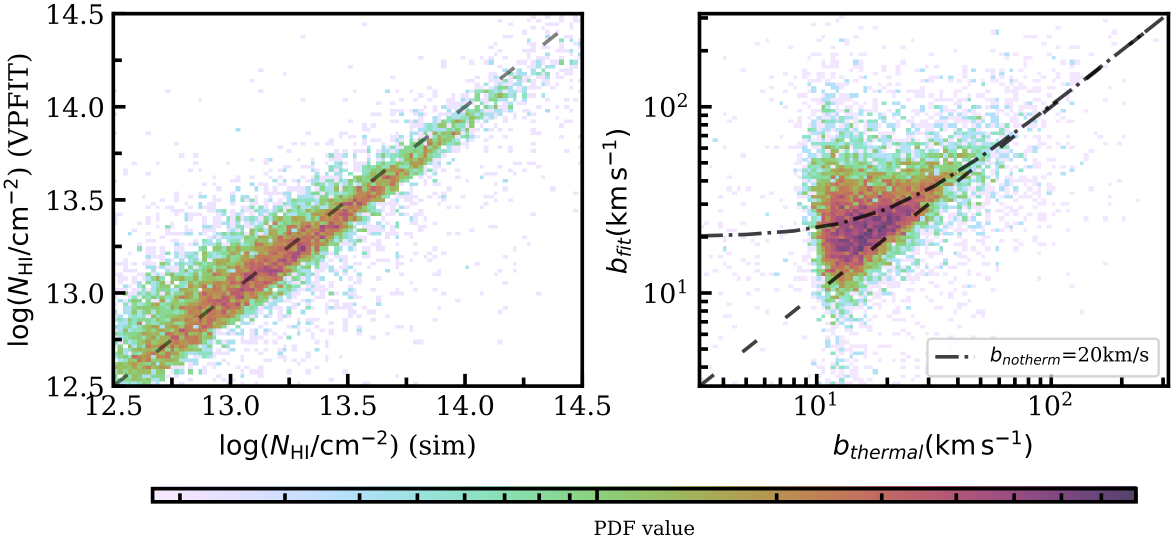

Figure 5.10: The observed variables $N_{\mathrm{HI,fit}}$, b_{fit} fitted by VPFIT compared with the physical quantities $N_{\mathrm{HI,sim}}$ and b_{thermal} of the simulated absorbers identified in the simulation skewers, where the $N_{\mathrm{HI,sim}}$ is calculated by integrating the n_{HI} of the absorbers along the LOS, and the b_{thermal} is computed by assuming the broadening of the Lyα lines are pure thermal. *left:* N_{HIfit} vs $N_{\mathrm{HI,sim}}$. *right:* b_{fit} vs b_{thermal}. The dash-dot line in the right panel represents the b-parameter resulting from the combination of the thermal component b_{thermal} and turbulence in the IGM with $b_{\mathrm{notherm}} = 20$ km/s.

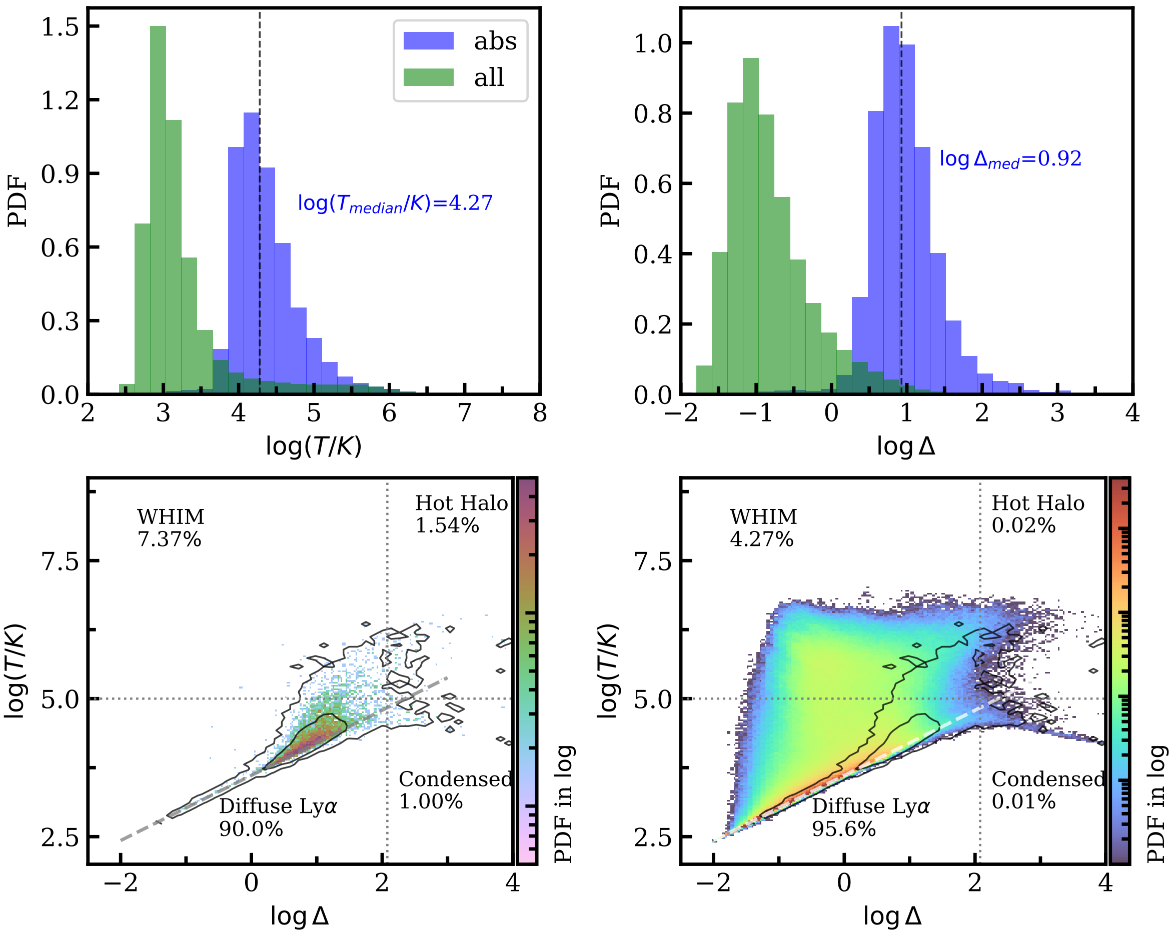

Figure 5.11: Distributions of Δ and T of the simulated Lyα absorbers in the Nyx simulation, compared with the full simulation. The ensemble consists of 10510 absorbers, all obtained from the 1000 spectra discussed in § 5.3.1. The top panels show the 1D distributions of T(right) and Δ(left) for the whole simulation (green) compared with simulated Lyα absorbers (blue). The medians of the T and Δ for the simulated absorbers are indicated by dashed black lines. The bottom panels plot the 2d T-Δ distributions for the Lyα absorbers(left) and for the whole simulation(right), while the contours for 1,3-σ (68% and 99.7%) of the T-Δ distribution of the absorbers are shown in both panels. The volume-weighted gas phases for absorbers and the whole simulation are given in the left panel and the right panel respectively. The best-fit power-law T-Δ relationships are given in the bottom panels as comparisons.

We summarize the (Δ, T) for the ensemble of simulated Lyα absorbers identified in the Nyx simulation in Fig. 5.11. Considering that we have established one-to-one correspondence between the simulated absorbers and observed (mock) absorption lines, we employ a consistent filter to both sets, which selects Lyα lines with, $12.5 \leq \log(N_{\rm HI}/{\rm cm}^{-2}) \leq 14.5$ and $0.5 \leq \log(b/{\rm km\ s}^{-1}) \leq 2.5$ (see § 2.1.3). In the upper panels, we plot the volume-weighted 1D marginal distributions of Δ and T for all simulation grid cells, juxtaposed with the 1D distributions of Δ and T for the simulated Lyα absorbers, showing that the simulated Lyα absorbers, in general, have higher temperature and over-density, compared with the full simulation. The peaks of the Δ and T distributions of the simulated Lyα absorber highlight the specific range of Δ and T to which the Lyα forest is sensitive at $z = 0.1$. More specifically, the Lyα forest is most sensitive to the IGM characterized by $\log\Delta = 0.92$ and $T = 10^{4.27}$K. It is worth mentioning that, the Lyα optical depth $\tau_{\rm Ly\alpha}$ is dependent on $\Gamma_{\rm HI}$. Consequently, the regions to which the Lyα forest is sensitive also depend on $\Gamma_{\rm HI}$. This point will be fully discussed later in § 5.3.3.

The bottom left panel of Fig. 5.11 shows the (volume-weighted) T-Δ distributions for simulated Lyα absorbers (left), and all grid cells in the simulation (right), while the volume-weighted gas fractions[3] are given in annotations for simulated absorbers and the whole simulation in the left and the right panel respectively. The black contours in both panels illustrate the 1 and 3 σ (68% and 99.7%) contours for the T-Δ distribution of the simulated Lyα absorbers. The T-Δ distribution of the simulated Lyα absorbers appears to be scattered at low-z, extending into the WHIM phase, due to the pervasive effects of shock heating. As per the gas phase fractions of the Lyα absorbers shown in the

---

[3] As previously mentioned, for each simulated Lyα absorber, we use the T and Δ at its $n_{\rm HI}$ peak, which dominates the Lyα absorption. To this end, when calculating the volume-weighted gas fractions, we do not take the physical size into account, but instead, only consider the one simulation cell where the $n_{\rm HI}$ reaches its maximum. This is reasonable since typical $n_{\rm HI}$ peaks seen in this study are so narrow that most of the $N_{\rm HI}$ comes from the peak cell. As a result, the so-called volume-weighted gas fractions for simulated absorbers are effectively unweighted. This approach is used for all gas fractions related to simulated Lyα absorbers throughout this paper.

bottom left panel, approximately 7% of the absorbers originate from the WHIM phase, suggesting that the low-z Lyα forest does probe the WHIM (see the right panel of Fig. 5.9 as an example), although its sensitivity is notably limited given the small fraction of lines arising from this phase. Such a result aligns with Tepper-García et al. (2012) regarding the detectability of the Broad Lyα Absorbers (BLAs) at low-z under realistic conditions.

5.3.2 Simulated Lyα Absorbers in IllustrisTNG and Illustirs

To further study the effects of the feedback mechanisms on the Lyα forest at $z = 0.1$, we identify the simulated Lyα absorbers in both the IllustrisTNG and Illustris simulations, and pair them to the corresponding absorption lines present in the mock spectra, following the method outlined in section 5.3.1. For each simulation,we carry out the identification process across 1000 mock spectra and summarize the physical properties of the simulated absorbers. It is worth mentioning that here the IllustrisTNG and Illustris simulations are tuned to have identical $\mathrm{d}N/\mathrm{d}z$, which requires different Γ_{HI} values (see § 5.1.3).

We plot the marginalized 1D distributions of the Δ and T for both IllustrisTNG (top) and Illustris (bottom) in Fig. 5.12. The plots show that the overall distributions of T and Δ for the two simulations are evidently different due to their different feedback recipes. For instance, the extreme feedback in Illustris simulation results in much more WHIM compared with IllustrisTNG, causing a secondary peak in its T distribution. However, the distributions of T and Δ for the absorbers in both simulations are comparable, with $\log(T_{\mathrm{med}}/\mathrm{K}) =4.33$, $\log\Delta_{\mathrm{med}} =0.97$ for IllustrisTNG, and $\log(T_{\mathrm{med}}/\mathrm{K}) = 4.16$, $\log\Delta_{\mathrm{med}} = 0.72$ for Illustris. Moreover, we discover that the differences in $\log(T_{\mathrm{med}}/\mathrm{K})$ and $\log\Delta_{\mathrm{med}}$ for the three simulations are actually caused by the different Γ_{HI} values used for the three simulations. The relevant discussion is presented in § 5.3.3.

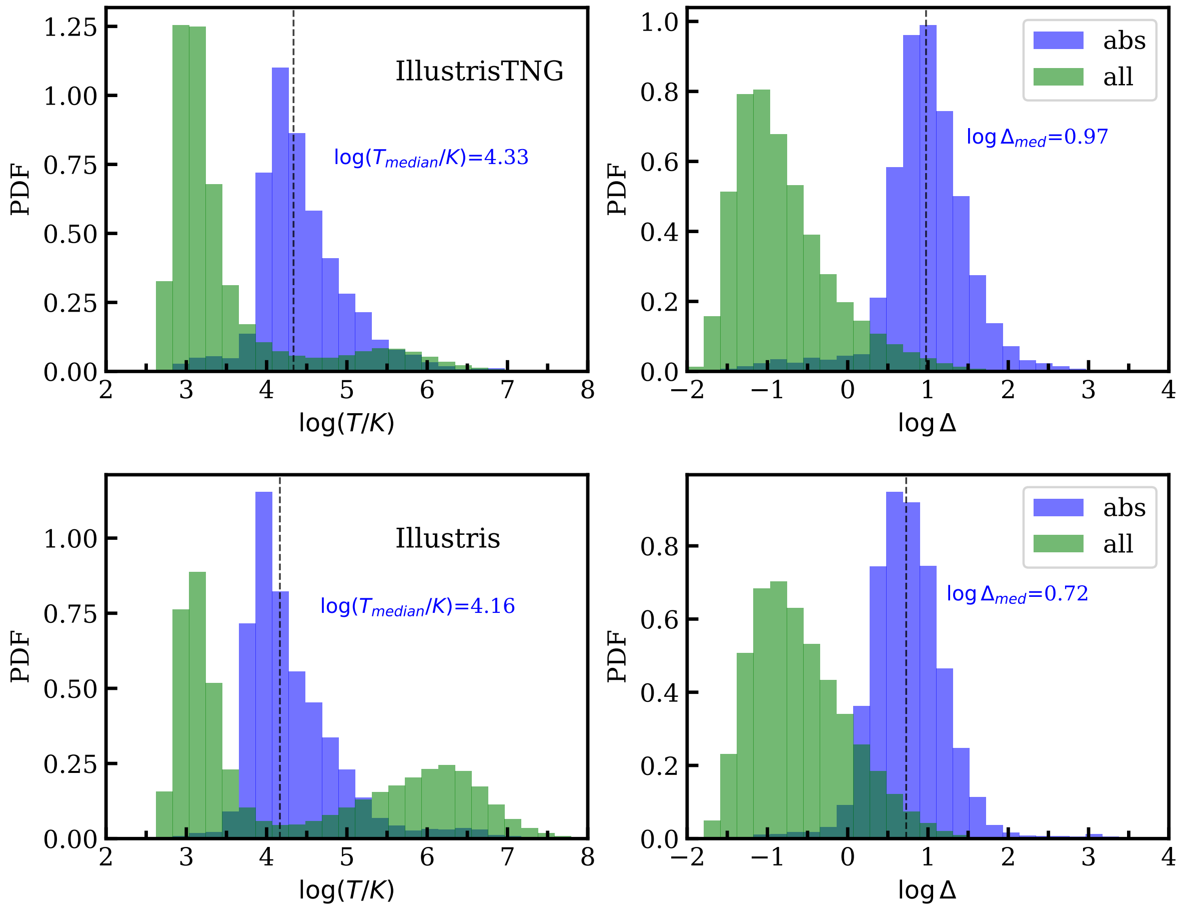

Figure 5.12: Marginalized 1D Δ and T distributions of the simulated Lyα absorbers in the IllustrisTNG (top) and Illustris (bottom) simulation. The medians of the T and Δ for the simulated absorbers are indicated by dashed black lines. The overall Δ and T distributions of the full simulations are plotted as comparisons. The two simulations share the same dN/dz, which is the same value observed in the D16 dataset.

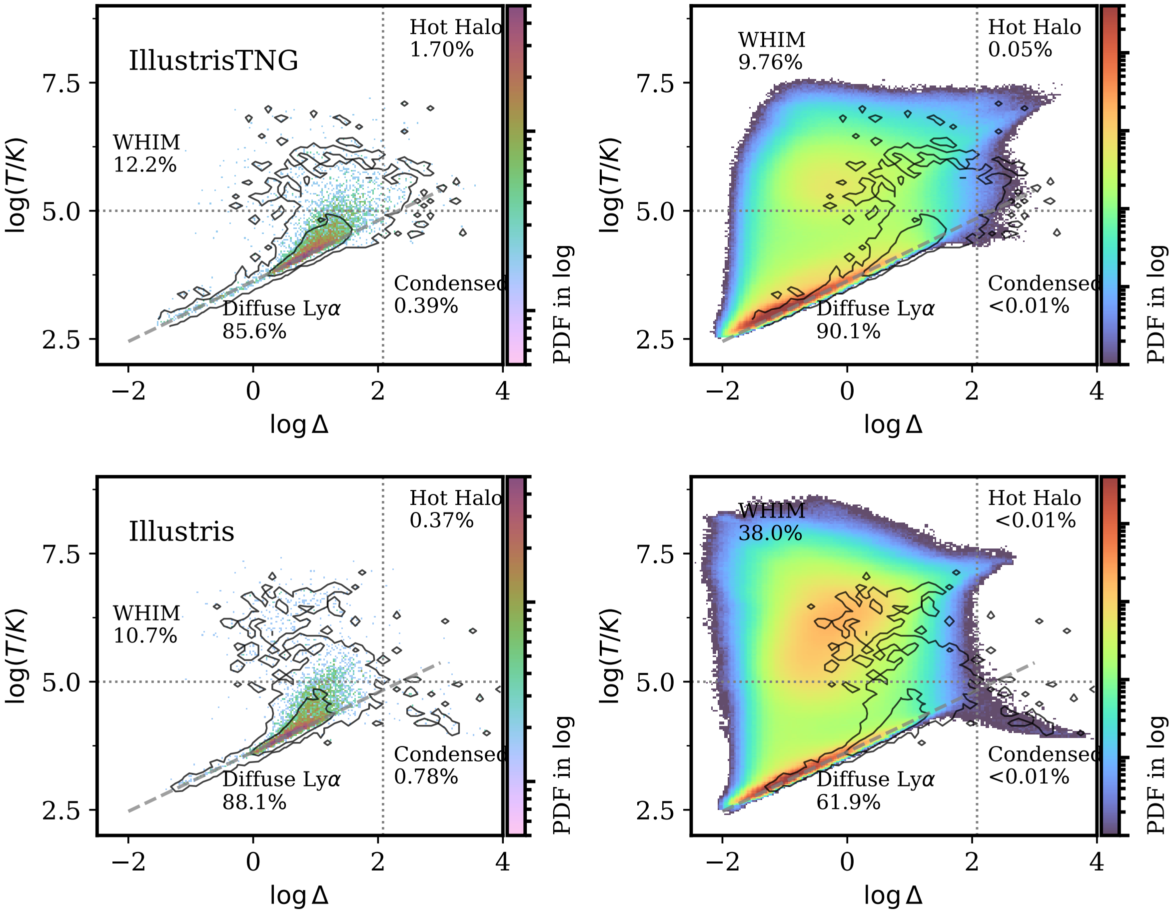

Figure 5.13: T-Δ distributions of the Lyα absorbers in the IllustrisTNG (top) and Illustirs (bottom) simulations, compared with the T-Δ distributions of full simulations (right). The contours for 1,3-σ (68% and 99.7%) of the T-Δ distribution of the absorbers are shown in both panels. The volume-weighted gas phases for absorbers and the whole simulation are given in the left panel and the right panel respectively. The best-fit power-law T-Δ relationships are given as comparisons. The two simulations are dN/dz matched.

The (volume-weighted) 2D T-Δ distributions for simulated Lyα absorbers in both IllustrisTNG (top) and Illustris (bottom) simulations are shown in the left column of Fig. 5.13, whereas the (volume-weighted) 2D T-Δ distributions for the whole simulations are given in the right column as comparisons. The volume-weighted gas fractions are given in the annotation for simulated absorbers and full simulation in the left and the right panels respectively. For the simulated Lyα absorbers, 12.2% (10.7%) of the Lyα absorbers arise from the WHIM for IllustrisTNG (Illustris), while the value for Nyx simulation is approximately 7%. The 1 and 3-σ (68% and 99.7%) contours for the T-Δ distributions for the simulated Lyα absorbers are also given in the Figure, showing that their T-Δ distributions are more scattered compared with these in Nyx simulation, especially for the WHIM phase absorbers. These differences are caused by stronger shock heating in IllustrisTNG and Illustris simulations compared with Nyx simulation, caused by their feedback mechanisms. However, while the (volume-weighted) WHIM fractions for the two simulations are remarkably different, 9.8% for IllustrisTNG and 38.0% for Illustris, the WHIM fractions for the Lyα absorbers are similar, both around 11%. Furthermore, in § 5.3.3 we demonstrate that the small difference in WHIM fractions for simulated absorbers in the two simulations actually arises from the different Γ_{HI} values used in the two simulations. Such a fact implies that the low-z Lyα forest does not probe the WHIM effectively under realistic conditions, which is consistent with the conclusion drawn by Khaire et al. (2023a).

5.3.3 Simulations under the same Γ_{HI}

Considering that the calculation of the Lyα optical depth $\tau_{\mathrm{Ly}\alpha}$ involves Γ_{HI}, and given that the observed absorption feature (i.e., the Lyα forest) consistently probes regions with $\tau_{\mathrm{Ly}\alpha} \sim 1$, it follows that the T and Δ of these regions probed by the Lyα forest, are influenced by the $\Gamma_{\rm HI}$ values. Such an argument can be qualitatively demonstrated by the fluctuating Gunn-Peterson approximation (FGPA, see Weinberg et al. 1997)

$$\tau_{\rm Ly\alpha} \propto x_{\rm HI} n_{\rm H} \propto \frac{n_{\rm H}^2 T^{-0.7}}{\Gamma_{\rm HI}} \propto \frac{\Delta^{2.7-\gamma}}{\Gamma_{\rm HI}} \propto \frac{T^{2/(\gamma-1)-0.7}}{\Gamma_{\rm HI}}, \tag{5.2}$$

where the $\tau_{\rm Ly\alpha}$ denotes the Lyα optical depth and the $n_{\rm H}$ is the hydrogen number density. Since the Lyα forest always probes the region with $\tau_{\rm Ly\alpha} \sim 1$, the last two terms in eq. 5.2 suggest that the $\Gamma_{\rm HI}$ is in positive correlation with Δ and T respectively, given the $\gamma \sim 1.6$ at $z = 0.1$.

In our analysis, the three simulations are tuned to match $\mathrm{d}N/\mathrm{d}z$. However, due to the degeneracy between $\Gamma_{\rm HI}$ and feedback mechanisms, each simulation ends up with a distinct $\Gamma_{\rm HI}$ value (refer to § 5.1.3). As a result, the T and Δ distributions of the simulated Lyα absorbers in these simulations are influenced not just by the feedback but also by the varying $\Gamma_{\rm HI}$ values. To isolate and examine solely the impact of feedback, we post-process the IllustrisTNG and Illustris simulations to align with the $\Gamma_{\rm HI}$ value used in Nyx, set at $\log(\Gamma_{\rm HI}/s^{-1}) = -13.093$. With this consistent $\Gamma_{\rm HI}$ across the three simulations, we re-perform the analysis from the prior section and summarize the results below. It is worth mentioning that the overall T-Δ distributions of simulations are determined by the cooling and heating processes during their evolution and are not altered by the post-processing of the $\Gamma_{\rm HI}$.

We plot the marginalized Δ and T distributions and their median values for Lyα absorbers in Nyx, IllustrisTNG, and Illustris simulations with the same $\mathrm{d}N/\mathrm{d}z$ in Fig. 5.14. Interestingly, for simulations with the same $\mathrm{d}N/\mathrm{d}z$, the T and Δ for absorbers are correlated with its $\Gamma_{\rm HI}$. More specifically, with $\Gamma_{\rm HI,Illustris} < \Gamma_{\rm HI,Nyx} < \Gamma_{\rm HI,IllustrisTNG}$ (see Fig. 5.3), we obtain $T_{\rm med,Illustris} < T_{\rm med,Nyx} < T_{\rm med,IllustrisTNG}$ and $\Delta_{\rm med,Illustris} < \Delta_{\rm med,Nyx} < \Delta_{\rm med,IllustrisTNG}$.

We plot the marginalized Δ and T distributions and their median values for Lyα absorbers in the three simulations under the same $\Gamma_{\rm HI}$ in Fig. 5.15. Under the same $\Gamma_{\rm HI}$, the T and Δ distributions for simulated absorbers in all three simulations become almost identical, having nearly the same median values for T and Δ respectively. Such a result suggests that while feedback evidently affects the overall T-Δ distributions of the low-z IGM (see Fig. 5.2), their impacts on the physical properties of the low-z Lyα forest (i.e., the T and Δ distributions) are not distinguishable under realistic conditions.

In Fig. 5.16, we plot the 2D T-Δ distributions of the simulated Lyα absorbers in IllustrisTNG (left) and Illustris (right), under the same $\Gamma_{\rm HI}$. While the overall IGM T-Δ distributions for the two simulations are evidently different (see Fig. 5.2), the T-Δ distributions of the simulated Lyα absorbers in these two simulations are similar, and the gas phase fractions for absorbers in both simulations are almost identical, suggesting that the small difference in the WHIM fractions of the simulated absorbers shown in Fig. 5.13 are caused by different $\Gamma_{\rm HI}$ values. Such results indicate that the $\Gamma_{\rm HI}$ has a much stronger impact on the Lyα forest compared with the feedback mechanisms implemented in IllustrisTNG and Illustris simulations.

5.4 Summary and Discussion

In this paper, we explore the effects of the WHIM, which causes significant dispersion in the IGM T-Δ distribution, on the low-z Lyα forest and the IGM thermal state $[T_0, \gamma]$ measured from it. We first evaluate the effectiveness of $[T_0, \gamma]$ as IGM parameters under the inference framework presented in Hu22, and compare its performance with the photoheating labels [A,B]. We discover that the thermal state [T_0, γ] still parameterizes the IGM effectively in spite of the dispersion in the IGM T-Δ distribution. We further apply the inference method to IllustrisTNG and Illustris simulations which implement

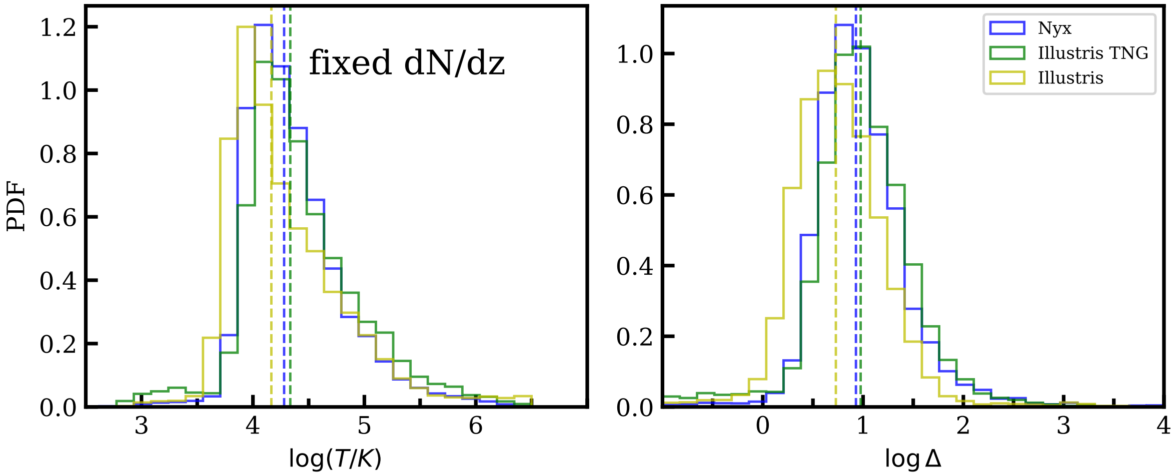

Figure 5.14: Marginalized 1d distribution of T (left), and Δ (right) of the simulated Lyα absorbers identified in Nyx (blue), IllustrisTNG (green), and Illustris (yellow) simulations. The medians of $\log T$ and $\log \Delta$ are indicated by vertical dashed lines. The three simulations are turned to have identical $\mathrm{d}N/\mathrm{d}z$.

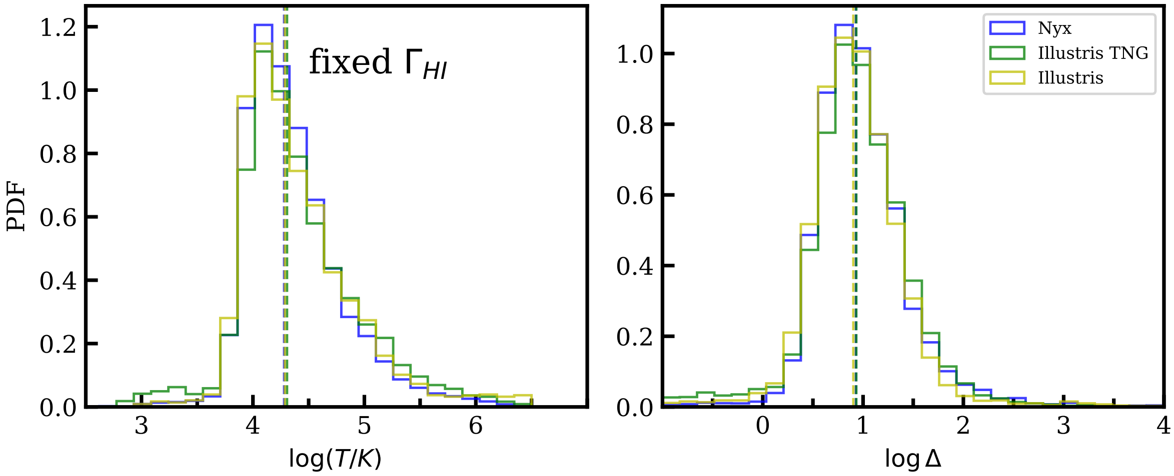

Figure 5.15: Marginalized 1d distribution of T (left), and Δ (right) of the simulated Lyα absorbers identified in Nyx (blue), IllustrisTNG (green), and Illustris (yellow) simulations. The medians of $\log T$ and $\log \Delta$ are indicated by vertical dashed lines. The three simulations used here are post-processed to have the same UV background photoionization rate, with Γ_{HI} =-13.093.

Figure 5.16: T-Δ distributions of the simulated Lyα absorbers in the IllustrisTNG (left) and Illustirs (right) simulations under the same Γ_{HI}. The contours for 1,3-σ (68% and 99.7%) of the T-Δ distribution of the absorbers are shown in both panels. The volume-weighted gas phase fractions for absorbers in both simulations are given respectively.

different variants of feedback, potentially making them better approximations to the real Universe. We discover that the [T_0,γ] of these two simulations can be recovered using the inference method within reasonable offsets. Considering the inference results and the huge difference across the three simulations in the IGM WHIM fractions, we conclude that the Lyα forest does not probe the WHIM effectively under realistic conditions, and the IGM thermal state [T_0,γ] is not affected by the shock heating caused by AGN feedback and other astrophysical processes significantly at $z = 0.1$. To further confirm our conclusion, we identified the Lyα absorbers in all three simulations at $z = 0.1$, and pair them to the corresponding absorption lines identified in the mock spectra. The physical properties of the simulated Lyα absorbers support our conclusion that the observable Lyα forest are not affected by the substantial WHIM in the low-z, and the thermal state $[T_0, \gamma]$ measured from the Lyα forest remains solid. In this section, we summarize our paper and present our discussion as follows.

- We compare the performance of $[T_0, \gamma]$ as neural network training labels against the photoheating labels [A,B], i.e. the photoheating rate rescaling factors used to generate the Nyx simulation suite with various thermal histories. Given that the $[A, B]$ parameters were actually used to generate the simulation outputs, one might expect that they would serve as a better set of labels than $[T_0, \gamma]$. However, the inference results show the efficacy of these two sets of labels are comparable, suggesting that the $[T_0,\ \gamma]$, which parameterize the power law T-Δ relationship, still effectively characterize the Lyα observables at low-z, notwithstanding the dispersion in the T-Δ distribution induced by shock heating at low-z.

- We explored the degree to which the presence of feedback can influence or bias the inference of the IGM thermal state parameters from the b-N_{HI} distribution. In the context of our inference framework, this question becomes: what would happen if we used a simulation grid without feedback to infer the thermal state of a Universe that has strong feedback? Would the feedback lead to biased inference? To address these questions, we apply our inference procedure trained on Nyx simulations without feedback to mock datasets from the IllustrisTNG and Illustris simulations which include feedback, whereby the latter serve as potential proxies for the real Universe. We find that the $[T_0,\gamma]$ of IllustrisTNG and Illustris can be recovered within small offset, where $\Delta \log(T_0/\mathrm{K}) = -0.022$ dex, $\Delta\gamma = 0.064$ for IllustrisTNG and $\Delta \log(T_0/\mathrm{K}) = 0.047$ dex, $\Delta\gamma = -0.094$ for Illustris. These offsets are smaller than the typical precision afforded by a realistic dataset, i.e., $\Delta \log T_0 \lesssim 0.5\ \sigma_{\log T_0}$, and $\Delta\gamma \lesssim \sigma_\gamma$.

- We developed a method to identify regions in the simulation responsible for the Lyα absorption lines identified via Voigt-profile fitting, allowing us to determine their temperature T and overdensity Δ from the simulation skewers. For the Nyx simulations, the simulated Lyα absorbers have a median density $\log \Delta_{\mathrm{median}} = 0.92$, a median temperature $T_{\mathrm{median}} = 10^{4.27}$K, and about 7% of the simulated Lyα absorbers have $T > 10^5$, making them outliers from the power-law T-Δ relationship. This low fraction is consistent with the previous study of Tepper-García et al. (2012) on the low-z BLAs.

- As pointed out in previous work (Bolton et al. 2022a, Khaire et al. 2023a, Tillman et al. 2023a), the Lyα forest is affected by the UV background, which impacts the $\mathrm{d}N/\mathrm{d}z$. Nevertheless, we observe that the temperature and overdensity of the region probed by the low-z Lyα forest are also affected by the UV background photoionization rate Γ_{HI} used in the simulation. For $\mathrm{d}N/\mathrm{d}z$ matched simulations, the T and Δ of the simulated Lyα absorbers are correlated with its Γ_{HI} respectively. Specifically, the Lyα forest probes regions with higher T and Δ given a higher Γ_{HI}.

This is because for Lyα absorbers with $\tau_{\mathrm{Ly}\alpha}$ $\sim$1, the fluctuating Gunn-Peterson approximation implies that $\Gamma_{\mathrm{HI}} \propto \Delta^{2.7-\gamma} \propto T^{2/(\gamma-1)-0.7}$, where $\gamma \sim 1.6$.

- We post-processing the three simulations to share the same Γ_{HI}, allowing us to explore the effects of different mechanisms. Under the same Γ_{HI}, the T and Δ distributions of the simulated Lyα absorbers across all three simulations become almost indistinguishable, converging to nearly identical median values, while the overall IGM T-Δ distributions remain different among the simulations, due to their distinct feedback mechanisms. For the WHIM fractions, the volume-weighted WHIM fractions for IllustrisTNG and Illustris stand at 9.8% and 38.0%, respectively, but the WHIM fractions for the simulated Lyα absorbers in both simulations are nearly identical, averaging around 11.6%. This suggests that while feedback significantly alters the low-z IGM T-Δ distribution, especially the WHIM phase gas, their impacts on the low-z Lyα forest is indistinguishable under realistic conditions. Such a conclusion aligns with the results derived from various statistics of the low-z Lyα forest by Khaire et al. (2023a).

Chapter 6

Conclusion and Outlook

In this thesis, I measure the thermal and ionization state of the IGM across seven redshift bins from $0 < z < 1.5$, spanning approximately 10 billion years in cosmic history. I utilize a novel inference method developed based on neural networks to jointly measure the thermal and ionization state of the IGM. This approach leverages the 2D b-N_{HI} distribution and Lyα line abundance $\mathrm{d}N/\mathrm{d}z$ to break parameter degeneracy and achieve high-precision measurements even with limited observational data.

I apply this method to 94 archival HST COS and STIS quasar spectra distributed across these seven redshift bins, providing a detailed evolutionary history of the low-z IGM thermal and ionization state. The results indicate that the IGM may be significantly hotter at low-z than previously expected, potentially reaching an isothermal state with mean density temperatures T_0 around $\sim 30,000$K and a power-law index $\gamma \sim 1.0$ at $z = 0.1$. In addition, the results also suggest that the unexpected thermal state of the IGM potentially emerges at $z \sim 1$.

Additionally, while the inferred Γ_{HI} align with theoretical models at $z \sim 1$, the Γ_{HI} measured at $z < 0.5$ lies substantially below the values given by the theoretical model, challenging the current understanding of the UVB models at $z < 0.5$.

As an alternative explanation of the observed higher-than-expected b-parameter, I model the micro-turbulence in the low-z IGM using a Gaussian velocity dispersion. I demonstrate that this approach increases the observed b-parameter. The inference results suggest that it would need velocity dispersion with $\sigma_v \sim 14$, 18, 11, and 10 km/s in small scale at $z = 0.1$, 0.2, 0.3 and 0.4. Such increasing turbulence towards low-z further suggests the discrepancy observed in the b-parameters, whether caused by micro-turbulence or not, gets stronger towards low-z.

Furthermore, I have rigorously verified that the observed discrepancies in the b-parameter are not due to feedback from AGN activities or galaxy formation processes, including extreme feedback models. This strengthens the argument that the observed discrepancies in the b-parameter are likely caused by unrecognized physical mechanisms, potentially involving additional heating or turbulence.

Looking forward, there are approximately 30 COS spectra and 20 STIS spectra not analyzed in this thesis due to the absence of metal identification. Once the metals are identified, these spectra can be readily used to measure the IGM thermal and ionization state. Meanwhile, we are preparing to submit a new HST proposal to acquire additional Lyα forest spectra at $z \sim 1$, a pivotal period in the thermal evolution history of the IGM. According to this thesis, this epoch is crucial as it marks the onset of the observed discrepancies in the b-parameter, which could eventaully lead to the discovery of new physics in the low-z IGM.

Furthermore, we plan to extend our inference methodology to simulations which incorporate more sophisticated and diverse feedback and heating mechanisms, such as the EAGLE simulation (Schaye et al. 2015) and CAMELS (Villaescusa-Navarro et al. 2021) projects. This expansion will enhance our understanding of how various feedback and heating processes influence the Lyα forest and the IGM, which may potentially resolve the discrepancies between simulations and observations regarding the b-parameter.

Appendix A

Appendix

A.1 Inference Without Absorber Density dN/dz

In this section we provide more details about the inference without using the absorber density. In such a case, the likelihood function would simply be the first term of Eq.(2.7), i.e.

$$\ln \mathcal{L} = \sum_{i=1}^{n} \ln P(b_i, N_{\mathrm{HI},i}). \tag{A.1}$$

Such likelihood function is evaluated based on our b-N_{HI} distribution emulator solely. To make better comparison, we use the same mock dataset and training dataset as used in §2.2.4. The MCMC posterior is given in Fig.A.1, where we obtain $\log(T_0\mathrm{K}) = 3.709^{+0.058}_{-0.073}$, $\gamma = 1.550^{+0.066}_{-0.068}$ and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1}) = 13.401^{+0.097}_{-0.090}$ from the marginalized distributions, whereas the true parameters are: $\log(T_0\mathrm{K}) = 3.643$, $\gamma = 1.591$ and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1}) = 13.458$. In comparison, the posterior obtained using Eq.(2.5), which takes into account the dN/dz, is shown in blue in Fig.A.1. As we show here, the two inference results are coherent, but our modified inference algorithm (green posteriors) perform better. By implementing the dN/dz feature, our modified inference algorithm provides more accurate results in both T_0 and Γ_{HI}, and reduce the uncertain in Γ_{HI} significantly. Furthermore, the inference without absorber density dose not pass the inference where the true model falls in the 1-σ (68%) interval for about 50% of the time.

In short, by employing the absorber density we not only evidently reduce the uncertainty in Γ_{HI} but also increase the accuracy in other parameters since the modification adds more information to the Bayesian analysis by matching the absorber density.

A.2 Inference test likelihood calculation

To calculate contours of cumulative probability distribution with high dimensionality is challenging in computation power. In our case, the parameter grid size is 100^3 and we have to compute the probability density function $P(b\ , N_{\mathrm{HI}} \mid T_0, \gamma, \Gamma_{\mathrm{HI}})$ many hundreds

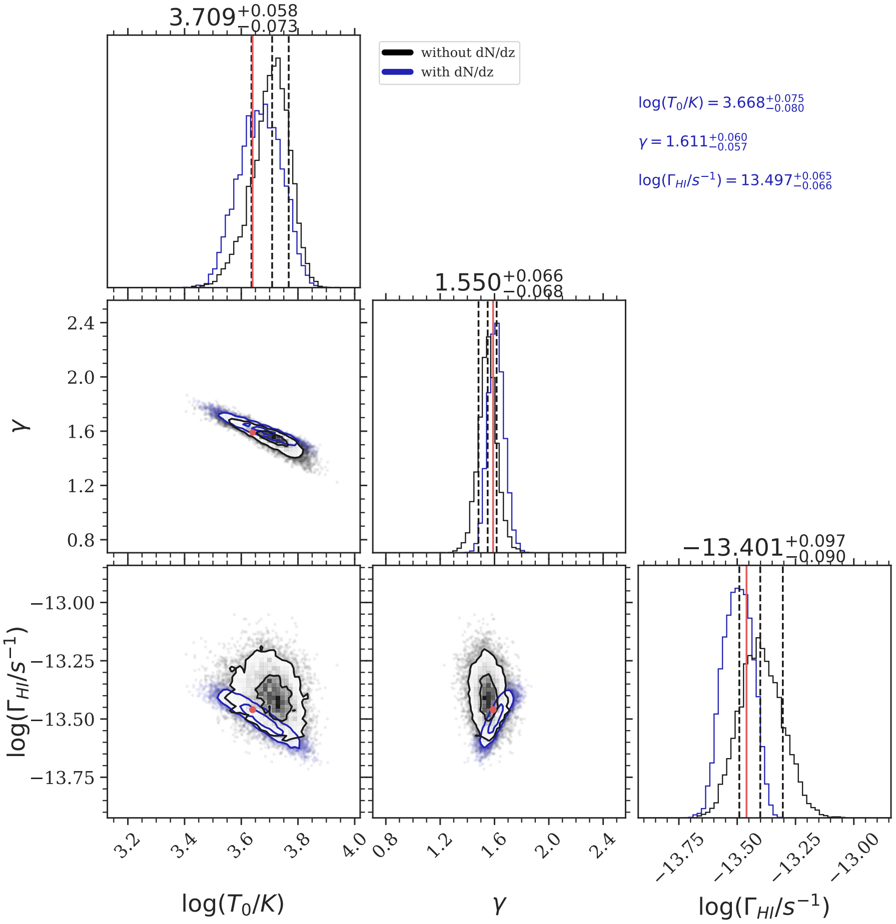

Figure A.1: MCMC posterior (black) for the Nyx model discussed in §2.2.4 based on the likelihood function without the absorber density Eq.(A.1). Projections of the true model is shown as red dot. Inner(outer) contours represents the projected 2D 1(2)-sigma interval. The parameters of true model are indicated by red lines in the marginal distributions, while the dashed black lines indicates the 16, 50, and 84 percentile values of the posterior. The true parameters are: $\log(T_0/\mathrm{K}) = 3.643$, $\gamma = 1.591$ and $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$ = -13.458. In comparison, the posterior obtained using Eq.(2.5), which takes into account the $\mathrm{d}N/\mathrm{d}z$, is shown in blue, while the medians of the posterior are shown in blue on the top right.

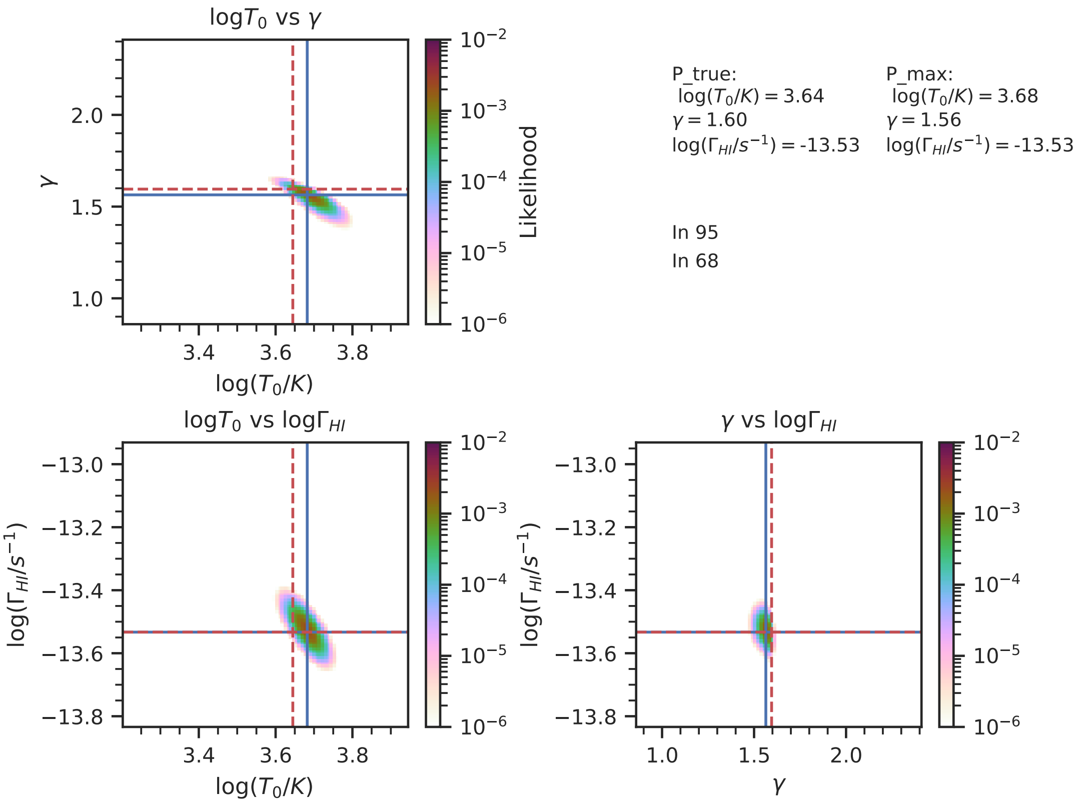

Figure A.2: Example of the distribution of the likelihood function sliced at the location of the true parameters (P_{true}, indicated by red dashed lines). The parameters corresponding to the maximum likelihood model $P_{\max}$ are indicated by blue solid lines. Values of both P_{true} and $P_{\max}$ are given in the up right. Calculation implies that the true parameters are in both the effective 1σ 68%) and 2σ (95%) intervals.

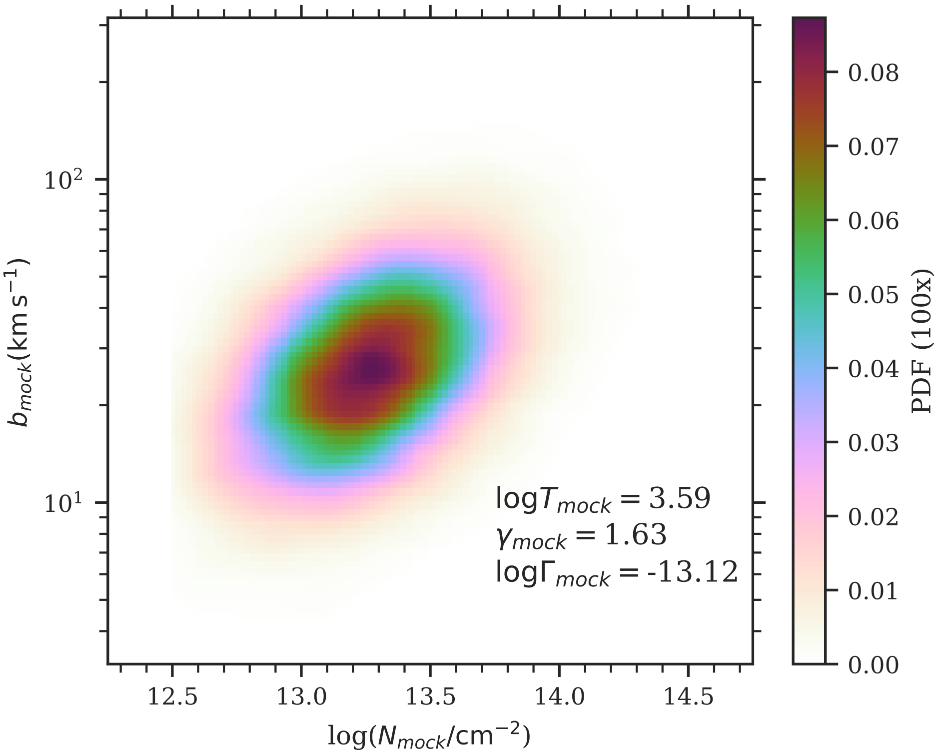

Figure A.3: The KDE based PDF of b-N_{HI} distribution of one of the toy models which is a 2D Gaussian distribution parameterized by T_{mock}, γ_{mock} and Γ_{mock} in analogy with thermal parameters T_0, γ , Γ_{HI}. The parameters of the toy model is shown in the right bottom corner of the plot. For illustration purposes, values of pdf are multiplied by 100 in the color bar.

times (i.e. the number of lines in the data set) to evaluate the likelihood function on a single point on the parameter grid (see Eq. 2.5). However, due to the structure of the b-N_{HI} PDF calculated by our DELFI emulator, we are able to save time by computing the likelihood function on the whole grid simultaneously, with help of vector operations implemented in `python`, though such treatment requires reconstruction of the likelihood function and needs extra amounts of memory. In comparison, our code is much faster than the MCMC prescription which would require a very long chain to interpolate the likelihood function on the whole grid to achieve the same precision. An example of the distribution of the likelihood function is shown in Fig.A.2.

A.3 Toy model

To verify the performance of our emulators in a clean environment, we build a toy model with a mock data set which roughly simulates the behavior of our real model. Here the toy b-N_{HI} distributions consist of 2D Gaussian distributions parameterized by

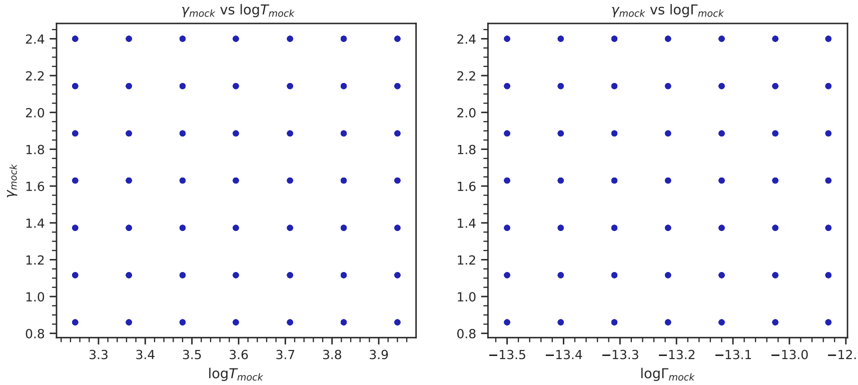

Figure A.4: The thermal grid used in our toy model. The left-hand panel is the γ - T_0 grid and the right-hand panel is γ - $\Gamma_{\rm HI}$ slice showing the 7 $\Gamma_{\rm HI}$ values we have for each point on the 2D γ - T_0 grid.

$T_{\rm mock}$ and $\gamma_{\rm mock}$, $\Gamma_{\rm mock}$ in analogy with thermal parameters T_0, γ and $\Gamma_{\rm HI}$. Here we follow the parameter dependence discussed in §2.2.3, i.e. both $T_{\rm mock}$ and $\gamma_{\rm mock}$ sets the y-axis location of the center of the Gaussian, while $\gamma_{\rm mock}$ also sets the tilted angle of the Gaussian, and the $\Gamma_{\rm mock}$ controls the density of data points for each model, in analogy with the $\Gamma_{\rm HI}$ which determines the absorber density dN/dz. For convenience, we set these mock parameters to be dimensionless. We tune these parameters in a way that the 'b-$N_{\rm HI}$ distribution' of our toy model falls roughly in the same range as the Nyx simulation, and we adopt absorber density emulated by our dN/dz emulator based on our Nyx simulations, so that the mock dN/dz follows the relationship between thermal parameters and absorber density in our Nyx simulation. We in total generate 7x7x7 = 343 (see Fig.A.4) models spanning the thermal grid. An example of the b-$N_{\rm HI}$ distribution of a toy model is shown in Fig.A.3, which is generated based on the Kernel Density Estimation (KDE) of the mock dataset using a smoothing bandwidth $(\sigma_{\log N_{\rm mock}}, \sigma_{\log b_{\rm mock}}) = (0.08, 0.32)$. Such choice of bandwidth is taken from Hiss et al. (2019).

For each toy model with different mock thermal parameters, we first generate a set of 2000 'imaginary' pathlength Δz_i, each of which equals to a randomly chosen observation spectra in Danforth et al. (2016) low-z Lyα dataset (i.e. for each model we generate a set of 2000 Δz_i but without actual spectra). For each 'imaginary' pathlength Δz_i we generate a set of mock 'b-N' pairs (lines), sampling from the b-$N_{\rm HI}$ distribution, while the number of lines N_i follows a Poisson distribution Pois(λ_i) with Poisson rate $\lambda_i = \Delta z_i \times ({\rm d}N/{\rm d}z)_{\rm model}$, where the $({\rm d}N/{\rm d}z)_{\rm model}$ is the absorber density of that model. The total number of lines for the model is thus $N_{\rm tot} = \sum_i^{2000} N_i$. At this point we obtain a training dataset with the same structure as the one described in §5.2.1, which consists of 'b-N' pairs labeled by thermal parameters. We then train the DELFI (b-$N_{\rm HI}$ distribution) and Gaussian (dN/dz) emulators based on the above dataset, and test our whole inference algorithm on the toy model following the prescription given in §2.2.4. An example of the

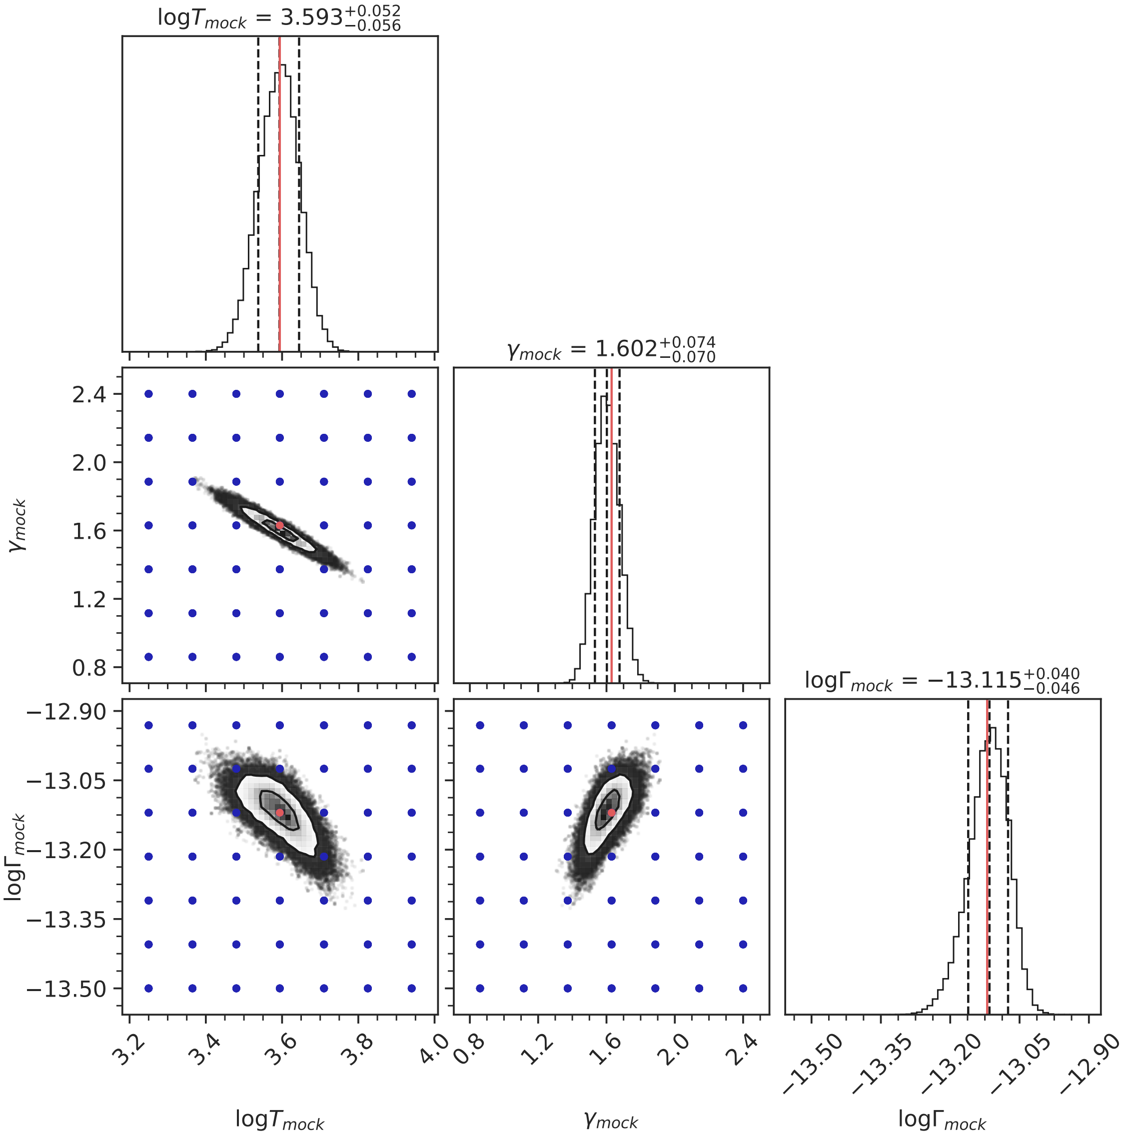

Figure A.5: MCMC posterior for the fit of the b-N_{HI} distribution from one of the toy models (absorbers shown as black points in Fig.A.6) using the likelihood function (Eq. 2.6) from DELFI and our Gaussian emulator (see § 2.2.2). Projections of the thermal grid used for generating models are shown as blue circles. Inner(outer) black contour represents the projected 2D 1(2)-sigma interval. The parameters of true model are indicated by red lines in the corner plot, while the dashed black lines indicates the 16, 50, and 84 percentile value of the posterior. The true parameters are: $\log T_{\mathrm{mock}} = 3.59$, $\gamma_{\mathrm{mock}} = 1.63$ and $\log \Gamma_{\mathrm{mock}} = -13.12$.

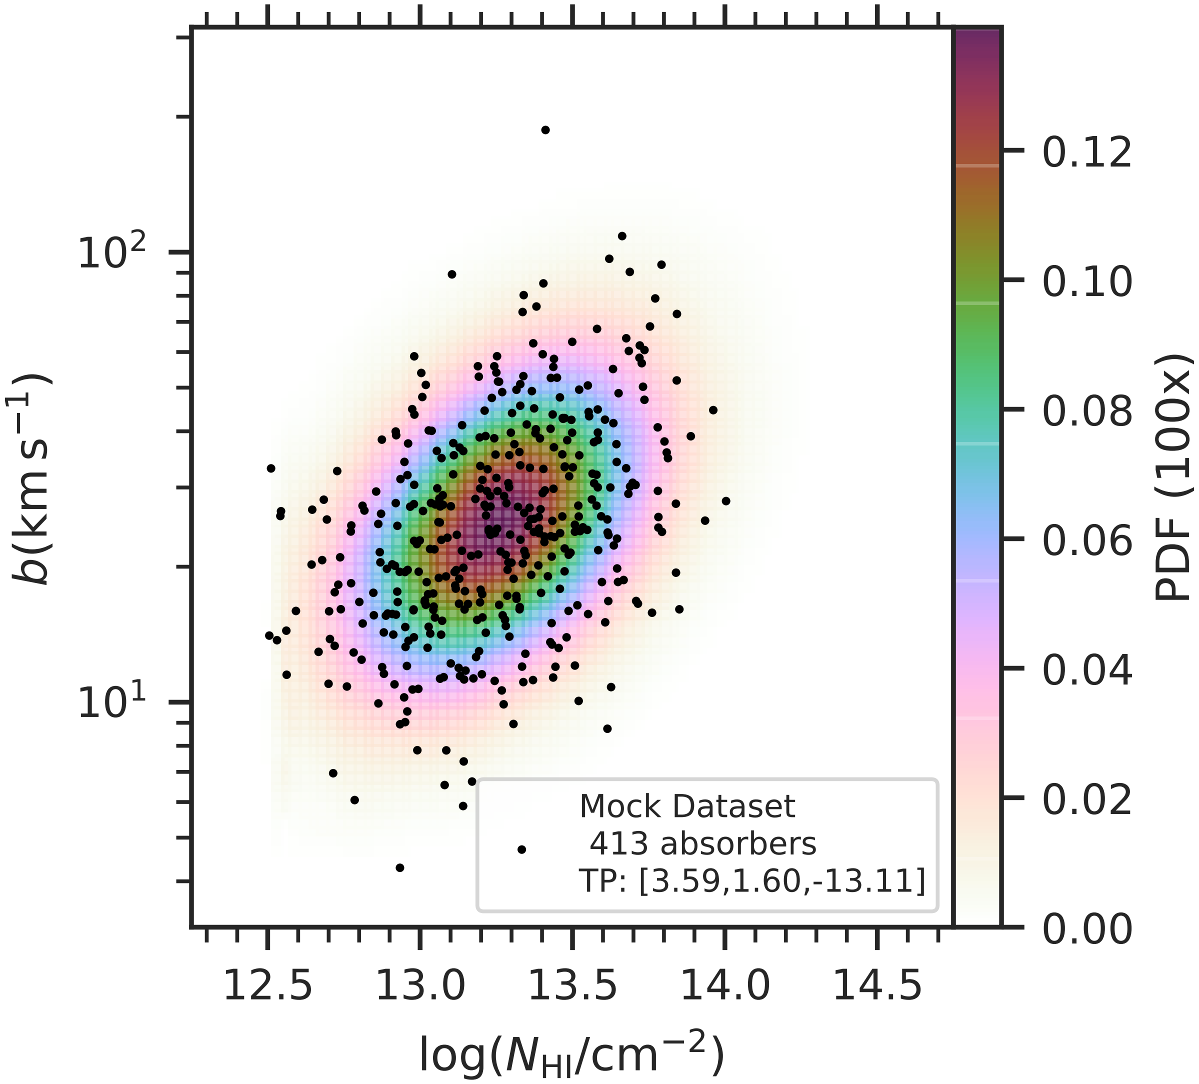

Figure A.6: The 'best fit' model b-N_{HI} distribution for the Gaussian toy model emulated by DELFI. It is emulated based on the best-fit parameters (median values of the marginalized MCMC posterior), which is shown in the right bottom corner of the plot. The true parameters are: $\log(T_{\mathrm{mock}} = 3.59$, $\gamma_{\mathrm{mock}} = 1.63$ and $\log(\Gamma_{\mathrm{mock}} =$ -13.12. For illustration purposes, values of pdf are multiplied by 100 in the color bar.

Table A.1: Table of results of the inference test for the toy model

models	Total	68(%)	95 (%)
3D toy model	240	165 ($68.75 \pm 2.92\%$)	225 ($93.75 \pm 1.67\%$)
2D toy model	300	199 ($66.33 \pm 2.67\%$)	284 ($94.67 \pm 1.33\%$)

inference result is shown below, including the MCMC posteriors (Fig.A.5) and the 'best fit' b-$N_{\rm HI}$ distribution recovered from mock dataset (Fig.A.6). As a comparison, the KDE based PDF of the b-$N_{\rm HI}$ distribution of the model is shown in Fig.A.3.

In the end, we perform inference test on our toy model for both 3D and 2D (without $\Gamma_{\rm HI}$) models to test the robustness of our whole inference pipeline following the method discussed in §2.2.5, and the results are given in table A.1, showing that our inference algorithm passes the inference test perfectly for an idealized model. Moreover, the inference on toy model of b-$N_{\rm HI}$ distribution performs slightly better than on Nyx simulation (see Appendix A.3). The reason could be that the toy model b-$N_{\rm HI}$ distributions are 2D Gaussian distributions that solely depends on the thermal parameters $T_{\rm mock}$ and $\gamma_{\rm mock}$, which is equivalent to say that the b-$N_{\rm HI}$ distribution fully preserved the thermal information of the IGM, however, in the Nyx simulation the b-$N_{\rm HI}$ distributions are affected by the complex astrophysical processes in the diffuse IGM, resulting in the loss of the thermal information.

Combining all results shown above, we conclude that our inference algorithm is able to recover the mock parameters with extraordinary accuracy under idealized condition, and our entire pipeline including b-$N_{\rm HI}$ distribution emulation, $\mathrm{d}N/\mathrm{d}z$ emulation, likelihood function and inference pipeline is robust.

A.4 Convergence test

Lukić et al. (2015) demonstrated that the b parameter of Lyα forest is sensitive to the simulation resolution, and its distribution converges for simulation finer than *L10N512* simulation (i.e., box size $L = 10h^{-1}$ Mpc and $N = 512^3$ dark matter particles and baryon grids which gives the resolution of 20 h^{-1} kpc) while the box size itself does not affect line parameters of the Lyα forest. Whereas above mentioned tests are done at redshift ~ 3, it is worthy to further investigate impact of the boxsize and resolution of the simulation on the Lyα forest at lower redshifts, since the nonlinear evolution at low redshift can affect the Lyα forest.

Here, we perform a convergence test at redshift $z = 0.5$ to to check if our results are independent of the simulation box-size at low redshift. To test the convergence we use two Nyx boxes; *L20N1024* (box-size = 20 h^{-1} Mpc, $N = 1024^3$ dark matter particles and baryon grids i.e resolution of 20 h^{-1} kpc), and *L100N4096* (box-size 100 h^{1-} Mpc and $N = 4096^3$ dark matter particles and baryon grids, resolution of 24 h^{-1} kpc). These two simulation boxes are ran following the same procedures given in section §2.1. In Fig. A.7, we plot the temperature T, overdensity Δ, and velocity along line-of-sight v_{los} of these

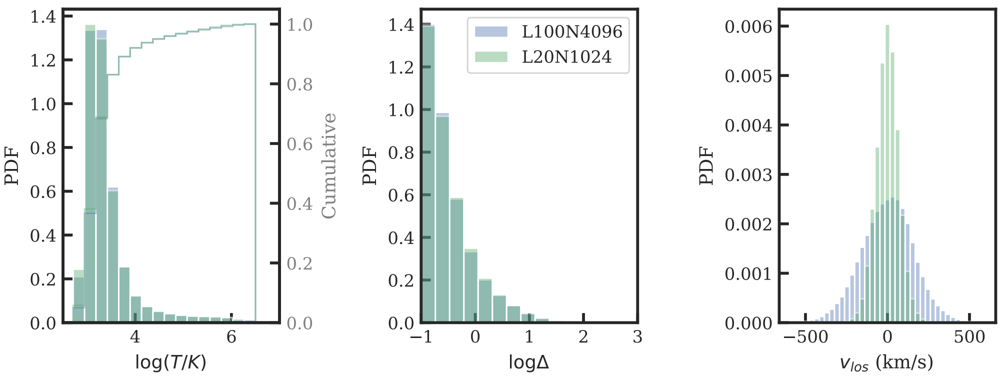

Figure A.7: From left to right, the 1D marginalized distribution of the temperature T, overdensity Δ, and velocity along line-of-sight $v_{\rm los}$. The unfilled histogram in the left most panel shows the CDF of the temperature distribution. The *L100N4096* box are shown in blue, while the *L20N1024* box are shown in green.

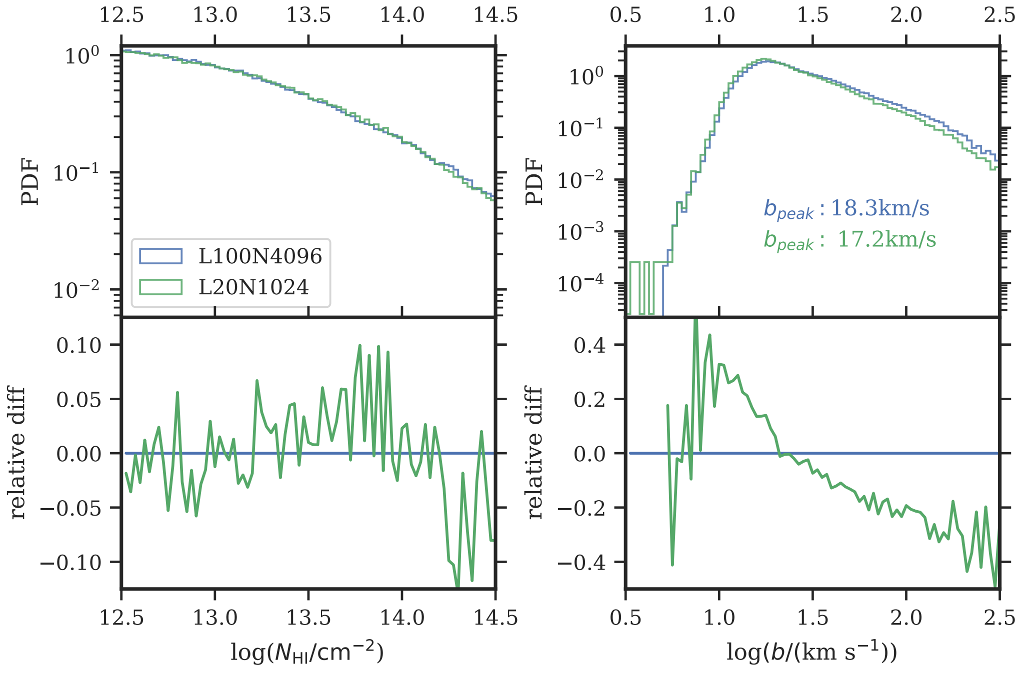

Figure A.8: The 1D marginalized $N_{\rm HI}$(left) and b(right) parameters of the two simulations. The relative differences are shown in the bottom panels. The *L100N4096* box is shown in blue, while the *L20N1024* box is shown in green. The peak of b parameters are given in the text.

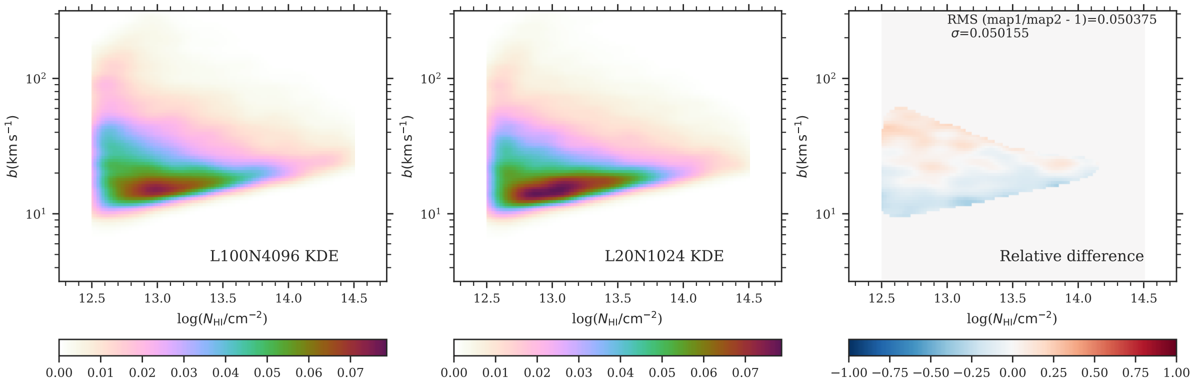

Figure A.9: The 2D KDE maps of b-$N_{\rm HI}$ distributions for simulations *L100N4096* (left-hand panel) and *L20N1024* (middle panel). The right-hand panel shows the relative difference $\Delta P = (P_{L100}/P_{L20} - 1)$. To avoid division by zero, we apply a threshold and only include regions integrating up to 75% for P_{L20}. The KDE maps are made from 20000 data points for each simulation, and the RMS and standard deviation of of the relative difference map are given in the right-hand panel. Details of the calculations are given in the text.

two simulations. We can see the distributions of T and Δ are alike for both while the small box *L20N1024* simulation has much smaller line-of-sight velocity. This is expected since line-of-sight velocities are dominated by the large scale modes that exist only in the large box simulations. However, these large velocities are because of bulk motion and therefore do not affect the parameters of the Lyα forest lines.

For both simulations, we follow the forward modeling and line fitting procedures discussed in Section §2.1, except that here we use a Gaussian LSF with fixed resolution R=3.5 km/s and assume a SNR=100. Such choices of resolution and SNR assure that the Lyα forest are fully resolved and the box-size effect are independent of resolution and instrument. For both simulations, we use the photoionization rate $\log \Gamma_{\rm HI}(\mathrm{s}^{-1}) = -13.308$. The 1D marginalized distributions of Doppler parameter b and column density $N_{\rm HI}$ of both simulations are presented in Fig. A.8. The $N_{\rm HI}$ distribution of the two simulations are in excellent agreement with each other, with the relative difference $\Delta_{P(N)} < 10\%$. The b parameter have very similar distributions for both simulations, where the two distributions agree with each other near the peak, with relative difference $\Delta_{P(b)} < 25\%$, and the difference increases as $\log b$ becomes smaller than 1.0 or larger than 2.0, which however have very small contribution in the total cumulative distribution. The peak values of the b parameter for both simulations are given in Fig. A.8, where the b distributions give $b_{\rm peak} = 18.3$ km/s and 17.2km/s for *L100N4096* and *L20N1024* simulations respectively. We count the dN/dz for both simulations, *L20N1024* gives dN/dz= 750, and *L100N4096* gives dN/dz= 700. The difference in dN/dz is about 7%. Furthermore, we plot the 2D b-$N_{\rm HI}$ distribution in Fig. A.9 for both simulations. These are 2D KDE maps each generated by 20000 data points collected from the $\{b, N_{\rm HI}\}$ dataset following the procedures described in Section §2.1. In the right most panel of Fig. A.9, we plot the relative difference of the KDE map, given by $\Delta P = (P_{L100}/P_{L20} - 1)$, where the P_{L100} and P_{L20} stand for the KDE for *L100N4096* and *L20N1024* simulations respectively. To avoid division by zero, we apply a small threshold and only include regions with $P_{L20} > P_{\rm TH}$, where $\int_{P_{\rm TH}}^{\infty} P {\rm d}P = 75\%$. We quantify the overall relative difference by calculating the root mean square and standard deviation of the ΔP. As shown in Fig. A.9, the relative differences in the 2D b-$N_{\rm HI}$ distribution are small and only about 5%. Therefore we conclude, even at $z \sim 0.5$ box-sizes do not affect the parameters of Lyα forest significantly.

Appendix B

Appendix

B.1 Observational data and Metal masks

In this Appendix, we present our observational spectra and the corresponding masks for the 12 archival HST STIS spectra. The original spectrum (normalized) is shown in grey, and the model based on VP-fitting is shown in blue. The noise vector of the original spectrum is shown in red, and the masked regions due to metal line detection are shown as green shaded regions. The Lyα lines used for our $\{b, N_{\mathrm{HI}}\}$ dataset (after all filters) are labelled by red vertical lines.

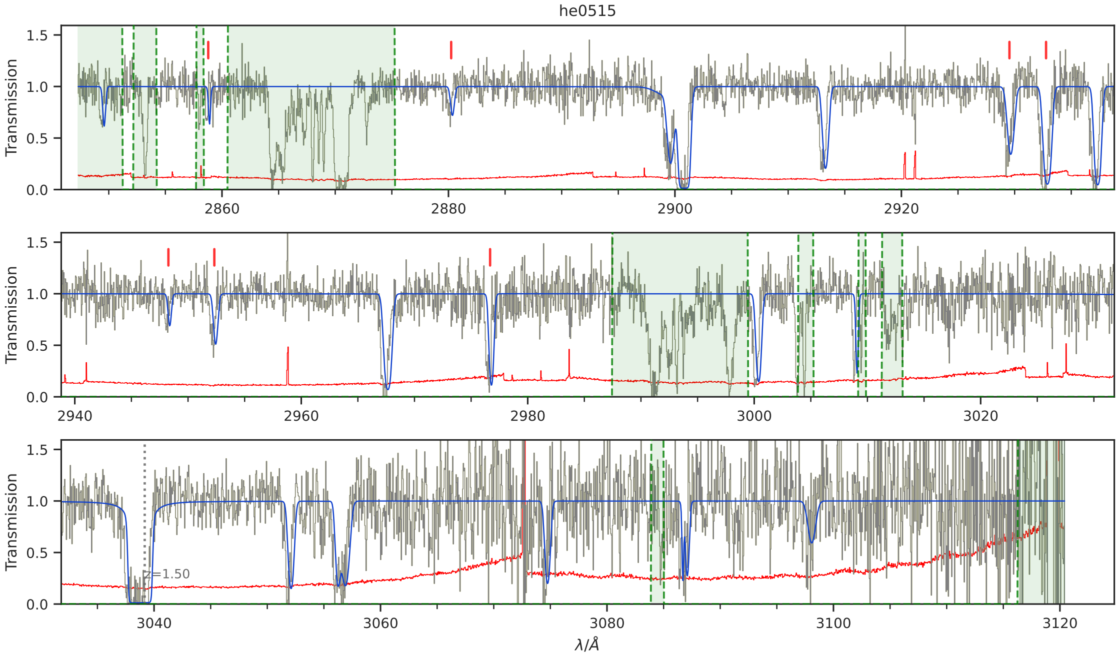

Figure B.1: Illustration of the processed STIS spectrum of HE0515-4414. The original spectrum is shown in gray, while a model spectrum based on line fitting (described in § 2.1.4) is shown in blue. The noise vector of the original spectrum is shown in red, and the masked regions are shown as a green shaded region. The Lyα lines used for our $\{b, N_{\rm HI}\}$ dataset are labelled by red vertical line.

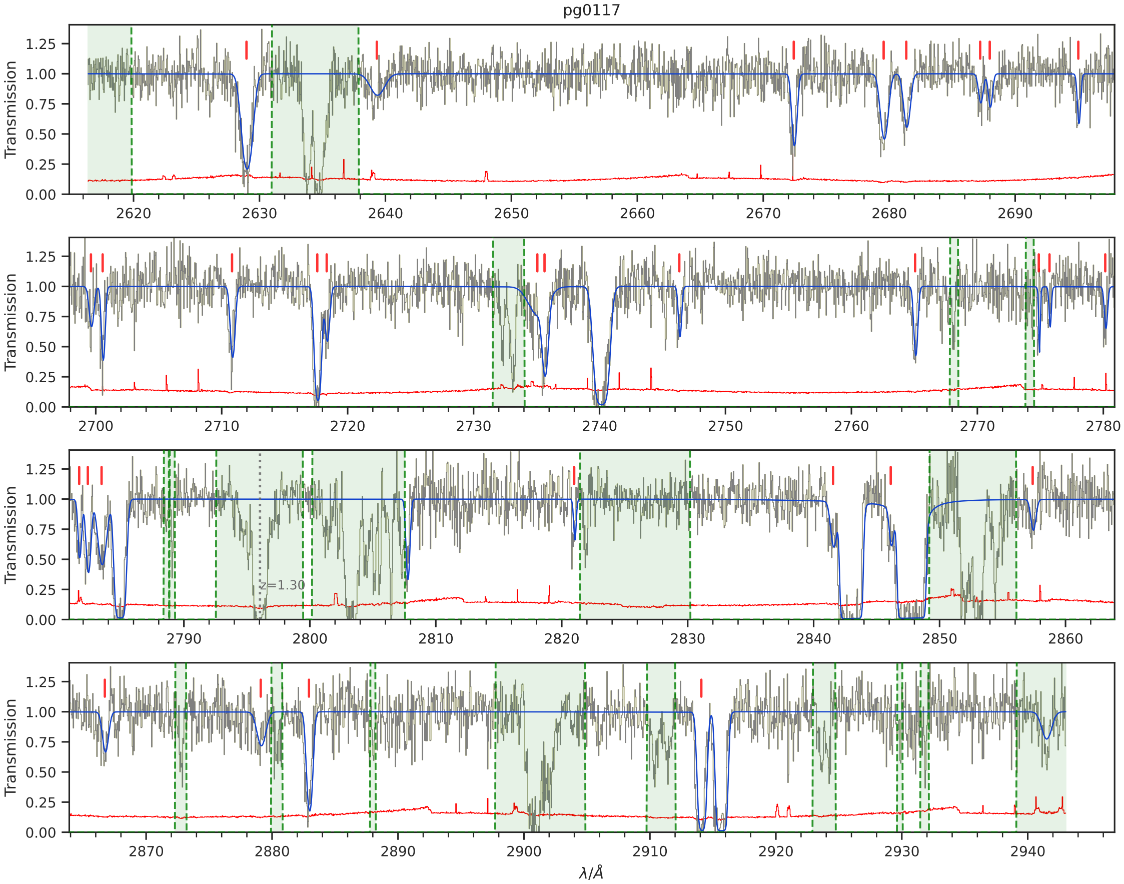

Fig. B.1 continued. Spectrum of PG0117+213.

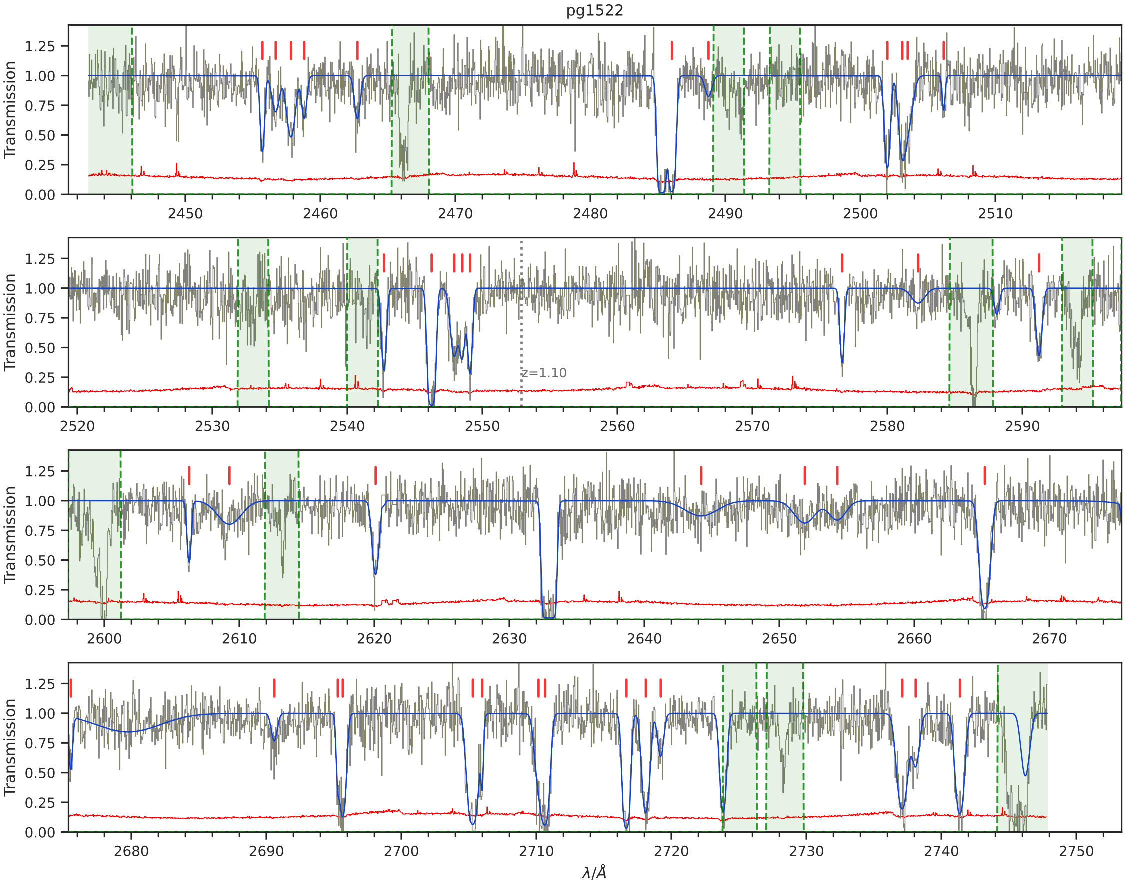

Fig. B.1 continued. Spectrum of PG1522+101.

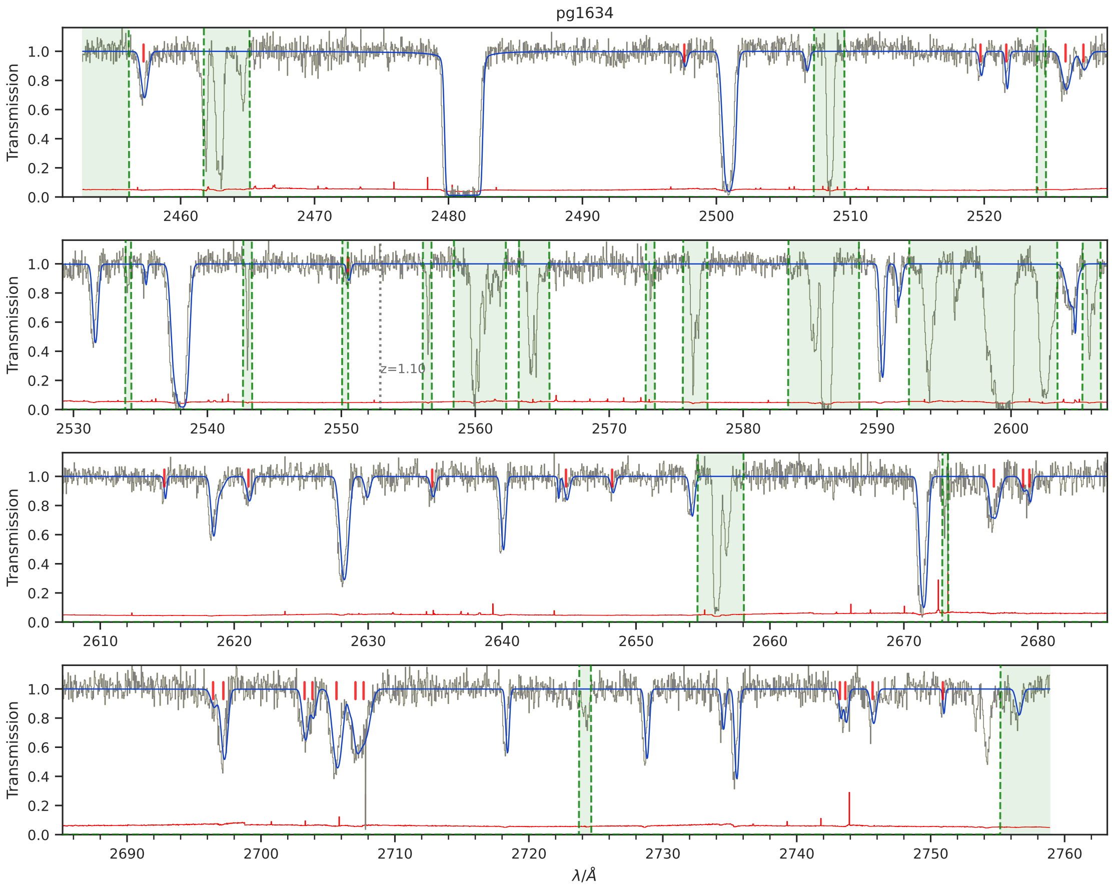

Fig. B.1 continued. Spectrum of PG1634+706.

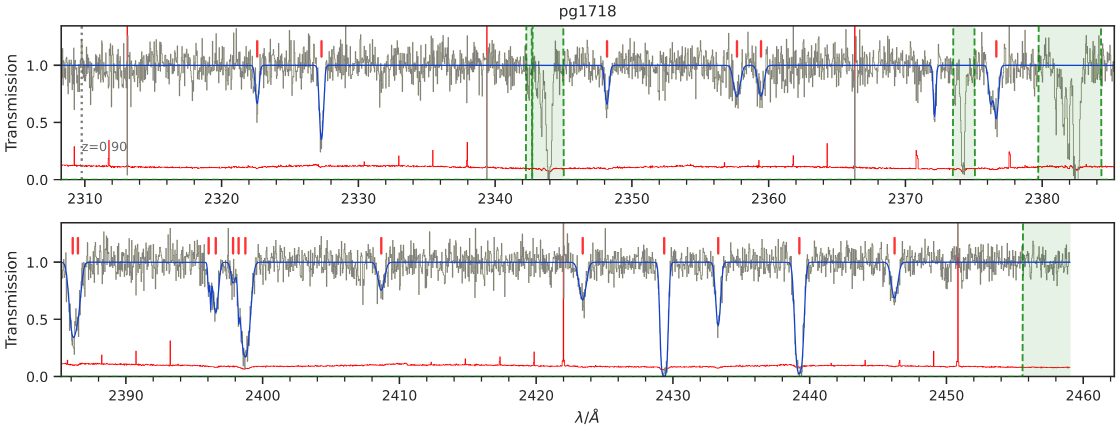

Fig. B.1 continued. Spectrum of PG1718+481.

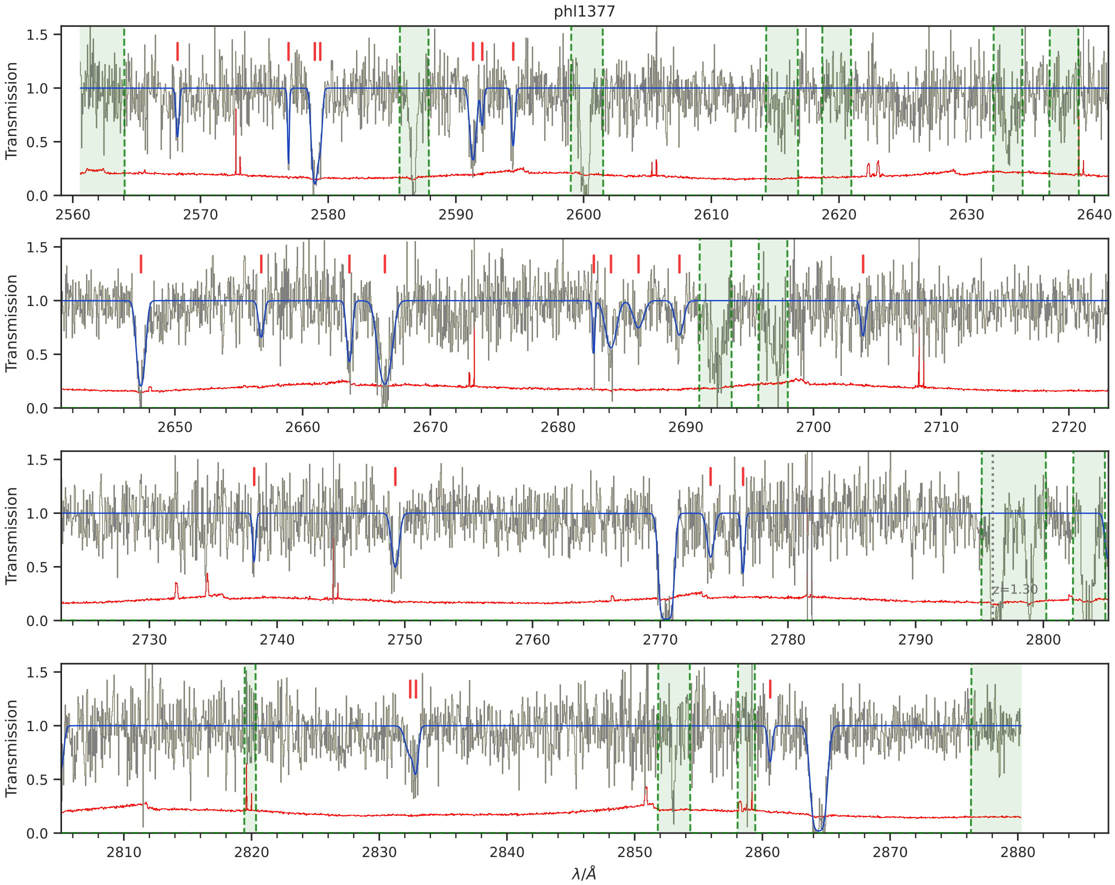

Fig. B.1 continued. Spectrum for PHL1377.

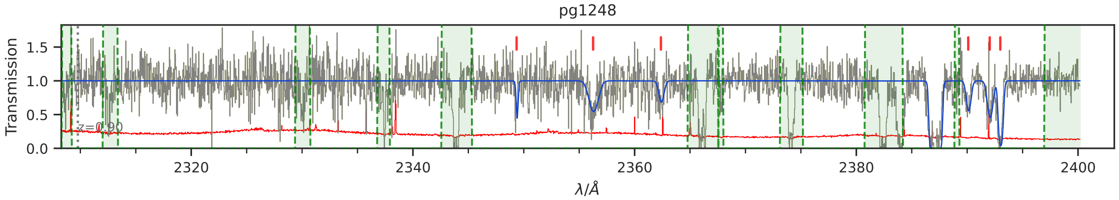

Fig. B.1 continued. Spectrum of PG1248+401.

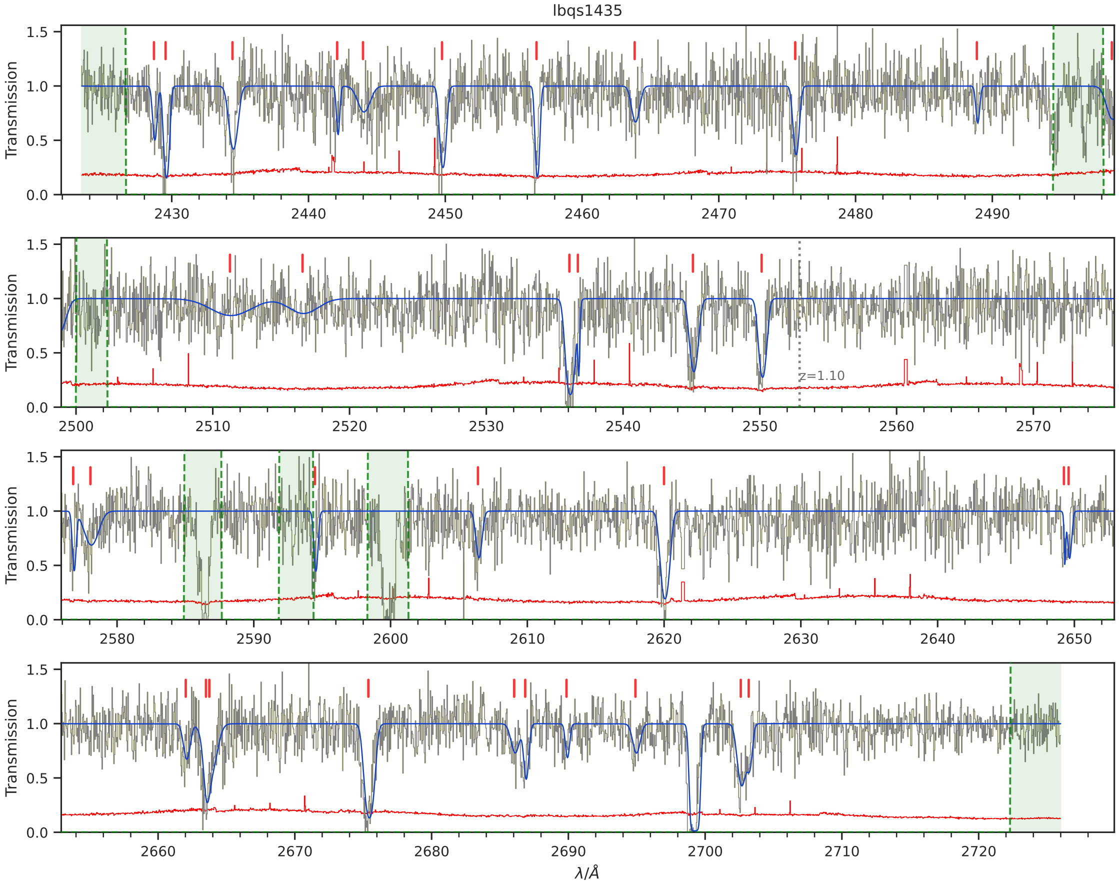

Fig. B.1 continued. Spectrum of LBQS1435-0134.

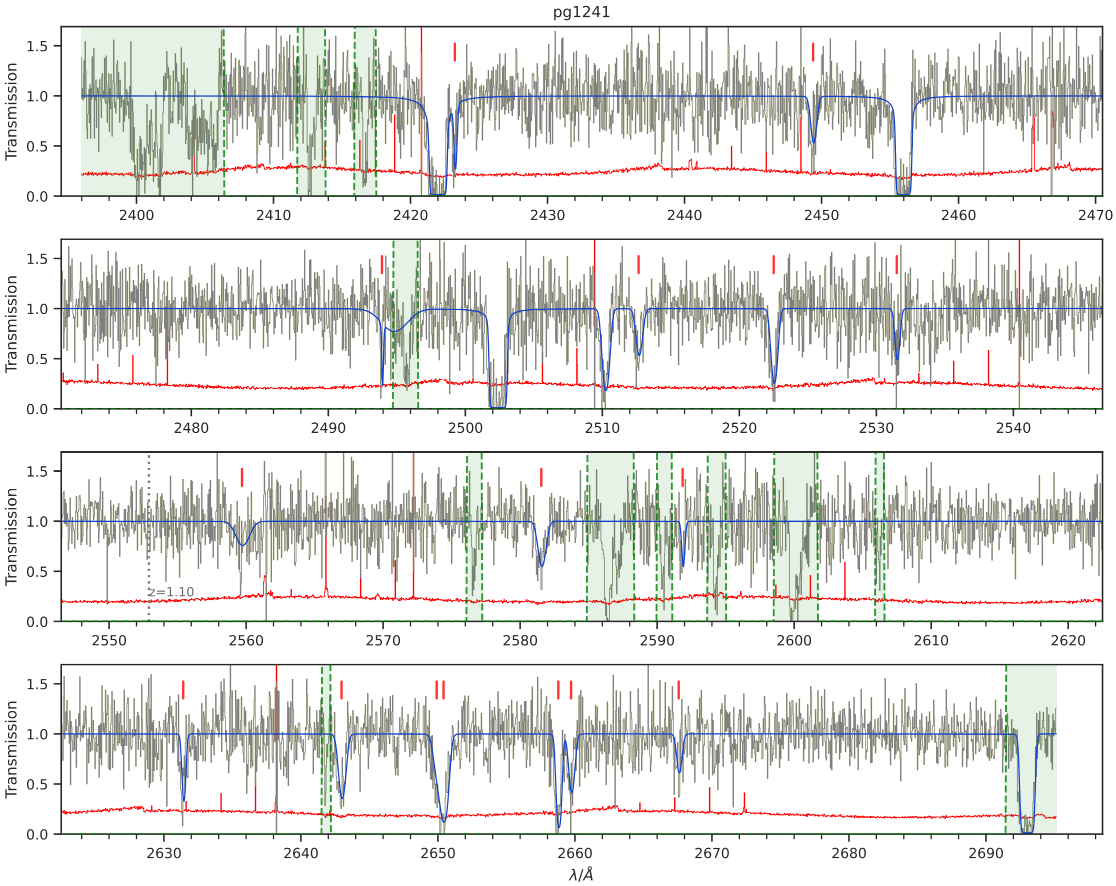

Fig. B.1 continued. Spectrum of PG1241+176.

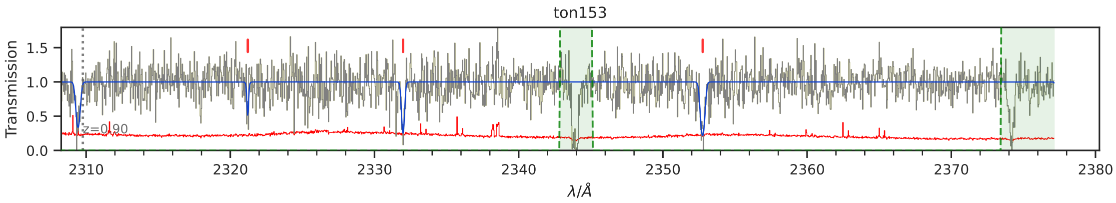

Fig. B.1 continued. Spectrum of TON153.

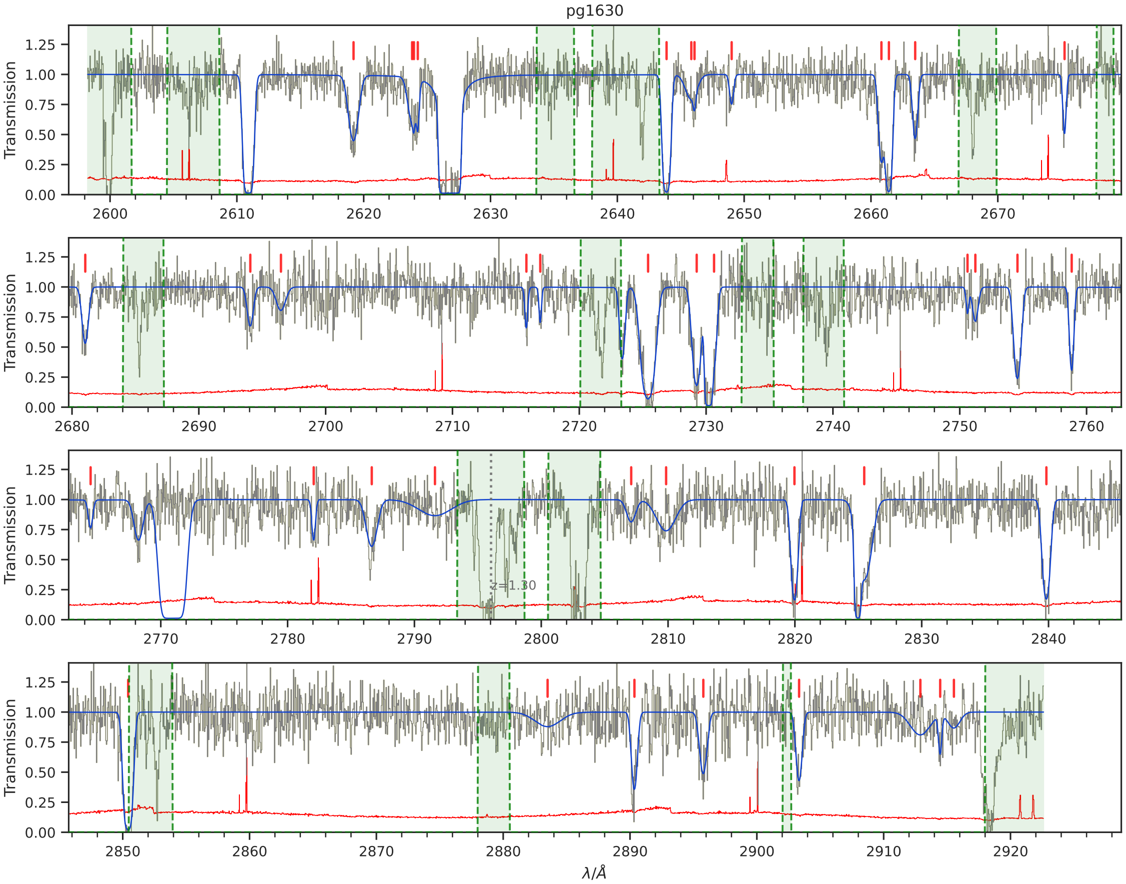

Fig. B.1 continued. Spectrum of PG1630+377.

Appendix C

Appendix

C.1 Inference based on the Photoheating labels [A,B]

In this section, we present our inference results using the framework where different Nyx models are labelled by the photoheating parameters [A,B] instead of the thermal state [T_0, γ], and the inference method returns [$A, B, \log\Gamma_{\rm HI}$]. The inference is conducted following the procedures described in § 5.2, based on the DELFI b-$N_{\rm HI}$ distribution emulator trained on training dataset labelled by [A,B, $\log\Gamma_{\rm HI}$], which returns $P(b\,, N_{\rm HI} \mid A, B, \log\Gamma_{\rm HI})$.

The simulation grid, parameterized by the photoheating labels [A,B], is given in Fig. C.3. An example of the MCMC posterior obtained based on the aforementioned likelihood function is given in Fig. C.1. The inference method returns A =1.321 (1.0), $B = -0.190$ (0.0), $\Gamma_{\rm HI} = -13.160$ (-13.093), whereas the true values are given in the parentheses. The posterior appears compact, with the medians of the marginalized posteriors landing within 1-σ errors for all three parameters. The b-$N_{\rm HI}$ distribution recovered from the mock dataset is presented in Fig. C.2, which is emulated by our DELFI b-$N_{\rm HI}$ distribution emulator based on the inferred parameters.

We perform an inference test following the § 5.2.2, in which we also exclude models that are too close to the parameter boundaries to avoid the truncation of the resulting posteriors. Specifically, we only use models with $3.3 < \log(T_0/{\rm K}) < 3.9$, $1.0 < \gamma < 2.3$, $-13.75 < \log(\Gamma_{\rm HI}/{\rm s}^{-1}) < -13.0$. The result of the inference test is shown in Fig. 5.4. The performance looks comparable to the one based on the thermal state [T_0, γ], suggesting that [T_0, γ] are still effective IGM parameters at low-z, notwithstanding the substantial dispersion in the IGM T-Δ distribution induced by pervasive shock heating at this redshift.

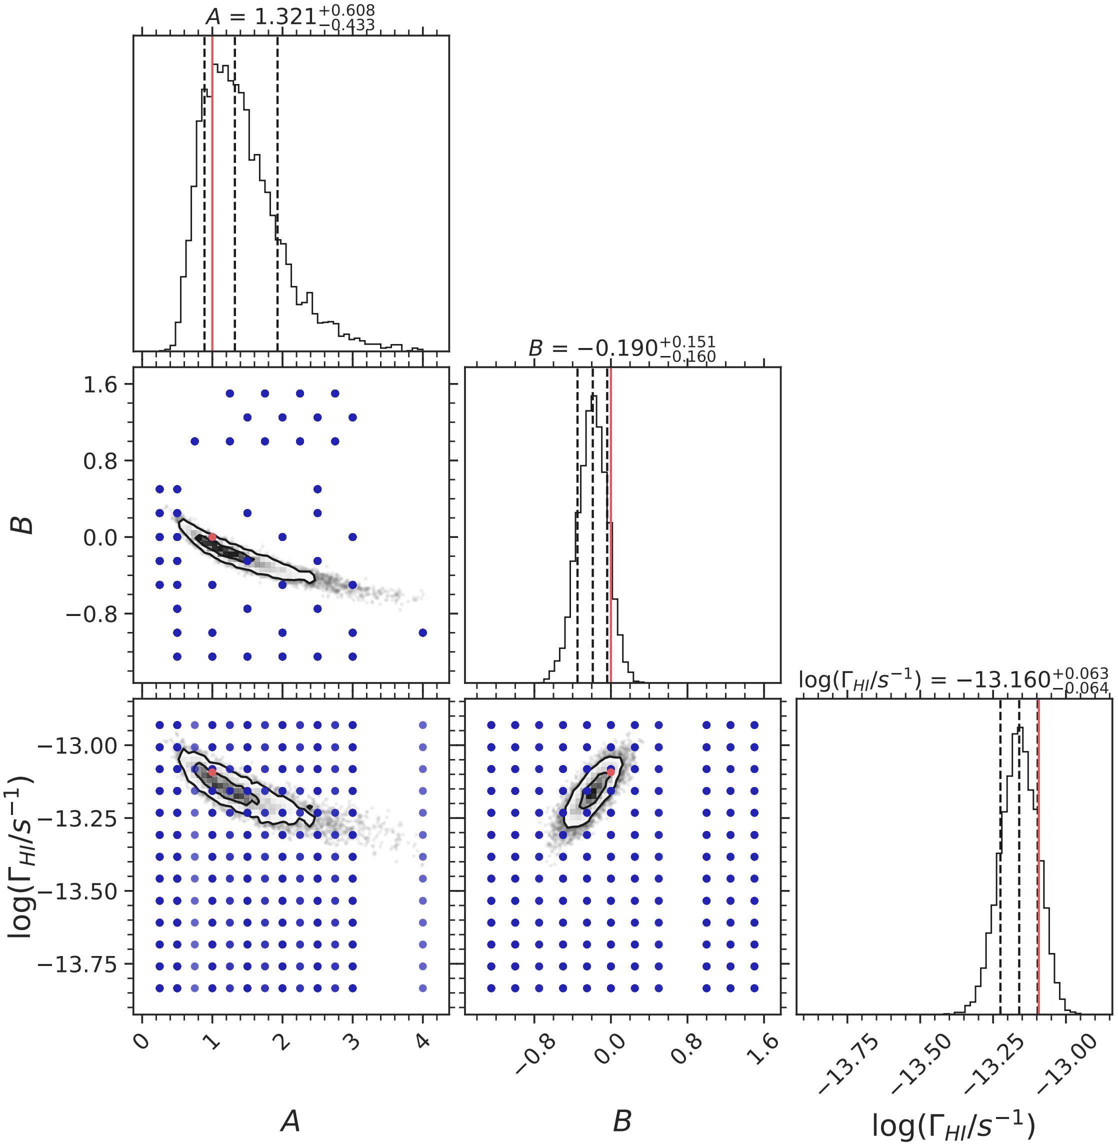

Figure C.1: An example of posterior obtained by our inference method based on [T_0, γ, $\Gamma_{\mathrm{HI()}}$]. Projections of the thermal grid used for generating models are shown as blue dots, while the true model is shown as red dots. The inner (outer) black contour represents the projected 2D 1(2)-sigma interval. The parameters of true models are indicated by red lines in the marginal distributions, while the dashed black lines indicate the 16, 50, and 84 percentile values of the posterior. The true parameters are: A = 1.0 and B =0.0, $\log(\Gamma_{\mathrm{HI}}/\mathrm{s}^{-1})$ = -13.093.

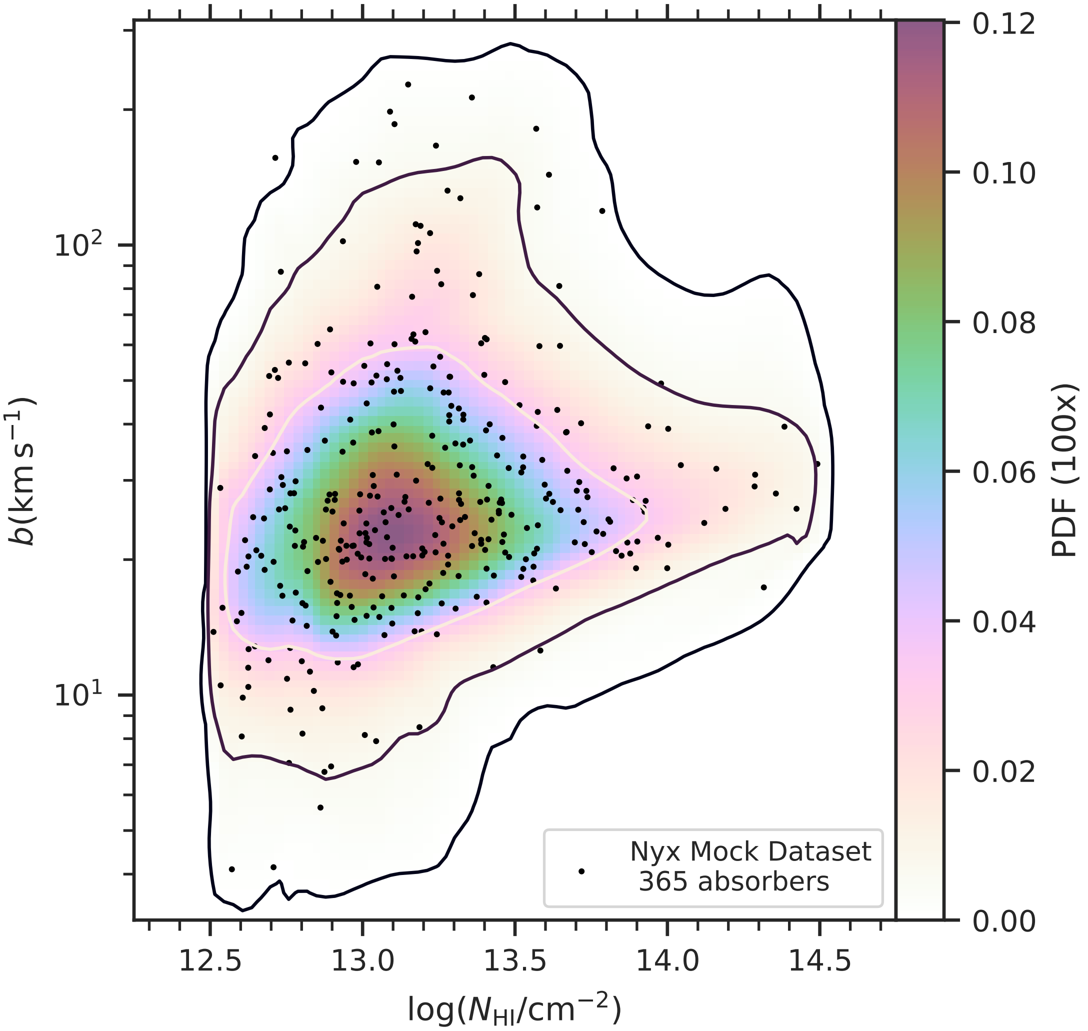

Figure C.2: The colour map is the full b-$N_{\rm HI}$ distribution recovered from the Nyx mock dataset, which is emulated by our DELFI emulator based on the best-fit parameters (median values of the marginalized MCMC posterior), where A = 3.695 (1.0) and B = 1.507 (0.0) and $\log(\Gamma_{\rm HI}/\mathrm{s}^{-1})$ =-13.237 (-13.093), the true parameters are given in parentheses. Black dots are the mock datasets we used in the inference. The contours correspond to cumulative probabilities of 68%, 95% and 99.7%. For illustration purposes, the values of pdf are multiplied by 100 in the colour bar.

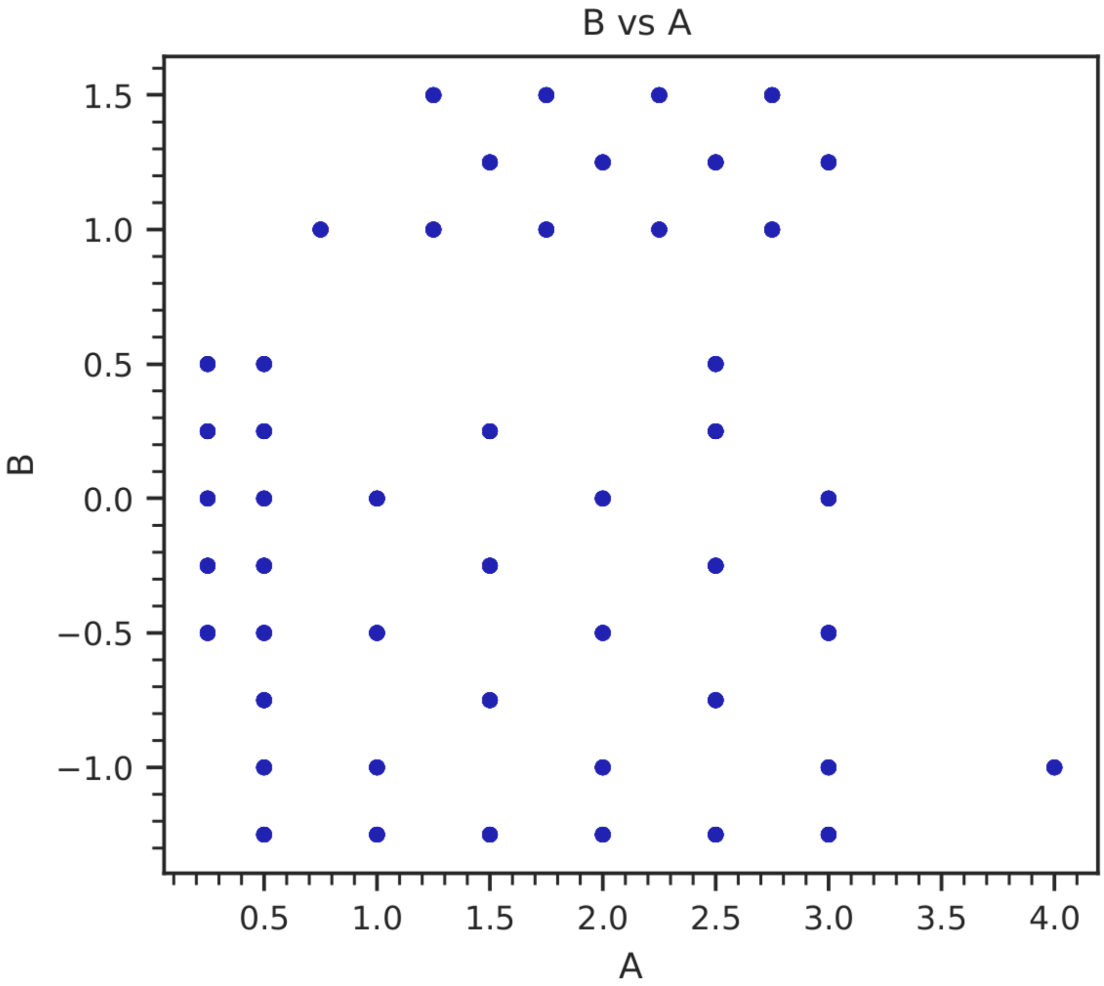

Figure C.3: Parameters grid (blue circles) from snapshots of hydrodynamic simulations of the THERMAL suite at $z = 0.1$ parameterized by the thermal state [A,B]

Bibliography

Acharya A., Khaire V., 2022, MNRAS, 509, 5559

Almgren A. S., Bell J. B., Lijewski M. J., Lukić Z., Van Andel E., 2013, ApJ, 765, 39

Alsing J., Wandelt B., 2018, Monthly Notices of the Royal Astronomical Society: Letters, 476, L60

Alsing J., Wandelt B., Feeney S., 2018, MNRAS, 477, 2874

Alsing J., Charnock T., Feeney S., Wandelt B., 2019, MNRAS, 488, 4440

Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O'Neil M., 2016, IEEE Transactions on Pattern Analysis and Machine Intelligence, 38, 252

Araya I. J., Padilla N. D., 2014, MNRAS, 445, 850

Armengaud E., Palanque-Delabrouille N., Yèche C., Marsh D. J. E., Baur J., 2017, MNRAS, 471, 4606

Becker G. D., Bolton J. S., 2013a, MNRAS, 436, 1023

Becker G. D., Bolton J. S., 2013b, MNRAS, 436, 1023

Becker G. D., Bolton J. S., Haehnelt M. G., Sargent W. L. W., 2011, MNRAS, 410, 1096

Boera E., Murphy M. T., Becker G. D., Bolton J. S., 2014, MNRAS, 441, 1916

Bolton J. S., 2007, The Observatory, 127, 262

Bolton J. S., Haehnelt M. G., 2007, MNRAS, 382, 325

Bolton J. S., Viel M., Kim T. S., Haehnelt M. G., Carswell R. F., 2008, MNRAS, 386, 1131

Bolton J. S., Becker G. D., Haehnelt M. G., Viel M., 2014, MNRAS, 438, 2499

Bolton J. S., Caputo A., Liu H., Viel M., 2022a, , 129, 211102

Bolton J. S., Gaikwad P., Haehnelt M. G., Kim T.-S., Nasir F., Puchwein E., Viel M., Wakker B. P., 2022b, MNRAS, 513, 864

Bolton J. S., Gaikwad P., Haehnelt M. G., Kim T.-S., Nasir F., Puchwein E., Viel M., Wakker B. P., 2022c, MNRAS, 513, 864

Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, MNRAS, 370, 645

Burchett J. N., et al., 2019, , 877, L20

Busca N. G., et al., 2013, , 552, A96

Carswell R. F., Webb J. K., 2014, VPFIT: Voigt profile fitting program (ascl:1408.015)

Chen Z., Oh S. P., 2024, MNRAS, 530, 4032

Chen H.-W., Johnson S. D., Zahedy F. S., Rauch M., Mulchaey J. S., 2017, , 842, L19

Croom S. M., et al., 2009, MNRAS, 399, 1755

Croton D. J., et al., 2006a, MNRAS, 365, 11

Croton D. J., et al., 2006b, MNRAS, 365, 11

Danforth C. W., et al., 2016, VizieR Online Data Catalog, p. J/ApJ/817/111

Davé R., Tripp T. M., 2001, ApJ, 553, 528

Davé R., Oppenheimer B. D., Katz N., Kollmeier J. A., Weinberg D. H., 2010, MNRAS, 408, 2051

Debuhr J., Quataert E., Ma C.-P., 2011, MNRAS, 412, 1341

Dixon K. L., Furlanetto S. R., 2009, ApJ, 706, 970

Fan X., et al., 2006, , 132, 117

Faucher-Giguère C.-A., 2020, MNRAS, 493, 1614

Faucher-Giguère C.-A., Lidz A., Hernquist L., Zaldarriaga M., 2008, ApJ, 688, 85

Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, ApJ, 703, 1416

Finkelstein S. L., et al., 2019, ApJ, 879, 36

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, , 125, 306

Fumagalli M., Haardt F., Theuns T., Morris S. L., Cantalupo S., Madau P., Fossati M., 2017, MNRAS, 467, 4802

Gaia Collaboration et al., 2016, , 595, A1

Gaia Collaboration et al., 2023, , 674, A1

Gaikwad P., Khaire V., Choudhury T. R., Srianand R., 2017a, MNRAS, 466, 838

Gaikwad P., Srianand R., Choudhury T. R., Khaire V., 2017b, MNRAS, 467, 3172

Gaikwad P., Choudhury T. R., Srianand R., Khaire V., 2018, MNRAS, 474, 2233

Gaikwad P., Srianand R., Haehnelt M. G., Choudhury T. R., 2021, MNRAS, 506, 4389

Garny M., Konstandin T., Sagunski L., Viel M., 2021, , 2021, 049

Garzilli A., Bolton J. S., Kim T. S., Leach S., Viel M., 2012, MNRAS, 424, 1723

Garzilli A., Theuns T., Schaye J., 2015, MNRAS, 450, 1465

Garzilli A., Theuns T., Schaye J., 2018, preprint, (`arXiv:1808.06646`)

Genel S., et al., 2014, MNRAS, 445, 175

Germain M., Gregor K., Murray I., Larochelle H., 2015, in International Conference on Machine Learning. pp 881–889

Ghavamian P., et al., 2009, Preliminary Characterization of the Post- Launch Line Spread Function of COS, COS Instrument Science Report 2009-01(v1), 23 pages

Goodman J., Weare J., 2010, CAMCoS, 5, 65

Green J. C., et al., 2012, ApJ, 744, 60

Gunn J. E., Peterson B. A., 1965, ApJ, 142, 1633

Haardt F., Madau P., 2012, ApJ, 746, 125

Haislmaier K. J., Tripp T. M., Katz N., Prochaska J. X., Burchett J. N., O'Meara J. M., Werk J. K., 2021, MNRAS, 502, 4993

Hiss H., 2019, PhD thesis, Dekanat der Fakultät für Physik und Astronomie, `https://doi.org/10.11588/heidok.00027299`

Hiss H., Walther M., Hennawi J. F., Oñ orbe J., O'Meara J. M., Rorai A., Lukić Z., 2018, ApJ, 865, 42

Hiss H., Walther M., Oñorbe J., Hennawi J. F., 2019, ApJ, 876, 71

Hopkins P. F., Hernquist L., Cox T. J., Kereš D., 2008a, , 175, 356

Hopkins P. F., Hernquist L., Cox T. J., Kereš D., 2008b, , 175, 356

Hu T., et al., 2022, MNRAS, 515, 2188

Hu T., Khaire V., Hennawi J. F., Onorbe J., Walther M., Lukic Z., Davies F., 2023a, arXiv e-prints, p. arXiv:2308.14738

Hu T., Khaire V., Hennawi J. F., Tripp T. M., Oñorbe J., Walther M., Lukic Z., 2023b, arXiv e-prints, p. arXiv:2311.17895

Hui L., Gnedin N. Y., 1997, MNRAS, 292, 27

Hussain T., Khaire V., Srianand R., Muzahid S., Pathak A., 2017, MNRAS, 466, 3133

Inoue A. K., Kamaya H., 2010, Earth, Planets and Space, 62, 69

Iršič V., et al., 2024, , 109, 043511

Khaire V., 2017, MNRAS, 471, 255

Khaire V., Srianand R., 2015, MNRAS, 451, L30

Khaire V., Srianand R., 2019a, MNRAS, 484, 4174

Khaire V., Srianand R., 2019b, MNRAS, 484, 4174

Khaire V., Srianand R., Choudhury T. R., Gaikwad P., 2016, MNRAS, 457, 4051

Khaire V., et al., 2019, MNRAS, 486, 769

Khaire V., Hu T., Hennawi J. F., Walther M., Davies F., 2023a, arXiv e-prints, p. arXiv:2306.05466

Khaire V., Hu T., Hennawi J. F., Burchett J. N., Walther M., Davies F., 2023b, arXiv e-prints, p. arXiv:2311.08470

Kimble R. A., et al., 1998, in Bely P. Y., Breckinridge J. B., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 3356, Space Telescopes and Instruments V. pp 188–202, doi:10.1117/12.324433

Kollmeier J. A., et al., 2014, , 789, L32

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

Lacki B. C., Charlton J. C., 2010, MNRAS, 403, 1556

Lee K.-G., et al., 2015, The Astrophysical Journal, 799, 196

Lehner N., et al., 2013, The Astrophysical Journal, 770, 138

Lidz A., Faucher-Giguère C.-A., Dall'Aglio A., McQuinn M., Fechner C., Zaldarriaga M., Hernquist L., Dutta S., 2010, ApJ, 718, 199

Lueckmann J.-M., Bassetto G., Karaletsos T., Macke J. H., 2018, arXiv preprint arXiv:1805.09294

Lukić Z., Stark C. W., Nugent P., White M., Meiksin A. A., Almgren A., 2015, MNRAS, 446, 3697

Lynds R., 1971, Apj, 164, L73

Madau P., Meiksin A., 1994, , 433, L53

Madau P., Pozzetti L., Dickinson M., 1998, ApJ, 498, 106

Marinacci F., et al., 2018, MNRAS, 480, 5113

Martizzi D., et al., 2019, MNRAS, 486, 3766

McDonald P., 2006, , 74, 103512

McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., 2001, ApJ, 562, 52

McGreer I. D., Mesinger A., D'Odorico V., 2015, MNRAS, 447, 499

McQuinn M., 2016, , 54, 313

McQuinn M., Upton Sanderbeck P. R., 2016, MNRAS, 456, 47

McQuinn M., Lidz A., Zaldarriaga M., Hernquist L., Hopkins P. F., Dutta S., Faucher-Giguère C.-A., 2009, Apj, 694, 842

Medallon S., Welty D., 2023, in , Vol. 22, STIS Instrument Handbook for Cycle 31 v. 22. p. 22

Ménard B., Scranton R., Fukugita M., Richards G., 2010, MNRAS, 405, 1025

Milutinović N., et al., 2007, MNRAS, 382, 1094

Miralda-Escudé J., Haehnelt M., Rees M. J., 2000, ApJ, 530, 1

Morrison J., Simon N., 2018, Journal of Computational and Graphical Statistics, 27, 648

Naiman J. P., et al., 2018, MNRAS, 477, 1206

Nasir F., Bolton J. S., Viel M., Kim T.-S., Haehnelt M. G., Puchwein E., Sijacki D., 2017, MNRAS, 471, 1056

Nelson D., et al., 2018, MNRAS, 475, 624

Nelson D., et al., 2019, Computational Astrophysics and Cosmology, 6, 2

Oñorbe J., Hennawi J. F., Lukić Z., 2017a, The Astrophysical Journal, 837, 106

Oñorbe J., Hennawi J. F., Lukić Z., Walther M., 2017b, The Astrophysical Journal, 847, 63

Palanque-Delabrouille N., et al., 2013, , 551, A29

Palanque-Delabrouille N., Yèche C., Schöneberg N., Lesgourgues J., Walther M., Chabanier S., Armengaud E., 2020, , 2020, 038

Papamakarios G., Murray I., 2016, in Advances in Neural Information Processing Systems. pp 1028–1036

Papamakarios G., Pavlakou T., Murray I., 2017, arXiv e-prints, p. arXiv:1705.07057

Papamakarios G., Sterratt D. C., Murray I., 2018, arXiv preprint arXiv:1805.07226

Penzias A. A., Wilson R. W., 1965, ApJ, 142, 419

Pillepich A., et al., 2018a, MNRAS, 473, 4077

Pillepich A., et al., 2018b, MNRAS, 475, 648

Planck Collaboration et al., 2014, , 571, A16

Prangle D., Blum M. G. B., Popovic G., Sisson S. A., 2014, Australian & New Zealand Journal of Statistics, 56, 309

Prochaska J. X., et al., 2019, , 243, 24

Puchwein E., Pfrommer C., Springel V., Broderick A. E., Chang P., 2012, MNRAS, 423, 149

Puchwein E., Haardt F., Haehnelt M. G., Madau P., 2019, MNRAS, 485, 47

Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, MNRAS, 430, 2427

Ricotti M., Gnedin N. Y., Shull J. M., 2000, ApJ, 534, 41

Robertson B. E., Ellis R. S., Furlanetto S. R., Dunlop J. S., 2015, ApJL, 802, L19

Rorai A., Hennawi J. F., White M., 2013, ApJ, 775, 81

Rorai A., et al., 2017, Science, 356, 418

Rorai A., Carswell R. F., Haehnelt M. G., Becker G. D., Bolton J. S., Murphy M. T., 2018, MNRAS, 474, 2871

Ross N. P., et al., 2013, ApJ, 773, 14

Rudie G. C., Steidel C. C., Pettini M., 2012a, ApJ, 757, L30

Rudie G. C., Steidel C. C., Pettini M., 2012b, , 757, L30

Ryden B., 2003, Introduction to Cosmology. Addison-Wesley, `https://books.google.com/books?id=z27vAAAAMAAJ`

Schaye J., Theuns T., Leonard A., Efstathiou G., 1999, MNRAS, 310, 57

Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, MNRAS, 318, 817

Schaye J., et al., 2015, MNRAS, 446, 521

Sellentin E., Starck J.-L., 2019, Journal of Cosmology and Astroparticle Physics, 2019, 021

Shull J. M., Moloney J., Danforth C. W., Tilton E. M., 2015, ApJ, 811, 3

Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007a, MNRAS, 380, 877

Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007b, MNRAS, 380, 877

Springel V., 2005, MNRAS, 364, 1105

Springel V., 2010, MNRAS, 401, 791

Springel V., et al., 2005, , 435, 629

Springel V., et al., 2018, MNRAS, 475, 676

Syphers D., Shull J. M., 2014, ApJ, 784, 42

Tanabashi M., et al., 2018, Phys. Rev. D, 98, 030001

Tepper-García T., 2006, MNRAS, 369, 2025

Tepper-García T., Richter P., Schaye J., Booth C. M., Dalla Vecchia C., Theuns T., 2012, MNRAS, 425, 1640

Theuns T., Zaroubi S., 2000, MNRAS, 317, 989

Theuns T., Schaye J., Haehnelt M. G., 2000, MNRAS, 315, 600

Theuns T., Schaye J., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell B., 2002, The Astrophysical Journal, 567, L103

Tillman M. T., et al., 2023a, arXiv e-prints, p. arXiv:2307.06360

Tillman M. T., Burkhart B., Tonnesen S., Bird S., Bryan G. L., Anglés-Alcázar D., Davé R., Genel S., 2023b, , 945, L17

Tripp T., 2014, The COS Absorption Survey of Baryon Harbors (CASBaH): Probing the Circumgalactic Media of Galaxies from z = 0 to z = 1.5, HST Proposal ID 13846. Cycle 22

Tripp T. M., Giroux M. L., Stocke J. T., Tumlinson J., Oegerle W. R., 2001, ApJ, 563, 724

Tripp T. M., Sembach K. R., Bowen D. V., Savage B. D., Jenkins E. B., Lehner N., Richter P., 2008, , 177, 39

Upton Sanderbeck P. R., D'Aloisio A., McQuinn M. J., 2016, MNRAS, 460, 1885

Viel M., Bolton J. S., Haehnelt M. G., 2009, Monthly Notices of the Royal Astronomical Society: Letters, 399, L39

Viel M., Schaye J., Booth C. M., 2013, MNRAS, 429, 1734

Viel M., Haehnelt M. G., Bolton J. S., Kim T.-S., Puchwein E., Nasir F., Wakker B. P., 2017, MNRAS, 467, L86

Villaescusa-Navarro F., et al., 2021, ApJ, 915, 71

Villasenor B., Robertson B., Madau P., Schneider E., 2021, ApJ, 912, 138

Walther M., Hennawi J. F., Hiss H., Oñorbe J., Lee K.-G., Rorai A., O'Meara J., 2017, The Astrophysical Journal, 852, 22

Walther M., Oñorbe J., Hennawi J. F., Lukić Z., 2019a, ApJ, 872, 13

Walther M., Oñorbe J., Hennawi J. F., Lukić Z., 2019b, ApJ, 872, 13

Weinberg D. H., Hernsquit L., Katz N., Croft R., Miralda-Escudé J., 1997, in Petitjean P., Charlot S., eds, Structure and Evolution of the Intergalactic Medium from QSO Absorption Line System. p. 133 (`arXiv:astro-ph/9709303`)

Weinberger R., et al., 2017, MNRAS, 465, 3291

Wolfson M., Hennawi J. F., Davies F. B., Oñorbe J., Hiss H., Lukić Z., 2021, MNRAS, 508, 5493

Wolfson M., Hennawi J. F., Davies F. B., Oñorbe J., 2022, arXiv e-prints, p. arXiv:2208.09013

Woodgate B. E., et al., 1998, , 110, 1183

Worseck G., et al., 2011, ApJl, 733, L24

Worseck G., Davies F. B., Hennawi J. F., Prochaska J. X., 2018, preprint, (`arXiv:1808.05247`)

Wotta C. B., Lehner N., Howk J. C., O'Meara J. M., Oppenheimer B. D., Cooksey K. L., 2019, The Astrophysical Journal, 872, 81

Yèche C., Palanque-Delabrouille N., Baur J., du Mas des Bourboux H., 2017, , 2017, 047

Zaldarriaga M., Hui L., Tegmark M., 2001, The Astrophysical Journal, 557, 519

Ziegel J. F., Gneiting T., 2014, Electronic Journal of Statistics, 8, 2619