



Universiteit
Leiden
The Netherlands

Les Lumières: probing the cosmic Epoch of Reionization with high-redshift quasars

Kist, T.

Citation

Kist, T. (2026, July 3). *Les Lumières: probing the cosmic Epoch of Reionization with high-redshift quasars*. Retrieved from <https://hdl.handle.net/1887/4307539>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4307539>

Note: To cite this publication please use the final published version (if applicable).

4

INFERRING LOCAL QUASAR IGM DAMPING WING CONSTRAINTS

Abstract

Lyman- α damping wings towards quasars are a highly sensitive probe of the neutral hydrogen (HI) content in the foreground intergalactic medium (IGM), not only constraining the global timing of reionization but also the *local* ionization topology near the quasar. Near-optimal extraction of this information is possible with the help of two recently introduced reionization model-independent summary statistics of the HI distribution in the IGM *before* the quasar started shining, complemented with the quasar's lifetime encoding the effect of its ionizing radiation as a third parameter. We introduce a fully Bayesian JAX-based Hamiltonian Monte Carlo (HMC) inference framework that allows us to jointly reconstruct the quasar's unknown continuum and constrain these local damping wing statistics. We put forward a probabilistic framework that allows us to tie these local constraints to any specific reionization model and obtain model-dependent constraints on the global timing of reionization. We demonstrate that we are able to constrain the (Lorentzian-weighted) HI column density in front of the quasar to a precision of $0.69^{+0.06}_{-0.30}$ dex and its original distance to the first neutral patch before the quasar started shining to $31.4^{+10.7}_{-28.1}$ cMpc (if a noticeable damping wing is present in the spectrum), extracting hitherto unused local information from the IGM damping wing imprint. Once tied to a specific reionization model, we find that the statistical fidelity of our constraints on the global IGM neutral fraction and the lifetime of the quasar improves, while retaining the same precision as achieved by pipelines that infer these parameters directly.

Work by: **Timo Kist**, Joseph F. Hennawi and Frederick B. Davies, submitted for publication in *Monthly Notices of the Royal Astronomical Society*, doi.org/10.48550/arXiv.2508.21812. Reprinted here in its entirety.

4.1 Introduction

Little doubt exists about the fact that the formation of the first stars and galaxies heralded the epoch of reionization, a major phase transition in the history of our universe driven by the highly energetic radiation emitted by these objects which progressively reionized all neutral hydrogen (HI) in the intergalactic medium (IGM). Most details about this landmark event, however, are yet to be determined, starting with its mere timing in cosmic history, through to its topological features, informing us about the nature of the sources and sinks of reionization. One of the most promising probes of reionization is the Lyman- α transition in the spectra of bright astrophysical sources such as quasars. Its remarkable sensitivity to even the smallest amounts of neutral hydrogen in the IGM, causing extended Gunn-Peterson absorption troughs already at global volume-averaged IGM neutral fractions of $\langle x_{\text{HI}} \rangle \gtrsim 10^{-4}$ (Gunn & Peterson 1965), and the red damping wing induced by quantum mechanical line broadening when $\langle x_{\text{HI}} \rangle$ becomes of order unity (Miralda-Escudé 1998), allow us to gain unique insights into how reionization proceeded over cosmic time.

While the first damping wing constraints have been reported for individual objects (Mortlock et al. 2011; Bolton et al. 2011; Greig et al. 2017b, 2019, 2022; Bañados et al. 2018; Davies et al. 2018a; Wang et al. 2020; Yang et al. 2020a; Āurovčíková et al. 2020; Reiman et al. 2020), Greig et al. (2024a) and Āurovčíková et al. (2024) recently analyzed the first statistical ensembles probing the final stages of reionization. Even tighter constraints, reaching much deeper into the core stages of reionization, will be achievable in the coming years thanks to the growing number of high-redshift quasar spectra available to us (D’Odorico et al. 2023; Onorato et al. 2025; Euclid Collaboration et al. 2026), most prominently driven by the unprecedented discoveries of new objects at ever-higher redshifts in the Euclid wide field survey (Euclid Collaboration et al. 2019; Bañados et al. 2025; Euclid Collaboration et al. 2026).

Moreover, smooth roll-offs in the spectra near individual Gunn-Peterson troughs in the foreground of specific quasars have been interpreted as damping wings, possibly pointing to the persistence of neutral islands in the IGM down to $5.5 \lesssim z \lesssim 6$ (Becker et al. 2024; Spina et al. 2024; Zhu et al. 2024; Sawyer et al. 2025). In addition to that, the advent of JWST has enabled the first damping wing constraints towards galaxies, pushing the highest-redshift frontiers (Curtis-Lake et al. 2023; Hsiao et al. 2024; Keating et al. 2024b; Park et al. 2025; Umeda et al. 2024, 2025; Mason et al. 2026). However, additional care needs to be taken in this context due to the presence of intrinsic damped Lyman- α absorbers (DLAs; Heintz et al. 2024, 2025) which can mimic the cosmological imprint from the IGM (Huberty et al. 2025),

introducing an additional nuisance process which needs to be marginalized out (Mason et al. 2026).

Quasars, on the other hand, as powerful sources of ionizing radiation, suffer considerably less from this issue as they ionize away all residual neutral gas within their Mpc-scale proximity zone. Any potentially remaining proximate DLA systems can be excluded based on the presence of associated weak metal absorption systems which can be identified with the help of supplementary high-quality spectra (Davies et al. 2025). The remaining complication is the reconstruction of the intrinsic quasar continuum that the cosmological damping wing imprint has to be disentangled from. As the properties of these continua do not appear to evolve notably with redshift (Shen et al. 2007), data-driven models can be built based on unabsorbed low-redshift continua, and a plethora of approaches have been developed in the past years, ranging from simple models based on principal component analysis (PCA) to non-linear, neural network-based ones (for a comprehensive overview, see Greig et al. 2024b), most recently as a unified framework to *jointly* reconstruct the quasar continuum and the IGM absorption imprint (Hennawi et al. 2025).

Kist et al. (2025b) demonstrated that approximately half the total error budget on the inferred IGM neutral fraction $\langle x_{\text{HI}} \rangle$ is due to this continuum reconstruction task, whereas the other half is sourced by the stochastic distribution of IGM density fluctuations, as well as the distribution of neutral patches during reionization. To eliminate the latter part of this stochasticity from the inference task itself, Kist et al. (2025c) introduced a new parameterization of IGM damping wings, probing the *local* ionization topology in front of a given quasar *before* it is altered by the quasar’s ionizing radiation. This three-parameter model comprises 1) the HI column density, weighted by a Lorentzian profile that accounts for the frequency dependence of the Lyman- α cross section, 2) the distance from the quasar to the first neutral patch, and 3) the quasar lifetime which encapsulates the effects of the quasar’s ionizing radiation.

Kist et al. (2025c) further demonstrated that these summary statistics are reionization model invariant, and argued that they, when tied to a specific reionization model, provide constraints not only on the timing but also the topology of reionization. This work is concerned with the practical realization of this inference task by leveraging a fully Bayesian framework introduced in Hennawi et al. (2025) which allows us to infer the first *local* IGM damping wing constraints. Our local summary statistics can be constrained freely from any prior assumptions about the underlying reionization model. We demonstrate how the topology information from a given reionization model can be folded into these constraints *subsequently* in a probabilistic manner, resulting in *topology-informed* local constraints, and, in addition, a *global* constraint on the IGM neutral fraction $\langle x_{\text{HI}} \rangle$. In a companion paper, we

apply this new approach to JWST/NIRSpec spectra of the two $z \sim 7.5$ quasars J1007+2115 and J1342+0928 (Kist et al. 2026).

We start by introducing this local parameter framework in Section 4.2, and proceed in Section 4.3 by describing our inference approach originally introduced in Hennawi et al. (2025) which is applicable both in the context of our new local IGM damping wing parameterization as well as the conventional global one, parameterized by the volume-averaged neutral fraction $\langle x_{\text{HI}} \rangle$ and the lifetime t_{Q} of the quasar. In Section 4.4 we test the statistical fidelity of the pipeline and quantify the precision of the local damping wing constraints. We also compare the precision of the resulting constraints on $\langle x_{\text{HI}} \rangle$ and t_{Q} to that of the directly inferred ones. We conclude in Section 4.5.

4.2 Theory: relating a local IGM damping wing model to the global timing of reionization

Due to its sensitivity to neutral hydrogen in the foreground IGM, the Lyman- α damping wing signature in the spectra of high-redshift sources is considered one of the key probes of the global volume-averaged IGM neutral fraction $\langle x_{\text{HI}} \rangle$ as a function of redshift. As has recently been noted (Chen 2024; Keating et al. 2024a; Kist et al. 2025c), considerable sightline-to-sightline variations are possible due to fluctuations in the cosmological density field as well as the patchy nature of reionization. Specifically, Kist et al. (2025c) identified a two-dimensional set of physical summary statistics that (along with the lifetime of the quasar as a third parameter) tightly parametrizes the characteristic shape of the IGM damping wing. These two statistics are informative not only about the global timing of reionization, but also the local ionization topology *before the quasar started shining*. Note that the quasar radiation inevitably modifies the topology that ultimately imprints the damping wing, but our physical interest is directed at the original topology which encapsulates the information about what we will henceforth refer to as the *pre-quasar* IGM, in contrast to the *post-quasar* one, impacted by the quasar’s ionizing radiation. In this section, we will provide a short but self-contained overview over this parameterization, and demonstrate subsequently how we can fold in the topological information that arises from the assumed reionization model through a prior on our local damping wing statistics, and how we can use this to relate back these summaries to the global timing of reionization, parameterized by the redshift evolution of the global IGM neutral fraction $\langle x_{\text{HI}} \rangle(z)$.

4.2.1 A local, topology-independent parameterization of quasar IGM damping wings

We start by defining the two summary statistics of the local ionization topology around quasars introduced in Kist et al. (2025c). We refer the reader to this work for all additional details. The authors demonstrated that the information contained in the damping wing optical depth τ_{DW} can largely be condensed into a single number, resulting in a near-optimal parameterization of the Lyman- α damping wing across the entire spectral range. For a quasar at redshift z_{QSO} with a post-quasar HI density field $n_{\text{HI}}^{\text{post}}$, the damping wing optical depth τ_{DW} as a function of Lyman- α rest-frame wavelength $\lambda_{\text{rest}} = \lambda_{\alpha}(1 + \frac{v}{c})$ is given by

$$\tau_{\text{DW}}(\lambda_{\text{rest}}) = \int_0^{R(z_{\text{QSO}})} n_{\text{HI}}^{\text{post}}(R) \cdot \sigma_{\alpha} \left(\frac{1 + z_{\text{QSO}}}{1 + z(R)} \lambda_{\text{rest}} \right) dR, \quad (4.1)$$

where σ_{α} is the Lyman- α cross section, dR the infinitesimal proper line-of-sight interval, and $R(z_{\text{QSO}})$ the corresponding proper distance from the observer to the quasar.

In essence, Eq. (4.1) constitutes a column density integral of the HI density field with an additional weighting kernel, governed by the Lyman- α cross section σ_{α} . To excellent approximation, σ_{α} at a given spectral velocity pixel v_{T} is of Lorentzian shape, i.e., $\sigma_{\alpha} \sim (v - v_{\text{T}})^{-2}$, whose impact we can capture by defining the Lorentzian line-of-sight average $\langle\langle \cdot \rangle\rangle_{\text{Lor}}$ of a field X (such as the HI density field n_{HI}) as

$$\langle\langle X \rangle\rangle_{\text{Lor}} \equiv \frac{1}{\mathcal{N}(u_{\text{min}}, u_{\text{max}})} \int_{u_{\text{min}}}^{u_{\text{max}}} \frac{X(u)}{(u+1)^2} du, \quad (4.2)$$

where $u = R/R_{\text{T}}$ is a dimensionless integration variable, and $\mathcal{N}(u_{\text{min}}, u_{\text{max}}) \equiv (u_{\text{max}} - u_{\text{min}})/((u_{\text{max}} + 1)(u_{\text{min}} + 1))$ a normalization factor ensuring $\langle\langle \mathbf{1} \rangle\rangle_{\text{Lor}} = 1$ for the identity field $\mathbf{1}$. Here we defined R_{T} as the (positive) proper distance value corresponding to the red-side velocity offset v_{T} via $R_{\text{T}} \equiv +v_{\text{T}}/H(z_{\text{QSO}})$.

Instead of directly taking the Lorentzian-weighted average of the post-quasar HI density field $n_{\text{HI}}^{\text{post}}$, we operate on its pre-quasar version $n_{\text{HI}}^{\text{pre}} = \langle n_{\text{H}} \rangle(z_{\text{QSO}}) \cdot x_{\text{HI}} \cdot \Delta$, unaffected by the ionizing quasar radiation and therefore directly informative about the pre-quasar ionization topology parameterized by the neutral fraction field x_{HI} , along with the dimensionless matter overdensity field Δ and the cosmic mean hydrogen density $\langle n_{\text{H}} \rangle(z_{\text{QSO}})$. Specifically, we define the *Lorentzian-weighted* HI column density

$$N_{\text{HI}}^{\text{DW}} \equiv 5.1 \times 10^{20} \text{ cm}^{-2} \times \left(\frac{\mathcal{N}\left(\frac{r_{\text{min}}}{r_{\text{T}}}, \frac{r_{\text{max}}}{r_{\text{T}}}\right)}{0.67} \right) \left(\frac{r_{\text{T}}}{18 \text{ cMpc}} \right) \times \left(\frac{1 + z_{\text{QSO}}}{1 + 7.54} \right)^2 \left(\frac{\langle\langle x_{\text{HI}} \cdot \Delta \rangle\rangle_{\text{Lor}}}{1} \right), \quad (4.3)$$

where the normalization factor is defined as above, and r_{\min} , r_{\max} and r_T are the comoving versions of the proper distances R_{\min} , R_{\max} and R_T .¹ With this definition at hand, it is then straightforward to show that we have identified an asymptotically optimal parameterization of the (pre-quasar) IGM damping wing, in the sense that at the velocity offset v_T itself, we find a direct proportionality between the optical depth and our summary statistic:

$$\tau_{\text{DW}}^{\text{pre}}(v = v_T) \simeq \frac{e^2}{m_e c^2} f_\alpha \gamma_\alpha \lambda_\alpha (c/v_T - 1)^2 \times N_{\text{HI}}^{\text{DW}} \quad (4.4)$$

in the limit where the integration limits approach $u_{\min} \rightarrow 0$ and $u_{\max} \rightarrow R(z_{\text{QSO}})/R_T$, and where the Lorentzian approximation of the Lyman- α cross section is valid.

To make the statistic adequate for the *post*-quasar optical depth despite the fact that we are operating on the *pre*-quasar HI density field, we have to account for the fact that quasars commonly ionize away all neutral material within the first few cMpc surrounding them, and therefore not contributing to the optical depth integral in Eq. (4.1). We do so by starting the integration in Eq. (4.2) at a common size of this ionized bubble, $r_{\min} = 4$ cMpc. We fix the upper integration limit to $r_{\max} = r_{\min} + 100$ cMpc since any neutral patches beyond this distance are down-weighted to such a high degree by the Lorentzian weighting kernel that they do not notably contribute to $N_{\text{HI}}^{\text{DW}}$ anymore. Lastly, we set the reference distance r_T with respect to which our summary statistic is defined, to $r_T = 18$ cMpc (corresponding to $v_T \simeq 2000$ km/s at $z_{\text{QSO}} = 7.54$), noting that we are not sensitive to this choice since the damping wing largely constitutes a one-parameter family, i.e., the imprint is so correlated that its value at one spectral pixel v_T closely determines its value at all other spectral pixels, as demonstrated in Kist et al. (2025c). As the parameter range for $N_{\text{HI}}^{\text{DW}}$ spans several orders of magnitude, we are in fact sensitive to the logarithmic quantity $\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}$, and for notational simplicity we henceforth adopt $\log N_{\text{HI}}^{\text{DW}}$ as short-hand notation for this.

We adopt as a second summary statistic the distance r_{patch} between the source and the first neutral patch in the *pre*-quasar topology. This quantity has been constrained by Mason et al. (2026) in the context of galaxy IGM damping wings, and has been shown in Kist et al. (2025c) to remain a meaningful summary for *quasar* IGM damping wings—despite the impact of the ionizing quasar radiation due to which the pre- and post-quasar distance to the first neutral patch do not necessarily agree (c.f. Schroeder et al. 2013, who inferred the post-quasar version of this distance). Most importantly, r_{patch} encapsulates information complementary to $\log N_{\text{HI}}^{\text{DW}}$ as demonstrated in Kist et al. (2025c).

¹Note that this is the convention we will adopt for all distances throughout this work.

4.2.2 Folding in the topology dependence and constraining the global IGM neutral fraction

Kist et al. (2025c) showed that the two local summary statistics $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} not only minimize the scatter of IGM transmission values (i.e., the amount of information that these statistics do not explain) in the damping wing region of the spectrum to $\lesssim 1\%$ across the entire range of physical parameter space, but that these two summaries are largely insensitive to the underlying reionization topology in the sense that the median and the 68-percentile scatter of the IGM transmission profiles at a given set of parameter values are identical, regardless of what reionization topology the underlying sightlines originate from. This implies that the specific distribution of neutral patches along a given sightline does not impact the damping wing shape, provided the $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} parameter values are fixed. All topological information is instead fully encoded in the statistical distribution of these statistics *within* a given topology. As a result, in a Bayesian sense, all assumptions about the ionization topology can be absorbed into the prior distribution imposed on $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$, while the IGM transmission likelihood, given $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} as model parameters, remains reionization model independent and can be determined by considering models of damping wing transmission profiles from *any* fiducial reference topology. We exploit this fact by constructing the likelihood based on sightlines generated according to the simplistic toy prescription introduced in Kist et al. (2025c) which augments the smoothness of the likelihood as a function of $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} , facilitating the practical task of sampling from the posterior distribution. The realistic topology only enters the analysis for determining the prior, and for converting the *local* $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ constraints into a *global* $\langle x_{\text{HI}} \rangle$ constraint. We do so by probabilistically relating global and local parameters to one another based on how the local parameters are distributed in a given global topology.

4.2.2.1 Simulating IGM transmission profiles

We simulate IGM transmission profiles following Davies et al. (2018a)’s hybrid approach of combining sightlines from cosmological hydrodynamical simulations with independent neutral fraction skewers and 1d radiative transfer. We extract 3600 density, velocity and temperature skewers originating at the 600 most massive halos with masses of $M_{\text{halo}} \geq 1.3 \times 10^{11} M_{\odot}$, and extending towards the six principal directions of the $z = 7.0$ snapshot of the Nyx hydrodynamical simulations (Almgren et al. 2013; Lukić et al. 2015). The simulation box measures 100 cMpc/h on a side and contains 4096^3 baryon and dark matter particles, respectively.

We simulate neutral fraction x_{HI} skewers in two different manners, 1) by extracting them from realistic semi-numerical reionization topologies, used

to determine our priors and convert our local constraints to a *global* $\langle x_{\text{HI}} \rangle$ constraint on the timing of reionization, or 2) by producing synthetic x_{HI} skewers according to a simple toy prescription which we use to determine the IGM transmission likelihood.

The former skewers originate from topologies generated using a modified version of the 21cmFAST code (Mesinger et al. 2011; Davies & Furlanetto 2022) at fixed global IGM neutral fractions of $\langle x_{\text{HI}} \rangle = 0.0, 0.05, \dots, 1.0$. Intermediate values between zero and one are achieved by tuning the ionizing efficiency ζ . From each simulation box (of size 400 cMpc on a 2048^3 initial and 512^3 output grid), we extract 20 randomly oriented skewers originating from each of the 500 most massive halos of masses $M_{\text{halo}} \geq 3 \times 10^{11} M_{\odot}$, giving a total of 10,000 x_{HI} skewers. We then combine each of our 3600 hydrodynamical Nyx sightlines with a random draw from these 10,000 neutral fraction skewers. By construction, these skewers vary smoothly as a function of the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$. We can also measure the local summary statistics $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} for each individual sightline,² and aggregate all sightlines according to these labels; however, this does not guarantee that these aggregated profiles vary smoothly as a function of $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} since the number of available sightlines can vary significantly in different regions of $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space.

Such small sample sizes have a particularly unfavorable impact on the noise level of the covariance matrices we have to estimate for determining the likelihood function (see Eq. (4.11)). Noisy covariance matrices can introduce discontinuities to this, and sampling from such a non-smooth distribution becomes practically impossible. We can avoid such issues by instead generating synthetic x_{HI} skewers according to the toy prescription introduced in Kist et al. (2025c) which by construction ensures smooth variations of the resulting skewers with respect to these local summaries. In short, we start from a given density sightline in its completely ionized state, and subsequently add in neutral patches of minimum size $\Delta r_{\text{min}} = 0.5$ cMpc, growing them out one-by-one to a maximum size $\Delta r_{\text{max}} = 5.0$ cMpc until a desired HI column density $\log N_{\text{HI}}^{\text{DW}}$ is reached. By definition, the position of the first neutral patch is set by r_{patch} , and the subsequent patches are added according to a fixed sequence of neutral patch locations for a given density sightline, ensuring continuity of the resulting skewers with respect to $\log N_{\text{HI}}^{\text{DW}}$. We emphasize that achieving continuity is the sole purpose of this particular approach, and that it does not seek to model any realistic physical processes occurring during reionization—in fact, Kist et al. (2025c) demonstrated that a physical reionization model is not necessary for our purposes of determining the IGM transmission likelihood, as the mean and

²Note that before determining r_{patch} , we smooth the x_{HI} field with a box-car filter of size 0.5 cMpc and define r_{patch} as the distance from the quasar where this smoothed field first exceeds a neutral fraction threshold of 50% to avoid being overly sensitive to extremely small-scale fluctuations in the x_{HI} field.

variance of the resulting IGM transmission profiles in our local damping wing parameterization only depend on the distribution of density fluctuations in the IGM but *not* on the specific distribution of neutral patches along the line of sight.

Further, it is important to note that the minimum and maximum column density values $\log N_{\text{HI}}^{\text{DW}}$ are sightline-dependent, set by the distribution of IGM density fluctuations along the line of sight for a completely neutral and a completely ionized sightline, respectively. At a given location in $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space, we can therefore end up with a number between 0 and 3600 sightlines, depending on how many sightlines allow for this specific parameter combination. We compute these synthetic x_{HI} skewers on an irregular $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter grid, reflecting the degree to which these parameters impact the resulting IGM transmission profiles. Specifically, our parameter grid consists of 21 column density values between $17.48 \leq \log N_{\text{HI}}^{\text{DW}} \leq 21.08$, and 18 neutral patch distances between $0.3 \text{ cMpc} \leq r_{\text{patch}} \leq 143.0 \text{ cMpc}$. A finer grid spacing is chosen towards higher column densities $\log N_{\text{HI}}^{\text{DW}}$ and shorter neutral patch distances r_{patch} where small parameter variations have a noticeable impact on the damping wing strength, whereas this is not necessary at low $\log N_{\text{HI}}^{\text{DW}}$ and high r_{patch} where no damping wing is present.

Finally, we combine hydrodynamical and neutral fraction skewers, adopting the Nyx temperature field for all ionized regions, and assuming an initially cold IGM with $T = 2000 \text{ K}$ for all neutral patches. We then perform one-dimensional radiative transfer along these sightlines to model the impact of the ionizing quasar radiation following [Davies et al. \(2016\)](#). Our model resembles the $z_{\text{QSO}} = 7.54$ quasar ULAS J1342+0928 with a simple light bulb light curve corresponding to an ionizing photon emission rate of $Q = 10^{57.14} \text{ s}^{-1}$ for quasar lifetimes t_{Q} on a logarithmic grid with 51 values between $t_{\text{Q}} = 10^3$ and 10^8 yr , and we henceforth adopt $\log t_{\text{Q}}$ as short-hand notation for $\log_{10} t_{\text{Q}}/\text{yr}$. To produce Lyman- α transmission profiles, we finally convolve the physical output fields with a Voigt profile ([Tepper-García 2006](#)).

As there is no star formation prescription in the Nyx simulations, strong proximate optically thick absorption line systems are not modeled realistically. This means we have to remove such sightlines from our full set of simulated IGM transmission profiles. This does not pose a major limitation to our approach as we do not aim to constrain the ionization state of the IGM based on targets where such strong absorption systems are present. As their absorption signature could easily be confused with the intergalactic one, this would introduce a significant additional modeling uncertainty when trying to disentangle it from the IGM damping wing, similar to what is required for galaxies ([Mason et al. 2026](#)). Observationally, such sightlines can be excluded a priori by identifying associated metal absorption lines in the spectrum of the source ([Davies et al. 2025](#)). In our simulated profiles, we identify

such sightlines by computing the HI column density within chunks of size 0.1 pMpc along the fully ionized realization of each sightline, and excluding all those sightlines containing at least one chunk with an (unweighted) HI column density of at least 10^{19} cm^{-2} within the first 5000 km/s from the source, leaving 2545 of the total of 3600 sightlines that can be used for the subsequent analysis.

We depict in Figure 4.1 the median profiles based on the synthetic x_{HI} skewers as a function of the three parameters of our local parameterization. From top to bottom, we show the variation with respect to $\log t_{\text{Q}}$, $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} , respectively. The parameters which are not varied in a given panel are fixed to the reference values of $t_{\text{Q}} = 10^6 \text{ yr}$, $\log N_{\text{HI}}^{\text{DW}} = 20.08$ and $r_{\text{patch}} = 8 \text{ cMpc}$ (black line in each panel). We observe the well-known trend of larger proximity zones for long-lived quasars due to the increasing size of the ionized bubble such objects have carved out around themselves (see e.g. Chen & Gnedin 2021). In line with this, the strength of the IGM damping wing decreases with increasing $\log t_{\text{Q}}$ due to the decreasing amount of neutral hydrogen along the line of sight. Similarly, the damping wing strength increases and the proximity zone size decreases with increasing HI column density $\log N_{\text{HI}}^{\text{DW}}$. Note, however, that this parameter captures the properties of the *pre*-quasar ionization topology, whereas t_{Q} encapsulates the effects of the quasar’s ionizing radiation.

The functional dependence of the profiles on the distance r_{patch} from the source to the first neutral patch is more complex. First, note that the median profiles in the two upper panels show a clear bump in transmission at $r_{\text{patch}} = 8 \text{ cMpc}$. This is because unlike the already ionized material closer to the quasar, this initially neutral patch at $r_{\text{patch}} = 8 \text{ cMpc}$ receives an additional amount of photoelectric heating when ionized by the quasar, leading to an enhancement in transmission in the corresponding region of the spectrum (c.f. Kist et al. 2025b). We can see in the bottom panel of Figure 4.1 that this position shifts with the distance the first neutral patch, with the $r_{\text{patch}} = 4 \text{ cMpc}$ profile showing a clear bump at that location instead. The remaining shape of the profiles depends highly non-trivially on the value of r_{patch} . For values of $0.3 \text{ cMpc} \leq r_{\text{patch}} < 4 \text{ cMpc}$, the first neutral patch is located *outside* the integration range of $\log N_{\text{HI}}^{\text{DW}}$ which starts at $r_{\text{min}} = 4 \text{ cMpc}$. As a result, r_{patch} is the only statistic capturing information about any *pre*-quasar neutral material in this region, and hence acts as a summary akin to $\log N_{\text{HI}}^{\text{DW}}$, with increasingly strong absorption for smaller values of r_{patch} .

This behavior turns around for $r_{\text{patch}} \geq 4 \text{ cMpc}$, where the damping wing shape is already captured to lowest order by $\log N_{\text{HI}}^{\text{DW}}$. Contrary to the case where $r_{\text{patch}} < 4 \text{ cMpc}$, we find a mild increase in the strength of the absorption signature with increasing r_{patch} . This is because we are concerned with *pre*-quasar statistics, and at fixed HI column density $\log N_{\text{HI}}^{\text{DW}}$, a neutral patch is more likely to be ionized away by the radiation of a quasar of a




Figure 4.1: Median IGM transmission profiles in our local damping wing parameterization. Variations with respect to t_Q , $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} are displayed in the upper, middle and lower panel, respectively. When fixed, these parameters are set to $t_Q = 10^6$ yr, $\log N_{\text{HI}}^{\text{DW}} = 20.08$ and $r_{\text{patch}} = 8$ cMpc (shown as black reference line in each panel). The dotted vertical line marks the spectral location corresponding to the reference choice of r_{patch} . All profiles are simulated following the prescription in Section 4.2.2.1 with synthetic x_{HI} profiles.

given lifetime $\log t_Q$ the closer it is located to the source, i.e., the smaller the value of r_{patch} . However, more nearby neutral patches are also the greatest contributors to the damping wing optical depth, and as such, the damping wing strength is lower the closer the first neutral patch was located to quasar, i.e., the more likely it was to get ionized away. When comparing to the upper two panels, however, we see that these effects are clearly subdominant to the impact of quasar lifetime $\log t_Q$ and HI column density $\log N_{\text{HI}}^{\text{DW}}$ on the IGM damping wing strength. However, the fact that the transmission profiles still change as a function of r_{patch} when $\log N_{\text{HI}}^{\text{DW}}$ and $\log t_Q$ are fixed shows that this parameter carries information complementary to that encapsulated by the other two parameters.

4.2.2.2 The prior on the local summary statistics induced by the reionization model

As demonstrated in Kist et al. (2025c), the mean and variance of the IGM transmission profiles at fixed $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter values do not depend on the reionization topology that the sightlines originate from. Instead, the topology exclusively determines *how common* a given parameter combination is. In other words, a statistical distribution on these parameters—be it prior or posterior—can directly be translated between the global $\langle x_{\text{HI}} \rangle$ parameterization and the local $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameterization.

When constraining the timing of reionization with quasar IGM damping wings, we do not want to impose any prior assumptions on the reionization state of the universe, i.e., a priori, we consider any global IGM neutral fraction value $\langle x_{\text{HI}} \rangle$ between 0 and 1 as equally likely. Formally speaking, we impose a flat prior $P(\langle x_{\text{HI}} \rangle) = \text{Unif}_{[0,1]}(\langle x_{\text{HI}} \rangle)$ on the global IGM neutral fraction, where $\text{Unif}_{[a,b]}(x)$ is a uniform distribution of the random variable x on the interval $[a, b]$.

But the assumed reionization topology determines the conditional distribution $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ of the local summaries given $\langle x_{\text{HI}} \rangle$, since clearly different local parameter configurations are differently common in different topologies with different global IGM neutral fractions, and henceforth we denote all distributions P_{top} that are affected by these topology assumptions with a corresponding subscript. Based on this, we can use the global prior $P(\langle x_{\text{HI}} \rangle)$ to write down the joint prior distribution on all three parameters:

$$\begin{aligned} P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) & \quad (4.5) \\ & = P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle) \times P(\langle x_{\text{HI}} \rangle). \end{aligned}$$

Marginalizing Eq. (4.5) over the global IGM neutral fraction, we see that the simple uniform prior on $\langle x_{\text{HI}} \rangle$ translates into a non-trivial, topology-informed prior on the local parameters $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$, given by

$$\begin{aligned} P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) \\ = \int d\langle x_{\text{HI}} \rangle P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle) \times P(\langle x_{\text{HI}} \rangle). \end{aligned} \quad (4.6)$$

It is this step where the assumptions about the reionization model affect our constraints: different reionization topologies can come with a different probabilistic mapping $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ between global and local labels. On the other hand, since the likelihood $L(\mathbf{t} | \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}, t_{\text{Q}})$ of the IGM transmission field \mathbf{t} is reionization model independent in our local parametrization, $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ is the distribution which encapsulates the *entire* topology dependence of our $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ constraints.

For the sake of quoting a bare set of local, topology-independent constraints on $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} , fully agnostic to the reionization model, we initially impose a constant prior distribution in two-dimensional $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space which we denote as $P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ (in contrast to the physical prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ resulting from Eq. (4.6)). This constant prior (marked by the dashed line of the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel of Figure 4.2) is only limited by the hard physical boundaries for these two quantities as we will discuss below. Based on this, we infer the topology-agnostic posterior distribution $P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ according to Bayes' theorem:

$$\begin{aligned} P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f}) &= L(\mathbf{f} | \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) \\ &\times P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) / P(\mathbf{f}). \end{aligned} \quad (4.7)$$

If we assume the full three-dimensional prior factorizes into $P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) = P(\log t_{\text{Q}}) \times P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$, folding in the topology dependence then simply amounts to a change of priors to the non-trivial prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ governed by the topology of interest as per Eq. (4.6):

$$\begin{aligned} P_{\text{top}}(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f}) \\ = \frac{P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})}{P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})} \times P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}), \end{aligned} \quad (4.8)$$

where the denominator is trivial since $P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ is constant by assumption, and it encloses the *entire* physical domain, implying that the support of any non-trivial prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ is guaranteed to be a subset of the support of $P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$.

We now proceed by practically determining $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ for the realistic semi-numerical reionization topology considered in this work.




Figure 4.2: Joint distribution $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ of the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ and the local summary statistics $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} , determined based on 21×2545 sightlines from the realistic semi-numerical reionization topology introduced in Section 4.2.2.1. Black dots denote individual sightline data which is arranged on a grid in the $\langle x_{\text{HI}} \rangle$ dimension since the sightlines were extracted from 21 distinct topologies with fixed global IGM neutral fractions between $0 \leq \langle x_{\text{HI}} \rangle \leq 1$. Contours correspond to the 68% and 95% percentile regions. The dashed line in the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ plane encloses the physically permitted domain for these parameters, determined based on simulations run with the synthetic prescription described in Section 4.2.2.1. Note that this boundary is constructed based on where we have at least 800 sightlines available to estimate means and covariances, and hence individual, rare realizations can lie outside of it.

Enforcing a uniform prior on $\langle x_{\text{HI}} \rangle$ is straightforward in the case at hand, as we are already provided with IGM transmission profiles at 21 regularly spaced $\langle x_{\text{HI}} \rangle$ parameter values between 0 and 1, well-representing a uniform prior distribution. By computing the local summary statistics $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} for all 21×2545 sightlines (i.e., 2545 distinct density sightlines combined with 21 different $\langle x_{\text{HI}} \rangle$ models each), and marginalizing over the underlying $\langle x_{\text{HI}} \rangle$ values, we are therefore immediately provided with a set of samples from the prior distribution $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ informed by this topology.

We show in Figure 4.2 a corner plot of the three-dimensional joint distribution $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ as given by Eq. (4.5) before performing the marginalization over $\langle x_{\text{HI}} \rangle$. For illustrational purposes, we depict 2d-marginals of a three-dimensional kernel density estimation (KDE) of the distribution rather than the actual samples, especially because of the discrete grid spacing in the $\langle x_{\text{HI}} \rangle$ dimension. Note that the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel of the plot can individually be understood as a contour plot of the prior distribution $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ given by Eq. (4.6) under the assumption of a flat prior on $\langle x_{\text{HI}} \rangle$.

For reference, the dashed line in the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel marks the two-dimensional constant prior $P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ which follows from our synthetic toy profiles used for constructing the IGM transmission likelihood. We immediately note that it covers a significantly more extended range of parameter space than the realistic ionization topology. This is because our toy prescription can produce profiles at *any* parameter combination that is not excluded physically. This particularly also includes sightlines that simultaneously show low HI column densities and short neutral patch distances, i.e. the bottom left region of the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel. To keep the column density low, such sightlines cannot contain many neutral patches in addition to the first one that is closest to the source demarcating the value of r_{patch} — a configuration that is rarely found in realistic reionization topologies due to the inside-out nature of reionization. On the other hand, high HI column densities with large neutral patch distances as would be seen in the top right corner of the panel are not found in *any* topology whatsoever. This is a hard physical constraint as the maximum HI column density, set by the number of sightlines which are *completely neutral* starting at r_{patch} (and, by definition, ionized at $r < r_{\text{patch}}$), clearly decreases as a function of r_{patch} . Note that even if all material further along the line of sight from r_{patch} is neutral, the density fluctuations limit the maximum value of $\log N_{\text{HI}}^{\text{DW}}$ that can be achieved. These maximum column density values are thus exclusively determined by the distribution of density fluctuations in the IGM.

In the same manner, we would expect the minimum column density to increase monotonically with decreasing r_{patch} , as these minimum values are set by those sightlines which are completely ionized except at a single, small neutral patch located at r_{patch} whose contribution to $\log N_{\text{HI}}^{\text{DW}}$ increases

according to the Lorentzian weighting with decreasing r_{patch} . The dashed line on the left-hand side of Figure 4.2 shows that this is indeed the case for all neutral patch distances $150 \text{ cMpc} \geq r_{\text{patch}} \geq 4 \text{ cMpc}$. Note that the sharp edges we see when following the dashed line from $r_{\text{patch}} = 150 \text{ cMpc}$ down to $r_{\text{patch}} = 4 \text{ cMpc}$ result from the discrete $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter grid whose spacing is chosen relatively coarsely in this region of parameter space.

However, this behavior turns around at $r_{\text{patch}} < 4 \text{ cMpc}$, where the minimum column density reduces back to $N_{\text{HI}}^{\text{DW}} \simeq 10^{18} \text{ cm}^{-2}$ (manifesting in the kink seen in the bottom left of the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel in Figure 4.2). This discontinuity arises from the fact that our lower integration limit r_{min} for the HI column density (see Eqs. (4.2) and (4.3)) is fixed to specifically this distance. As a result, any neutral material closer to the quasar than r_{min} will not contribute to $\log N_{\text{HI}}^{\text{DW}}$. To further understand the seemingly peculiar shape, now imagine two different scenarios: first, a sightline with $r_{\text{patch}} < r_{\text{min}}$ which contains a single neutral patch at the location r_{patch} closer to the quasar than r_{min} but which is ionized from r_{min} onward. Second, a sightline with $r_{\text{patch}} \geq r_{\text{min}}$ which only contains a neutral patch at the location r_{patch} more distant from the quasar than r_{min} . Then the neutral patch in the second scenario with $r_{\text{patch}} \geq r_{\text{min}}$ certainly contributes to $\log N_{\text{HI}}^{\text{DW}}$ while that of the first sightline does not since it is located outside of the integration range of $\log N_{\text{HI}}^{\text{DW}}$. Despite its small r_{patch} value, the sightline in the first scenario therefore has a smaller HI column density than the sightline in the second, whose r_{patch} value is higher but whose neutral patch inevitably contributes to $\log N_{\text{HI}}^{\text{DW}}$. This behavior generalizes as all sightlines with $r_{\text{patch}} \geq r_{\text{min}}$ by construction have at least one neutral patch contributing to $\log N_{\text{HI}}^{\text{DW}}$, and as a result, the minimum column density of sightlines with $r_{\text{patch}} < r_{\text{min}}$ is *smaller* than that of sightlines with $r_{\text{patch}} \geq r_{\text{min}}$, giving rise to the kink we observe at $r_{\text{patch}} = 4 \text{ cMpc}$ in Figure 4.2.

Note further that the realizations of the realistic topology (black dots in Figure 4.2) seemingly extend beyond the hard physical boundary enclosed by the dashed line. A number of effects conspire to cause this behavior. First, and most importantly, the dashed line respects the additional requirement of having available at least 800 toy profiles to estimate smooth covariance matrices for the IGM transmission likelihood at any given point in $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space, excluding certain configurations that are rare but not entirely impossible to achieve physically. Second, as already pointed out above, the dashed line is a linear interpolation across our simulation grid which we chose rather coarsely towards lower $\log N_{\text{HI}}^{\text{DW}}$ and higher r_{patch} values where the transmission profiles cease to vary with these parameters. Third, this behavior looks even more drastic for the probability contours shown in this figure where smoothing effects make the distribution appear wider than it in reality is.

Aside from the limited support in the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ plane, the shape of the realistic prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ differs significantly from a uniform




Figure 4.3: Conditional distribution $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ of the two local summary statistics $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} given the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$, determined based on 2545 sightlines from the realistic semi-numerical reionization topology introduced in Section 4.2.2.1 at six different global IGM neutral fractions between $0 \leq \langle x_{\text{HI}} \rangle \leq 1$. Dots denote individual sightline data, and contours the 68% scatter. For visibility, we smoothed the one-dimensional marginal distributions with a kernel density estimation (KDE). The dashed line in the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ plane encloses the physically permitted domain for these parameters, determined based on the simulations run with the synthetic prescription described in Section 4.2.2.1.

distribution. We observe a bimodal distribution with a more pronounced peak towards higher HI column densities of $N_{\text{HI}}^{\text{DW}} \gtrsim 10^{20} \text{ cm}^{-2}$ and smaller neutral patch distances $r_{\text{patch}} \lesssim 20 \text{ cMpc}$, and a lower peak at low column densities $N_{\text{HI}}^{\text{DW}} \lesssim 10^{18.5} \text{ cm}^{-2}$ with large distances $r_{\text{patch}} \approx 143 \text{ cMpc}$ to the first neutral patch. The axis of degeneracy between the peaks naturally arises due to the fact that by definition, sightlines with a more nearby neutral patch can accommodate higher HI column densities. However, the fact that this is not a perfect one-to-one relation implies that our second summary statistic r_{patch} captures additional information not contained in $\log N_{\text{HI}}^{\text{DW}}$.

Comparing to Figure 4.3 where we show the *conditional* distribution $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ of the two local summaries $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} given six different values of the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$, we immediately see that the former peak is due to sightlines originating from more neutral topologies, while the latter one corresponds to the most ionized

ones. An axis of degeneracy connecting the two peaks is also apparent in these panels; however, there is significant scatter about this axis due to the fact that the distribution of neutral patches along a given line of sight can vary significantly even at fixed global IGM neutral fraction $\langle x_{\text{HI}} \rangle$. It is this amount of scatter which we exclude from the primary inference task by adopting our local parameterization rather than the global one.

Note that the high- r_{patch} peak is artificial in the sense that our simulated skewers have a finite length of 100 cMpc/h. Strictly speaking, r_{patch} is ill-defined for all skewers that do not contain *any* neutral patch along their entire range. In such cases, we set r_{patch} to the maximum distance of 100 cMpc/h $\simeq 143$ cMpc, noting that this value really only constitutes a lower limit for the actual distance to the first neutral patch.³ However, due to the Lorentzian decline of the Lyman- α cross section σ_{α} , such distant neutral patches hardly leave any imprint on the observed IGM transmission profile anyways, and therefore this artificial bound will not bias our constraints.

4.2.2.3 Converting local measurements to a global $\langle x_{\text{HI}} \rangle$ constraint

Due to their local nature, our two summary statistics encapsulate information not only about the global timing of reionization, but also the local ionization topology itself. The former bit is conventionally quoted in terms of the global volume-averaged IGM neutral fraction $\langle x_{\text{HI}} \rangle$. Here we elaborate on how such a global constraint can easily be obtained from our local summary statistics by again invoking the stochastic relation $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ between global and local parameters. Apart from setting the prior, the topology dependence therefore naturally comes in a second time when converting the local $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ constraints to a global constraint on $\langle x_{\text{HI}} \rangle$.

Specifically, let us assume that, given an observed quasar spectrum \mathbf{f} , we inferred the two local summaries $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ as well as the lifetime of the quasar $\log t_{\text{Q}}$, i.e., we obtained samples from the local, topology-agnostic posterior distribution $P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ via Eq. (4.7), but we are interested in the full topology-informed posterior distribution $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ of both global and local parameters. By decomposing the latter via Bayes' theorem, we can see that

$$\begin{aligned} & P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f}) \\ &= L_{\text{top}}(\mathbf{f} | \langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) \\ &\quad \times P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) / P(\mathbf{f}). \end{aligned} \tag{4.9}$$

We can simplify this expression by harnessing the topology-independence of the IGM transmission likelihood in our local parameterization, which implies $L_{\text{top}}(\mathbf{f} | \langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) = L(\mathbf{f} | \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$

³Technically speaking, we are thus inferring the parameter $r_{\text{patch}} \equiv \min(r_{\text{patch}}^{\text{phys}}, 143 \text{ cMpc})$, where $r_{\text{patch}}^{\text{phys}}$ is the true physical distance to the first neutral patch.

as the shape of the transmission profiles is not affected by what topology of what global neutral fraction $\langle x_{\text{HI}} \rangle$ they originate from, provided that the values of the local summary statistics $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} are fixed. By further assuming an independent lifetime prior $P(\log t_{\text{Q}})$ such that $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) = P(\log t_{\text{Q}}) \times P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$, and decomposing the remaining joint distribution $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ as per Eq. (4.5), we can straightforwardly combine Eqs. (4.7) and (4.9) to replace the likelihood $L(\mathbf{f} | \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ with the topology-agnostic posterior $P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ and arrive at

$$P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f}) = \frac{P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})}{P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})} \times P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle) \times P(\langle x_{\text{HI}} \rangle). \quad (4.10)$$

Note that the denominator is trivial since it only consists of the constant topology-agnostic prior $P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ that covers the entire physical domain (dashed line in Figures 4.2 and 4.3). In essence, folding in the topology dependence to obtain the full four-dimensional topology-informed posterior distribution $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ is thus simply a matter of multiplying together the local, topology-agnostic posterior $P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ that we originally inferred, and the conditional probability distribution $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ which encapsulates *all* information about the assumed ionization topology, along with the prior $P(\langle x_{\text{HI}} \rangle)$. In Eq. (4.10), we are therefore 1) imposing a non-trivial, physically motivated prior distribution on our local parameters, and 2) obtaining a constraint on the global timing of reionization.

In practice, we have to keep in mind that we are only provided with *samples* from the two aforementioned distributions. In order to perform the multiplication in Eq. (4.10), we introduce a parameter grid to evaluate the distributions on. We choose the $(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}})$ grid spacing in accordance with our parameter grid for the IGM transmission profiles discussed in Section 4.2.2.1, and likewise for $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ based on the irregular parameter grid on which our toy transmission profiles used to determine the IGM transmission likelihood are simulated.⁴ We then represent the two distributions through histograms with bins defined by these grids. As soon as $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ is determined according to Eq. (4.10), we can arbitrarily marginalize the distribution via numerical integration over our parameter grid in order to obtain, e.g., the conventionally quoted posterior $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}} | \mathbf{f})$ on the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$

⁴Adopting this irregular binning ensures that we make use of the enhanced sensitivity for profiles with strong damping wings, i.e. at high $\log N_{\text{HI}}^{\text{DW}}$ and low r_{patch} , where we chose a finer grid spacing. We formally investigate the precision of our measurements as a function of astrophysical parameter space in Section 4.4.3.

and quasar lifetime $\log t_Q$, or the associated one-dimensional marginals $P_{\text{top}}(\langle x_{\text{HI}} \rangle | \mathbf{f})$ or $P_{\text{top}}(\log t_Q | \mathbf{f})$.

4.3 Methods: inference pipeline

We described in the previous section a set of summary statistics that encodes the local information encapsulated in the IGM damping wing imprint in a topology-independent fashion. We now proceed by introducing an inference framework that can be used to constrain these parameters based on observed high-redshift quasar spectra. The formalism is based on the one established in Hennawi et al. (2025) and Kist et al. (2025b) as a fully Bayesian framework to infer the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ and the lifetime $\log t_Q$ of a quasar based on its observed spectrum \mathbf{f} . We start this section with a short summary providing an overview over all major modeling components, focusing especially on the adaptations made to constrain our local summary statistics rather than directly inferring the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$. For specific details, we point the reader to Hennawi et al. (2025).

The damping wing signature is imprinted upon the spectra of high-redshift sources (such as quasars in our case) by the foreground IGM. In order to extract the information it encapsulates about the global—or local—ionization topology, we therefore first have to disentangle this imprint from the unabsorbed, intrinsic spectrum of the source. A two-step method has been the conventional approach to address this task (see e.g. Greig et al. 2017a; Davies et al. 2018a): first, the intrinsic continuum of the quasar is reconstructed based on its correlation with the emission lines in the unabsorbed region of the spectrum redward of the Lyman- α line. *Subsequently*, the resulting continuum-normalized spectrum can be used to draw conclusions about the ionization state of the surrounding IGM. Hennawi et al. (2025) introduced for the first time a fully Bayesian framework that *jointly* performs these two tasks while operating on the entire spectral range, thus accounting for the full covariance resulting from the continuum reconstruction and the IGM transmission stochastic process.

The heart of the framework is a full generative model for high-redshift quasar spectra \mathbf{f} . Starting from a quasar continuum \mathbf{s} , based on a dataset of ~ 45000 unabsorbed, low-redshift ($1.878 < z < 3.427$) continua, we fold in the IGM absorption imprint based on simulated IGM transmission profiles \mathbf{t} (see Section 4.2.2.1), parameterized as a function of $\boldsymbol{\theta} \equiv (\log t_Q, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$. After forward-modeling instrumental effects by convolving these profiles with the line-spread function of a hypothetical spectrograph, as well as adding a realistic, heteroscedastic spectral noise vector $\boldsymbol{\sigma}$, we are equipped with a realistic forward model for high-redshift quasar spectra \mathbf{f} .

To tackle the inverse problem of constraining the astrophysical parameters $\boldsymbol{\theta}$ based on an observed quasar spectrum \mathbf{f} , we construct a data-

driven low-dimensional parametric PCA model for the quasar continuum based on the aforementioned set of low-redshift spectra. Together with the simulated IGM transmission profiles \mathbf{t} , we use this to determine the full likelihood $L(\mathbf{f}|\boldsymbol{\theta}, \boldsymbol{\eta}, \boldsymbol{\sigma})$ of the observed quasar spectrum \mathbf{f} as a function of the low-dimensional vector $\boldsymbol{\eta}$ of (latent) PCA coefficients describing the full continuum \mathbf{s} and the astrophysical parameter vector $\boldsymbol{\theta}$, conditioned on the observational noise vector $\boldsymbol{\sigma}$. Hennawi et al. (2025) demonstrated that we can approximate this likelihood as

$$L(\mathbf{f}|\boldsymbol{\sigma}, \boldsymbol{\theta}, \boldsymbol{\eta}) = \mathcal{N}(\mathbf{f}; \langle \mathbf{t} \rangle \circ \langle \mathbf{s} \rangle, \boldsymbol{\Sigma} + \langle \mathbf{S} \rangle \mathbf{C}_t \langle \mathbf{S} \rangle + \langle \mathbf{T} \rangle \mathbf{C}_s \langle \mathbf{T} \rangle), \quad (4.11)$$

where $\mathcal{N}(\mathbf{f}; \boldsymbol{\mu}, \mathbf{K})$ is the multivariate normal distribution of the random variable \mathbf{f} with mean $\boldsymbol{\mu}$ and covariance matrix \mathbf{K} . Further, $\mathbf{t} \circ \mathbf{s}$ denotes the element-wise (Hadamard) product of the two mean vectors $\langle \mathbf{t} \rangle$ and $\langle \mathbf{s} \rangle$, and we defined the matrices $\boldsymbol{\Sigma} \equiv \text{diag}(\boldsymbol{\sigma})$, $\mathbf{T} \equiv \text{diag}(\mathbf{t})$, $\mathbf{S} \equiv \text{diag}(\mathbf{s})$, as well as the covariance matrices \mathbf{C}_t and \mathbf{C}_s of \mathbf{t} and \mathbf{s} , respectively.

Note that this likelihood operates on the entire spectral range, both redward and blueward of the Lyman- α line, and hence covers the smooth IGM damping wing as well as the Lyman- α forest region with the quasar proximity zone. This allows us to *jointly* infer the astrophysical parameters $\boldsymbol{\theta}$ and the continuum nuisance parameters $\boldsymbol{\eta}$.

4.3.1 Continuum dimensionality reduction model

We now proceed with a short summary of our parametric model for the quasar continuum. In short, we are using an updated version of the principal component analysis (PCA) model described in Hennawi et al. (2025) which additionally accounts for luminosity variations in the spectra. This allows us to construct the PCA decomposition based on a ~ 3 times larger dataset of low-redshift continua, and extend the spectral coverage redward of Lyman- α out to the Mg II line. A detailed description of this new model will be provided in Hennawi et al. (2026).

In short, the dataset we use comprises 44,587 low-redshift ($1.878 < z < 3.427$) spectra from the SDSS-III Baryon Oscillation Spectroscopic Survey (BOSS) and SDSS-IV Extended BOSS (eBOSS). These spectra cover a rest-frame wavelength range of 1175 – 3000 Å and have a resolution of $R \sim 2000$ with a median signal-to-noise ratio of $S/N > 10$ within a 5 Å region around the rest-frame wavelength 1285 Å. For a small subset of these sources, we also obtained near-infrared spectra from the Gemini Near-Infrared Spectrograph Distant Quasar Survey (GNIRS-DQS; Matthews et al. 2021). The spectra were taken with the Gemini Near-Infrared Spectrograph (GNIRS; Elias et al. 2006) at the Gemini North Observatory. These spectra cover the $\sim 0.8\text{--}2.5\mu\text{m}$ range to encompass the H β and [O III] region. They have a resolution of $R \sim 1100$ and were obtained to reach sensitivities comparable to

the SDSS spectra at $\lambda_{\text{obs}} \sim 5000 \text{ \AA}$. All spectroscopic data from GNIRS-DQS were re-reduced with PypeIt (Prochaska et al. 2020).

We compute the PCA decomposition based on 42,854 (i.e., $\sim 96\%$) of the full set of continua and keep the remaining 1733 apart as test set for estimating the reconstruction error and drawing mock continua. The PCA decomposition is computed based on the logarithmic continua, and we keep the 7 PCA vectors accounting for the largest variance as basis vectors of our model. Note that since we perform the decomposition in log-space, the first PCA vector captures the normalization of the continua, and hence the luminosity dependence of the spectral shape is intrinsically accounted for. Note further that we opted for a weighted PCA that allows us to also include continua with missing pixels. For convenience, however, we chose our test set such that it exclusively contains continua with full spectral coverage. This provides us with a dimensionality-reduced description $\mathbf{s}_{\text{DR}}(\boldsymbol{\eta})$ of the true continuum \mathbf{s} with the model parameters corresponding to the seven PCA coefficients $\boldsymbol{\eta}$.⁵

Hennawi et al. (2025) showed that the associated relative continuum reconstruction error $\boldsymbol{\delta} = (\mathbf{s} - \mathbf{s}_{\text{DR}})/\mathbf{s}$ is well-approximated by a Gaussian distribution, and, as a result,

$$P(\mathbf{s}|\boldsymbol{\eta}) = \mathcal{N}(\mathbf{s}; \langle \mathbf{s}(\boldsymbol{\eta}) \rangle, \mathbf{C}_{\mathbf{s}}(\boldsymbol{\eta})). \quad (4.12)$$

We estimate the mean $\langle \mathbf{s}(\boldsymbol{\eta}) \rangle \equiv \mathbf{s}_{\text{DR}}(\boldsymbol{\eta}) \circ (1 + \langle \boldsymbol{\delta} \rangle)$ and the covariance $\mathbf{C}_{\mathbf{s}}(\boldsymbol{\eta}) \equiv \text{diag}(\mathbf{s}_{\text{DR}}(\boldsymbol{\eta})) \mathbf{C}_{\boldsymbol{\delta}} \text{diag}(\mathbf{s}_{\text{DR}}(\boldsymbol{\eta}))$ of this distribution based on the test set of 1733 continua, where $\langle \boldsymbol{\delta} \rangle$ and $\mathbf{C}_{\boldsymbol{\delta}}$ are the mean and covariance of the continuum reconstruction error $\boldsymbol{\delta}$.

4.3.2 IGM transmission likelihood

All patches of neutral hydrogen in the surrounding IGM imprint an absorption signature on the continuum of the quasar. This is represented by IGM transmission profiles \mathbf{t} which we obtain from numerical simulations (see Section 4.2.2.1), parameterized by a set of astrophysical parameters $\boldsymbol{\theta}$. In the conventional approach, one adopts $\boldsymbol{\theta} \equiv (\langle x_{\text{HI}} \rangle, \log t_{\text{Q}})$ which we henceforth refer to as the *global* parameterization, whereas we here propose to consider the *local* parameterization $\boldsymbol{\theta} \equiv (\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$.

⁵As demonstrated in Kist et al. (2025b), a five-dimensional PCA model plus normalization factor is sufficiently flexible to represent the full quasar continuum \mathbf{s} for the purpose of astrophysical parameter inference based on IGM damping wings. As we perform the PCA in log-space without the need for an additional normalization factor, and due to the extended red-side spectral coverage of 3000 \AA (as compared to 2000 \AA in the analysis by Kist et al. 2025b), we increase the dimensionality of our PCA model to seven.

In either case, to go forward, we follow [Hennawi et al. \(2025\)](#) and approximate the distribution of \mathbf{t} given the astrophysical parameters $\boldsymbol{\theta}$ as a multivariate Gaussian distribution:

$$P(\mathbf{t}|\boldsymbol{\theta}) = \mathcal{N}(\mathbf{t}; \langle \mathbf{t}(\boldsymbol{\theta}) \rangle, \mathbf{C}_{\mathbf{t}}(\boldsymbol{\theta})). \quad (4.13)$$

It is a well-known fact that the assumption of Gaussianity is in fact not fully justified when using $\langle x_{\text{HI}} \rangle$ as label ([Lee et al. 2015](#); [Davies et al. 2018a](#)), and neither in the case where $\langle x_{\text{HI}} \rangle$ gets replaced with our local summary statistics as we find in this analysis. The approximation, however, is vital to our formalism as it allows us to derive a closed-form analytical expression for the likelihood of the full quasar spectrum. To make sure that this does not result in an overestimation of our constraining power, we apply a principled procedure to broaden the posteriors by the degree required to ensure statistically faithful constraints (see also [Hennawi et al. 2025](#)). In future, we aim to address these issues while retaining the full constraining power by learning the true non-Gaussian shape of $P(\mathbf{t}|\boldsymbol{\theta})$ with the help of simulation-based inference.

As far as this work is concerned, we simply estimate the mean transmission $\langle \mathbf{t}(\boldsymbol{\theta}) \rangle$ and the covariance $\mathbf{C}_{\mathbf{t}}(\boldsymbol{\theta})$ based on our simulated IGM transmission profiles. Note that sampling from the posterior distribution via Hamiltonian Monte Carlo (HMC) becomes challenging in cases where the posterior distribution is not a smooth function of the parameters $\boldsymbol{\theta}$ as this can prevent sufficient mixing of the HMC chains. The most critical part of our framework are the covariance matrices $\mathbf{C}_{\mathbf{t}}(\boldsymbol{\theta})$ which are prone to becoming noisy if estimated based on too low a number of sightlines. This potential pitfall is addressed by our synthetic procedure for generating x_{HI} sightlines that vary smoothly as a function of $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} , as described in Section 4.2.2.1 (see also [Kist et al. 2025c](#)). This allows us to measure means and covariances based on a large number of sightlines at any desired location in $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space with continuity with respect to those parameters intrinsically built in. An additional virtue of this procedure is that it allows us to cover the *entire* physically accessible domain in $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space with a sufficient number of simulated sightlines for estimating covariance matrices.

Note also that before computing $\langle \mathbf{t}(\boldsymbol{\theta}) \rangle$ and $\mathbf{C}_{\mathbf{t}}(\boldsymbol{\theta})$, we forward-model instrumental effects for each individual transmission profile by convolving it with a Gaussian line-spread function (LSF) of width FWHM = 100 km/s and rebinning it onto a coarse velocity grid with a pixel spacing of 500 km/s, covering a rest-frame wavelength range of 1175–3000 Å. Even though strictly speaking this should not be done before multiplying in the quasar continuum, we perform these operations in *pre-processing* on the bare IGM transmission profiles, as doing this downstream would make the inference computationally infeasible without recognizable precision gains.

4.3.3 Inference procedure

With a full framework at hand for the likelihood $L(\mathbf{f}|\boldsymbol{\sigma}, \boldsymbol{\theta}, \boldsymbol{\eta})$ of an observed quasar spectrum \mathbf{f} given astrophysical parameters $\boldsymbol{\theta}$ and continuum nuisance parameters $\boldsymbol{\eta}$, we are now in the position to sample from the associated posterior distribution once we specified our priors. We adopt a uniform prior on the logarithmic quasar lifetime $P(\log t_Q) = \text{Unif}(3, 8)$, covering a wide range of physically plausible values (Khrykin et al. 2021). In the context of our local parameter framework, we impose a two-dimensional constant prior $P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ on our local summary statistics, enclosed by the (non-trivial) boundaries set by the distribution of IGM density fluctuations as discussed in Section 4.2.2.2. For reference, we also perform the inference in the conventional, global $(\langle x_{\text{HI}} \rangle, \log t_Q)$ framework, in which case we impose an uninformative, uniform prior on the IGM neutral fraction $P(\langle x_{\text{HI}} \rangle) = \text{Unif}(0, 1)$ in addition to the aforementioned log-uniform lifetime prior. In practice, to aid the sampling procedure, we apply sigmoid transformations to all bounded parameters to make them fully unbounded. To account for the non-trivial two-dimensional prior boundary of $P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$, we first transform r_{patch} , and subsequently $\log N_{\text{HI}}^{\text{DW}}$, conditioned on r_{patch} . This is required since the minimum and maximum HI column densities $\log N_{\text{HI}}^{\text{DW}}$ change as a function of r_{patch} . The PCA coefficients η_i parameterizing the shape of the quasar continuum are already allowed to vary over the entire real axis $P(\eta_i) = \text{Unif}(-\infty, \infty)$, so here we only remove their mean and rescale them by their standard deviation.

We then deploy Hamiltonian Monte-Carlo (HMC) to sample from the resulting posterior distribution, using the HMC implementation with a No U-Turn Sampler (NUTS) from the NumPyro probabilistic programming library based on the machine learning and autograd framework JAX (Bradbury et al. 2018; Bingham et al. 2018; Phan et al. 2019). For a given spectrum, we run four HMC chains with 1000 warm-up steps each, and an additional 1000 steps for sampling. Furthermore, we perform sigma-clipping on the spectrum before performing the inference itself in order to make sure our results do not get biased by a small number of outlier pixels. To that end, we start by iteratively fitting for the maximum-likelihood model with the help of an Adam optimizer with a learning rate of 0.5 and 1000 optimization steps. We compute the χ^2 statistic between the data and this best-fit model, masking up to five $> 4\sigma$ outlier pixels on the red side of the spectrum ($\lambda_{\text{rest}} > 1260 \text{ \AA}$) per iteration. We repeat the same optimizing and clipping procedure on the remaining unmasked part of the spectrum for up to 5 iterations before we start sampling from the posterior distribution via HMC.

4.3.4 Generating mock spectra

To verify the performance of our pipeline, we apply it to mock spectra generated in the following way. We start by drawing an unabsorbed continuum \mathbf{s} from the test set of 1733 low-redshift spectra used in Section 4.3.1 to estimate the continuum reconstruction error. Given a desired astrophysical parameter vector $\boldsymbol{\theta}$, we further draw an IGM transmission profile \mathbf{t} from the simulated profiles described in Section 4.2.2.1, with instrumental effects folded in as discussed in Section 4.3.2 via convolution with a Gaussian LSF of width FWHM = 100 km/s and rebinning to a 500 km/s velocity grid. We interpolate the continuum \mathbf{s} onto the same wavelength grid (extending from 1175 Å to 3000 Å), and multiply it together with the IGM transmission profile \mathbf{t} . Finally, we generate realistic heteroscedastic noise realizations including telluric absorption as well as contributions from object photons, sky background and detector read noise with help of the `SkyCalc_ipy` package (Leschinski 2021); see Hennawi et al. (2025) for details. The exposure time of the hypothetical instrument is adjusted such that the median signal-to-noise ratio per 100 km/s velocity interval is $S/N = 10$ within a 5 Å interval centered at 1285 Å in the rest-frame. Two different ensembles of such mock spectra are considered in this work:

1. A set of 10×10 mock spectra on a parameter grid spanned by the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ and the quasar lifetime $\log t_{\text{Q}}$ with values of $\langle x_{\text{HI}} \rangle = 0.05, 0.15, \dots, 0.95$ and $\log t_{\text{Q}} = 3.25, 3.75, \dots, 7.75$. The underlying IGM transmission profiles originate from the realistic semi-numerical reionization topology described in Section 4.2.2.1. For each sightline, we can also uniquely determine the values of our two local pre-quasar summary statistics $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} . By construction of this sample, the distribution of these statistics follows the topology-informed prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$. We run both the global and the local version of our inference pipeline on these spectra, inferring either the global parameters $(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}})$ or the local ones $(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$. We use this mock ensemble to compare the fits to these spectra in the context of the two different parameterizations, to perform coverage tests, and to compute the corrections required to pass them, as will be discussed in the subsequent section.
2. A set of 360 mock spectra on a $3 \times 4 \times 4$ parameter grid in the local $(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameterization with values of $\log t_{\text{Q}} = 4, 6, 8$, $\log N_{\text{HI}}^{\text{DW}} = 20.1, 20.3, 20.5, 20.7$ and $r_{\text{patch}} = 2, 6, 10, 14$ cMpc. At each location in parameter space, we simulate 10 distinct mocks, and the keep the continuum and noise realization of these 10 mock spectra fixed when varying $\log t_{\text{Q}}$, $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} to isolate the impact of these three parameters from those additional sources of stochasticity. Note that realizations of IGM transmission profiles only exist at 12 out

of the 16 locations in $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space since some parameter combinations are physically excluded by the distribution of density fluctuations in the IGM (see e.g. Figure 4.2), and hence we end up with 360 rather than 480 spectra. The underlying IGM transmission profiles originate from synthetic neutral fraction sightlines generated according to our analytical prescription described in Section 4.2.2.1. As such, they do not have an associated global IGM neutral fraction value $\langle x_{\text{HI}} \rangle$. We therefore only run the local version of our inference pipeline on these spectra to determine the precision with which we can infer $\log t_{\text{Q}}$, $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} .

4.3.5 Coverage tests

A reliable inference framework guarantees that the inferred posterior distribution is statistically faithful. For example, when inferring parameters from 100 different mock spectra, and considering the 68% highest-density region (HDR) of each inferred posterior, we statistically expect the true value to be contained in this 68% HDR for 68 of these 100 mocks, and likewise for any other credibility level α . More generally speaking, given a set of N_{mock} mock inferences, we can determine the *expected coverage probability* C_α corresponding to the *credibility level* α as the number of mocks whose true parameter value is indeed contained in the α -th HDR of the posterior distribution, divided by the total number of mocks in the sample.

We can perform a *coverage test* by explicitly computing C_α for values of α in the entire range $\alpha \in [0, 1]$. This test is passed if we find $C_\alpha = \alpha$ for all $\alpha \in [0, 1]$, implying that the inferred posteriors contain the truth exactly as many times as they are statistically expected to. On the other hand, finding $C_\alpha > \alpha$ indicates that the inferred posteriors are *underconfident*, i.e., less constraining than they could be. *Overconfident*, or too narrow posteriors where $C_\alpha < \alpha$ pose an even larger problem as they can lead us to draw conclusions that are not actually backed by the data. This is caused by flaws in the inference pipeline, such as bugs, inappropriate priors, or, as in the case at hand, an approximate likelihood prescription.

Eliminating such imperfections is essential before quoting any physical parameter constraints. While it would exceed the scope of this work to restructure our pipeline with the help of simulation-based inference to avoid the approximate analytical likelihood expression in Eq. (4.11), Hennawi et al. (2025) introduced a principled way of retrospectively broadening the inferred posteriors in such a way that they are guaranteed to pass a coverage test. In essence, the idea is to relabel the α -th HDR of a given posterior distribution as the C_α -th HDR according to the mapping $\alpha \mapsto C_\alpha$ determined in the coverage test. Doing so for all credibility levels α by construction ensures perfect coverage for the new coverage-corrected posterior distribution. While Hennawi et al. (2025) outlined the practical procedure for the case where

we are provided with MCMC samples from the posterior distribution, we here elaborate on an analogous strategy applicable in the case where the posterior is evaluated on a parameter grid such as the one we are faced with when converting our local parameter constraints to global ones.

Specifically, suppose that the full parameter space spanned by $\boldsymbol{\theta}$ is divided into grid cells i with central value $\boldsymbol{\theta}_i$ and volume V_i . Then the probability mass contained in the i -th pixel is

$$P_i \equiv \int_{V_i} d\boldsymbol{\theta} P(\boldsymbol{\theta}) \simeq V_i \cdot P(\boldsymbol{\theta}_i), \quad (4.14)$$

where for clarity we suppressed all other variables in the argument of the posterior distribution $P(\boldsymbol{\theta})$. To determine the highest-density regions, we can now order all grid cells according to the probability mass they contain, resulting in a permutation $\pi(i)$ of the original grid cells i , with $\pi^{-1}(j)$ mapping back the sorted grid cells j to their original index. We then define the α_N -th HDR as the region enclosed by the N highest-probability cells, i.e.,

$$\sum_{j=1}^N P_{\pi^{-1}(j)} = \alpha_N. \quad (4.15)$$

Note that this can be seen as the cumulative distribution function of the reordered cells. This also underlines that summing over all N_{cell} cells, we obtain $\sum_{j=1}^{N_{\text{cell}}} P_{\pi^{-1}(j)} = \alpha_{N_{\text{cell}}} = 1$.

To perform a coverage test, we can now easily determine for each mock the index N_0 as the smallest index N for which the true parameter is still contained in its α_N -th HDR. The coverage probability corresponding to the credibility level α is then simply given by the relative number of objects whose value of α_{N_0} (i.e., the size of the smallest HDR that still contains the truth) is less or equal to α . In other words, by rank-ordering the N_{mock} values of α_{N_0} , and dividing their rank by the total number of mocks N_{mock} , we are immediately provided with the corresponding coverage probabilities $C_{\alpha_{N_0}}$. By interpolating C_{α} among intermediate α values, we can obtain a smooth approximation of the coverage curve across the entire range $\alpha \in [0, 1]$.

Once this curve is determined, we can also make use of it to correct for any potential imperfections in the coverage behavior. This can be straightforwardly accomplished in a backward approach where we again loop over all grid cells N ordered by their probability mass $P_{\pi^{-1}(N)}$, and instead assign them the *correct* mass $P_{\pi^{-1}(N)}^{\text{rew}}$ dictated by their expected coverage probability C_{α_N} . Specifically, when considering the N -th cell, we have to set

$$P_{\pi^{-1}(N)}^{\text{rew}} \equiv C_{\alpha_N} - \sum_{j=1}^{N-1} P_{\pi^{-1}(j)}^{\text{rew}}. \quad (4.16)$$

Starting at the cell $N = 1$ containing the highest probability mass, and subsequently moving to lower-probability cells j , this uniquely defines the

values of the coverage-corrected posterior distribution $P^{\text{rew}}(\boldsymbol{\theta})$ at each grid cell.

4.4 Results: global vs local parameter inference

Equipped with our new local parameterization of quasar IGM damping wings (Section 4.2), as well as an inference framework to constrain them (Section 4.3), we now proceed by testing the statistical fidelity of this framework on large ensembles of mock spectra and quantifying the precision with which we can constrain both the local and global astrophysical parameters. Importantly, we also demonstrate how we can relate our local constraints back to the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$, obtaining the full four-dimensional, topology-informed posterior $P(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ on both global and local parameters.

4.4.1 Inference for a single mock

We start with an illustration of our approach by applying it to the mock spectrum of a quasar which has been shining for $t_{\text{Q}} = 10^{6.25}$ yr, observed at the mid-stages of reionization (at $\langle x_{\text{HI}} \rangle = 0.55$). The underlying IGM transmission profile is based on a sightline from the semi-numerical reionization topology corresponding to that global neutral fraction value, and we measure $N_{\text{HI}}^{\text{DW}} = 10^{20.25} \text{ cm}^{-2}$ and $r_{\text{patch}} = 12.9 \text{ cMpc}$ for the local pre-quasar summary statistics of this sightline, corresponding to a typical sightline in an intermediately neutral IGM (c.f. Figure 4.3). We feed in the same mock spectrum to both our local parameter inference framework introduced in Section 4.3, and, for comparison, the global version of the pipeline that directly infers the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ rather than the local summaries $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} . Figure 4.4 depicts the mock spectrum with/without noise in black/purple, with the underlying continuum and noise vector shown in green and yellow, respectively. The upper/lower panels depict the reconstructed spectrum (red) and continuum (blue) inferred in the context of the local/global parameterization. While the median inferred models are remarkably similar, the 68% scatter (marked by the shaded regions), reflecting parameter uncertainty, continuum reconstruction errors, as well as spectral noise of the inferred model spectrum is significantly smaller in the local parameterization as compared to the global one. This is a direct consequence of the fact that the stochasticity of reionization is an intrinsic part of the model in the global parameterization, but has been removed in the local one as demonstrated in Kist et al. (2025c).

The black contours in the corner plot in Figure 4.5 depict the coverage corrected local ($\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}$) constraints that give rise to the model spectrum shown in the upper panel of Figure 4.4. Here we already




Figure 4.5: Posterior distributions inferred from the mock spectrum depicted in Figure 4.4 in the local IGM damping wing parameterization (black) and the global one (yellow). In both cases, the distribution is marginalized over 7 nuisance parameters describing the shape of the quasar continuum and coverage corrected according to the procedure described in Section 4.3.5. Additionally depicted in green is the topology-informed version of the local constraints, entailing the non-trivial prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$, explicitly shown in purple in the extra panel, and additionally providing a constraint on the global IGM neutral fraction (x_{HI}) in good agreement with the directly inferred one.

marginalized out the seven nuisance parameters associated to the quasar continuum. We can see that the constraints extend over a vast range in the $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ plane with a mild preference towards higher HI column densities and shorter neutral patch distances. The contours in the $(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}})$ panel hint at a similar degeneracy between these two parameters in the range $10^{19.5} \text{ cm}^{-2} \lesssim N_{\text{HI}}^{\text{DW}} \lesssim 10^{21} \text{ cm}^{-2}$ as is often observed between the parameters $\langle x_{\text{HI}} \rangle$ and $\log t_{\text{Q}}$. This degeneracy arises because both a higher HI column density $\log N_{\text{HI}}^{\text{DW}}$ (or a higher global IGM neutral fraction $\langle x_{\text{HI}} \rangle$), as well as a shorter quasar lifetime $\log t_{\text{Q}}$ are causing a stronger IGM damping wing. At $N_{\text{HI}}^{\text{DW}} \lesssim 10^{19.5} \text{ cm}^{-2}$, however, the axis of degeneracy becomes perfectly vertical. This is because at such low column densities, there ceases to be any identifiable damping wing imprint present in the spectrum, and as such, all constraining power with respect to $\log N_{\text{HI}}^{\text{DW}}$ is lost. The size of the proximity zone, on the other hand, places a lower limit on the lifetime of the quasar, since a $t_{\text{Q}} \simeq 10^3 \text{ yr}$ quasar can never carve out a proximity zone of this size, regardless of the reionization state of the pre-quasar IGM. The lifetime of this object therefore cannot be significantly smaller than $t_{\text{Q}} \simeq 10^4 \text{ yr}$ such that the lifetime constraint extends along the vertical direction. A similar effect is observed in the $(\log t_{\text{Q}}, r_{\text{patch}})$ plane where neutral patch distances $r_{\text{patch}} \lesssim 25 \text{ cMpc}$ are preferred at most lifetimes $t_{\text{Q}} \gtrsim 10^5 \text{ yr}$, while the contours become vertical at $t_{\text{Q}} \simeq 10^4 \text{ yr}$ as our sensitivity to more distant neutral patches vanishes. This is because at lifetimes of $t_{\text{Q}} \gtrsim 10^5 \text{ yr}$, a drop in the proximity zone and absorption redward of Lyman- α pushes the constraints toward smaller patch sizes. The reason we can constrain this along with $\log N_{\text{HI}}^{\text{DW}}$ is that at fixed $\log N_{\text{HI}}^{\text{DW}}$, large values of r_{patch} would require rare density fluctuations at $r > r_{\text{patch}}$, high enough to produce these signatures in the spectrum.

Recall that by assuming the constant prior $P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ that is only limited by the physical boundaries for these parameters, we initially sampled from the topology-agnostic posterior $P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$. Binning these HMC samples into a histogram now allows us to multiply them with the conditional distribution $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ to fold in the topology dependence via Eq. (4.10) and obtain the full four-dimensional topology-informed posterior distribution $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ on both global *and* local parameters, shown as green contours in Figure 4.5. Since the data was not overly constraining with regards to our local summary statistics, the prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ —shown explicitly in purple in the extra panel of the same figure—has a clear impact on the shape of the topology-informed posterior which is shifted notably towards higher HI column densities due to the higher prevalence of such sightlines in the realistic ionization topologies (see Figure 4.2). The effect on the r_{patch} posterior is somewhat weaker, but here too we can see that the high- r_{patch} tail of the distribution gets suppressed. As these low $\log N_{\text{HI}}^{\text{DW}}$ and high r_{patch} values go along with lifetimes of $t_{\text{Q}} \simeq 10^4 \text{ yr}$, the long-lifetime mode of the converted

posterior obtains significantly more probability mass than in the unconverted case.

Turning to the resulting $\langle x_{\text{HI}} \rangle$ constraints, we see that the posterior extends over the entire range of $0 \leq \langle x_{\text{HI}} \rangle \leq 1$. In essence, this is because almost any ionization topology contains a number of sightlines with local parameter values similar to the one at hand, as seen in Figure 4.3. The probability mass cumulates around intermediate neutral fractions with a peak very close to the true value of $\langle x_{\text{HI}} \rangle = 0.55$. Note that the $(\langle x_{\text{HI}} \rangle, \log N_{\text{HI}}^{\text{DW}})$ and $(\langle x_{\text{HI}} \rangle, r_{\text{patch}})$ panels of the corner plot can be understood as a reweighted version of the joint distribution $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ of our local summary statistics shown in Figure 4.2, reweighted according to the posterior distribution $P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ as per Eq. (4.10).⁶ The $(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}})$ panel recovers the well-known degeneracy between these two parameters, reflecting the fact that both a higher IGM neutral fraction and a shorter quasar lifetime increase the strength of the damping wing signature. We can now also compare the marginal $(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}})$ constraints to the ones obtained by inferring these two parameters directly, without invoking our local parameterization, shown as yellow contours in Figure 4.5. These constraints are in excellent agreement with the green contours obtained by converting our local constraints.

This demonstrates that our local parameterization allows us to recover the consistent global $\langle x_{\text{HI}} \rangle$ constraints on the timing of reionization as conventionally quoted in the literature. However, our local parameterization comes with a number of advantages: first, the global neutral fraction $\langle x_{\text{HI}} \rangle$ is not the only parameter constrained by the statistics $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} . Analyzing an ensemble of sightlines in our local parameterization, and determining the distribution of the summaries $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} provides us with information about the underlying ionization topology, in addition to the constraints on the timing of reionization. Second, our local framework facilitates model comparison of different reionization models which can come with different ionization topologies. Their effect can easily be folded into our constraints via the conditional distribution $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ that provides the stochastic mapping from local to global parameters. Third, the statistical fidelity of our pipeline with respect to the parameters $\langle x_{\text{HI}} \rangle$ and $\log t_{\text{Q}}$ improves if we constrain them based on our local parameter constraints rather than infer them directly. We quantify this effect in the subsequent section using the concept of expected coverage probabilities introduced in Section 4.3.5.

⁶As per Eq. (4.5), Eq. (4.10) is nothing else but the the product of the joint distribution $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ and the topology-agnostic posterior distribution $P(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$.




Figure 4.6: Coverage test results in the form of expected coverage probability C_α as a function of the credibility level α for different sets of model parameters. The coverage is always computed after marginalizing over 7 nuisance parameters describing the shape of the quasar continuum. The coverage probability of the three parameters ($\log t_Q, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}$) of our local framework is depicted in black, the coverage after folding in the topology dependence and adding $\langle x_{\text{HI}} \rangle$ as an additional parameter in white. The resulting coverage on $\langle x_{\text{HI}} \rangle$ and $\log t_Q$ after marginalizing over the two local parameters $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} is depicted in green, and is compared to the coverage curve we obtain when directly inferring $\langle x_{\text{HI}} \rangle$ and $\log t_Q$ (yellow).

4.4.2 Coverage test

The mock spectrum discussed in the previous section is part of an ensemble of 10×10 mock spectra spanning the entire $(\langle x_{\text{HI}} \rangle, \log t_Q)$ parameter grid (see Section 4.3.4). Each sightline contains an IGM transmission profile drawn from the realistic reionization topology, and hence we can also determine true $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} values based on the underlying pre-quasar HI density field of each sightline. Having available the ground truth of all four parameters $\langle x_{\text{HI}} \rangle, \log t_Q, \log N_{\text{HI}}^{\text{DW}}$ and r_{patch} , we can readily perform coverage tests with respect to the unconverted posterior distributions $P(\log t_Q, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$, as well as the full four-dimensional converted ones $P_{\text{top}}(\langle x_{\text{HI}} \rangle, \log t_Q, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$, and arbitrary marginals thereof.⁷

We show the results of these coverage tests in Figure 4.6, where we depict the expected coverage probability C_α as a function of the credibility level

⁷Note that before performing any of these coverage tests, we marginalize over all nuisance parameters $\boldsymbol{\eta}$ related to the quasar continuum since we are mainly interested in the coverage on the astrophysical parameters $\boldsymbol{\theta} = (\langle x_{\text{HI}} \rangle, \log t_Q, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$.

α , which, in the optimal case, should follow the black dashed one-to-one relation. Firstly, we see that the coverage probability of our local parameter inference (solid black curve) stays consistently below this line, and is therefore overconfident. This is likely because of the Gaussian approximation of the IGM transmission likelihood (Eq. 4.13), which Hennawi et al. (2025) identified as the main source of overconfidence in their global parameter inference pipeline. We reproduce their findings by performing a second coverage test on the directly inferred parameters $\langle x_{\text{HI}} \rangle$ and $\log t_{\text{Q}}$, based on the same ensemble of mock spectra. The overconfidence of this coverage curve, depicted in yellow in Figure 4.6 (see also the left-hand panel of Figure 11 in Hennawi et al. 2025), is somewhat more mild than in the local parameterization (black). The reason for this is likely that the local parameterization has removed a notable degree of scatter from the IGM transmission profiles. While this additional amount of scatter acts as a major additional source of uncertainty for the final parameter constraints (Kist et al. 2025b), it also aids the inference by gaussianizing the likelihood function according to the central limit theorem. Our Gaussian approximation to the IGM transmission likelihood (which applies in both cases) is therefore slightly better justified in the global parameterization than in the local one, leading to the somewhat better (but still suboptimal) coverage behavior.

The coverage of the four-dimensional, converted posterior distribution $P(\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \mathbf{f})$ (white curve) looks almost identical to the unconverted, three-dimensional one since it remains affected by the overconfidence with respect to the local parameters $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} . On the other hand, after marginalizing out these two parameters, we are left with a nearly perfect coverage curve on the remaining two parameters $\langle x_{\text{HI}} \rangle$ and $\log t_{\text{Q}}$ (green curve). Notably, the coverage behavior is even better than if we had inferred these two parameters directly. This shows that our local parameter inference pipeline can be advantageous even if all interest is geared towards global constraints on the timing of reionization, disregarding the local summaries $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} as nuisance parameters. This result is somewhat surprising given the overconfidence when inferring the local parameters on their own. However, it is important to realize that the coverage on other parameters, subsequently obtained from those, is not necessarily similarly bad.

In the case at hand, note that the mapping from local to global parameters is probabilistic, and therefore in particular highly non-injective, i.e., many local parameter combinations can get mapped to the same global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ (c.f. Figure 4.3). Even if there are imperfections in the local $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} constraints, these parameters can still get mapped to the 'correct' global $\langle x_{\text{HI}} \rangle$ value. Specifically, this appears to be the case here to such a high degree that the converted ($\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}$) constraints are statistically more faithful than the directly inferred ones. An additional factor that likely contributes is that binning the conditional distribution

$P(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ in order to convert the constraints via Eq. (4.10) smooths out the contours and thus improves the coverage to a certain extent, whereas such a binning step is not required when inferring $\langle x_{\text{HI}} \rangle$ and $\log t_{\text{Q}}$ directly.

4.4.3 Inference precision

Having investigated the coverage behavior of our inference framework, we now proceed by quantifying its precision. We do so by considering the one-dimensional marginal posterior distribution of a given astrophysical parameter, and computing its width

$$\Delta_{68\%}(\theta) = \frac{1}{2}(p_{84\%}(\theta) - p_{16\%}(\theta)) \quad (4.17)$$

as a measure of the 1σ error on the parameter $\theta \in \{\langle x_{\text{HI}} \rangle, \log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}\}$, where $p_q(\theta)$ is the q -th percentile of the respective marginal posterior distribution. Before determining these percentiles, we perform a coverage test and reweigh the inferred posterior distribution according to the prescription discussed in Section 4.3.5 and Hennawi et al. (2025) to ensure that we do not quote overly optimistic precision values.

4.4.3.1 Local parameterization

We start by investigating the precision we can achieve on $\log t_{\text{Q}}$, $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} when performing the inference in our new local parameterization. Here we only determine the inference precision in the region of parameter space where a non-negligible damping wing imprint is present, i.e., where we can expect at least some degree of sensitivity to these parameters. In other words, we choose to concentrate our computational resources for this exercise to regions where our constraints are not fully prior-dominated. These are sightlines with comparably high HI column densities $\log N_{\text{HI}}^{\text{DW}}$ and shorter distances r_{patch} to the first neutral patch. We therefore consider a 4×4 grid in $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space with values of $20.1 \leq \log N_{\text{HI}}^{\text{DW}} \leq 20.7$ and $2 \text{ cMpc} \leq r_{\text{patch}} \leq 14 \text{ cMpc}$, i.e., mock ensemble (ii) described in Section 4.3.4. In what follows, the configuration with $\log N_{\text{HI}}^{\text{DW}} = 20.1$ and $r_{\text{patch}} = 14 \text{ cMpc}$ can be seen as representative of the remaining parts of parameter space where $\log N_{\text{HI}}^{\text{DW}} < 20.1$ and/or $r_{\text{patch}} > 14 \text{ cMpc}$. Note that due to the physical constraints on the possible $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ combinations, models only exist for 12 out of the 16 points on this grid. Recall that it does *not* take a topology-informed prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ to exclude the remaining 4 parameter combinations. These are *physically excluded* based on the topology-independent distribution of density fluctuations in the IGM.

We add a third grid dimension by considering three different lifetime values of $t_{\text{Q}} = 10^4, 10^6$ and 10^8 yr . At each location in $(\log t_{\text{Q}}, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$




Figure 4.7: Inference precision with respect to the three parameters of our local damping wing parameterization in different regions of parameter space. Precision on $\log t_Q$, $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} is shown in the top, middle and bottom row, respectively. The lifetime is fixed to $t_Q = 10^4$, 10^6 , and 10^8 yr in the left, middle and right column, respectively, and each panel depicts the precision in the region of $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space where damping wings are present. Each pixel shows the average precision values of 10 different mock objects from mock ensemble (ii) described in Section 4.3.4, where the same 10 continua are used at each location in parameter space to isolate the effects of the astrophysical parameter variations. Annotated values are the mean of all precision values in a given panel.

parameter space, we consider 10 distinct mock spectra, and we average the inferred precision values over these 10 objects. To further suppress variations throughout parameter space due to other sources of stochasticity, we consider the *same* 10 continuum draws and noise realizations at each location in parameter space. Overall, this amounts to a set of $12 \times 3 \times 10 = 360$ mock spectra. Note that due to the fact that these parameter values do not span the full prior range of the three astrophysical parameters ($\log t_Q, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}$), we do not determine the coverage behavior on this set of mocks but rather on the one considered in Section 4.4.2. This is possible since all that is required for the reweighting procedure outlined in Section 4.3.5 is the functional relationship between credibility levels α and expected coverage probability C_α . Practically, this means we consider the unconverted local parameter coverage (black curve) from Figure 4.6, and use it to reweigh our $(\log t_Q, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ samples obtained from the new mock ensemble introduced in this section. We then compute the inference precision based on the percentiles of the *reweighted* samples according to Eq. (4.17).

The averages of these precision values as a function of the astrophysical parameter values are plotted in Figure 4.7 where we show from top to bottom the inference precision on $\log t_Q, \log N_{\text{HI}}^{\text{DW}}$ and r_{patch} , respectively. The three columns represent the three different mock quasar lifetimes of $t_Q = 10^4, 10^6$ and 10^8 yr, respectively, and each panel itself depicts the precision in $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space. The annotated values are the mean precision of all mocks in the respective panel.

We begin by analyzing the lifetime precision $\Delta_{68\%}(\log t_Q)$ for which we do not identify a notable dependence on the local parameters $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} . The mean precision values are 0.66 dex at $t_Q = 10^4$ yr and 0.64 dex at $t_Q = 10^6$ yr, decreasing to 0.44 dex at $t_Q = 10^8$ yr, as seen in the upper row of Figure 4.7. This enhanced precision at long lifetimes is in line with the trend identified in Kist et al. (2025b) in the context of the global parameterization, attributed to the thermal proximity effect due to Helium II ionization. This becomes apparent for quasars with long lifetimes and thus a significantly extended Helium ionization front where the associated thermal heating has enhanced the strength of the IGM transmission peaks, making such lifetimes easier to identify through this distinct feature. We refrain from a quantitative comparison of our lifetime precision values to those quoted in Kist et al. (2025b) due to the fact that due to the different parameterizations, an apples-to-apples comparison of the same regions of parameter space is challenging, and, more importantly, we are investigating here the precision of our *topology-agnostic* constraints, whereas the global constraints in Kist et al. (2025b) are inevitably topology-informed. Instead, we proceed in the subsequent section by performing a comparison of the precision values we can obtain in global $(\langle x_{\text{HI}} \rangle, \log t_Q)$ parameter space after conversion of our local constraints versus directly inferring $\langle x_{\text{HI}} \rangle$ and $\log t_Q$.

For now, proceeding to the HI column density inference precision $\Delta_{68\%}(\log N_{\text{HI}}^{\text{DW}})$, we observe similar trends as found by Kist et al. (2025b) for the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$, that is, an increasingly high precision the stronger the IGM damping wing. This is unsurprising given that the strongest imprints can most easily be disentangled from the intrinsic continuum of the quasar. This trend is clearly seen when focusing on rows of fixed r_{patch} at short to intermediate quasar lifetimes in the first two panels of the middle row of Figure 4.7. Vice versa, at fixed $\log N_{\text{HI}}^{\text{DW}}$ (i.e., in a given column of one these panels), the inference precision on this parameter appears to deteriorate the closer the location r_{patch} of the first neutral patch to the quasar. A possible reason for this is the somewhat increased scatter of IGM transmission profiles in the local parameterization at small r_{patch} values (see Figures 8 and 9 in Kist et al. 2025c) which makes it harder to reconstruct the true column density of a given sightline.

Both aforementioned trends with $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} cease to exist when the lifetime of the quasar is long ($t_{\text{Q}} = 10^8 \text{ yr}$) where we find a uniform $\log N_{\text{HI}}^{\text{DW}}$ inference precision across the entire range of parameter space considered in this figure. This is likely because the damping wing imprint is so weak at these lifetimes that it can hardly be disentangled from the intrinsic spectrum of the quasar, even at the highest HI column densities $\log N_{\text{HI}}^{\text{DW}}$.

The most significant trends are seen for the inference precision on the distance r_{patch} of the quasar to the first neutral patch, depicted in the bottom row of Figure 4.7. At the shortest lifetime of $t_{\text{Q}} = 10^4 \text{ yr}$, we can identify a highly pronounced trend of enhanced precision at higher HI column densities $\log N_{\text{HI}}^{\text{DW}}$ and shorter neutral patch distances r_{patch} , improving from $\sim 40 \text{ cMpc}$ down to $< 5 \text{ cMpc}$. These high precision values can be explained by the fact that the location of the ionization front for low-lifetime objects coincides well with the pre-quasar neutral patch location since such quasars have not yet carved out an ionized patch that extends significantly beyond this location. As a result, r_{patch} can easily be reconstructed from the spectrum of such objects. However, if the patch location is too far away, the decrease of the quasar's photoionization rate $\Gamma_{\text{QSO}} \sim 1/r^2$, suppresses all transmission already at a distance closer to the quasar than r_{patch} , and hence we lose sensitivity to this parameter (see also Figure 8 in Kist et al. 2025c).

At longer quasar lifetimes, inference precision on r_{patch} deteriorates since most such objects have ionized away the first pre-quasar neutral patch, making it significantly harder (if not even entirely impossible) to reconstruct its location. At $t_{\text{Q}} = 10^6 \text{ yr}$, a relatively high precision is retained for the $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}) = (10^{20.7} \text{ cm}^{-2}, 2 \text{ cMpc})$ and $(10^{20.5} \text{ cm}^{-2}, 10 \text{ cMpc})$ models. These are parameter combinations where the location of the ionization front after 10^6 yr still coincides well with the location of the first pre-quasar neutral patch (compare also Figure 9 in Kist et al. 2025c), and as such,

is still reconstructible at a reasonable precision. On the other hand, after shining for $t_Q = 10^8$ yr, the quasar has ionized away so much neutral material that virtually any information about the location of the pre-quasar neutral patch is lost, regardless of the location in $(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space.

In summary, we have found that in certain regions of parameter space, we are sensitive to all three parameters of our local IGM damping wing parameterization, demonstrating its value for astrophysical parameter inference. Specifically, we saw that we can reconstruct the quasar lifetime $\log t_Q$ to $\sim 0.4 - 0.7$ dex, the HI column density to ~ 0.7 dex, improving down to ~ 0.2 dex when the damping wing is strong, and the neutral patch distance r_{patch} up to < 5 cMpc at short quasar lifetimes of $t_Q \simeq 10^4$ yr, whereas sensitivity gets lost the longer the quasar has been shining.

4.4.3.2 Global parameterization

As pointed out in the previous section, an apples-to-apples comparison of the inference precision achieved in the local parameterization to that quoted in Kist et al. (2025b) in the context of the conventional global parameterization is not possible. To perform such a comparison, we instead consider mock ensemble (i) described in Section 4.3.4 where the x_{HI} sightlines originate from realistic semi-numerical ionization topologies and we can compare the precision achieved on $\langle x_{\text{HI}} \rangle$ and $\log t_Q$ by converting the local $(\log t_Q, \log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ constraints to these global parameters versus inferring them directly.

Figure 4.8 depicts these precision values obtained on $\langle x_{\text{HI}} \rangle$ in the upper row and on $\log t_Q$ in the lower row, comparing the converted local constraints in the left column to the directly inferred global ones in the right one. Each panel by itself shows the precision values in the $(\langle x_{\text{HI}} \rangle, \log t_Q)$ parameter plane, where each pixel corresponds to the average precision of 4 mocks from our 10×10 mock object ensemble. Note that there is a direct correspondence between the two right-hand panels and the bottom-right panels of Figures 9 and 10 in Kist et al. (2025b). We average the precision over 4 mocks in each pixel to suppress sightline-to-sightline stochasticity as far as possible. Note that to that end, our mock ensemble considered in the previous section had fixed continua and noise draws across the entire range of parameter space. This is *not* the case in this section as we are only interested in the differences in inference precision based on the two parameterizations, so a one-to-one comparison between fixed pixels in the left and right panels is still possible.

What we find when performing this comparison is a near perfect agreement between the precision values obtained in the two parameterizations. Small differences only become apparent in $\langle x_{\text{HI}} \rangle$ precision towards the (high- $\langle x_{\text{HI}} \rangle$, low- $\log t_Q$) end where the global constraints become somewhat more precise. This is a consequence of the binning we adopt when converting the local to the global constraints which are limited by the $\langle x_{\text{HI}} \rangle$ grid spacing




Figure 4.8: Inference precision with respect to the two parameters $\langle x_{\text{HI}} \rangle$ and $\log t_Q$ of the global damping wing parameterization in different regions of parameter space, compared between our local and global inference framework. Precision on $\langle x_{\text{HI}} \rangle$ is shown in the upper row, precision on $\log t_Q$ in the lower row. The left column shows precision values of constraints converted from the local to the global parameterization, whereas the ones in the right column correspond to $(\langle x_{\text{HI}} \rangle, \log t_Q)$ inferred directly. Each pixel shows the average precision values of 4 different mock objects originating from realistic ionization topologies, i.e., based on the mock ensemble (i) described in Section 4.3.4.

of 0.05. On the other hand, the converted constraints are slightly more precise in regions of low precision (e.g. $\langle x_{\text{HI}} \rangle$ precision at the long-lifetime end $t_Q \simeq 10^8$ yr, or $\log t_Q$ precision at intermediate lifetimes $t_Q \simeq 10^{5-6}$ yr). This likely results from the stronger coverage corrections that have to be applied in the context of the global parameterization where the inference is somewhat overconfident.

Besides these minor differences, our main conclusion is that the inference precision on $\langle x_{\text{HI}} \rangle$ and $\log t_Q$ coincides remarkably well, regardless of the parameterization that the inference was originally performed in. This demonstrates that we are not losing any information on the global timing of reionization by first inferring local constraints and subsequently tying them to a global reionization model. Note that on the other hand, we also cannot expect significantly tighter constraints than from the direct global parameter inference since the conversion step necessarily implies that we are again

folding in the stochasticity of reionization that our local parameterization has removed from the inference task itself. However, our new parameterization allows us to gain hitherto unused information about the local ionization topology in front of a quasar while still retaining the same precision on the parameters $\langle x_{\text{HI}} \rangle$ and $\log t_{\text{Q}}$ when inferring them directly in the context of the global parameterization.

4.5 Conclusions

IGM damping wings towards quasars are a unique probe not only of the history of reionization but also its topology, and [Kist et al. \(2025c\)](#) introduced a new set of local summary statistics encapsulating all this information. Here we presented a fully Bayesian inference framework that allows us to extract this information in a topology-independent fashion. These constraints are informative about the local ionization topology before the quasar started shining, and, when tied to a specific reionization model, also about the global timing of reionization.

To establish this connection, we introduced a probabilistic framework that allows us to map our local damping wing statistics, consisting of a Lorentzian-weighted HI column density $\log N_{\text{HI}}^{\text{DW}}$ as well as the distance r_{patch} of the quasar to the first neutral patch in the pre-quasar topology, to the global volume-averaged IGM neutral fraction $\langle x_{\text{HI}} \rangle$. The virtue of our approach is that the two local summaries $\log N_{\text{HI}}^{\text{DW}}$ and r_{patch} , along with the lifetime $\log t_{\text{Q}}$ of the quasar, can be inferred independently of any assumptions about the reionization model. We demonstrated that all these assumptions can be encoded in a non-trivial prior $P_{\text{top}}(\log N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ which can be folded in *subsequently*, facilitating the comparison of different reionization models.

In addition, we found that, after marginalizing over the local constraints, this two-step procedure improves the statistical fidelity of our inference pipeline. The inference of the local damping wing statistics themselves shows a certain degree of overconfidence, but we correct for this by broadening the inferred posteriors in a principled way before quoting actual constraints. In the future, simulation-based inference approaches will be key to also capture non-Gaussianities in the IGM transmission likelihood, and hence make use of the full information contained in the spectra while removing the need for coverage correction altogether ([Chen 2024](#)).

Based on a large sample of mock spectra, we quantified that precision at which we can infer the three astrophysical model parameters, and we found that we can constrain the quasar lifetime to $0.58_{-0.13}^{+0.13}$ dex, the HI column density to $0.69_{-0.30}^{+0.06}$ dex, and the distance to the first neutral patch to $31.4_{-28.1}^{+10.7}$ cMpc for model parameter combinations that would imprint a significant damping wing upon the spectrum of the quasar. Furthermore, we

found that after tying our local constraints to a given reionization model, our precision on the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ and the quasar lifetime $\log t_{\text{Q}}$ is on par with the precision achieved when inferring these parameters directly in the context of the conventional, global parameterization, while removing the need for coverage corrections. This advantage comes at the price of a somewhat increased computational cost due to the increased dimensionality of parameter space which also increases the number of required simulation models.

The biggest virtue of adopting our approach, though, is that it provides model-independent, physical information about the local ionization topology in front of a given quasar. If concerned with multiple sightlines at a similar redshift, setting up a Bayesian hierarchical model will allow us to combine these topology-agnostic local constraints to a topology-informed constraint not only on the global timing of reionization but also its topology (Sharma et al. 2025). In addition, the local information we gain through these statistics will allow us to draw novel connections to the 21 cm power spectrum, and the ionized bubble sizes measured in studies of Lyman- α emission from galaxies. We will further explore these connections in future work.

Acknowledgements

We acknowledge helpful conversations with the ENIGMA group at UC Santa Barbara and Leiden University, especially with Da-Ming Yang about the low-redshift continuum data. This work made use of NumPy (Harris et al. 2020), SciPy (Virtanen et al. 2020), JAX (Bradbury et al. 2018), NumPyro (Bingham et al. 2018; Phan et al. 2019), sklearn (Pedregosa et al. 2011), Astropy (Astropy Collaboration et al. 2013, 2018, 2022), Pytype (Prochaska et al. 2020), SkyCalc_ipy (Leschinski 2021), h5py (Collette 2013), Matplotlib (Hunter 2007), corner.py (Foreman-Mackey 2016), and IPython (Pérez & Granger 2007). TK and JFH acknowledge support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No 885301). JFH acknowledges support from NSF grant No. 2307180.

Data Availability

The derived data generated in this research will be shared on reasonable requests to the corresponding author.