



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).

Les Lumières

Probing the cosmic Epoch of Reionization
with high-redshift quasars

Proefschrift

ter verkrijging van
de graad van doctor aan the Universiteit Leiden,
op gezag van rector magnificus prof.dr. S. de Rijcke,
volgens besluit van het college voor promoties
te verdedigen op vrijdag 3 juli 2026
klokke 11.30 uur

door

Timo Kist

geboren te Zeven, Duitsland
in 1997

Promotores:

Prof.dr. J.F. Hennawi

Prof.dr. J. Schaye

Promotiecommissie:

Prof.dr. I.A.G. Snellen

Prof.dr. K.H. Kuijken

Prof.dr. H. Hoekstra

Prof.dr. M.G. Haehnelt (University of Cambridge)

Dr. F.B. Davies (Max Planck Institute of Astronomy)

Dr. E. Sellentin

Each cow on the cover represents a different reionisation topology whose average neutral fraction $\langle x_{\text{HI}} \rangle$ is computed within its respective coordinate system, assuming that white regions are neutral and black ones ionised. The line in the background is an artistic impression of a quasar spectrum with Lyman- α damping wing and proximity zone, generated according to the procedure outlined in Chapter 2 of this thesis.

Illustration by Annika Dahlkemper, scientific add-ons by the author.

An electronic version of this dissertation is available at
<https://scholarlypublications.universiteitleiden.nl/>.

"Sapere aude! Habe Muth, dich deines eigenen Verstandes zu bedienen!"
– *Beantwortung der Frage: Was ist Aufklärung?*, Immanuel Kant

"Sapere Aude! Have courage to use your own understanding!"
– *An Answer to the Question: What is Enlightenment?*, Immanuel Kant

CONTENTS

1	Introduction	1
1.1	The first billion years of our Universe	3
1.1.1	The first seconds	3
1.1.2	The first millennia	5
1.1.3	The first billion years	6
1.2	Reionisation – Les Lumières cosmologiques	6
1.2.1	The integrated effect of reionisation	7
1.2.2	The global timing of reionisation	8
1.2.3	The topology of reionisation	9
1.3	Modelling Reionisation	11
1.3.1	Reionisation equations	12
1.3.2	Reionisation simulations	14
1.4	Quasars – Les Lumières astrophysiques	15
1.4.1	The Lyman- α transition	16
1.4.2	The Lyman- α forest	17
1.4.3	Quasars during the Epoch of Reionisation	17
1.4.3.1	The Gunn-Peterson effect	18
1.4.3.2	Dark pixel fractions	19
1.4.3.3	Effective optical depth distributions	20
1.4.3.4	The mean free path of ionising photons	20
1.4.3.5	Damping wings	21
1.4.3.6	Proximity zones	23
1.5	Other observational probes of Reionisation	24
1.5.1	The Cosmic Microwave Background	24
1.5.2	Galaxies	26
1.5.3	Gamma-ray bursts	27
1.5.4	The 21 cm line	28
1.5.5	The patchy kSZ effect	28
1.6	Inference – Le bon sens réduit au calcul	29
1.6.1	Bayes' theorem	29
1.6.2	Sampling	30
1.7	This Thesis – Les Lumières	32
1.8	Outlook – Les nouvelles Lumières	34

2	Quantifying the precision of IGM damping wing measurements towards quasars	37
2.1	Introduction	38
2.2	Methods	41
2.2.1	Training Data and Dimensionality Reduction	43
2.2.2	Simulating Damping Wing Observations	44
2.2.2.1	Full-Simulation Model	44
2.2.2.2	Analytical Simulator	45
2.2.3	Generating Mock Spectra	49
2.2.4	Inference Procedure	51
2.3	Sensitivity Analysis	54
2.3.1	Inference Tests	55
2.3.2	Inference Precision	56
2.3.3	Parameter Scan	57
2.3.3.1	Latent Dimension	57
2.3.3.2	Wavelength coverage	64
2.3.3.3	Signal to Noise Ratio	66
2.3.3.4	Spectral resolution	67
2.3.3.5	The optimal hyper-parameter setting	69
2.3.4	The Total Error Budget	69
2.4	Conclusions	75
2.A	Correction to the Analytical Damping Wing Formula	77
2.B	Inference Tests	79
2.C	Parameter Scan for the Full-Simulation Model	81
3	A local, topology-independent parametrization of quasar IGM damping wings	85
3.1	Introduction	86
3.2	Towards a three-parameter model of quasar IGM damping wings	88
3.2.1	Introducing a local HI column density label	88
3.2.2	Choice of the $N_{\text{HI}}^{\text{DW}}$ integration limits	92
3.2.3	The damping wing as a one-parameter family	94
3.2.4	A second label: the distance to the first neutral patch	96
3.2.5	Relation to previous definitions	97
3.3	Quantifying the IGM transmission scatter	99
3.3.1	Simulating IGM transmission fields	99
3.3.2	Global parameterization	101
3.3.3	Local parameterization	104
3.4	Towards a topology independent local IGM damping wing parameterization	108

3.4.1	An analytic toy prescription for generating neutral fraction profiles	110
3.4.1.1	Description of the prescription	111
3.4.1.2	Generating IGM transmission profiles	112
3.4.2	Comparison to a realistic reionization topology	113
3.4.2.1	Physically accessible regions of parameter space	114
3.4.2.2	Agreement between realistic reionization model and toy prescription	117
3.4.2.3	Parameter dependence of the IGM transmission profiles at short quasar lifetimes	119
3.4.2.4	Parameter dependence of the IGM transmission profiles at intermediate to long quasar lifetimes	120
3.5	Conclusions	121
3.A	On the relation between τ_{DW} and $N_{\text{HI}}^{\text{DW}}$	124
3.B	On the hyperparameter dependence of $N_{\text{HI}}^{\text{DW}}$	125
4	Inferring local quasar IGM damping wing constraints	131
4.1	Introduction	132
4.2	Theory: relating a local IGM damping wing model to the global timing of reionization	134
4.2.1	A local, topology-independent parameterization of quasar IGM damping wings	135
4.2.2	Folding in the topology dependence and constraining the global IGM neutral fraction	137
4.2.2.1	Simulating IGM transmission profiles	137
4.2.2.2	The prior on the local summary statistics induced by the reionization model	142
4.2.2.3	Converting local measurements to a global $\langle x_{\text{HI}} \rangle$ constraint	148
4.3	Methods: inference pipeline	150
4.3.1	Continuum dimensionality reduction model	151
4.3.2	IGM transmission likelihood	152
4.3.3	Inference procedure	154
4.3.4	Generating mock spectra	155
4.3.5	Coverage tests	156
4.4	Results: global vs local parameter inference	158
4.4.1	Inference for a single mock	158
4.4.2	Coverage test	163
4.4.3	Inference precision	165
4.4.3.1	Local parameterization	165

4.4.3.2	Global parameterization	169
4.5	Conclusions	171
5	First constraints on the local ionization topology in front of two quasars at $z \sim 7.5$	173
5.1	Introduction	174
5.2	Methods	175
5.3	Local IGM damping wing constraints at $z \sim 7.5$	177
5.3.1	J1007+2115	178
5.3.2	J1342+0928	182
5.4	Summary and Conclusions	186
	Bibliography	188
	English summary	201
	Nederlandse samenvatting	205
	List of Publications	209
	Curriculum Vitae	211
	Acknowledgements	213

1 | INTRODUCTION

The lights on the night sky have fascinated us humans since time immemorial, and it might be hard to find civilisations that have not developed their own myths and legends trying to make sense of the phenomena that they would observe. Given the sky's intangibility, it is maybe not surprising that these explanations were often related to stories of creation, involving all sorts of deities, higher powers or other mythical creatures. Perhaps all the more remarkable is the fact that *science* has convinced us that looking at the night sky indeed means nothing less than *looking into the very past of our own Universe*. As if this was not enough, the last century came with the rapid consolidation of a theory – the Λ CDM model – that allows us to tell our Universe's *entire* evolution history – based on the variety of lights that we are observing in the sky.

Instrumental to that end was the unprecedented number of paradigm shifts that our understanding of physics has gone through in the past century, including the development of general relativity (Einstein 1916) and quantum mechanics (e.g. Heisenberg 1925; Born & Jordan 1925; Schrödinger 1926a,b,c,d), subsequently extended to quantum field theory (Tomonaga 1946; Feynman 1948; Schwinger 1948; Dyson 1949), as our most fundamental physical theories to date. Even more so in the field of cosmology, essentially our *entire* base of knowledge has been generated in that century. The rise of modern cosmology perhaps started looming in 1920 when Harlow Shapley and Heber Curtis in their *Great Debate* discussed the question if the then observed *spiral nebulae* would be of intra- or extragalactic origin, and therefore, more fundamentally, if the Universe would extend *beyond our own galaxy, the Milky Way*. The question ultimately got settled with Edwin Hubble's discovery that these nebulae's recessional velocities would linearly increase with their distance to Earth (Hubble 1929), confirming earlier observations by Vesto Slipher (Slipher 1913, 1917) and suggesting that they were indeed nothing else but separate *galaxies* other than the Milky Way.

Two conflicting theories arose in the subsequent decades as physical frameworks to explain the expansion of the Universe implied by these observations: first, the *steady-state theory* where the Universe would eternally expand, retaining its appearance over cosmic time by accommodating the spatial expansion via the concomitant creation of new matter that would help keep the matter density constant over time (Bondi & Gold 1948; Hoyle 1948). On the other hand, the *Friedmann equations*, derived as a homogeneous and isotropic solution to the Einstein field equations (Friedmann 1922; Lemaitre 1927), got by without invoking the creation of new matter. This alternative idea, later dubbed the *Big Bang theory* by Fred Hoyle, instead postulated that the Universe started out in an initial hot dense state. Subsequently,

as the Universe spatially expanded, matter would dilute and the Universe would cool until it reached its present state. Irrefutable evidence for the latter theory ultimately arrived in 1964 with the discovery of the *cosmic microwave background* (CMB) radiation (Penzias & Wilson 1965), a nearly ideal blackbody spectrum, observed isotropically across the entire radio sky, that could convincingly be explained only as the result of a *Big Bang*.

While the discovery of the CMB had settled the fundamental debate between steady-state and Big Bang theory, many details were still yet to be explained. In the following years, the hypothesis of the existence of a *dark*, unseen type of matter gained new attention. Originally postulated significantly earlier (Thomson 1904; Poincare 1906; Kapteyn 1922; Oort 1932; Zwicky 1933), new observations of galaxy rotation curves (Rubin & Ford 1970) and other probes suggested (Ostriker & Peebles 1973; Einasto et al. 1974) that galaxies would have to be surrounded by *halos* of dark matter – an order of magnitude larger in extent than the galaxies themselves – and that this unknown type of matter would even make up the largest fraction of all matter in the Universe. At the present day, the existence of dark matter has become an integral assumption for our entire understanding of cosmology (e.g. Davis et al. 1985), and the Standard Model of Particle Physics has been extended in numerous directions to predict theoretical candidate particles that could account for this unseen type of matter. However, despite considerable efforts to detect dark matter on the particle-level, so far all searches have remained fruitless (e.g. Abdalla et al. 2022).

A further paradigm shift followed no more than 30 years ago, in 1998, when observations of Type Ia Supernovae (Riess et al. 1998; Perlmutter et al. 1999) suggested that our Universe is currently expanding at an *accelerated* rate. Explaining this behaviour requires the postulation of a further *dark* contributor to the Universe’s total energy density: *dark energy* – whose nature is by no means better understood than that of dark matter. In the simplest interpretation, dark energy is nothing else but a *cosmological constant* Λ , originally introduced by Albert Einstein in his field equations of general relativity to keep the Universe static (Einstein 1917). Together with baryonic matter, dark matter, photons and neutrinos, dark energy forms the main ingredient of the Λ CDM *model* that has seen an unprecedented triumph in the recent decades, fitting observations at a remarkable level of precision (e.g. Planck Collaboration et al. 2020). The current cosmological best-fit value for Λ , however, is not less than 120 orders of magnitude away from the quantum field theory-based prediction from first principles, suggesting fundamental limitations in our understanding of the physics behind Λ (Weinberg 1989). At the same time, cosmologists keep challenging the simple picture of a cosmological constant from an observational perspective, conducting ever better surveys to search for evidence that the dark energy equation of state must be evolving over time (e.g. Abdul Karim et al. 2025), leaving vast space for new discoveries even after the revolutionary past century.




Figure 1.1: Schematic timeline of Universe’s evolution history according to the Λ CDM model. Major events that took place over the past 13.8 billion years are labelled on the timeline, and the schematic drawings show the qualitative state of cosmic structure at the respective point in time. The Epoch of Reionisation is heralded by the formation of the first stars and galaxies, approximately within the first billion years. Credit: ESA.

1.1 The first billion years of our Universe

Despite the fact that its most prominent components are far from being understood, and that the Λ CDM model itself might not even be the final word in our understanding of cosmology, it provides an excellent fit to present observations and allows us to tell in remarkable detail the story of our Universe’s past 13.8 billion years, following the timeline sketched in Figure 1.1. This thesis will contribute to telling the tale of its first billion years, and we start this chapter by telling the prologue.

1.1.1 The first seconds

The very beginnings of the Universe, or more specifically, its first $t_{P1} \sim 5 \times 10^{-44}$ s will remain untold for the time being, at least for as long as quantum field theory and the theory of general relativity have not been merged into a unified theory of quantum gravity that would allow us to describe this regime. However, once the Planck time t_{P1} has passed and general relativity enters its reign of validity, the evolution of spacetime – defined by the metric tensor $g_{\mu\nu}$ – is governed by the *Einstein field equations*

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad (1.1)$$

relating the curvature of spacetime, quantified by the Einstein tensor $G_{\mu\nu}$, to the matter components of the Universe, summarised in the stress energy tensor $T_{\mu\nu}$ (Einstein 1916). Despite the fact that we can technically describe their *evolution* from first principles, we are still faced with an *initial conditions* problem: only a highly fine-tuned choice of parameters governing the Universe’s initial curvature and matter content would result in a Universe as flat, homogeneous and isotropic as the one that we are living in today.

Cosmic Inflation, an early phase of *accelerated* spatial expansion of the Universe, has been proposed as a solution to these issues (Guth 1981; Linde 1982; Albrecht & Steinhardt 1982). Driven by one or more scalar fields, their contribution comes to dominate over all other stress energy components, smoothing and flattening the Universe. During such an inflationary epoch, physical distances grow exponentially faster than the Hubble radius whose size governs which scales are in causal contact at a given point in time. As a result, the previously smoothed modes exit the horizon until eventually entering again in the subsequent hot expansion phase, explaining the smoothness of the Universe at the present day, even on its largest scales that would otherwise never have been in causal contact before. At the same time, small inflationary quantum fluctuations provide the seeds for the subsequent formation of the cosmic large-scale structure that shapes our Universe today. The excellent agreement of these inflationary predictions with observations of the CMB (Planck Collaboration et al. 2020) are seen as strong evidence that our Universe indeed went through an inflationary phase in its first $\sim 10^{-33} - 10^{-32}$ s. Inflation has to end with a *reheating* process where the inflationary scalar field(s) decay and deposit their stress energy into a thermal bath of standard model particles.

A vast range of models can drive an inflationary epoch, sharing the above generic predictions. The fact that many of the simplest models, requiring the least amount of fine tuning, have already been ruled out has sparked new debate (Ijjas et al. 2013) and led to the development of alternative theories such as *bouncing models* (Ijjas & Steinhardt 2016) where a period of *slow contraction* would precede the standard phase of hot expansion and similarly solve the problems that inflationary scenarios are aiming to address (e.g. Cook et al. 2020; Kist & Ijjas 2022). *B*-mode polarisation patterns detected in future CMB experiments (e.g. Ade et al. 2019), predicted by many inflationary but not by bouncing models, are expected to provide new evidence on this matter in the coming years.

In either case, in what follows – with its initial conditions set – the Universe enters its hot Big Bang expansion phase, starting in a *radiation-dominated* epoch where relativistic components of the primordial plasma (i.e., photons and neutrinos) dominate the total energy density. The initially enormous temperature begins to decline as the Universe expands, inducing the successive breaking of Standard Model symmetries into the four known fundamental forces, and the emergence of massive elementary particles via the Higgs mechanism (Higgs 1964; Englert & Brout 1964; Guralnik et al. 1964) during the *electroweak phase transition* after 10^{-12} s, and later that of hadrons such as protons and neutrons during the *quantum chromodynamics (QCD) phase transition* (10^{-5} s after the Big Bang). After 1 second, the point is reached where neutrinos *decouple* from the rest of the primordial plasma, and freely stream through the Universe until the present day, forming the so-called *cosmic neutrino background (CνB)*. As the Universe cools further,

particles and their anti-particles annihilate once their respective mass (or, equivalently, energy) scale is reached, starting with hadronic particles and later continuing with leptons until at around 10s after the Big Bang, electrons and positrons have annihilated as well. A major remaining conundrum about these (a priori symmetric) *baryogenesis* and *leptogenesis* processes is how the present-day matter-antimatter asymmetry could come into place.

More clarity exists about the subsequent nuclear fusion of light chemical elements (deuterium, helium, lithium and other light isotopes) during *Big Bang nucleosynthesis* (BBN, [Alpher et al. 1948](#)). This process began once the universe's temperature fell below the binding energies of the lightest elements and continued as long as a sufficient temperature for nuclear fusion was retained, resulting in a characteristic abundance of these primordial elements by a time of around 20 minutes after the onset of the hot Big Bang.

1.1.2 The first millennia

The evolution that follows upon these turbulent first minutes takes place on significantly longer – cosmological – timescales. After around 47 000 years, at a redshift of $z \sim 3400$, *matter-radiation equality* was reached where matter superseded radiation as dominant component of the Universe's total energy density. At that point, photons, protons and electrons were still in thermal equilibrium, forming an ionised plasma opaque to radiation due to frequent Thomson scattering events with the numerous free electrons. This changed around 370 000 years after the Big Bang, at a redshift of $z \sim 1100$, when the binding energy of neutral hydrogen was reached such that protons and electrons could *recombine* and form the first atoms of neutral hydrogen.

Photons, due to the sudden absence of free electrons to scatter off, now had a mean free path that exceeded the horizon distance such that they *decoupled* and started travelling freely for the first time in cosmic history. The Universe had become *transparent*. Since these photons had been in thermal equilibrium with the rest of the primordial plasma right until these stages, and they travelled freely ever since, they still form a nearly perfect blackbody spectrum, reflecting the state of the Universe at this *surface of last scattering*. Detected for the first time in 1965 by [Penzias & Wilson \(1965\)](#), this Cosmic Microwave Background (CMB) radiation – now redshifted to microwave frequencies – can be seen as an early echo of the Big Bang, reflecting in minute detail the state of the Universe at $z \sim 1100$ and acting as our single most important source of information not only about these early stages of our Universe but also many later epochs, based on the marks they leave on the CMB photons. Most prominently, besides its overall isotropy and smoothness, the CMB contains small temperature variations, reflecting the state of density perturbations in the early Universe that would keep growing in the matter-dominated epoch and serve as seeds for the formation of cosmic structure that we observe today.

1.1.3 The first billion years

Despite its newly found transparency, the Universe subsequently entered what we call its *dark ages*, owing to the fact that the recently released CMB photons quickly redshifted away from visible frequencies into the infrared. The only other source of radiation at that time were microwave photons emitted via the 21 cm spin flip transition of the recently recombined neutral hydrogen atoms. As such, the Universe remained devoid of visible light until eventually the first stars would form out of the gradually growing primordial density perturbations. Indirect evidence for the existence of these pristine, metal-poor *population III stars* exists in form of traces of otherwise inexistent heavy elements in quasar emission spectra, but direct observational proof of their existence is still a long time coming (see e.g. [Bromm 2013](#)).

The cosmic dark ages gradually came to an end when cosmic structure formation gathered pace, resulting in a sharp increase in the presence of luminous structure. This process is decisively supported by the presence of *dark matter*, outnumbering ordinary matter by a factor of 5 – 6 and only interacting gravitationally with standard model particles. As such, it forms a hierarchical structure of filaments and halos that we call the *cosmic web*, enshrouding vast low-density cosmic voids. These *halos* of dark matter act as supporting structures where ordinary matter can collapse further into *stars and galaxies*. Indeed, the James Webb Space Telescope (JWST) has already enabled the discovery of galaxies out to redshifts of $z \sim 14$ (e.g. [Carniani et al. 2024](#); [Witstok et al. 2025](#)), no more than 300 million years after the Big Bang. The light from these new luminous sources gradually ended the cosmic dark ages and initiated the cosmic era that is the subject of this thesis, the *Epoch of Reionisation*.

1.2 Reionisation – Les Lumières cosmologiques

Many cosmic epochs have been termed with distinctively pictorial names, some of them even in remarkable analogy to eras of the more recent *human* history. Coincidence or not, it is in any case worthwhile to fully appreciate the metaphor of the cosmic *dark ages* and the period that logically has to follow. In the previous millennium, the *Age of Enlightenment* ended the *medieval* dark ages when individuals emerged from their self-imposed immaturity ([Kant 1784](#)). In a similar vein, the Epoch of Reionisation – the ‘Age of *Cosmic* Enlightenment’ – is shaped by the formation of the first *astrophysical* individuals. These first stars, galaxies and quasars quite literally illuminated – or enlightened – our Universe, not only ending the cosmic dark ages but also causing a profound impact on the Universe’s subsequent evolution by connecting the properties of galaxies and other astrophysical sources of ionising photons to the large-scale history of our

Universe and its fundamental physical properties. In other words, the Epoch of Reionisation links astrophysics to cosmology, and neither the former nor the latter can be truly understood without a firm grasp of what happened in this era.

The estates where the reionisation process unfolds are what we call the *intergalactic medium* (IGM), composed of the vast voids between galaxies – the space where most cosmic baryons actually reside. The evolution of the IGM’s ionisation state is what informs us about the way reionisation proceeded and its interplay with the growth and physical properties of galaxies that drove it. The relevant physical field to that effect is the number density field of cosmic hydrogen n_{H} , separated into the two components

$$n_{\text{H}} = n_{\text{HI}} + n_{\text{HII}}, \quad (1.2)$$

where n_{HI} is the number density field of *neutral* hydrogen atoms and n_{HII} that of *ionised* ones. The ionisation state of the IGM is then commonly quantified in terms of the *fractional fields* $x_{\text{HI}} \equiv n_{\text{HI}}/n_{\text{H}}$ and $x_{\text{HII}} \equiv n_{\text{HII}}/n_{\text{H}}$. Of further relevance is the number density n_e of free electrons, given by counting the remaining ions, where – besides hydrogen – singly (He II) and doubly (He III) ionised helium are non-negligible contributors:

$$n_e = n_{\text{HII}} + n_{\text{HeII}} + 2n_{\text{HeIII}}. \quad (1.3)$$

Before reviewing past and ongoing modelling efforts aiming to trace the full evolution of these fields, we start by describing how the reionisation process proceeded on a *global* level.

1.2.1 The integrated effect of reionisation

Most fundamentally, from a cosmologist’s perspective, the impact of reionisation is summarised in a single free parameter, the reionisation optical depth τ_e – notably, one of *the* six free parameters of the Λ CDM model. This is because the Epoch of Reionisation leaves a characteristic imprint on the CMB due to Thomson scattering of CMB photons by the free electrons released during reionisation. Over the course of cosmic time, starting from the decoupling of the CMB at $z_{\text{CMB}} \simeq 1100$, this results in an integrated optical depth

$$\tau_e = \sigma_{\text{T}} \int_0^{z_{\text{CMB}}} n_e \frac{c(1+z)^2}{H(z)} dz, \quad (1.4)$$

where $\sigma_{\text{T}} = 8\pi/3 \cdot (e^2/4\pi\epsilon_0 m_e c^2)^2$ is the Thomson scattering cross section of the electron. The result of this is a suppression of the amplitude of the CMB temperature power spectrum proportional to a factor of $e^{-2\tau_e}$, in nearly perfect degeneracy with a reduction of the primordial amplitude of scalar perturbations A_s as a further free parameter of the Λ CDM model.

However, scattering of the large-scale CMB quadrupole radiation field also introduces a linear polarization pattern that can be identified in the E -mode polarisation power spectrum of the CMB as the so-called *reionisation bump* at low multipoles l , scaling as τ_e^2 (see e.g. [Reichardt 2016](#)). Due to these effects, an understanding of reionisation is indispensable for precision constraints on other cosmological parameters based on observations of the CMB ([Elbers 2025](#); [Jhaveri et al. 2025](#); [Sailer et al. 2026](#)).

1.2.2 The global timing of reionisation

Moving beyond its overall integrated effect, the logical next step consists in understanding how reionisation proceeded as a function of cosmic time. Constraining this *global timing* of reionisation is the most fundamental goal of all astrophysical reionisation studies and also a main objective of this thesis. We measure its progression in terms of the global *volume-averaged* neutral and ionised fractions $\langle x_{\text{HI}} \rangle$ and $\langle x_{\text{HII}} \rangle$ in the IGM, conventionally defined as $\langle x_{\text{HI}} \rangle \equiv \langle n_{\text{HI}} \rangle / \langle n_{\text{H}} \rangle$ and $\langle x_{\text{HII}} \rangle \equiv \langle n_{\text{HII}} \rangle / \langle n_{\text{H}} \rangle$. An evolution equation for these quantities can then be derived via simple photon counting arguments in a comoving patch of the Universe with volume V (see e.g. [Gnedin & Madau 2022](#)). Let $N_{\text{ion}}(t)$ be the number of ionising photons at cosmic time t in that given patch, then the number of ionised hydrogen atoms $N_{\text{HII}}(t)$ is equal to $N_{\text{ion}}(t)$ minus the number of those photons that recombine (on a given recombination time scale t_{rec}):

$$N_{\text{HII}}(t) = N_{\text{ion}}(t) - \int_0^t N_{\text{HII}} \frac{dt'}{t_{\text{rec}}}. \quad (1.5)$$

Averaging over a cosmologically representative volume V to take the limit $N_i/V \rightarrow \langle n_i \rangle$ for any of the aforementioned indices i , and differentiating with respect to t , we can derive the well-known Madau reionisation equation

$$\frac{d\langle x_{\text{HII}} \rangle}{dt} = \frac{\dot{n}_{\text{ion}}}{n_{\text{H},0}} - \frac{\langle x_{\text{HII}} \rangle}{t_{\text{rec}}} \quad (1.6)$$

that describes, on a global level, the transition of cosmic hydrogen from neutral to completely ionised ([Madau et al. 1999](#)).

It also emphasises the prominent distinction between *sources* and *sinks* of reionisation, embodied by the two distinct terms of the equation, whose nature is still subject to active investigation. Current consensus is that galaxies – as opposed to more powerful sources such as active galactic nuclei (AGN) and quasars – are the dominating sources of ionising photons, driving the reionisation process ([Madau & Dickinson 2014](#); [Bouwens et al. 2015](#); [Robertson et al. 2015](#); [Finkelstein et al. 2019](#)) (although JWST observations of early galaxies and AGN have been sparking new debate; see e.g. [Madau](#)

et al. 2024; Muñoz et al. 2024; Simmonds et al. 2024). Under this assumption, the source term in Eq. (1.6) is commonly parameterised into the three factors

$$\dot{n}_{\text{ion}} = f_{\text{esc}} \xi_{\text{ion}} \rho_{\text{SFR}}, \quad (1.7)$$

where ρ_{UV} is the UV luminosity density, ξ_{ion} the ionising photon production efficiency, i.e., the number of ionising photons emitted per unit UV luminosity, and f_{esc} is the fraction of these photons escaping the interstellar medium (ISM) and circumgalactic medium (CGM) and thus actually ionising hydrogen atoms in the diffuse IGM. While UV luminosity functions are well-determined (Bouwens et al. 2021; Adams et al. 2024; Donnan et al. 2024; Robertson et al. 2024), ξ_{ion} and f_{esc} are significantly harder to constrain and still remain an active research subject (Endsley et al. 2023; Atek et al. 2024; Muñoz et al. 2024; Saxena et al. 2024; Simmonds et al. 2024).

Besides these sources of ionising photons, we also have to account for recombinations via the sink term in Eq. (1.6). Their timescale t_{rec} is impacted by the actual gas density distribution whose effect we can absorb in the *clumping factor* that is commonly defined as

$$C_{\text{HII}} \equiv \frac{\langle n_{\text{HII}}^2 \rangle}{\langle n_{\text{HII}} \rangle^2}. \quad (1.8)$$

Its value has to be estimated externally via hydrodynamical simulations, commonly resulting in values of $C_{\text{HII}} \lesssim 3$ at $z \gtrsim 6$ (Gnedin & Ostriker 1997; Madau et al. 1999; Pawlik et al. 2009; Kaurov & Gnedin 2014; So et al. 2014). With a recombination rate $\dot{n}_{\text{rec}} = \alpha_{\text{HII}}^{\text{A}} \langle n_e n_{\text{HII}} \rangle \sim \alpha_{\text{HII}}^{\text{A}} \langle n_{\text{HII}}^2 \rangle = \alpha_{\text{HII}}^{\text{A}} C_{\text{HII}} \langle n_{\text{HII}} \rangle^2$, this then amounts to a recombination timescale

$$t_{\text{rec}} = \frac{\langle n_{\text{HII}} \rangle}{\dot{n}_{\text{rec}}} = \frac{1}{\alpha_{\text{HII}}^{\text{A}} \langle n_{\text{HII}} \rangle C_{\text{HII}}}, \quad (1.9)$$

where $\alpha_{\text{HII}}^{\text{A}}$ is the Case A recombination coefficient, accounting for recombinations into all excited states *and* the ground state (as appropriate for reionisation modelling; see Kaurov & Gnedin 2014).

1.2.3 The topology of reionisation

While Eq. (1.6) provides us with a basic analytical prescription of the temporal progression of reionisation on a global level, it is far from explaining reionisation to a full degree. This is because reionisation proceeds in a highly inhomogeneous, *patchy* way. In other words, the x_{HI} field is largely binary, with either a neutral or an ionised value at a given point in space. The composition of such neutral and ionised patches determines the *topology* of reionisation which describes the way reionisation proceeded on a spatial level and is thus also highly informative about the nature of the sources and sinks of ionising radiation.




Figure 1.2: Two-dimensional slice of a simulated ionisation field x_{HI} at redshift $z = 7.5$ with a global volume-average neutral fraction of $\langle x_{\text{HI}} \rangle = 0.5$, generated with the semi-numerical 21cmFAST code (Mesinger et al. 2011). White regions are neutral, black ones ionised. Dots denote the locations of the most massive halos within a 4 cMpc thick slice, color- and size-coded by their mass. Credit: Davies et al. (2018a).

Historically, two options have been considered for the overall progression of reionisation. First, *outside-in* scenarios, where photons – if they managed to escape from the dense host halos where the ionising sources are residing – would firstly ionise underdense regions of the IGM due to their modest need for ionising photons, whereas higher-density regions would first remain self-shielded from the ionising background until they would gradually get ionised as well (Miralda-Escudé et al. 2000).

On the other hand, motivated by the first numerical simulations of reionisation, Furlanetto et al. (2004) and Furlanetto & Oh (2005) developed a new model for the growth of ionised regions, instead favouring an *inside-out* scenario. Based on excursion-set theory (Bond et al. 1991) – in analogy to the Press-Schechter formalism for halo formation (Press & Schechter 1974) – they came to the conclusion that ionised bubbles would first form around clustered sources in overdense regions. These bubbles would then grow as time progresses, and eventually overlap and merge with other bubbles until the IGM has been ionised in its entirety. An example of a simulated topology at $\langle x_{\text{HI}} \rangle = 0.5$, building upon this model, is depicted in Figure 1.2.

Improved numerical simulations – as well as observational hints that Lyman- α emitting galaxies (LAEs) at high redshift appear to be highly clustered in what would correspond to overdense, highly ionised regions (Hu et al. 2021; Trapp et al. 2023) – have led us to conclude that reionisation indeed proceeds largely in an inside-out fashion. Cosmic filaments, on the other hand, due to their lack of ionising sources at increased densities compared to cosmic voids, might be the last regions to reionise in accordance with the original outside-in argument (Kaurov 2016). Further details about the sizes of ionised bubbles and the topology of reionisation are yet to be determined, and as such, a large body of research is still concerned with these topics (e.g. Mason & Gronke 2020; Hayes & Scarlata 2023; Lu et al. 2024b; Witstok et al. 2024).

1.3 Modelling Reionisation

Modelling reionisation is a highly non-trivial task due to the plethora of different physical effects and scales involved. As such, *numerical simulations* of reionisation have been developed in many different flavours of varying complexity, owing to the difficulty of modelling the radiative processes occurring in this era jointly with the gravitational evolution of dark and baryonic matter, forming the large-scale structure of the Universe.

Efforts to simulate the cosmological evolution of our Universe started out with pure *N-body simulations* focussing exclusively on the dynamics of collisionless dark matter that would collapse into what we call the *cosmic web* (e.g. Davis et al. 1985; Springel et al. 2005). As computational resources grew, it became possible to additionally include the fluid-like dynamics of baryonic matter, as well as a growing number of *sub-grid prescriptions* for varying types of *feedback* effects from physical processes occurring below the resolution limit of these simulations, for example due to star formation or active galactic nuclei (AGN), matching the observed populations of galaxies at a remarkable level of detail and enabling unique new insights into the physical processes behind their formation and evolution (e.g. Vogelsberger et al. 2014; Schaye et al. 2015; Pillepich et al. 2018; Davé et al. 2019; Pakmor et al. 2023; Schaye et al. 2023).

Nonetheless, accounting for *all* physical effects that are at play – even if not modelled from first principles – would greatly exceed the computational resources available to us. As such, the selection of the physical effects included in the modelling has to be tailored to the cosmological and astrophysical observables to compare the simulations with. In the context of reionisation, this prominently requires accounting for the radiative effects of the ionising sources. The required *radiative transfer* calculations are expensive, especially when performing full *radiation-hydrodynamical simulations* where the radiative effects are self-consistently coupled to a conventional

hydrodynamical simulation tracing the evolution of dark and baryonic matter (e.g. Gnedin 2014; Ocvirk et al. 2016, 2020; Garaldi et al. 2022; Kannan et al. 2022; Lewis et al. 2022). As a result, many approximate, *semi-numerical* schemes have been developed with the aim of reducing the computational complexity while still modelling reionisation at a reasonable level of precision (e.g. Mesinger et al. 2011). Before providing a quick overview over these approaches, we build intuition by discussing the basic underlying evolution equations that need to be solved.

1.3.1 Reionisation equations

To model reionisation, we turn back to the neutral fraction field x_{HI} introduced in Section 1.2. Rather than describing the global evolution of its volume-averaged version $\langle x_{\text{HI}} \rangle$, governed by Eq. (1.6), we now aim to construct an evolution equation for the field itself. We restrict this section to a basic set of simplified modelling equations, illustrating the most important contributions, and point the reader to Gnedin & Madau (2022) for a more comprehensive review of past and ongoing modelling approaches. We start by following the same reasoning as in Section 1.2, counting ionisation and recombination events to construct a simple rate equation with source and sink terms analogous to Eq. (1.6).

The atoms in a gas of neutral hydrogen can be ionised in two ways, via *photoionisation* and *collisional ionisation*, characterised by the ionisation rates Γ_{HI} and $\Gamma_{e\text{HI}}$. *Photoionisation* occurs if a Lyman continuum photon, i.e., a photon with wavelength shorter than the Lyman limit $\lambda_{\text{HI}} = 911.3 \text{ \AA}$ (corresponding to an energy greater than $h\nu_{\text{HI}} = 13.6 \text{ eV}$), encounters a neutral hydrogen atom. We characterise the underlying radiation field by its (space- and angle-averaged) specific intensity J_ν whose evolution is described by the *radiative transfer equation*

$$(\partial_t - \nu H(t) \partial_\nu) J_\nu + 3H(t) J_\nu = -c\kappa_\nu J_\nu + \frac{c}{4\pi} \epsilon_\nu, \quad (1.10)$$

where ϵ_ν is the proper space-averaged volume emissivity and k_ν is the absorption coefficient (see e.g. Peebles 1993; Haardt & Madau 1996, 2012).

For a given radiation field, we can then obtain the *photoionisation rate* Γ_{HI} of neutral hydrogen by integrating the product of the number $J_\nu/h\nu$ of photons at frequency ν and the HI absorption cross section $\sigma_{\text{HI}}(\nu)$ over all Lyman continuum frequencies:

$$\Gamma_{\text{HI}} = 4\pi \int_{\nu_{\text{HI}}}^{\infty} \frac{J_\nu}{h\nu} \sigma_{\text{HI}}(\nu) d\nu. \quad (1.11)$$

Collisional ionisation of neutral hydrogen atom is caused by thermal electrons, and hence its rate is given by

$$\Gamma_{e\text{HI}} = n_e \int_{\nu_{\text{HI}}}^{\infty} \sigma_{e\text{HI}}(v) v f(v) dv, \quad (1.12)$$

where $f(v)$ is the Maxwell–Boltzmann distribution of the electron velocity v , and $\sigma_{e\text{HI}}(v)$ is the collisional ionization cross section.

Recombination, i.e., the capture of a free electron by a proton, takes place at a rate given by the product of the free electron density n_e and the Case A recombination coefficient $\alpha_{\text{HII}}^{\text{A}}$ as introduced in Section 1.2.2, given by the sum

$$\alpha_{\text{HII}}^{\text{A}} = \sum_{i=1}^{\infty} \sum_{l=0}^{n-1} \int_{\nu_{\text{HI}}}^{\infty} \sigma_{nl}(v) v f(v) dv, \quad (1.13)$$

where $\sigma_{nl}(v)$ is the recombination cross section into the state characterised by the quantum numbers n and l .

Putting together all ionisation and recombination contributions, the full rate equation for the evolution of the neutral (or ionised) fraction field reads

$$\frac{dx_{\text{HI}}}{dt} = -x_{\text{HI}} \Gamma_{\text{HI}} - x_{\text{HI}} \Gamma_{e\text{HI}} + \alpha_{\text{HII}}^{\text{A}} n_e x_{\text{HII}}. \quad (1.14)$$

Although we omit them here for brevity, a full set of modelling equations would include analogous evolution equations for the helium fractions x_{HeI} , x_{HeII} and x_{HeIII} (see e.g. Bolton & Haehnelt 2007a).

An additional important consequence of the photoionisation process is that the IGM gets *photoheated* to temperatures of several thousands of Kelvin at a photoheating rate of

$$\mathcal{H}_{\text{HI}} = 4\pi n_{\text{HI}} \int_{\nu_{\text{HI}}}^{\infty} \frac{J_{\nu}}{h\nu} (h\nu - h\nu_{\text{HI}}) \sigma_{\text{HI}}(\nu) d\nu, \quad (1.15)$$

summarised along with all other sources of heating in the total heating rate \mathcal{H} . This means that we simultaneously have to trace the *thermal evolution* of the IGM based on all relevant heating and cooling sources, commonly modelled via the equation

$$\frac{dT}{dt} = -2H(t)T + \frac{2T}{3\Delta} \frac{d\Delta}{dt} + \frac{T}{\mu} \frac{d\mu}{dt} + \frac{(\gamma - 1)\mu}{k_{\text{B}}n} (\mathcal{H} - \Lambda), \quad (1.16)$$

where Δ is the cosmic overdensity field, μ the mean molecular weight and n the number density of all free baryonic particles (i.e., including ions *and* electrons), $\gamma = 5/3$ is the adiabatic index, and \mathcal{H} and Λ are the total heating and cooling rates, respectively (see e.g. Hui & Gnedin 1997).

This completes the set of the most important modelling equations of the radiative transfer problem during the reionisation era. Note that we are faced with a highly non-trivial set of coupled partial differential equations. Specifically, we have to solve the radiative transfer equation (1.10), governing the intensity of the ionising sources, jointly with the evolution equations of all ionised species (see Eq. (1.14) for hydrogen), coupled to each other via the photoionisation rate Γ_{HI} and the absorption coefficient

κ_ν . These equations, in return, are also coupled to the thermal evolution equation (1.16) via the heating and cooling rates \mathcal{H} and Λ . Conversely, the temperature evolution back-reacts on the collisional ionisation rate $\Gamma_{e\text{HI}}$ and the recombination coefficient α_{HII}^A by changing the thermal distribution of electrons $f(v)$. On top of this, everything is coupled to the evolution of dark and baryonic matter, governing the overdensity field Δ , as well as other astrophysical feedback effects that may have an impact on any of these fields. The solution of this coupled system of PDEs therefore poses a highly challenging numerical task.

1.3.2 Reionisation simulations

Self-consistent, fully coupled *radiation-hydrodynamical simulations* are considered the holy grail of these attempts, jointly modelling all of the aforementioned effects. Important examples of such simulation suites are CROC (Gnedin 2014), Thesan (Garaldi et al. 2022; Kannan et al. 2022) and CoDa (Ocvirk et al. 2016, 2020; Lewis et al. 2022), or the recently developed Nyx-RT code (Marshak et al. 2025) that is poised to extend the hydrodynamical Nyx simulations (Almgren et al. 2013; Lukić et al. 2015) to include radiative transfer. However, the computational cost of these simulations is considerable. As such, the underlying radiative transfer schemes often have to remain approximate, and the box sizes (~ 100 cMpc) and resolutions ($\sim 0.1 - 1$ pMpc at $z = 6$) of these simulations are not able to compete with those of their purely hydrodynamical counterparts. In addition, and most crucially, this means that most simulation suites are limited to at most a handful of high-resolution flagship runs, hindering a systematic exploration of the effects of variations in cosmological parameters and/or the timing and topology of reionisation.

A less costly alternative to such simulations are *partially coupled simulations* where the aforementioned modelling components are still included but not all of them are fully coupled to each other, for example by performing radiative transfer in post-processing on previously completed hydrodynamical simulations such as done in the C2-Ray simulations (Iliev et al. 2006; Mellema et al. 2006), or by Aubert & Teyssier (2010) with their Aton code, recently revised to also include helium (Asthana et al. 2024). A further example are the Sherwood-Relics simulations (Puchwein et al. 2023), building upon the Sherwood simulations (Bolton et al. 2017) and including a multi-step hybrid technique for treating radiative transfer effects to capture the effects of patchy reionisation on the IGM, allowing for a wider exploration of astrophysical and cosmological parameter variations.

Semi-numerical models are an even more efficient category of techniques that do not require cosmological (hydrodynamical or N-body) simulations to model the full non-linear distribution of cosmic matter but rather evolve the Gaussian field of cosmic initial conditions directly via perturbation

theory. The prime example of these techniques is the 21cmFAST code (Mesinger et al. 2011) which predicts the distribution of ionised gas at a given redshift via an excursion set-based formalism (Furlanetto et al. 2004; Furlanetto & Oh 2005) from the Zel'dovich approximation (Zel'dovich 1970). An alternative example is the Abundance Matching Box for the Epoch of Reionization (AMBER) that instead determines the reionisation-redshift field from a certain base field such as the angle-averaged specific intensity J_ν via an abundance matching approach (Trac et al. 2022). While being approximate, these methods come with the decisive advantage of allowing for cheap simulations of multiple large boxes with varying reionisation histories as needed for cosmological parameter inference.

This also motivates the hybrid approach for the simulation models adopted in this thesis: in order to simulate large cosmological volumes including the rarest overdensities where quasars are assumed to reside, and to explore variations with respect to the volume-averaged neutral fraction (x_{HI}) of the IGM, we make use of the efficient semi-numerical 21cmFAST code to simulate an array of different ionisation topologies. However, for an optimal modelling of the diffuse IGM and its impact on the Lyman- α forest observed in high-redshift quasar spectra, we combine these semi-numerical topologies with density, velocity and temperature fields from the Nyx hydrodynamical simulations that were designed specifically for this purpose. Finally, as we only have to model single quasar sightlines, the radiative transfer problem can significantly be simplified to a single dimension. Hence, we perform radiative transfer in post-processing on the sightlines that have been extracted from the aforementioned hydrodynamical simulations and ionisation topologies. This illustrates how the complexity of modelling reionisation is often broken down via custom-built approaches tailored to the problem at hand.

1.4 Quasars – Les Lumières astrophysiques

Having obtained an overview over the cosmological epoch that we would like to study, as well as the state of our modelling efforts to that end, we now turn to the observational probes – *the astrophysical lights* – that provide a means to gain this understanding. Yet another time, the first relevant historical developments date back to the previous century, to the late 1950s, when radio astronomers detected unusually bright, point-like *quasi-stellar radio sources* (see e.g. Edge et al. 1959), dubbed *quasars* (Chiu 1964), whose nature remained unclear at first. A breakthrough followed in 1963 in the form of the first detection of a – still point-like – visual counterpart to one of these sources which suggested an unusually high redshift of $z = 0.158$, placing the object *at a cosmological distance* despite its star-like appearance (Hazard et al. 1963; Schmidt 1963). Within the next few decades, this interpretation

was actively debated, until eventually consensus was reached that quasars were indeed as distant and powerful sources as originally suggested, making them *the brightest (non-transient) astrophysical lights in the entire Universe*, appearing on the sky as *quasi-stellar objects* (QSOs) despite their distance of millions of parsecs from our own galaxy.

An explanation for their extraordinary power was provided one year later by [Salpeter \(1964\)](#) by invoking the (at that point by no means proven) existence of *black holes*. In that picture, the quasar radiation would be released by gas in an accretion disc of gas falling onto a supermassive black hole (SMBH) in the centre of a galaxy. While highly controversial in these years, especially the first images of *shadows* of black holes from the Event Horizon Telescope ([Event Horizon Telescope Collaboration et al. 2019, 2022](#)) have convinced us of their existence, and that such a supermassive black hole should reside in the centre of any larger galaxy, many of which are shining as active galactic nuclei (AGN) and the most powerful ones as quasars (e.g. [Padovani et al. 2017](#)).

1.4.1 The Lyman- α transition

To date, quasars can be considered as our main source of information about the Epoch of Reionisation, owing to the fact that these objects already existed at sufficiently early cosmic stages, and to the exquisite sensitivity of their spectra to imprints from reionisation.

Particularly sensitive in this regard is the highly resonant *Lyman- α transition*. A Lyman- α photon corresponds to the transition of a neutral hydrogen atom between its 1s ground state and the 2p excited state at a wavelength of $\lambda_{\text{Ly}\alpha} = 1215.67 \text{ \AA}$. Its cross section can be calculated from first principles (see e.g. [Heitler 1954](#); [Bethe & Salpeter 1957](#); [Karzas & Latter 1961](#); [Berestetskii et al. 1971](#)), resulting in

$$\sigma_{\alpha}(\nu) = \frac{\pi e^2}{4\pi\epsilon_0 m_e c} f_{\alpha} \phi_{\alpha}(\nu), \quad (1.17)$$

where the oscillator strength $f_{\alpha} = 2^{13}/3^9 \simeq 0.416$ follows from the corresponding quantum-mechanical matrix element, and $\phi_{\alpha}(\nu)$ is the normalised line profile function, accounting for *line broadening*.

Contributions from two distinct line broadening mechanisms are relevant in the case at hand, resulting in the common expression of ϕ_{α} as a *Voigt profile* ϕ_{Voigt} ([Voigt 1912](#)). First, the line is *broadened naturally* due to the Heisenberg uncertainty principle ([Heisenberg 1927](#)) since energy and lifetime of the excited state cannot both be sharply determined at the same time. The consequence of this is a Lorentzian profile

$$\phi_{\text{Lor}}(\nu) = \frac{R_{\alpha}/(\pi\nu_{\alpha})}{(\nu/\nu_{\alpha} - 1)^2 + R_{\alpha}^2} \quad (1.18)$$

with the factor $R_\alpha = \Gamma_\alpha/4\pi\nu_\alpha$ determined by the decay constant $\Gamma_\alpha = 6.265 \times 10^8 \text{ s}^{-1}$ of the Lyman- α transition that is inversely related to the lifetime of the 2p state.

Second, thermal motion of the gas described by a Maxwell-Boltzmann velocity distribution causes *Doppler broadening* according to a Gaussian profile

$$\phi_{\text{Gauss}}(\nu) = \frac{1}{\sqrt{2\pi\sigma_{\nu_\alpha}^2}} e^{-(\nu-\nu_\alpha)^2/2\sigma_{\nu_\alpha}^2}, \quad (1.19)$$

where the gas temperature T sets the profile's width $\sigma_{\nu_\alpha} = \sqrt{\frac{k_B T}{m_e c^2}} \nu_\alpha$. The full Voigt profile ϕ_{Voigt} is then obtained by convolving ϕ_{Lor} and ϕ_{Gauss} .

1.4.2 The Lyman- α forest

For quasars, the Lyman- α transition is relevant in both emission and absorption: a strong Lyman- α emission line is produced due to the highly energetic, ionising UV radiation in the accretion disk, and the resulting recombination and resonant scattering events. As the radiation from the quasar propagates into the IGM, a characteristic absorption signature is imprinted upon its spectrum, the *Lyman- α forest*. This is because any cloud of neutral hydrogen that intersects our line of sight to the quasar absorbs those photons that have been redshifted to the Lyman- α resonance by the time they reach that cloud. The result is a *forest* of absorption and transmission spikes tracing the locations where there were and where there were no such clouds of neutral hydrogen present. This makes the Lyman- α forest an exceptional tool of modern cosmology, probing the cosmic large-scale structure via baryon acoustic oscillations (BAO) (McDonald 2003; Busca et al. 2013; Slosar et al. 2013; Font-Ribera et al. 2014; Delubac et al. 2015; du Mas des Bourboux et al. 2017, 2020; Alam et al. 2021; Adame et al. 2025) and the flux power spectrum (Croft et al. 1998; McDonald et al. 2000; Croft et al. 2002; Seljak et al. 2005, 2006; McDonald et al. 2006; Palanque-Delabrouille et al. 2013, 2015; Chabanier et al. 2019; Ivanov et al. 2025; Karaçaylı et al. 2025; Ravoux et al. 2025) whose smaller scales are also sensitive to the thermal state of the IGM (Hui & Gnedin 1997; Lidz et al. 2010; Becker et al. 2011; Boera et al. 2019; Walther et al. 2019; Gaikwad et al. 2020, 2021; Villasenor et al. 2022) and the free-streaming of dark matter (Viel et al. 2004, 2005, 2013; Garzilli et al. 2017; Iršič et al. 2017a,b; Villasenor et al. 2023; Garcia-Gallego et al. 2025b).

1.4.3 Quasars during the Epoch of Reionisation

Towards higher redshifts, when the Epoch of Reionisation has not fully ended yet, the presence of neutral hydrogen starts attenuating the Lyman- α forest




Figure 1.3: Two-dimensional representation of 527 quasar spectra at redshifts $z > 5.3$. The Lyman- α forest transmission seen blueward of the Lyman- α line gradually gets suppressed by increasingly large Gunn-Peterson troughs at higher redshifts. Credit: [Fan et al. \(2023\)](#).

as compared to how it appears at lower redshifts. To quantify this effect, we consider the Lyman- α transmission field

$$t(\lambda_{\text{obs}}) = e^{-\tau(\lambda_{\text{obs}})}, \quad (1.20)$$

where $\tau(\lambda_{\text{obs}})$ is the Lyman- α optical depth at the observed-frame wavelength λ_{obs} . The total optical depth $\tau(\lambda_{\text{obs}})$ is obtained by integrating an infinitesimal optical depth element $d\tau = \langle n_{\text{HI}} \rangle \sigma_{\alpha}(\nu) dl$ along the line-of-sight interval dl (e.g. [Miralda-Escudé 1998](#); [Mortlock 2016](#)), resulting in

$$\tau_{\text{DW}}(\lambda_{\text{obs}}) = \int n_{\text{HI}}(l) \sigma_{\alpha}(\nu) dl. \quad (1.21)$$

1.4.3.1 The Gunn-Peterson effect

As Eq. (1.21) suggests, any enhancement of the density of neutral hydrogen along the line of sight causes increased Lyman- α attenuation. Remarkably, the Lyman- α resonance is strong enough that even the last remainders of neutral hydrogen during the final stages of reionisation are sufficient to

cause deep absorption troughs adjacent to the Lyman- α line where we would otherwise observe Lyman- α forest spikes, as first noted by Gunn & Peterson (1965). The strength of these *Gunn-Peterson (GP) troughs* can be derived via the lowest-order approximation of the Lyman- α line profile as a Dirac delta distribution

$$\phi_\alpha(\nu) = \delta(\nu - \nu_\alpha), \quad (1.22)$$

resulting in the Gunn-Peterson optical depth

$$\tau_{\text{GP}}(z_{\text{QSO}}) = \frac{\pi e^2 f_\alpha \lambda_\alpha \langle n_{\text{HI}} \rangle(z_{\text{QSO}})}{4\pi \varepsilon_0 m_e c H(z_{\text{QSO}})} \quad (1.23)$$

$$\begin{aligned} &= 5.3 \times 10^5 \left(\frac{h}{0.68} \right) \left(\frac{1-Y}{0.76} \right) \left(\frac{\Omega_b}{0.049} \right) \left(\frac{\Omega_m}{0.31} \right)^{-\frac{1}{2}} \\ &\times \left(\frac{1+z_{\text{QSO}}}{1+7.54} \right)^{\frac{3}{2}} \left(\frac{H_0 \sqrt{\Omega_m (1+z_{\text{QSO}})^3}}{H(z_{\text{QSO}})} \right) \left(\frac{\langle x_{\text{HI}} \rangle(z_{\text{QSO}})}{1.0} \right) \end{aligned} \quad (1.24)$$

that acts on all photons with $\lambda_{\text{rest}} \leq \lambda_\alpha$ as they get redshifted into the Lyman- α resonance. This shows that even global volume-averaged neutral fractions of $\langle x_{\text{HI}} \rangle(z_{\text{QSO}}) = O(10^{-4})$ are sufficient to produce substantial GP absorption troughs. The emergence of these troughs at higher redshifts can be observed in Figure 1.3, representing 527 quasar spectra at redshifts $z > 5.3$. Measurements of τ_{GP} as well as length distributions of these absorption troughs have been among the earliest probes to place constraints on reionisation (Fan et al. 2000, 2002, 2006; Zhu et al. 2022).

1.4.3.2 Dark pixel fractions

The most straightforward – and in fact also the only model-independent – way of utilising this effect to constrain $\langle x_{\text{HI}} \rangle$ as a function of redshift z is counting the fraction of *dark pixels* in high-redshift quasar spectra as proposed by Mesinger (2010). Given that any non-extinguished pixel must be ionised, whereas, given the high sensitivity of the Lyman- α transition, dark regions could still be neutral or already ionised as well, this places upper limits on $\langle x_{\text{HI}} \rangle$ at a given redshift (McGreer et al. 2011, 2015; Jin et al. 2023) which can be complemented with their analogues from higher-order Lyman-series transitions (Davies et al. 2026). However, performing actual measurements of $\langle x_{\text{HI}} \rangle$ rather than placing upper limits is not possible with this technique, and the derived limits quickly approach unity as soon as the absorption troughs start extending over large parts of the spectral range, limiting the constraining power of dark pixel fractions to the end stages of reionisation.

1.4.3.3 Effective optical depth distributions

More quantitative measurements are possible by not only estimating *if* – but *how much* – Lyman- α flux is effectively transmitted at a given redshift z . This, however, requires a model for the intrinsic spectrum $s(\lambda_{\text{obs}})$ of the quasar, allowing us to determine the Lyman- α transmission fraction from the observed flux $f(\lambda_{\text{obs}})$ as

$$t(\lambda_{\text{obs}}) = \frac{f(\lambda_{\text{obs}})}{s(\lambda_{\text{obs}})}. \quad (1.25)$$

By evaluating this expression at a given position in the spectrum, we can obtain $t(z) = t(\lambda_{\text{obs}} = (1+z)\lambda_{\alpha})$ at the redshift corresponding to the respective observed-frame wavelength λ_{obs} , conventionally averaged over length or redshift intervals of a certain width. A given spectrum thus provides us with an array of measurements of $t(z)$, extending along the redshift range that the respective line of sight is probing. We can then convert these values to an *effective Gunn-Peterson optical depth*

$$\tau_{\text{eff}}(z) = -\ln t(z), \quad (1.26)$$

and the cumulative distribution function (CDF) of such $\tau_{\text{eff}}(z)$ measurements along many lines of sight can then be compared to the CDF obtained from simulations (Becker et al. 2015; Bosman et al. 2018; Eilers et al. 2018; Yang et al. 2020b). This approach has recently led to the conclusion that reionisation must have ended *late*, not before $z \sim 5.3$ (Bosman et al. 2022) because the $\tau_{\text{eff}}(z)$ distributions become incompatible with those obtained from homogeneous UV background (UVB) models starting at $z \gtrsim 5.3$. This suggests the presence of remaining neutral patches as late as $z \sim 5.3$ that could account for inhomogeneities in the UVB as favoured by the data.

1.4.3.4 The mean free path of ionising photons

As reionisation proceeded and ionised bubbles grew and merged, Lyman continuum photons could travel increasingly large distances before they would encounter an atom of neutral hydrogen that they would then ionise. In other words, the *mean free path* λ_{mfp} of ionising photons grew over time. Measurements of λ_{mfp} thus provide a further way of placing constraints on the timing of the reionisation era. In fact, the mean free path even forges a link between the IGM and the sources of ionising radiation since we can approximate the specific intensity of the ionising background as

$$J_{\nu} = \lambda_{\text{mfp}} (1+z)^2 f_{\text{esc}} \frac{\epsilon_{\nu}}{4\pi}, \quad (1.27)$$

assuming that the comoving mean free path λ_{mfp} , emitted by a population of Poisson distributed sources, is significantly smaller than the horizon

$\lambda_{\text{mfp}} \ll c/H(z)$ such that cosmological expansion does not matter when solving Eq. (1.10) (see e.g. Bolton & Haehnelt 2007b; Haardt & Madau 2012). Hence the photoionisation rate follows $\Gamma_{\text{HI}} \sim \epsilon \lambda_{\text{mfp}}$ as per Eq. (1.11) for a given spectral shape of the radiation field with total ionising emissivity ϵ .

Various techniques have been deployed to measure the mean free path, such as inferring Γ_{HI} and λ_{mfp} based on models that relate these quantities to measurements of the Lyman- α optical depth (Miralda-Escudé & Ostriker 1990; Meiksin & Madau 1993; Haardt & Madau 1996; Bolton & Haehnelt 2007b; Faucher-Giguère et al. 2008; Davies & Furlanetto 2016; D’Aloisio et al. 2018; Choudhury et al. 2021; Gaikwad et al. 2023). An alternative consists in directly measuring λ_{mfp} as the distance over which the transmitted ionising flux in mean quasar spectra has declined by a fraction $1/e$ (Prochaska et al. 2009; Fumagalli et al. 2013; O’Meara et al. 2013; Worseck et al. 2014; Becker et al. 2021). Lastly, the mean free path has also been constrained via the distribution of absorbers along lines of sight towards *individual* quasars (Songaila & Cowie 2010; Rudie et al. 2013; Romano et al. 2019; Bosman 2021). All these methods have lately converged on a rapid evolution of the mean free path between $5 \lesssim z \lesssim 6$ (Becker et al. 2021; Bosman 2021; Gaikwad et al. 2023), corroborating the picture of a rapid end to reionisation between these redshifts.

1.4.3.5 Damping wings

In order to use quasar spectra to obtain quantitative constraints on the Epoch of Reionisation even beyond its final stages, we can make use of the aforementioned natural broadening of the Lyman- α line which manifests as a so-called *damping wing* imprint in a quasar’s spectrum when the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ reaches order unity (Miralda-Escudé 1998) – and which is the subject of this thesis.

The signature is characterised by a suppression of the transmission *redward* of the Lyman- α line as depicted in Figure 1.4 for a simulated example sightline. As per Eq. (1.17), its shape and strength is determined by the Lyman- α cross section σ_{α} and the exact amount of neutral hydrogen along the line of sight. While technically the full Voigt profile would have to be considered, its shape is well-determined by the Lorentzian component ϕ_{Lor} . To build analytical intuition, we therefore restrict ourselves in this section to substituting Eq. (1.18) into Eq. (1.17), obtaining the damping wing optical depth

$$\tau_{\text{DW}}(\lambda_{\text{obs}}) = \frac{4R_{\alpha}\tau_{\text{GP}}(z_{\text{ion}})}{\pi} \left(\frac{\lambda_{\text{obs}}}{(1+z_{\text{ion}})\lambda_{\alpha}} \right)^{3/2} \times \left(I \left[\frac{(1+z_{\text{ion}})\lambda_{\alpha}}{\lambda_{\text{obs}}} \right] - I \left[\frac{(1+z_{\text{end}})\lambda_{\alpha}}{\lambda_{\text{obs}}} \right] \right), \quad (1.28)$$




Figure 1.4: Quasar IGM Damping Wings. *Upper:* Neutral fraction field x_{HI} along the sightline from a quasar shining for 10^6 yr into three different environments with global neutral fractions of $\langle x_{\text{HI}} \rangle = 0, 0.5$ and 1 . *Lower:* The more neutral the IGM, the stronger the damping wing imprinted upon the quasar’s intrinsic spectrum (black), and the smaller its proximity zone.

where we integrated from the redshift z_{ion} of the quasar’s ionisation front to the end of reionisation at z_{end} , and where $I(x)$ is a dimensionless factor given by

$$I(x) = \frac{x^{1/2}}{4(1-x)} + \frac{1}{8} \ln \left(\frac{1-x^{1/2}}{1+x^{1/2}} \right). \quad (1.29)$$

Mortlock et al. (2011) reported the first observational evidence for such an imprint in the spectrum of the quasar ULAS J1120+0641 at $z = 7.09$. This finding sparked a first wave of interest in probing reionisation with IGM damping wings, where broadly consistent constraints were derived from the aforementioned object and ULAS J1342+0928 at $z = 7.54$ (Bolton et al. 2011; Greig et al. 2017b; Bañados et al. 2018; Davies et al. 2018a; Greig et al. 2019; Āurovčíková et al. 2020; Reiman et al. 2020), as well as DES J0252-0503 at $z = 7.00$ and J1007+2115 at $z = 7.51$ (Wang et al. 2020; Yang et al. 2020a; Greig et al. 2022), and a couple of other lower-redshift quasars (Schroeder et al. 2013). However, due to the the scarcity of quasars at sufficiently high redshifts, the analyses remained limited to this handful of objects until Āurovčíková et al. (2024) and Greig et al. (2024a) presented the first studies of larger quasar samples, probing the end stages of reionisation between $5.8 \lesssim z \lesssim 7$ with quasar damping wings. To date, four largely independent analysis pipelines have been developed and applied to observational data (Greig et al. 2017a; Davies et al. 2018a; Āurovčíková et al. 2020; Hennawi et al. 2025). The last of these four pipelines also forms the basis of the analyses presented in this thesis. Establishing such a new analysis framework is particularly timely in the light of the now ongoing Euclid mission whose wide-field survey is anticipated to identify up to hundreds of new quasars at

$z \gtrsim 7$ suitable for damping wing analysis (Euclid Collaboration et al. 2019), with the first objects already having been identified (Euclid Collaboration et al. 2026).

In addition to that, damping wing absorption may also provide evidence for the persistence of the last neutral islands at the end stages of reionisation. As pointed out by Malloy & Lidz (2015), if such a neutral island were to intersect a quasar sightline, it could cause a damping wing in the Lyman- α forest region of the quasar spectrum, adjacent to the island's Gunn-Peterson absorption trough. The first hints of such signatures have recently been identified in individual (Becker et al. 2024) and stacked (Spina et al. 2024; Zhu et al. 2024) quasar spectra at $5.5 \lesssim z \lesssim 6$, but their interpretation still requires careful simulation modelling (Sawyer et al. 2025; Gnedin & Zhu 2025).

1.4.3.6 Proximity zones

The ionisation state of the IGM in the foreground of a quasar is not only influenced by the topology of reionisation but also by the quasar's own ionising radiation. The consequence of this radiation is a Mpc-scale ionised bubble around the quasar whose size grows with its (UV-luminous) lifetime t_Q . To first approximation, assuming that each emitted photon ionises exactly one hydrogen atom in a homogeneous surrounding IGM, its radius R_{ion} can be approximated as

$$R_{\text{ion}} = \left(\frac{3Q t_Q}{4\pi \langle n_{\text{HI}} \rangle} \right)^{1/3}, \quad (1.30)$$

where Q is the quasar's emission rate of ionising photons (Cen & Haiman 2000). The lifetime-dependence of R_{ion} establishes an important connection to the growth of the supermassive black holes (SMBHs) that are powering these quasars. SMBH growth mechanisms are actively debated (e.g. Bogdán et al. 2024; Maiolino et al. 2024a,b), trying to explain how these objects could grow so massive even at the earliest redshifts where quasars and AGN are observed. In combination with mass constraints, measurements of how long these objects have been shining as UV-luminous quasars are therefore essential in constraining their growth history.

Observationally, the ionised bubble around a quasar manifests as a *proximity zone* (or *nearzone*) R_{prox} in its spectrum blueward of the Lyman- α line. Due to the excess ionisation from the quasar, the transmission stays enhanced in this region of the spectrum, as seen in Figure 1.4. However, due to residual neutral gas or density clumps within an ionised bubble, as well as the observational measurement procedure¹, R_{prox} is not always identical

¹ R_{prox} is commonly defined as the location where the continuum-normalised flux, smoothed with a 20 Å boxcar filter, first drops below 10%.

to the radius of the ionisation front R_{ion} . As such, detailed proximity zone modelling is required (e.g. Davies et al. 2020; Satyavolu et al. 2023a; Zhou et al. 2024), but nonetheless, R_{prox} remains an important summary statistic, informative about the lifetime of a quasar, and measured extensively in the literature (e.g. Eilers et al. 2017, 2020; Ishimoto et al. 2020; Satyavolu et al. 2023b). However, due to the lifetime’s degeneracy with the global ionisation state of the IGM, a joint measurement of the two quantities – as conducted in this thesis – is inevitable when probing quasars at redshifts where reionisation is ongoing.

1.5 Other observational probes of Reionisation

Quasars are only a single class of objects in a diverse collection of observational probes that are used to place constraints on the Epoch of Reionisation. A selection of recent constraints derived from this variety of probes is depicted in Figure 1.5. First, a largely model-independent integral constraint can be derived from the *Cosmic Microwave Background* as measured by the Planck satellite (Planck Collaboration et al. 2020). *Galaxies* provide a considerably more numerous alternative to quasars, probing reionisation through a number of distinct features such as the abundance and the clustering of Lyman- α emitters (LAEs) (e.g. Ouchi et al. 2010; Sobacchi & Mesinger 2015; Ouchi et al. 2018; Umeda et al. 2025), equivalent width distributions of Lyman-break galaxies (LBGs) (e.g. Stark et al. 2010, 2011; Mason et al. 2018, 2019; Tang et al. 2024), or damping wing features in their spectra (e.g. Curtis-Lake et al. 2023; Mason et al. 2026; Umeda et al. 2026), although at a significantly reduced constraining power per object. Damping wings can further be identified in the afterglow spectra of *gamma-ray bursts* (e.g. Lamb & Reichardt 2000; Totani et al. 2006; Fausey et al. 2025a). Future radio observations of the *21 cm spin-flip transition* of neutral hydrogen (Furlanetto et al. 2006; Pritchard & Loeb 2012; Liu & Shaw 2020) and the *patchy kinetic Sunyaev-Zel’dovich (pkSZ) effect* (Sunyaev & Zeldovich 1980) targeted by current (Reichardt et al. 2021; Beringue et al. 2025) and next-generation (Matsumura et al. 2014; Abazajian et al. 2016; Ade et al. 2019) CMB surveys will set new standards in the coming decade. In this section, we provide a brief overview over these other probes.

1.5.1 The Cosmic Microwave Background

As discussed in Section 1.2.1, CMB photons carry information about reionisation as they get Thomson scattered off of the free electrons that get released during this epoch. As such, large CMB surveys such as the Wilkinson Microwave Anisotropy Probe (WMAP) and the Planck satellite have placed constraints on τ_e as defined in Eq. (1.4). While the first constraint from




Figure 1.5: Selected constraints on the history of reionisation obtained from a variety of probes. Grey regions are based on CMB measurements by [Planck Collaboration et al. \(2020\)](#) combined with dark pixel fractions from [McGreer et al. \(2015\)](#). Quasar-based constraints include Lyman- α effective optical depth distributions ([Bosman et al. 2022](#); [Gaikwad et al. 2023](#)), dark pixel fractions ([McGreer et al. 2015](#); [Jin et al. 2023](#); [Davies et al. 2026](#)), dark gap statistics ([Zhu et al. 2022](#)) and IGM damping wings ([Davies et al. 2018a](#); [Wang et al. 2020](#); [Yang et al. 2020a](#); [Ďurovčiková et al. 2020](#); [Reiman et al. 2020](#); [Greig et al. 2022, 2024a](#); [Ďurovčiková et al. 2024](#); [Kist et al. 2026](#)). Galaxy-based probes are luminosity functions and clustering of Lyman- α emitters ([Morales et al. 2021](#); [Kageura et al. 2025](#); [Umeda et al. 2025](#)), equivalent width distributions of Lyman-break galaxies ([Bolan et al. 2022](#); [Nakane et al. 2024](#); [Tang et al. 2024](#); [Jones et al. 2025](#)) and IGM damping wings ([Mason et al. 2026](#); [Umeda et al. 2026](#)). Damping wing constraints from afterglow spectra of gamma ray bursts have been obtained by [Totani et al. \(2006\)](#); [Patel et al. \(2010\)](#); [Chornock et al. \(2013\)](#); [Totani et al. \(2014\)](#); [Hartoog et al. \(2015\)](#); [Melandri et al. \(2015\)](#); [Totani et al. \(2016\)](#); [Fausey et al. \(2025a,b\)](#). For better visibility, the upper summary panel scales linearly with cosmic time.

WMAP was still as high as $\tau_e = 0.17 \pm 0.04$ (Kogut et al. 2003), it reduced to $\tau_e = 0.089 \pm 0.014$ (Hinshaw et al. 2013) based on the final 9-year data release. Planck then successively reduced the inferred value further, starting with a value of $\tau_e = 0.097 \pm 0.038$ (Planck Collaboration et al. 2014) consistent with WMAP that then decreased via $\tau_e = 0.079 \pm 0.017$ (Planck Collaboration et al. 2016) down to $\tau_e = 0.0544 \pm 0.00755$ (Planck Collaboration et al. 2020) as the final legacy Planck result.

Recent baryon acoustic oscillation (BAO) measurements from the Dark Energy Spectroscopic Instrument (DESI) have led to claims of a $\sim 3\sigma$ tension with the basic Λ CDM model as supported by the Planck results, interpreted as a possible hint towards dynamical dark energy or neutrino masses smaller than the upper bound derived from neutrino oscillation experiments (Abdul Karim et al. 2025). This has sparked a new interest in the reionisation optical depth τ_e as a possible means to alleviate these tensions by absorbing them into a higher value for τ_e than preferred by Planck’s legacy result, and more in line with the older Planck and WMAP constraints (Jhaveri et al. 2025; Sailer et al. 2026). This underpins the need for complementary constraints from other astrophysical probes of reionisation, independent of the Planck CMB data (Elbers 2025; Garcia-Gallego et al. 2025a; Kageura et al. 2026).

1.5.2 Galaxies

Galaxies – as the main drivers of reionisation (Madau & Dickinson 2014; Bouwens et al. 2015; Robertson et al. 2015; Finkelstein et al. 2019) – naturally carry a wealth of information about the progression of this epoch. The challenge in extracting this information consists in complex gas and star formation physics that simultaneously affect their evolution, rendering the derivable constraints highly model dependent. Their biggest advantage, on the other hand, is their sheer number density as compared to significantly rarer sources such as quasars. Moreover, JWST has allowed us to detect galaxies at unprecedented redshift frontiers, far beyond those of the most distant quasars known, reaching deep into the early stages of reionisation, with spectroscopically confirmed objects out to $z \sim 14$ (e.g. Bunker et al. 2023; Curtis-Lake et al. 2023; Carniani et al. 2024; Witstok et al. 2025).

Of particular interest for constraining reionisation is the class of *Lyman- α emitting galaxies* (LAEs). These typically young, low-mass and highly star-forming galaxies – with their characteristic Lyman- α emission line that gets attenuated with the increased presence of neutral hydrogen – therefore become harder to detect the more neutral hydrogen is present in the IGM. As a result, number counts of LAEs in the form of luminosity functions have extensively been used to constrain the reionisation state of the IGM as a function of redshift (e.g. Ouchi et al. 2010; Konno et al. 2014; Zheng et al. 2017; Inoue et al. 2018; Konno et al. 2018; Goto et al. 2021; Morales et al.

2021; Ning et al. 2022; Wold et al. 2022; Kageura et al. 2025; Umeda et al. 2025). In a similar manner, the LAEs that we do detect are more likely to reside in large ionised regions whereas those in neutral regions are more prone to be too attenuated. Measuring the clustering of LAEs therefore provides an alternative to simple number counts and can simultaneously yield information about bubble sizes and the topology of reionisation (Ouchi et al. 2010; Sobacchi & Mesinger 2015; Ouchi et al. 2018; Umeda et al. 2025).

Lyman-break galaxies, identified by the characteristic drop-off in their spectrum at the Lyman limit, are another class of galaxies often used to constrain reionisation by measuring their *Lyman- α equivalent widths*, i.e., the brightness of their Lyman- α line relative to their UV continuum. The distribution of equivalent width measurements then serves as a probe of $\langle x_{\text{HI}} \rangle$ because the fraction of objects with high equivalent widths, showing strong Lyman- α emission, is expected to decrease in a more neutral IGM as already argued above (e.g. Stark et al. 2010, 2011; Ono et al. 2012; Schenker et al. 2012; Mesinger et al. 2015; Mason et al. 2018; Hoag et al. 2019; Mason et al. 2019; Jung et al. 2020; Bolan et al. 2022; Bruton et al. 2023; Nakane et al. 2024; Tang et al. 2024; Jones et al. 2025).

In addition to these probes, recent JWST observations have led to the first claims of *damping wing absorption* in individual (Curtis-Lake et al. 2023; Hsiao et al. 2024; Park et al. 2025; Mason et al. 2026) and stacked (Umeda et al. 2024, 2026) spectra of high-redshift galaxies. However, it has been pointed out that these signatures are in many cases dominated by very similar imprints from localised damped Lyman- α absorption systems (DLAs) in the foreground of these galaxies Heintz et al. (2024, 2025); Huberty et al. (2025). As such, the existing constraints need to be treated with caution and only marginalisation over the potential presence of DLAs can be expected to yield robust constraints (Mason et al. 2026).

1.5.3 Gamma-ray bursts

Damping wings have not only been reported towards galaxies and quasars, but also towards *gamma-ray bursts* (GRBs) – the brightest transient events seen on the sky, most likely to be associated to superluminous supernova events. Their brightness allows them to be observed out to extremely high redshifts, potentially out to $z \sim 20$ (Lamb & Reichart 2000). The fact that their afterglow spectra that can be modelled as a simple power law spectrum as well as the absence of significant own ionising radiation makes GRBs a conceptually very clean probe of the ionisation state of the IGM. The caveat is that such events are extremely rare and not more than a handful of such events have been detected at $z > 6$ to date (Totani et al. 2006; Patel et al. 2010; Chornock et al. 2013; Melandri et al. 2015; Fausey et al. 2025a), and, as a matter of fact, the derived constraints on $\langle x_{\text{HI}} \rangle$ have not always been

consistent (Chornock et al. 2013; Totani et al. 2014; Hartoog et al. 2015; Totani et al. 2016; Fausey et al. 2025b).

1.5.4 The 21 cm line

The 21 cm spin flip transition of neutral hydrogen, lending its name from its rest-frame wavelength of $\lambda_{21\text{ cm}} = 21.106\text{ cm}$, is the transition between the two hyperfine levels of the 1s ground state of the neutral hydrogen atom. Due to the ubiquitous presence of neutral hydrogen in the IGM, the hope is that observations of this 21 cm signal – at present redshifted to radio frequencies – would allow us to create detailed *intensity maps* of neutral hydrogen at any position in the sky, probing the progression of reionisation at a level of spatial detail unachievable with any of the aforementioned astrophysical probes (see e.g. Furlanetto et al. 2006; Pritchard & Loeb 2012; Liu & Shaw 2020). Despite this conceptual simplicity and promise, the practical detection of the 21 cm signal has proven as a considerable challenge due to significant foregrounds and instrumental systematics, orders of magnitude larger than the signal itself (see e.g. Cheng et al. 2018).

Nonetheless, sensitivities are steadily increasing and ongoing experiments have already placed first upper limits on the strength of the signal (Mertens et al. 2020; Trott et al. 2020; Abdurashidova et al. 2022). A boost in sensitivity is expected in the coming decade with the advent of the low-frequency array of the Square Kilometre Array (SKA-low) (Koopmans et al. 2015).

1.5.5 The patchy kSZ effect

The Cosmic Microwave Background carries another signature from the reionisation era that could potentially be utilised by the next generation of CMB surveys. The *Sunyaev-Zel'dovich (SZ) effect* arises when CMB photons inverse-Compton scatter off of high-energy electrons, causing characteristic anisotropies in the CMB (Sunyaev & Zeldovich 1980). The effect exists in two flavours: the *thermal SZ* (tSZ) effect caused by energy transfer from hot electrons to colder CMB photons, and the *kinetic SZ* (kSZ) effect due to scattering off of electrons that are in bulk motion relative to the CMB, resulting in a Doppler shift that then produces small temperature anisotropies in the CMB. While the thermal SZ effect is dominated by the signal from galaxy clusters, the patchy nature of reionisation causes a non-negligible contribution to the kSZ effect. This is due to the free electrons within the ionised bubbles that are growing during that epoch. The consequence of the bulk motion of these bubbles is a *patchy kSZ* (pkSZ) imprint on the CMB (Gruzinov & Hu 1998; Knox et al. 1998).

However, this imprint has to be disentangled from other post-reionisation contributions to the total kSZ signal. First tentative detections have been

reported by the South Pole Telescope (SPT) (Reichardt et al. 2021) and the Atacama Cosmology Telescope (ACT) (Beringue et al. 2025), and more conclusive measurements are expected from the next stage of CMB surveys such as LiteBIRD (Matsumura et al. 2014), CMB-S4 (Abazajian et al. 2016) and the Simons Observatory (Ade et al. 2019).

1.6 Inference – Le bon sens réduit au calcul

Having reviewed the theory and modelling efforts behind the cosmic epoch that we are trying to constrain, as well as its observational signatures, the only missing ingredient is a framework that allows us to connect the former and the latter and draw conclusions from the observational data on our models of reionisation. A fundamental limitation inherent to cosmology is that we cannot set up and conduct experiments in the same vein as experimentalists in other fields of physics are able to. Instead, all we are provided with is *a single realisation of our Universe* – the one that we are living in – and we have to constrain our models based exclusively on observations of this single realisation, making the problem naturally suited for *Bayesian statistics*.

1.6.1 Bayes’ theorem

Perhaps it is more than a coincidence that the development of the theory behind this brings us back from the *cosmic* to the *human* Age of Enlightenment. It was in 1763 when Richard Price published the manuscript of his friend Thomas Bayes who had passed away two years earlier, containing the first application of what we now know as *Bayes’ theorem* (Bayes & Price 1763). Re-derived and extended by Pierre-Simon Laplace (Laplace 1812, 1814, 1820), probability in its Bayesian notion constitutes our subjective measure of belief of a given hypothesis or model. Bayesian logic is therefore sometimes interpreted as an extension of the Boolean logic of *true* and *false* to continuous probability values in between these two extremes, or, as Laplace famously put it himself, as *essentially just the extension of common sense to calculus* (Laplace 1814).

More formally speaking, we in the situation where we are provided with a theoretical model \mathcal{M} whose free parameters $\boldsymbol{\theta}$ we are aiming to constrain based on a set of data \boldsymbol{x} that we collected. In other words, we are trying to determine the *posterior* probability distribution $P(\boldsymbol{\theta}|\boldsymbol{x})$, quantifying our degree of belief in the parameters $\boldsymbol{\theta}$ after having seen the data \boldsymbol{x} . The law of conditional probabilities then allows us to write down Bayes’ theorem, relating the *posterior* probability to three other distributions:

$$P(\boldsymbol{\theta}|\boldsymbol{x}) = \frac{L(\boldsymbol{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{P(\boldsymbol{x})}. \quad (1.31)$$

Here $L(\mathbf{x}|\boldsymbol{\theta})$ is the *likelihood* of observing the data \mathbf{x} , given a fixed set of parameters $\boldsymbol{\theta}$, encoding the *predictions of our theory*. The *prior* probability $\pi(\boldsymbol{\theta})$ encodes our *prior beliefs* about what values of $\boldsymbol{\theta}$ we deem likely or unlikely. These beliefs could, for example, be based on physical arguments or previous experiments. Finally, the *evidence* $P(\mathbf{x})$ quantifies the probability of observing the data \mathbf{x} *at all* (in the framework of our model \mathcal{M}). Computing $P(\mathbf{x})$ in the context of different models therefore also provides a way of performing *model comparison* by determining which model maximises the chance of collecting the observed data whatsoever. However, computing $P(\mathbf{x}) = \int d^n\boldsymbol{\theta} L(\mathbf{x}|\boldsymbol{\theta}) \pi(\boldsymbol{\theta})$ can be highly challenging since it involves an integral over the parameter vector $\boldsymbol{\theta}$ whose dimensionality can be high. On the other hand, from the standpoint of evaluating the posterior $P(\boldsymbol{\theta}|\mathbf{x})$ as a function of the parameters $\boldsymbol{\theta}$, the evidence $P(\mathbf{x})$ does not constitute more than a normalisation constant and therefore does not warrant explicit computation when no comparison of different models is required.

1.6.2 Sampling

After setting up the posterior distribution $P(\boldsymbol{\theta}|\mathbf{x})$ by specifying likelihood and prior, we are still faced with the difficulty of practically evaluating Eq. (1.31). A naive evaluation on a parameter grid quickly becomes infeasible in high dimensions, and as such, the common approach is to resort to *sampling techniques*, amongst which *Markov Chain Monte Carlo* (MCMC) methods hold the most prominent position. The many different flavours of MCMC samplers that have been implemented to date all share the common idea of constructing a *Markov Chain* – a sequence of samples $\{\boldsymbol{\theta}_i\}_{i \in \mathbb{N}}$ where a given sample $\boldsymbol{\theta}_i$ only depends on its immediate precursor $\boldsymbol{\theta}_{i-1}$ – whose density approximates the distribution of interest, i.e., $P(\boldsymbol{\theta}|\mathbf{x})$. In essence, this is achieved via a random walk through the posterior distribution $P(\boldsymbol{\theta}|\mathbf{x})$. The first incarnation of this idea is the *Metropolis-Hastings algorithm* (Metropolis et al. 1953; Hastings 1970) which, at a given location in the chain, draws a new sample from a given proposal distribution, and either accepts or discards the proposed sample according to a specific acceptance criterion, ensuring that the principle of *detailed balance*, i.e., the reversibility of each transition, is satisfied.

The original Metropolis-Hastings algorithm has been extended in various ways in order to improve its efficiency and allow for a more favourable scaling to high dimensions. The specific technique that we will be applying in this thesis is the *Hamiltonian Monte Carlo* (HMC) algorithm (Duane et al. 1987; Neal 1996) that combines the Metropolis-Hastings algorithm with *Hamiltonian dynamics* (Hamilton 1834, 1835) to allow for a more efficient exploration of the probability landscape than a naive random walk

allows. Treating the parameter vector $\boldsymbol{\theta}$ as position variable, its dynamics is governed by the Hamiltonian

$$\mathcal{H}(\boldsymbol{\theta}, \mathbf{p}) = U(\boldsymbol{\theta}) + \frac{1}{2} \mathbf{p}^T M^{-1} \mathbf{p}, \quad (1.32)$$

where the potential energy is chosen to be the negative logarithmic probability distribution of interest

$$U(\boldsymbol{\theta}) \equiv -\ln P(\boldsymbol{\theta}|\mathbf{x}), \quad (1.33)$$

and the kinetic term is defined by a given symmetric and positive definite *mass matrix* M for a set of artificially introduced conjugated momenta \mathbf{p} . After drawing a set of random momenta \mathbf{p}_i at a given location in parameter space $\boldsymbol{\theta}_i$, *Hamilton's equations*

$$\frac{d\boldsymbol{\theta}}{dt} = \vec{\nabla}_{\mathbf{p}} \mathcal{H} = M^{-1} \mathbf{p} \quad (1.34)$$

$$\frac{d\mathbf{p}}{dt} = -\vec{\nabla}_{\boldsymbol{\theta}} \mathcal{H} = \vec{\nabla}_{\boldsymbol{\theta}} \ln P(\boldsymbol{\theta}|\mathbf{x}) \quad (1.35)$$

can then be solved via a symplectic integrator such as the Leapfrog algorithm to propose a new sample $\boldsymbol{\theta}_{i+1}$ that is accepted or rejected according to the standard Metropolis-Hastings acceptance criterion. The advantage of the HMC algorithm is that it makes an informed guess for $\boldsymbol{\theta}_{i+1}$ by following a Hamiltonian trajectory through the potential landscape defined by the probability distribution, rather than drawing the new sample from a fixed proposal distribution, maximising the acceptance rate.

The number L of integration steps as well as the step size ϵ of the integrator are critical hyperparameters of the algorithm. Choosing these parameters too small would result in samples that are close to one another, effectively just recovering the inefficient random walk behaviour that the algorithm is designed to improve upon. A too large step size, on the other hand, would result in inaccurate integrations and hence a lower acceptance rate, while too many steps (of an adequate size ϵ) would eventually lead to a point where the trajectory starts to repeat itself in an oscillatory fashion, and hence the repeated computation of the same dynamics. The *No-U-Turn Sampler* (Hoffman & Gelman 2011) avoids these issues by successively constructing a binary tree of forward and backward leapfrog steps until the point where the trajectory would start to loop around, and then choosing a new sample from that tree.

As Eq. (1.35) suggests, the Hamiltonian Monte Carlo algorithm requires – and harnesses – gradient information. As such, it has only been able to gather pace in the past few years with the rise of *automatic differentiation* techniques as an alternative to symbolic and numerical differentiation. Unlike these approaches, automatic differentiation frameworks directly accumulate derivatives of composite functions via the repeated application of the chain

rule to all of their components. This has been a decisive factor that facilitated the current success of machine learning techniques by enabling efficient gradient-based decent and backpropagation. As such, the Hamiltonian Monte Carlo sampler takes advantage of the latest developments in computational infrastructure.

1.7 This Thesis – Les Lumières

Les Lumières – in the broader or the stricter sense, the title of this thesis unites all its main ingredients: the cosmological epoch we are aiming to constrain, the astrophysical lights we are using for that purpose, and the Bayesian way of reasoning that allows us to connect the former and the latter. Most of the chapters that will follow are based on or will extend the Bayesian inference pipeline for quasar damping wing analysis that has recently been introduced by Hennawi et al. (2025).

In **Chapter 2**, we leverage this pipeline to perform a detailed analysis of the precision of quasar IGM damping wing measurements. Several factors are contributing to the total error budget and need to be accounted for. 1) The (a priori unknown) *intrinsic continuum* of the quasar has to be reconstructed in order to disentangle it from the intergalactic absorption signature, introducing a continuum reconstruction error that will contribute to the overall uncertainty. 2) Since reionisation is a patchy process, sightline-to-sightline variations are a crucial factor. Especially during the mid-stages of reionisation, certain individual sightlines might still be mostly neutral while others are significantly ionised. A detailed modelling of the associated scatter is only possible via numerical simulations and introduces significant uncertainties. 3) The ionisation front of a given quasar – whose size is largely determined by the quasar’s lifetime – has a significant impact on the damping wing strength. As such, the lifetime of the quasar has to be inferred simultaneously in order to be able to marginalise over its degeneracy with $\langle x_{\text{HI}} \rangle$. 4) The observational setup determines the inference precision that we are ultimately able to achieve – and sets requirements on the sophistication of our likelihood prescription.

In this chapter, we perform a detailed study of all these contributions. The first part of the chapter is devoted to determining the optimal observational setup in terms of wavelength coverage, signal-to-noise ratio and spectral resolution. Crucially, we find that our approximate likelihood prescription prevents us from benefitting from higher-resolution spectra since the precision gains are evened out by the corrections we have to apply as to not quote overconfident (say: unjustly narrow) posterior distributions. We also determine the optimal dimensionality of our model for the quasar continuum and then proceed by breaking down the total error budget of IGM damping wing measurements as a function of astrophysical parameter

space. We find that we are able to constrain the IGM neutral fraction to $28.0^{+8.2}_{-8.8}$ % and the quasar lifetime to $0.80^{+0.22}_{-0.55}$ dex, where approximately half of the associated error budget is due to the continuum reconstruction task and the other half due to the stochastic nature of reionisation.

In **Chapter 3**, we further address that second contribution and introduce two new summary statistics that are robust against cosmic variance in the IGM and informative not only about the global timing of reionisation as quantified by $\langle x_{\text{HI}} \rangle$ but also the *local* ionisation topology. We can extract this local line-of-sight information by defining 1) the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ whose weighting mimics the frequency dependence of the Lyman- α cross section, and 2) the quasar’s distance r_{patch} to the first neutral patch. Both these statistics are defined on the original topology *before* it is altered by the quasar’s own ionising radiation, and we propose to use them in combination with the lifetime t_{Q} of the quasar, summarising the effects of its ionising radiation, as an improved parameterisation of quasar IGM damping wings. We demonstrate that this parameterisation reduces the scatter of the IGM transmission in the damping wing region of the spectrum down to $\lesssim 1$ %.

In the final section of the chapter, we show that these statistics are independent of the chosen model of reionisation. In other words, when selecting sightlines that share the same values of our summary statistics, all scatter in the resulting IGM transmission profiles solely arises from density fluctuations in the IGM, independently of what ionisation topology the sightlines originated from. We demonstrate the latter by introducing a simple numerical procedure for generating neutral fraction skewers, showing that the statistical properties of the resulting IGM transmission profiles match those derived from realistic semi-numerical ionisation topologies.

In **Chapter 4**, we put this new IGM damping wing parameterisation into action by incorporating it into our previously introduced Bayesian inference pipeline. The aforementioned topology-independence allows us to constrain our new local summary statistics in a *model-independent* fashion. We show that we can tie these constraints to any given reionisation model in a separate subsequent step. In our framework, the assumption of a specific reionisation model amounts to imposing a specific prior on our local summary statistics, given by their distribution within that assumed topology. This then allows us to also quote a *model-dependent* constraint on the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ – tied to the specific reionisation model that was assumed. By performing inference on a large set of mock spectra, we demonstrate that our local, model-independent measurements constrain the Lorentzian-weighted HI column density to $0.69^{+0.06}_{-0.30}$ dex and the quasar’s original distance to the first neutral patch to $31.4^{+10.7}_{-28.1}$ cMpc if a noticeable damping wing is present in the spectrum.

Finally, in **Chapter 5**, we apply our new pipeline to JWST/NIRSpec spectra of two of the highest-redshift quasars known to date: J1007+2115 at

$z = 7.51$ and J1342+0928 at $z = 7.54$. After folding in the model-dependent topology information, we find that J1007+2115 is most likely to reside in a $\langle x_{\text{HI}} \rangle = 0.32^{+0.22}_{-0.20}$ neutral IGM while shining for a remarkably short time of $\log_{10} t_{\text{Q}} = 4.14^{+0.74}_{-0.18}$. The constraints on our local summary statistics turn out to be largely prior-dominated for this object, as is common for a sightline without a clear damping wing imprint in its spectrum, suggesting that the sightline is largely ionised. For J1342+0928, we infer a neutral fraction of $\langle x_{\text{HI}} \rangle = 0.58^{+0.23}_{-0.23}$ along with an intermediately long lifetime of $\log_{10} t_{\text{Q}} = 5.64^{+0.25}_{-0.43}$. This source does exhibit a non-negligible damping wing signature, and we are able to constrain its column density to $N_{\text{HI}}^{\text{DW}} = 20.24^{+0.25}_{-0.22}$, along with a distance $r_{\text{patch}} = 10.9^{+5.6}_{-5.9}$ cMpc to the first neutral patch. We find our constraints on $\langle x_{\text{HI}} \rangle$ and t_{Q} derived from both these objects to be consistent with those quoted in the literature.

In addition, we also demonstrate how our new parameterisation naturally allows us to quantify the impact of potential local absorption systems in the foreground of a source on the observed damping wing signal. The possibility of the presence of such absorbers in front of J1342+0928 has recently been pointed out by [Davies et al. \(2025\)](#). In our parameterisation, this can simply be understood as an additional contribution to the Lorentzian-weighted column density.

1.8 Outlook – Les nouvelles Lumières

More conclusive constraints on the Epoch of Reionisation are looming very near at the horizon – across all observational probes, and for quasar IGM damping wings in particular. The launch of JWST has allowed us to find galaxies at redshifts entirely unheard of, reaching deep into the core stages of reionisation (e.g. [Bunker et al. 2023](#); [Curtis-Lake et al. 2023](#); [Carniani et al. 2024](#); [Witstok et al. 2025](#)). While significant efforts will be necessary to robustly extract information from the damping wing imprints towards these sources ([Mason et al. 2026](#)) or the spatial distributions of LAEs ([Maitra et al. 2026](#)), these probes will be complemented in the coming decade with new observing facilities that are poised to reveal first robust measurements of the cosmic 21 cm signal and the patchy kSZ effect. Both these signals are currently still hidden behind foregrounds, systematics and other nuisances, but the next generation of facilities – SKA-low ([Koopmans et al. 2015](#)), as well as LiteBIRD ([Matsumura et al. 2014](#)), CMB-S4 ([Abazajian et al. 2016](#)) and the Simons Observatory ([Ade et al. 2019](#)) – have the potential of constraining reionisation at an unprecedented level of precision.

The next leap forward in the context of quasar IGM damping wings is even closer ahead on the horizon. The Euclid wide-field survey has already drastically increased the number of known quasars at $z \gtrsim 7$ ([Euclid Collaboration et al. 2026](#)), and is anticipated to identify up to hundreds

of more objects in the coming years (Euclid Collaboration et al. 2019). Spectroscopic follow-up of these new discoveries with telescopes such as JWST will provide us with unprecedented samples for quasar damping wing analysis, with a first sample of 27 new JWST/NIRSpec spectra at $6.5 < z < 7.7$ already being taken at this writing.

This thesis aims to establish a solid statistical groundwork for their analysis. The new parameterisation introduced in Chapters 3 and 4 will make it possible for future analyses to separate out the model-independent extraction of constraints on the local ionisation topology from their model-dependent conversion to constraints on the global timing of reionisation. Larger sample sizes will eventually also open up possibilities for the comparison of different reionisation models by inferring the actual distribution of these local summaries. Further directions include their extension to the case of galaxies as background sources. Due to the higher prevalence of localised absorption systems in the foreground of galaxies as compared to quasars (Huberty et al. 2025), incorporating the absorber column density as an additional contribution to the inferred total Lorentzian-weighted column density as proposed in Chapter 5 will be particularly useful for jointly constraining the intergalactic damping wing signal and that due to localised absorption systems (Mason et al. 2026).

A further caveat identified in Chapter 2 remains unaddressed by this thesis. The approximate nature of our likelihood prescription introduces increasingly large biases when analysing higher-resolution data. In order to quote statistically faithful constraints, we are thus forced to retrospectively broaden our posteriors by an amount that makes them comparable in width to those inferred from low-resolution data. A natural approach to improve upon this issue are *simulation-based inference* methods (Alsing et al. 2019; Cranmer et al. 2020) that do not require any assumptions about the underlying likelihood by deploying a neural network model to learn the likelihood function directly from the simulated data. The application of simulation-based inference techniques therefore has the prospect of significant precision gains from high-resolution data. Addressing these points, and applying our techniques to larger samples of present (D’Odorico et al. 2023; Onorato et al. 2025) and upcoming (Euclid Collaboration et al. 2026) quasar spectra at ever higher redshifts (Euclid Collaboration et al. 2019) will soon shed bright new light on our understanding of the Epoch of Reionisation.



2 | QUANTIFYING THE PRECISION OF IGM DAMPING WING MEASUREMENTS TOWARDS QUASARS

Abstract

We investigate the precision with which the Lyman- α damping wing signature imprinted on the spectra of high-redshift quasars (QSOs) by the foreground neutral intergalactic medium (IGM) can measure the history of cosmic reionization. We leverage a novel inference pipeline based on a generative probabilistic model for the entire spectrum (both red- and blueward of the Lyman- α line), accounting for all relevant sources of uncertainty – the stochasticity caused by patchy reionization, the impact of the quasar’s ionizing radiation on the IGM, its unknown intrinsic spectrum, and spectral noise. Performing fast JAX-based Hamiltonian Monte-Carlo (HMC) parameter inference, we precisely measure the underlying global IGM neutral fraction as well as the lifetime of the quasar. Running a battery of tests on over a thousand mocks, we find optimal precision when running the pipeline with a six parameter PCA continuum model (five coefficients and a normalization) on $S/N \sim 10$ spectra, binned to a ~ 500 km/s velocity pixel scale, and extending at least out to the C IV $\lambda 1549$ Å emission line. After marginalizing out nuisance parameters associated with the quasar continuum, a single spectrum constrains the IGM neutral fraction to $28.0^{+8.2}_{-8.8}$ % and the quasar lifetime to $0.80^{+0.22}_{-0.55}$ dex, improving notably towards spectra with a stronger IGM damping wing imprint. Higher precision can be achieved by averaging over statistical quasar samples. We identify two primary sources of uncertainty that contribute approximately equally to the total error budget: the uncertain quasar continuum model and the stochastic distribution of neutral regions arising from both the reionization topology and the location of the quasar’s ionization front.

Work published in: **Timo Kist**, Joseph F. Hennawi and Frederick B. Davies, *Monthly Notices of the Royal Astronomical Society*, Volume 538, Issue 4, April 2025, Pages 2704–2728, doi.org/10.1093/mnras/staf460. Reprinted here in its entirety.

2.1 Introduction

The Lyman- α damping wing signature of high-redshift quasars in the intergalactic medium (IGM) provides a unique way of probing the history of reionization. Being highly sensitive to the presence of intergalactic neutral hydrogen, Lyman- α absorption saturates at global volume-averaged IGM neutral fractions as small as $\langle x_{\text{HI}} \rangle \sim 10^{-4}$ (Gunn & Peterson 1965). Even more importantly, at $\langle x_{\text{HI}} \rangle = O(0.1)$, the absorption signature also starts to extend reward of the Lyman- α line in the form of the quantum-mechanical line-broadening that is well-described by the Lorentzian component of the Voigt profile (Miralda-Escudé 1998). This *damping wing* absorption profile is a direct hallmark of the IGM optical depth and thus the reionization state of the IGM at the redshift of the observed object.

Predicted by Miralda-Escudé (1998), a quasar IGM damping wing was first claimed for the $z = 7.09$ QSO ULAS J1120+0641 by Mortlock et al. (2011). Re-analysis of the same object (as well as the $z = 7.54$ QSO ULAS J1342+0928) with different approaches has resulted in varying constraints on $\langle x_{\text{HI}} \rangle$ (Bolton et al. 2011; Greig et al. 2017b; Bañados et al. 2018; Davies et al. 2018a; Greig et al. 2019; Āurovčíková et al. 2020; Reiman et al. 2020). To date, three distinct analysis pipelines (Greig et al. 2017a; Davies et al. 2018a; Āurovčíková et al. 2020) have been applied to an additional number of QSOs at the highest redshifts currently accessible such as DES J0252-0503 at $z = 7.00$ and J1007+2115 at $z = 7.51$ (Wang et al. 2020; Yang et al. 2020a; Greig et al. 2022). Lately, Āurovčíková et al. (2024) and Greig et al. (2024a) presented the first analyses of larger QSO samples, probing the end stages of reionization between $5.8 \lesssim z \lesssim 7$. Moreover, a number of studies has recently suggested the presence of damping wing absorption adjacent to individual Gunn-Peterson troughs in the foreground of a quasar due to neutral islands in the $5.5 \lesssim z \lesssim 6$ IGM (Becker et al. 2024; Zhu et al. 2024; Spina et al. 2024). The number of available objects will increase drastically in the next years with the Euclid wide-field survey detecting $\gtrsim 100$ new quasars at $z > 7$, and even objects out to $z \sim 9$ (or higher, depending on the assumed quasar luminosity function (QLF); Euclid Collaboration et al. 2019). Observations of their spectra will provide us with the first statistical QSO sample reaching deep into the heart of reionization.

The interest in damping wings has gained additional momentum recently with the first galaxy damping wings claimed in (stacks of) ultra high-redshift galaxy spectra observed by the James Webb Space Telescope (JWST) (Curtis-Lake et al. 2023; Hsiao et al. 2024; Umeda et al. 2024; Keating et al. 2024b). However, Heintz et al. (2024) pointed out a possibly significant contamination by intrinsic damped Lyman- α absorbers (DLAs) in some of these recently detected high-redshift galaxies, and Heintz et al. (2025) indeed identified strong integrated DLAs in a large fraction of galaxies (60% at $z \sim 6$ and up

to 65 – 90 % at $z > 8$). Originating in the interstellar medium (ISM) or the circumgalactic medium (CGM), such proximate DLAs can interfere with or even dominate over the cosmologically relevant IGM damping wing imprint, potentially obstructing robust conclusions about the ionization state of the IGM. This underlines the importance of quasars as a complementary probe in the context of IGM damping wing analysis.

Though fewer in number as compared to galaxies, quasars come with the advantage that their strong ionizing radiation carves out a Mpc-scale H II-region, the *quasar proximity zone*, where all residual neutral CGM gas is ionized away, sidestepping the aforementioned problems related to intrinsic DLAs. Avoiding such issues with certainty, however, might require a strict exclusion of QSOs with proximate DLAs as recently pointed out in [Davies et al. \(2025\)](#) with regards to ULAS J1342+0928. The task of identifying such objects will indeed be feasible for quasars (unlike in the context of much fainter objects such as galaxies) by using the capabilities of JWST or ELT to identify weak metal absorption systems in their spectra ([Davies et al. 2025](#)).

An additional virtue of quasars as compared to galaxies is the simplicity in modelling their continuum that the IGM damping wing is imprinted upon. The reconstruction of this continuum acts as a nuisance stochastic process when constraining $\langle x_{\text{HI}} \rangle$ based on the IGM damping wing imprint. Owing to the absence of evolution between $2 \lesssim z \lesssim 5$ as evidenced by [Shen et al. \(2007\)](#), it stands to reason to construct empirical quasar continuum models based on large datasets of $O(1000 - 10\,000)$ low-redshift objects from surveys such as SDSS (see [Greig et al. 2024b](#), for a comprehensive overview and comparison of methods).

The strength of the IGM damping wing imprint is (to leading order) not only affected by the global volume-averaged IGM neutral fraction $\langle x_{\text{HI}} \rangle$ at the quasar’s redshift but also by its light curve: the longer a quasar has been shining, the larger the ionized region it has carved out in its vicinity, manifesting in a weaker damping wing imprint and a more extended proximity zone. Under the assumption of a light-bulb light curve, this is determined by a single nuisance parameter, the lifetime of the quasar. While acting as a nuisance parameter when constraining the history of reionization, the quasar lifetime also bears physical information on its own about the formation and growth of supermassive black holes (SMBHs) whose understanding is increasingly challenged by JWST observations of moderately massive SMBHs at ultra-high redshifts ([Übler et al. 2024](#); [Maiolino et al. 2024b](#); [Kokorev et al. 2023](#); [Larson et al. 2023](#); [Bogdán et al. 2024](#)). Proximity zone sizes as simple summary statistics informing about the lifetime of a quasar have been measured extensively in the literature ([Eilers et al. 2017, 2020](#); [Ishimoto et al. 2020](#); [Satyavolu et al. 2023b](#)), and remain the subject of ongoing modelling efforts (e.g. [Davies et al. 2020](#); [Satyavolu et al. 2023a](#); [Zhou et al. 2024](#)).

Further complications arise due to the fact that reionization is a patchy process: the local size of reionization bubbles at a given value of $\langle x_{\text{HI}} \rangle$ can vary significantly (Davies et al. 2018a), especially during the mid-stages of reionization, hence acting as an additional source of stochasticity when inferring $\langle x_{\text{HI}} \rangle$. Besides that, fluctuations in the density field source stochasticity in the size of the ionized bubble carved out by the quasar (Davies et al. 2018a). These effects can only be accounted for with cosmological hydrodynamical simulations including ionizing radiative transfer.

Hennawi et al. (2025) introduced a fully Bayesian approach to jointly infer the global volume-averaged IGM neutral fraction $\langle x_{\text{HI}} \rangle$, quasar lifetime t_{Q} and the underlying continuum of an observed quasar spectrum, accounting for the full covariance arising from continuum reconstruction as well as the Lyman- α transmission stochastic process. The pipeline entails a parametric model for the quasar continuum based on principal component analysis (PCA), constructed from a large sample of $\sim 15\,000$ low-redshift ($2.149 < z < 4$) continua. IGM transmission fields describing the IGM damping wing are constructed by combining skewers from the hydrodynamical Nyx simulations (Almgren et al. 2013; Lukić et al. 2015) and semi-numerical reionization topologies (Mesinger et al. 2011; Davies & Furlanetto 2022) via 1d radiative transfer (Davies et al. 2016). Hennawi et al. (2025) introduced a likelihood prescription that allows us to *jointly* fit the continuum and the IGM damping wing imprint, accounting for all uncertainties due to the stochasticity of the reionization topology, the unknown quasar lifetime, continuum reconstruction errors, and spectral noise. In addition, the likelihood operates on the entire spectrum and hence overcomes the limitations of the red-blue split commonly performed in previous approaches (e.g. Greig et al. 2017a; Davies et al. 2018a; Durovčiková et al. 2020; Reiman et al. 2020) which by construction cannot use the information on the blue side of the spectrum to recover the continuum.

A rigorous quantification of the associated error budget is key to deriving competitive cosmological and astrophysical constraints on $\langle x_{\text{HI}} \rangle$ and t_{Q} from observed high-redshift quasar spectra. The scope of our work in this context is twofold: 1) we determine the optimal observational setup to obtain maximally tight yet statistically faithful constraints, and 2) we break down the error budget across parameter space into the individual contributions of all model components.

Inference precision will inevitably be affected by hyperparameters of the underlying models such as the dimension and hence the flexibility of our parametric continuum model. Additional impact comes from the specifics of the observational setup such as the covered wavelength range, signal-to-noise ratio or spectral resolution. The inference performance with respect to these parameters informs us about the optimal observational strategy and the cost-benefit ratio of modern-day instruments for the purposes of our analysis. Observational considerations aside, the aforementioned parameters can also have a profound impact on the computational complexity of the

inference problem at hand. Given the cubic scaling of the evaluation of Gaussian likelihood terms with the number of spectral pixels, a too ambitious choice of the aforementioned hyper-parameters can easily make the problem computationally expensive or even intractable. This is particularly true for the highest spectral resolutions that are already achievable with modern-day instruments. Our analysis explores these dependencies and makes practical suggestions for the optimal parameter setting for the inference machinery introduced in Hennawi et al. (2025).

Secondly, inference precision intrinsically depends on the values of the physical parameters that are being inferred. A lower IGM neutral fraction $\langle x_{\text{HI}} \rangle$, for example, or a longer quasar lifetime t_{Q} lead to a weaker damping wing imprint, making the parameter values harder to infer. Additional uncertainty arises from the two major stochastic processes that are part of our pipeline: the parametric quasar continuum model and the Lyman- α transmission model accounting for the patchiness of reionization and stochasticity in the position of the quasar ionization front due to density fluctuations. We rigorously quantify their contributions to the overall uncertainty on the inferred IGM neutral fraction $\langle x_{\text{HI}} \rangle$ and quasar lifetime t_{Q} , and show how these uncertainties vary as a function of astrophysical parameter space.

Section 2.2 provides a short but largely self-contained summary of the inference pipeline introduced in Hennawi et al. (2025). We present our hyper-parameter sensitivity analysis in Section 2.3, culminating in the full breakdown of the total error budget in Section 2.3.4. We conclude in Section 2.4. In Appendix 2.A we correct a typo in the analytical damping wing optical depth equation provided by Mortlock (2016). Appendix 2.B summarizes the inference tests we performed to demonstrate the statistical fidelity of our pipeline, and Appendix 2.C extends the sensitivity analysis presented in Section 2.3.3 to our full forward model based on cosmological simulations.

2.2 Methods

We employ the inference procedure introduced in Hennawi et al. (2025) to infer the global volume-averaged neutral fraction $\langle x_{\text{HI}} \rangle$ and quasar lifetime t_{Q} with help of the damping wing signature that appears in high-redshift quasar spectra. We will restrict this section to a short summary of the formalism, and refer the reader to Hennawi et al. (2025) for additional details.

The observable in this context is the spectrum \mathbf{f} of the quasar that comes with a noise vector $\boldsymbol{\sigma}$. Inferring the above mentioned two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} given a measurement of \mathbf{f} and $\boldsymbol{\sigma}$ requires a forward model of IGM damping wing absorption and an expression for the likelihood of the observed spectrum given the parameters of our model. To that end, we assume that the full observed flux \mathbf{f} arises from a quasar continuum \mathbf{s} that is

subject to an IGM transmission field \mathbf{t} and spectral noise $\boldsymbol{\sigma}$. In practice, we do not have perfect knowledge of \mathbf{s} and \mathbf{t} as we cannot observe each of them separately. In other words, they constitute *latent variables* of our model. To move forward, we assume that they are governed by stochastic processes: we introduce a new low-dimensional latent variable $\boldsymbol{\eta}$ that represents the full quasar continuum \mathbf{s} , and we generate IGM transmission fields \mathbf{t} based on simulations that take as input a set of astrophysical parameters $\boldsymbol{\theta}$. Hennawi et al. (2025) showed that under the assumption of Gaussianity for the stochastic processes governing \mathbf{s} and \mathbf{t} , the full likelihood of measuring the flux \mathbf{f} can be approximated as

$$L(\mathbf{f}|\boldsymbol{\sigma}, \boldsymbol{\theta}, \boldsymbol{\eta}) = \mathcal{N}(\mathbf{f}; \langle \mathbf{t} \rangle \circ \langle \mathbf{s} \rangle, \boldsymbol{\Sigma} + \langle \mathbf{S} \rangle \mathbf{C}_t \langle \mathbf{S} \rangle + \langle \mathbf{T} \rangle \mathbf{C}_s \langle \mathbf{T} \rangle), \quad (2.1)$$

where $\mathcal{N}(\mathbf{f}; \boldsymbol{\mu}, \mathbf{K})$ is the standard normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix \mathbf{K} . Throughout, $\mathbf{x} \circ \mathbf{y}$ denotes the element-wise (Hadamard) product of two vectors \mathbf{x} and \mathbf{y} , and bold lowercase symbols denote vectors, bold uppercase symbols matrices. Here we defined $\boldsymbol{\Sigma} \equiv \text{diag}(\boldsymbol{\sigma})$, $\mathbf{S} \equiv \text{diag}(\mathbf{s})$, and $\mathbf{T} \equiv \text{diag}(\mathbf{t})$, with \mathbf{C}_s and \mathbf{C}_t as the corresponding covariance matrices.

We emphasize that a red-blue split of the spectrum is not required at any stage of the derivation of Eq. (2.1). Our forward model is a generative process for the full flux across the entire spectral range, and hence Eq. (2.1) also takes into account information from the blue side of the spectrum, most particularly the region of the smooth IGM damping wing signature and the quasar proximity zone. This implies that we are not restricted to the commonly adopted approach of predicting the blue-side continuum *before* fitting for the optimal parameters $\boldsymbol{\theta}$ on the resulting continuum-normalized spectrum (Greig et al. 2017a,b; Davies et al. 2018a,b; Bañados et al. 2018; Greig et al. 2019; Ďurovčiková et al. 2020; Wang et al. 2020; Yang et al. 2020a). Instead, we infer the astrophysical parameters $\boldsymbol{\theta}$ *jointly* with the nuisance parameters $\boldsymbol{\eta}$ representing the continuum across the entire spectral range.

The approach of considering the full spectrum only gained momentum very recently in the context of low-redshift quasar continuum prediction (Sun et al. 2023), and has been introduced for high-redshift quasar damping wings in Hennawi et al. (2025). It comes at the price of errors and bias resulting from the fact that Eq. (2.1) is approximate due to the non-Gaussianity of the IGM transmission stochastic process (Lee et al. 2015; Davies et al. 2018a). The approximate nature of our likelihood prescription can result in misplaced or too narrow posterior distributions, and we correct for the concomitant loss of information when determining inference precision. This paper rigorously quantifies the errors resulting down the line in the inference process, and explores the optimal parameter setting to maximize inference precision with a view on observational and computational cost of the inference task at hand.

2.2.1 Training Data and Dimensionality Reduction

Lacking a principled physical model for quasar continua, we adopt a data-driven approach to construct a lower-dimensional parametric model $\mathbf{s}_{\text{DR}}(\boldsymbol{\eta})$ of the full continuum \mathbf{s} . Due to the strong correlations among individual pixels, it is not favorable nor necessary to infer the value of \mathbf{s} at every single spectral pixel. Instead, we make use of these inter-pixel correlations to reduce the dimensionality of the continuum from the total number of spectral pixels n_λ to a lower-dimensional latent representation, fully specified by the new n_{latent} -dimensional latent variable $\boldsymbol{\xi}$ (with $n_{\text{latent}} \ll n_\lambda$). A simple and well-established approach to do so is based on principal component analysis (PCA) (e.g. Suzuki et al. 2005; Suzuki 2006; Pàris et al. 2011; Davies et al. 2018b), resulting in the expression

$$\mathbf{s}_{\text{DR}}(\boldsymbol{\xi}) = \langle \mathbf{s} \rangle + \boldsymbol{\xi}^T \mathbf{A}. \quad (2.2)$$

Here $\langle \mathbf{s} \rangle$ denotes the mean of the n_{train} training continua \mathbf{s} ; \mathbf{A} is the $(n_{\text{latent}} \times n_\lambda)$ -dimensional matrix of PCA basis vectors, and $\boldsymbol{\xi}$ the n_{latent} -dimensional vector of PCA coefficients. Along with the amplitude s_{norm} of the continuum at $\lambda = 1450 \text{ \AA}$, we end up with a $(n_{\text{latent}} + 1)$ -dimensional parametric model $\mathbf{s}_{\text{DR}}(\boldsymbol{\eta})$ for the quasar continuum, where $\boldsymbol{\eta} = (s_{\text{norm}}, \boldsymbol{\xi})$. Again, we stress that our model represents the continuum across the entire spectral range without a division into a blue-side and a red-side part.

Our choice of a standard PCA decomposition as dimensionality reduction (DR) method is motivated by the fact that Hennawi et al. (2025) showed it to perform comparably well to more complex machine learning-based techniques using Gaussian processes or variational autoencoders. To construct the PCA model, we assume that we can always find an (unabsorbed) low-redshift quasar spectrum that is a close representative of a given high-redshift continuum.¹ More specifically, we make use of 15 559 low-redshift quasar spectra from the SDSS-III Baryon Oscillation Spectroscopic Survey (BOSS) and SDSS-IV Extended BOSS (eBOSS) with SDSS autofits. The resolution of the spectra is $R \sim 2000$ and their redshifts are in the range $2.149 < z < 4$. The spectra are chosen such that each of them covers the full rest-frame wavelength range between 1170 \AA and 2040 \AA and that their median signal-to-noise ratio is $S/N > 10$ within a 5.0 \AA region centered around the rest-frame wavelength 1285 \AA . In order to allow for a joint PCA decomposition of all training spectra, we are rebinning the spectra onto a common rest-frame wavelength grid linearly spaced in velocity with pixels of size $dv_{\text{pix}} = 140 \text{ km/s}$. For further details, we refer the reader to

¹Note that we are *not* assuming that the intrinsic shape of quasar continua does not change as a function of redshift. It suffices to make the weaker assumption that we see examples of high-redshift quasars at low z . The opposite, i.e., that the full plethora of shapes we observe at low redshift also has to be present at high z , does not necessarily have to be true.

Hennawi et al. (2025). We divide the set of continua into $n_{\text{train}} = 14781$ training objects to determine the PCA basis \mathbf{A} and $n_{\text{test}} = 778$ test objects to draw mock continua from and estimate the continuum reconstruction error. This corresponds to a 95% – 5% training-test split. We account for the reconstruction error stochastic process by defining the relative continuum reconstruction error as

$$\boldsymbol{\delta} = \frac{\mathbf{s} - \mathbf{s}_{\text{DR}}(\boldsymbol{\eta})}{\mathbf{s}}, \quad (2.3)$$

where the division by \mathbf{s} is to be understood as an element-wise operation. As shown in Hennawi et al. (2025), $\boldsymbol{\delta}$ closely follows a Gaussian distribution $\mathcal{N}(\boldsymbol{\delta}; \langle \boldsymbol{\delta} \rangle, \Delta)$, and we estimate the mean $\langle \boldsymbol{\delta} \rangle$ and the covariance matrix Δ by applying our DR formalism to the $n_{\text{test}} = 778$ test objects. Eq. (2.3) implies that $\mathbf{s} \simeq \mathbf{s}_{\text{DR}}(\boldsymbol{\eta}) \circ (\mathbf{1} + \boldsymbol{\delta})$ up to linear order in $\boldsymbol{\delta}$, and therefore we obtain

$$P(\mathbf{s}|\boldsymbol{\eta}) = \mathcal{N}(\mathbf{s}; \langle \mathbf{s}(\boldsymbol{\eta}) \rangle, \mathbf{C}_s(\boldsymbol{\eta})) \quad (2.4)$$

if we define $\langle \mathbf{s}(\boldsymbol{\eta}) \rangle \equiv \mathbf{s}_{\text{DR}}(\boldsymbol{\eta}) \circ (\mathbf{1} + \langle \boldsymbol{\delta} \rangle)$, and $\mathbf{C}_s(\boldsymbol{\eta}) \equiv \text{diag}(\mathbf{s}_{\text{DR}}(\boldsymbol{\eta})) \Delta \text{diag}(\mathbf{s}_{\text{DR}}(\boldsymbol{\eta}))$. Note that $\langle \boldsymbol{\delta} \rangle$ and Δ (and hence also $\langle \mathbf{s}(\boldsymbol{\eta}) \rangle$ and $\mathbf{C}_s(\boldsymbol{\eta})$) depend on the choice of latent dimension n_{latent} of the DR model, and are thus recomputed whenever we vary n_{latent} .

2.2.2 Simulating Damping Wing Observations

To model IGM damping wing absorption profiles, we generate IGM transmission fields \mathbf{t} , considering their variation with respect to two astrophysical parameters $\boldsymbol{\theta} = (\langle x_{\text{HI}} \rangle, t_{\text{Q}})$, where $\langle x_{\text{HI}} \rangle$ is the global volume-averaged IGM neutral hydrogen fraction, and t_{Q} is the lifetime of the quasar. We explore two different approaches of generating \mathbf{t} for a given set of parameters $\boldsymbol{\theta}$: 1) Following Davies et al. (2018a), we post-process sightlines from the cosmological hydrodynamical Nyx simulations (Almgren et al. 2013; Lukić et al. 2015) and 21cmFAST reionization topologies (Mesinger et al. 2011) with 1d ionizing radiative transfer (Davies et al. 2016), and 2), for the purpose of breaking up the error budget into the separate contributions of all individual model components, we eliminate the stochasticity between $\boldsymbol{\theta}$ and the simulated IGM damping wing strength by constructing a simple analytical model that relates them deterministically. We introduce these two approaches in the following sections.

2.2.2.1 Full-Simulation Model

To obtain maximally realistic IGM transmission fields, we follow the procedure introduced in Davies et al. (2018a): we extract 1200 density, velocity, and temperature skewers from a snapshot of the Nyx hydrodynamical simulations (Almgren et al. 2013; Lukić et al. 2015) at redshift $z = 7.0$. The simulations have a box size of 100 cMpc/h, and contain 4096^3 baryon and

dark matter particles each. A realistic reionization topology is generated with help of a modified version (Davies & Furlanetto 2022) of the 21cmFast code (Mesinger et al. 2011), tuning the ionizing efficiency ζ to obtain 21 different ionization fields with global volume-averaged neutral fractions between $\langle x_{\text{HI}} \rangle = 0$ and $\langle x_{\text{HI}} \rangle = 1$ in steps of $\Delta \langle x_{\text{HI}} \rangle = 0.05$. From each run, we extract 10000 randomly oriented x_{HI} skewers originating at the 500 most massive halos. We combine each of the 1200 Nyx density and temperature skewers with a random x_{HI} skewer by performing one-dimensional radiative transfer along the sightlines (Davies et al. 2016), integrating up to 51 different quasar lifetimes between $t_{\text{Q}} = 10^3$ yr and $t_{\text{Q}} = 10^8$ yr in steps of $\Delta \log_{10}(t_{\text{Q}}/\text{yr}) = 0.1$ (Davies et al. 2019). For this, we assume a flat Λ CDM cosmology with $h = 0.685$, $\Omega_b = 0.047$, $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$, and $\sigma_8 = 0.8$. After excluding 17 of the resulting skewers because of strong proximate DLA absorption, this yields a 21×51 grid in $\theta = (\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ with 1183 distinct Lyman- α transmission fields at each point in parameter space.

Generally, the IGM transmission stochastic process $P(\mathbf{t}|\theta)$ is well known to be a non-Gaussian (Lee et al. 2015; Davies et al. 2018a). Nevertheless, for the sake of deriving an analytical expression (Eq. (2.1)) for the full likelihood $L(\mathbf{f}|\sigma, \theta, \eta)$, we approximated it as such, i.e.,

$$P(\mathbf{t}|\theta) = \mathcal{N}(\mathbf{t}; \langle \mathbf{t}(\theta) \rangle, \mathbf{C}_{\mathbf{t}}(\theta)). \quad (2.5)$$

In this context, we estimate the mean $\langle \mathbf{t}(\theta) \rangle$ and the covariance $\mathbf{C}_{\mathbf{t}}(\theta)$ empirically from the aforementioned simulated realizations. As shown in Hennawi et al. (2025), this is a reasonably close approximation at sufficiently low spectral resolution. The tests in Appendix 2.B reveal that this is not the case at higher resolutions where Eq. (2.5) can lead to notable biases in our parameter inference that we have to take into account when evaluating inference precision. We describe a structured way to do so in Section 2.3.1.

2.2.2.2 Analytical Simulator

While the aforementioned numerical simulator captures the relevant physics for modeling quasar proximity zones and IGM damping wings, one potential drawback is that the size of the proximity zone and strength of the IGM damping wing are stochastically related to the two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and $\log_{10}(t_{\text{Q}}/\text{yr})$. For example, for a fixed global volume averaged neutral fraction $\langle x_{\text{HI}} \rangle$, the reionization topology results in a broad distribution of distances to the nearest patch of neutral hydrogen (see Figure 2 of Davies et al. 2018a). In addition, small-scale overdensities corresponding to Lyman limit systems or partial Lyman limit systems (LLS or PLLS) can significantly attenuate the quasar radiation, and the presence or absence of such fluctuations produce scatter in the location of the quasar ionization front (Chen & Gnedin 2021). These physical effects, which ultimately arise

from density fluctuations on a hierarchy of scales, source stochasticity in the proximity zone size and the strength of the IGM damping wing.

It is conceivable that an alternative parameterization exists that would allow one to fit for the strength of the IGM damping wing as set by the total column density of neutral gas in the foreground of the quasar, which could potentially remove much of the stochasticity between damping wing strength and the parameters $\langle x_{\text{HI}} \rangle$ and $\log_{10} t_{\text{Q}}$. From the standpoint of the goals of this paper, a more deterministic relationship between damping wing strength and model parameters would be very valuable, as it would allow us to isolate how factors like continuum reconstruction errors, latent space dimension, S/N ratio, and spectral resolution contribute to the overall error budget; whereas, the relative contributions of these factors will be concealed for the more stochastic numerical simulator described in the previous section, once the intrinsic scatter in damping wing strength at fixed $\theta = \{\langle x_{\text{HI}} \rangle, \log_{10} t_{\text{Q}}\}$ dominates the overall noise budget. To this end, we also explore a more deterministic analytical IGM damping wing simulator, which we now describe.

For simplicity, to generate the IGM damping wing we consider an isotropically emitting quasar in an IGM of uniform density and uniform neutral fraction. In a second step, we then allow for density fluctuations in the proximity zone via a lognormal Lyman- α forest model.

We start by writing the mean neutral hydrogen number density at the redshift z_{QSO} of the quasar as

$$\langle n_{\text{HI}} \rangle = n_{\text{H},0} (1 + z_{\text{QSO}})^3 \langle x_{\text{HI}} \rangle, \quad (2.6)$$

where $n_{\text{H},0}$ is the present mean hydrogen number density and $\langle x_{\text{HI}} \rangle$ is the global volume-averaged IGM neutral fraction at z_{QSO} . Here $n_{\text{H},0} = 3H_0^2(1-Y)\Omega_b/8\pi Gm_p \simeq 1.9 \times 10^{-7}$ atoms/cm³ with big bang nucleosynthesis (BBN) hydrogen fraction $1 - Y = 0.76$. All cosmological parameters for the analytical simulator described in this section are set according to a [Planck Collaboration et al. \(2020\)](#) cosmology. Assuming that each photon emitted by the quasar ionizes precisely one hydrogen atom in the quasar's immediate vicinity and ignoring recombinations as well as Hubble expansion, the radius R_{ion} of the resulting ionization front evolves with the lifetime t_{Q} of the quasar as

$$\begin{aligned} R_{\text{ion}} &= \left(\frac{3Qt_{\text{Q}}}{4\pi \langle n_{\text{HI}} \rangle} \right)^{1/3} \\ &= 12.0 \text{ cMpc} \left(\frac{Q}{10^{57.1} \text{ s}^{-1}} \right)^{1/3} \left(\frac{t_{\text{Q}}}{10^6 \text{ yr}} \right)^{1/3} \left(\frac{\langle x_{\text{HI}} \rangle}{1.0} \right)^{-1/3} \end{aligned} \quad (2.7)$$

(Cen & Haiman 2000), where Q is the quasar's emission rate of ionizing photons, obtained by integrating the luminosity L_ν per photon energy $h\nu$ over all frequencies larger than the Lyman-limit frequency ν_{LL} :

$$Q = \int_{\nu_{\text{LL}}}^{\infty} \frac{L_\nu}{h\nu} d\nu. \quad (2.8)$$

Here we assume that the luminosity follows a power-law shape $L_\nu = L_{\text{LL}} (\nu/\nu_{\text{LL}})^{-\alpha_s}$ blueward of the Lyman limit frequency ν_{LL} with a spectral index $\alpha_s = 1.7$ and L_{LL} set by the J -band magnitude and redshift of the quasar which we choose as $z_{\text{QSO}} = 7.54$ and $J = 20.3$, resembling the quasar ULAS J1342+0928. Altogether, this amounts to $Q = 10^{57.1} \text{ s}^{-1}$.

The IGM transmission field $t(\lambda_{\text{obs}})$ as a function of observed-frame wavelength $\lambda_{\text{obs}} = (1 + z_{\text{QSO}}) \lambda_{\text{rest}}$ is governed by the Lyman- α optical depth $\tau(\lambda_{\text{obs}})$ as

$$t(\lambda_{\text{obs}}) = e^{-\tau(\lambda_{\text{obs}})}. \quad (2.9)$$

We obtain the optical depth $\tau(\lambda_{\text{obs}})$ by integrating the infinitesimal optical depth $d\tau = \langle n_{\text{HI}} \rangle \sigma_\alpha(\nu) dl$ along a line-of-sight interval dl (Miralda-Escudé 1998; Mortlock 2016):

$$\tau_{\text{DW}}(\lambda_{\text{obs}}) = \int \langle n_{\text{HI}} \rangle \sigma_\alpha(\nu) dl. \quad (2.10)$$

Here $\sigma_\alpha(\nu)$ is the scattering cross section of the Lyman- α line which is given as

$$\sigma_\alpha(\nu) = \frac{\pi e^2}{m_e c} f_\alpha \phi_\alpha(\nu), \quad (2.11)$$

where $f_\alpha \simeq 0.416$ is the Lyman- α oscillator strength and $\phi_\alpha(\nu)$ the line profile function. We assume that $\phi_\alpha(\nu)$ is of Lorentzian shape

$$\phi_\alpha(\nu) = \frac{R_\alpha / (\pi \nu_\alpha)}{(\nu/\nu_\alpha - 1)^2 + R_\alpha^2} \quad (2.12)$$

with Lyman- α frequency $\nu_\alpha = c/\lambda_\alpha$, Lyman- α wavelength $\lambda_\alpha = 1215.67 \text{ \AA}$, and $R_\alpha = \Gamma_\alpha/4\pi\nu_\alpha$ determined by the decay constant $\Gamma_\alpha = 6.265 \times 10^8 \text{ s}^{-1}$ of the Lyman- α transition. Bach & Lee (2015) showed this to be a more accurate approximation to the full quantum-mechanical cross section than the commonly used two-level model by Miralda-Escudé (1998). Integrating Eq. (2.10) from the redshift z_{ion} of the ionization front to the end of reionization (that we assume at $z_{\text{end}} = 6.0$; our results are not sensitive to this choice) yields the damping wing optical depth

$$\tau_{\text{DW}}(\lambda_{\text{obs}}) = \frac{4R_\alpha \tau_{\text{GP}}(z_{\text{ion}})}{\pi} \left(\frac{\lambda_{\text{obs}}}{(1 + z_{\text{ion}})\lambda_\alpha} \right)^{3/2} \times \quad (2.13)$$

$$I \left[\frac{(1 + z_{\text{ion}})\lambda_\alpha}{\lambda_{\text{obs}}} \right] - I \left[\frac{(1 + z_{\text{end}})\lambda_\alpha}{\lambda_{\text{obs}}} \right].$$

Here $\tau_{\text{GP}}(z)$ is the Gunn-Peterson optical depth

$$\tau_{\text{GP}}(z) = 5.3 \times 10^5 \left(\frac{1+z}{1+7.54} \right)^{\frac{3}{2}} \left(\frac{H_0 \sqrt{\Omega_m (1+z)^3}}{H(z)} \right) \left(\frac{\langle x_{\text{HI}} \rangle}{1.0} \right), \quad (2.14)$$

and $I(x)$ is a dimensionless integral factor given by

$$I(x) = \frac{x^{1/2}}{4(1-x)} + \frac{1}{8} \ln \left(\frac{1-x^{1/2}}{1+x^{1/2}} \right), \quad (2.15)$$

following the derivation in [Mortlock \(2016\)](#)².

Finally, to model the Lyman- α forest fluctuations in the quasar proximity zone, we start by generating realizations of a lognormal model using a modified version of the approach adopted by [Karaçaylı et al. \(2020\)](#), which is itself a modification of [McDonald et al. \(2006\)](#). In this approach, the non-linear and non-Gaussian HI column density field (and as a result, the optical depth field) is generated by applying a squared lognormal transformation to a Gaussian baryon field, reproducing the theoretically expected mean flux \bar{F} evolution and power spectrum of the Lyman- α forest. We draw realizations of the lognormal model at a reference redshift of $z_{\text{ref}} = 4.95$ and scale them to the quasar redshift $z_{\text{QSO}} = 7.54$ assuming a metagalactic UVB photoionization rate of $\Gamma_{\text{HI}} = 10^{-13}$ 1/s at that redshift. Combining these high-redshift Lyman- α forest profiles with the aforementioned damping wing profiles provides us with a full Lyman- α transmission model that we can compare to the (more realistic) full-simulation model introduced in Section 2.2.2.1.

We show three random skewers produced by each of the two models in Figure 2.1. All skewers were generated at a fixed IGM neutral fraction of $\langle x_{\text{HI}} \rangle = 0.5$ and a fixed quasar lifetime of $t_{\text{Q}} = 10^6$ yr, yet the three skewers based on the full-simulation model show significant differences. In particular, the sizes of their proximity zones vary between $\sim 1.0 - 2.5$ pMpc, reflective of the stochastic bubble size distribution in the underlying semi-numerical reionization topology. By construction, the three analytical skewers look almost identical as they only differ by the draw of the lognormal Lyman- α forest fluctuations. This illustrates how the analytic formalism introduced in this section eliminates the stochasticity of reionization from our forward modelling, allowing us to measure the impact of other sources of uncertainty more accurately, and quantify the extent as to which the intrinsic stochasticity of reionization contributes to the total error budget.

²Note that due to a typo in [Mortlock \(2016\)](#), the prefactor of the second term in Eq. (2.15) differs with respect to that in Eq. (40) in [Mortlock \(2016\)](#). We elaborate on the details and consequences of this correction in Appendix 2.A.




Figure 2.1: Three random examples of IGM transmission skewers generated with the full-simulation model (see Section 2.2.2.1) and the analytical model (see Section 2.2.2.2). All skewers were generated assuming a global IGM neutral fraction of $\langle x_{\text{HI}} \rangle = 0.5$ and a quasar lifetime of $t_{\text{Q}} = 10^6$ yr, convolved with a Gaussian line-spread function with FWHM = 100 km/s and rebinned to a 50 km/s pixel scale in the blue part of the spectrum. The size of the ionization front varies significantly between the three full-simulation skewers due to the stochasticity of the ionized bubble size distribution that is incorporated in this model but not in the analytical one.

2.2.3 Generating Mock Spectra

While our inference scheme can be readily applied to real observational data, we restrict this work to mock spectra in order to quantify the error budget of the inference pipeline. As in Hennawi et al. (2025), our models resemble the quasar ULAS J1342+0928 at $z = 7.54$ with $J = 20.3$ and realistic forward modelling of (hypothetical) instrument systematics and heteroscedastic spectral noise.

We generate mock spectra \mathbf{f} by combining unabsorbed continua \mathbf{s} with IGM transmission fields \mathbf{t} and a noise realization $\boldsymbol{\sigma}$. To that end, we draw real low-redshift continua \mathbf{s} from the set of 778 autofit test objects introduced in Section 2.2.1, and random spectral amplitudes s_{norm} from a uniform distribution, normalizing the continuum flux at 1450 Å to the

drawn value of s_{norm} . The range of the spectral amplitude distribution is chosen in accordance with the minimum and maximum amplitudes among 100 calibration spectra (with magnitudes in the range $J = 20.3 \pm 0.1$).

To obtain IGM transmission fields \mathbf{t} , we draw realizations of the astrophysical parameters $\boldsymbol{\theta} = (\langle x_{\text{HI}} \rangle, t_{\text{Q}})$, and then apply one of the forward models introduced in Sections 2.2.2.1 and 2.2.2.2 to obtain a realization of \mathbf{t} at the drawn parameter values. The parameter ranges covered by these models also define implicit priors for the inference. We adopt uniform distributions for both astrophysical parameters (uniform in log-space for t_{Q}) and choose the limits such that they cover a broad range of values that might be observationally relevant, noting that previous analysis of lower-redshift data has hinted towards a significantly more narrow log-normal quasar lifetime distribution centered around $t_{\text{Q}} \sim 10^6$ yr (Khrykin et al. 2021).

We consider global IGM neutral fractions in the full range $\langle x_{\text{HI}} \rangle \in [0, 1]$. When working with the analytical IGM transmission model, we have to restrict this to $\langle x_{\text{HI}} \rangle \in [10^{-4}, 1]$ to avoid divergences in our model (see Eq. (2.7)). The quasar lifetime range in the full-simulation model is $t_{\text{Q}} \in [10^3 \text{ yr}, 10^8 \text{ yr}]$ in accordance with our simulation grid. For the analytical model, we allow for lifetimes between $t_{\text{Q}} \in [10^5 \text{ yr}, 10^{8.84} \text{ yr}]$, where the upper lifetime limit corresponds to the age of the universe at redshift $z_{\text{QSO}} = 7.54$. Whenever we directly compare the two models, we adjust the lifetime range of the analytical model to that of the full-simulation model.

Note also that each of the 1183 skewers governing the IGM transmission field \mathbf{t} in our full-simulation model is pre-computed on a discrete 21×51 grid of parameter values $\boldsymbol{\theta} = (\langle x_{\text{HI}} \rangle, t_{\text{Q}})$. We randomly choose one of the 1183 available skewers and linearly interpolate the IGM transmission field \mathbf{t} along the parameter grid to the drawn value of $\boldsymbol{\theta}$.

To simulate instrumental effects, we subsequently convolve the IGM transmission field \mathbf{t} with a Gaussian line-spread function (LSF) of a given width (by default FWHM = 100 km/s) and rebin it onto the final velocity grid. Since the transmission field is defined only on a rest-frame wavelength grid of 1185 – 2000 Å, whereas our PCA model operates on the range 1170 – 2040 Å, we have to interpolate the continuum onto the smaller grid in order to combine it with the transmission skewers. To keep the computational cost of the inference procedure at a tractable minimum, we work on a hybrid wavelength grid by splitting up the spectrum at $\lambda_{\text{blue-red}} = 1218.10$ Å into a high-resolution blue part and a low-resolution red part as described in Hennawi et al. (2025). High resolution is only relevant around the Lyman- α line and blueward in the Lyman- α forest region, while a coarse wavelength grid is sufficient to resolve the smooth part of the spectrum redward of $\lambda_{\text{blue-red}}$. Throughout the analysis, we restrict the red part to a relatively coarse velocity pixel spacing ($dv_{\text{red}} = 500$ km/s) and only investigate the effects of resolving the Lyman- α region with varying resolution. When we change spectral resolution, we always keep the spectral sampling factor

between the FWHM of the LSF and the velocity spacing dv_{blue} in the blue part of the spectrum fixed at $\text{FWHM}/dv_{\text{blue}} = 2$.

Lastly, we use `SkyCalc_ipy` (Leschinski 2021) to generate realistic noise realizations including telluric absorption and contributions from object photons, sky background, and detector read noise. We tune the exposure time of our hypothetical instrument such that we achieve a given median signal-to-noise ratio S/N per 100 km/s velocity interval (by default S/N = 10).

For the purposes of our sensitivity analysis, we work with a fixed sample of 100 mock spectra drawn from the priors introduced above. Our pipeline is implemented in such a way that we can generate the exact same set of mock objects given different (hyper-)parameter settings such as signal-to-noise ratio S/N or spectral resolution FWHM. Note that this includes the draw of the continuum as well as the Lyman- α transmission skewers, and, in all cases besides spectral resolution, also the noise realization. This allows us to isolate the impact of the aforementioned hyper-parameters on inference precision from the stochasticity of the process itself by considering a moderate number of 100 mock spectra.

2.2.4 Inference Procedure

The likelihood derived in Section 2.2 in combination with the forward models from Sections 2.2.2.1 and 2.2.2.2 allows us to perform Bayesian inference to compute the posterior distribution of parameters $\Theta = (\langle x_{\text{HI}} \rangle, \log_{10} t_{\text{Q}}/\text{yr}, s_{\text{norm}}, \boldsymbol{\eta})$ given a (mock) observational quasar spectrum \mathbf{f} . Again, $(\langle x_{\text{HI}} \rangle, \log_{10} t_{\text{Q}}/\text{yr})$ are the two astrophysical parameters of interest that we infer jointly with the continuum nuisance parameters $\boldsymbol{\eta} = (s_{\text{norm}}, \boldsymbol{\xi})$, i.e., the continuum amplitude and the best-fitting PCA coefficients. For the purposes of this paper, we adopt uniform priors on the astrophysical parameters $\langle x_{\text{HI}} \rangle$ and $\log_{10} t_{\text{Q}}/\text{yr}$ that arise from the fixed parameter grid in the full-simulation model or the analytical model domain. The respective parameter ranges are summarized in Section 2.2.3. Our prior on the spectral amplitude s_{norm} is $\text{Uniform}(0, \infty)$, assuring positivity of this parameter, and we allow each PCA coefficient α_i of the continuum model to vary over the full range $\text{Uniform}(-\infty, \infty)$. Practically, in our numerical pipeline, we apply variable transformations to all (half-)bounded parameters to make them fully unbounded.

We sample from the posterior distribution using a Hamiltonian Monte-Carlo (HMC) algorithm with a No U-Turn Sampler (NUTS) kernel implemented in the `NumPyro` probabilistic programming package. The `NumPyro` package is based on the machine learning and autograd framework `JAX`, allowing us to implement our scheme in a fully differentiable and GPU-accelerated way. For each inference task, we run 4 MCMC chains for 2000 MCMC steps each, 1000 of which are discarded as warm-up steps.

An example inference run is shown in Figure 2.2. The mock object is part of the ensemble described and analyzed in more detail in Section 2.3.4. We

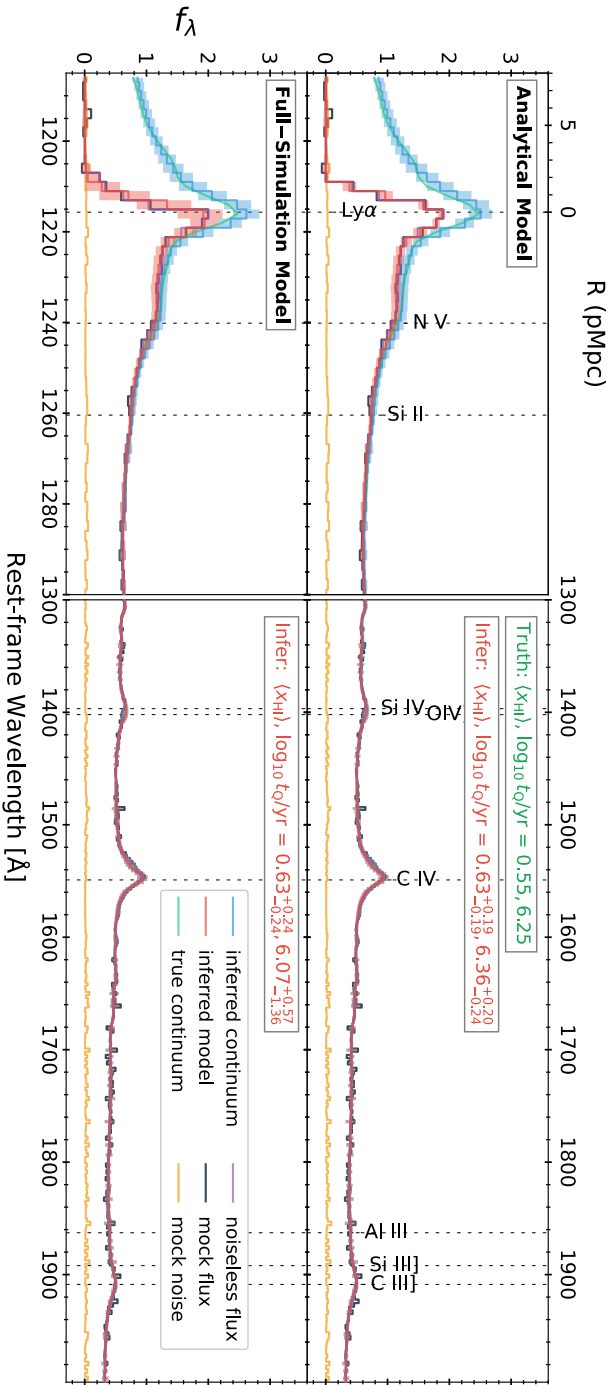


Figure 2.2: Model components of the true and the inferred spectrum of one of our 1000 mock quasars considered in Section 2.3.4 with a global IGM neutral fraction of $\langle x_{\text{HI}} \rangle = 0.55$ and a quasar lifetime of $\log_{10} t_Q/\text{yr} = 6.25$. In the upper panel, IGM transmission skewers are generated by the analytical model introduced in Section 2.2.2.2, in the lower panel by the full-simulation model from Section 2.2.2.1. The PCA continuum model has a latent dimension of $n_{\text{latent}} = 5$ and the observational setup is fixed to $\lambda_{\text{cut}} = 2000 \text{ \AA}$, $S/N = 10$, and $\text{FWHM} = 100 \text{ km/s}$ rebinned to $d\nu_{\text{blue}} = 500 \text{ km/s}$. For better visibility of the damping wing region, the rest-frame wavelength axis is stretched on the left side of the figure between 1185 \AA and 1300 \AA . The full mock spectrum of the quasar is depicted in black and consists of the true continuum (green) combined with IGM transmission (purple) and mock spectral noise (yellow). The inferred model spectrum is depicted in red, consisting of the inferred continuum (blue) combined with the inferred IGM transmission field. Solid lines represent the median inferred models, shaded regions denote the 16% and the 84% percentile variations reflecting parameter uncertainty, continuum reconstruction errors, as well as spectral noise. All inferred model spectra are simulated based on reweighted HMC samples, assuming that each ensemble passes a marginal coverage test with respect to the two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_Q . The 1σ -uncertainties are notably larger for the full-simulation model due to the additional stochasticity that it encompasses.




Figure 2.3: Marginal $(\langle x_{\text{HI}} \rangle, t_Q)$ -posterior distributions inferred from the two mock spectra depicted in Figure 2.2. All HMC samples are reweighted, assuring that each ensemble passes a marginal coverage test with respect to the two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_Q as described in Section 2.3.1. The contours of the full-simulation model (light blue) extend notably farther along the axis of degeneracy between global IGM neutral fraction and quasar lifetime than those of the deterministic analytical model (dark blue).

compare in the upper and the lower panel the same object, combined with full-simulation and analytical IGM transmission skewers as introduced in Sections 2.2.2.1 and 2.2.2.2, respectively. Note that the shaded 1σ regions need to entail the uncertainty contributions from all individual model components to allow for a meaningful assessment of the goodness-of-fit of the inferred model. To that end, we compute for each HMC sample $(\boldsymbol{\theta}, \boldsymbol{\eta})$ a draw of the continuum $\mathbf{s} = \mathbf{s}_{\text{DR}}(\boldsymbol{\eta}) \circ (\mathbf{1} + \boldsymbol{\delta})$, where we obtain $\mathbf{s}_{\text{DR}}(\boldsymbol{\eta})$ by evaluating our PCA model and the spectral amplitude at $\boldsymbol{\eta}$, and obtain $\boldsymbol{\delta}$ by drawing a realization of the associated continuum reconstruction error from Eq. (2.4). We further draw a random IGM transmission skewer \mathbf{t} for the given value of $\boldsymbol{\theta}$, using the analytical or the full-simulation IGM transmission model, respectively, giving $\mathbf{f} = \mathbf{s} \circ \mathbf{t}$. Lastly, we draw a noise realization $\tilde{\boldsymbol{\sigma}}$ from a Gaussian distribution with noise vector $\boldsymbol{\sigma}$. The depicted uncertainties on the inferred model and continuum are the 16% and 84% percentile variations of the realizations $\mathbf{f} + \tilde{\boldsymbol{\sigma}}$ and $\mathbf{s} + \tilde{\boldsymbol{\sigma}}$, hence reflecting parameter uncertainty, continuum reconstruction errors, as well as spectral noise. Note also that we use the coverage-corrected HMC samples for this procedure as discussed in Section 2.3.1.

While we are able to reconstruct the continuum reasonably well with either IGM transmission model, the error bars of the inferred model (red shaded regions) are significantly smaller when working with the analytical one. This is a direct consequence of the additional stochasticity present in the full-simulation model. As a result, also the marginal posterior distribution in Figure 2.3 extends significantly farther along the axis of degeneracy between $\langle x_{\text{HI}} \rangle$ and $\log_{10} t_{\text{Q}}/\text{yr}$.

This degeneracy arises because both parameters affect the neutral hydrogen content in the surrounding IGM and thus have a direct impact on the strength of the IGM damping wing signature. As a result, we can not always tell with certainty if we are observing a young quasar residing in a predominantly ionized IGM, or an older object in a globally more neutral IGM. This is because by carving out a larger ionized region, older quasars have significantly decreased the HI content along their sightlines as compared to the global volume-averaged value $\langle x_{\text{HI}} \rangle$, making their transmission profiles appear more akin to that seen for younger objects which have carved out a shorter ionization front but reside in a globally more ionized IGM.

Since Eqs. (2.7), (2.13) and (2.14) directly encode these dependencies, Figure 2.3 shows a tight axis of degeneracy for the posterior obtained with the analytical IGM transmission model. In case of the full-simulation model, this degeneracy extends significantly farther across parameter space and entails notably more scatter due to the additional sources of stochasticity that are treated in this model: firstly, because of the patchiness of reionization, the HI content along a given sightline can differ significantly from the global volume-averaged value $\langle x_{\text{HI}} \rangle$ which is approximated as uniform in the analytical IGM transmission model. Secondly, fluctuations in the density field are giving rise to additional stochasticity in the location of the ionization front in our full-simulation model. Altogether, this amounts to the notable degeneracy seen in Figure 2.3.

In Section 2.3.4 we provide a quantitative analysis of how the differences between the two models impact the overall inference precision we can achieve.

2.3 Sensitivity Analysis

Before applying our inference scheme to real-world data, it is essential to understand its sensitivity with respect to both model and hyper-parameters, as well as the associated error budget. The ultimate goal of the inference procedure is to obtain a 1) *faithful*, and 2) *precise* estimate of the posterior distribution of the astrophysical parameters of interest, i.e., the global volume-averaged IGM neutral fraction $\langle x_{\text{HI}} \rangle$ and the lifetime t_{Q} of the quasar at hand. In the next section, we will review the concept of coverage probability as a metric of overconfidence in the inferred posterior distribution and describe how to correct for potential imperfections if required. Subsequently, we can

introduce a *precision* metric that we can use to understand the sensitivity of the inference scheme with respect to various hyper-parameters related to the continuum DR model and the observational setup of the input spectra.

In order to isolate their impact from additional sources of stochasticity, we consider the same set of 100 mock spectra generated with different hyper-parameter choices as described in Section 2.2.3 throughout the entire analysis. Further, we eliminate the stochastic dependence of proximity zone size and IGM damping wing strength on the astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} by adopting the analytical IGM transmission model introduced in Section 2.2.2.2 as our default forward model of IGM transmission. Its deterministic relations (2.7) and (2.13) allow us to identify the impact of the aforementioned hyper-parameters on inference precision as clearly as possible. In Appendix 2.C, we show that the same trends, though weakened by the additional source of stochasticity, remain true when using realistic IGM transmission skewers based on cosmological simulations.

In the final part of this section, we will then, informed about the optimal hyper-parameter setting, provide a rigorous break-down of the contributions from all individual model components to the total error budget as a function of astrophysical parameter space.

2.3.1 Inference Tests

We show that the inferred posterior distributions are statistically faithful by running inference tests for ensembles of mock objects drawn from the full prior range. By computing expected coverage probabilities, we confirm that the inferred posterior distribution faithfully represents the true distribution of mock objects, i.e., that the inference is not under- or overconfident with regards to the inferred parameter values, or, in other words, that the inferred posterior distributions are not too wide and not too narrow. To this end, we introduce the concept of coverage probabilities: given a credibility level $\alpha \in [0, 1]$, the coverage probability C_α quantifies how frequently the true probability is contained in the α -th credibility level of the inferred posterior (see e.g. Sellentin & Starck 2019).

More practically speaking, for a given inference run we compute the credibility level α as the α -th highest-density region (HDR) of the inferred posterior distribution by sorting all MCMC samples with respect to posterior probability and choosing the α -th highest fraction. We then test if the true posterior value is contained in this α -th HDR and repeat the procedure for each object in our ensemble of mock spectra. The fraction of objects for which this is the case is called the expected coverage probability C_α . In the ideal case, we should find $C_\alpha = \alpha$ at all coverage levels. In this case, we say that we *pass* the coverage test. Otherwise, if $C_\alpha > \alpha$, we are said to be *underconfident* or *conservative* about the significance of our results, i.e., the inferred posterior distribution is wider than it could be according

to the true distribution. On the other hand, and more crucially, if $C_\alpha < \alpha$, we are *overconfident* about the significance of our results and our posterior distribution is shifted or narrower than the true distribution would allow.

In the latter case, it is necessary to broaden the posteriors to make them compatible with the true distribution. Hennawi et al. (2025) (see also Wolfson et al. 2023) described a principled mathematical way of assigning unique weights to the draws of a MCMC chain to weigh them in such a way that the inferred coverage probability C_α exactly fulfills $C_\alpha = \alpha$ at all credibility levels α . We apply these corrections to all runs considered in this paper as many runs show some degree of overconfidence (see Appendix 2.B). For consistency, we also apply the same technique to runs that are consistent with perfect coverage as this does not affect the results and obviates the need to introduce a precise threshold for overconfidence.

Inference tests can be performed by considering the full-dimensional posterior distribution, or, alternatively, a lower-dimensional marginalized version of it. In our case, the only physical parameters of interest are the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ and the quasar lifetime t_{Q} , whereas the remaining $n_{\text{latent}} + 1$ parameters are nuisance parameters related to the shape of the continuum. In order to make sure that our posteriors are compatible not only with the distribution in full parameter space but also in the astrophysical subspace, we perform coverage tests with respect to both the full $(2 + n_{\text{latent}} + 1)$ -dimensional posterior distribution and the 2-dimensional marginal $(\langle x_{\text{HI}} \rangle, \log_{10} t_{\text{Q}}/\text{yr})$ -posterior density. It turns out that the full coverage probability is always on par with or more conservative than the marginal one. We therefore adopt the marginal coverage probability as the decisive measure of overconfidence and use it to compute the above mentioned importance weights. We confirmed that in all cases considered that the marginally coverage corrected models also pass full coverage tests. We show the (marginal) inference test results for all mock ensembles treated in Section 2.3.3 in Appendix 2.B.

2.3.2 Inference Precision

After assuring that the inferred posteriors are unbiased in the sense defined above, our metric for inference precision is the width of the posterior distributions after marginalizing over all nuisance parameters $\boldsymbol{\eta}$ related to the quasar continuum, i.e., the PCA coefficients $\boldsymbol{\xi}$ and the amplitude of the spectrum s_{norm} . We adopt as a measure for this width the (symmetrized) 1σ -error $\Delta_{68\%}(x)$ of the parameter $x \in \{\langle x_{\text{HI}} \rangle, \log_{10} t_{\text{Q}}/\text{yr}\}$, i.e., half the distance between the 84% and 16% percentile of the respective 1d-marginal posterior density

$$\Delta_{68\%}(x) = \frac{1}{2}(P_{84\%}(x) - P_{16\%}(x)), \quad (2.16)$$

where $P_q(x)$ is the q -th percentile of the marginal posterior with respect to the parameter x .

We quantify precision in a statistical sense by averaging over the inference results of the mock ensemble with $\langle x_{\text{HI}} \rangle$ and t_{Q} drawn from the full prior volume as defined in Section 2.2.3. Given that our method will mainly target the transition from a neutral to a reionized universe at intermediate neutral fractions, and that current measurements favor a significantly more narrow quasar lifetime distribution than that covered by our prior (Khrykin et al. 2021), it is instructive to investigate how well our inference scheme performs in this particular region of parameter space. To that end, we define a subset of *fiducial* objects with neutral fractions in the range $0.25 \leq \langle x_{\text{HI}} \rangle \leq 0.75$ and lifetimes between $5.25 \leq \log_{10} t_{\text{Q}}/\text{yr} \leq 6.75$. In addition to the overall median precision, we always quote the median of these 16 (11) objects for the analytical (full-simulation) IGM transmission model to confirm that the trends we identify are not the result of an atypical region of parameter space.

2.3.3 Parameter Scan

In this section, we will systematically investigate the effects of various hyper-parameters related to the continuum model and the observational setup on inference precision as measured in terms of $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ and $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$ defined in Eq. (2.16). We do so by subsequently varying 1) the latent dimension n_{latent} of our PCA continuum model, 2) the maximum wavelength coverage λ_{cut} on the red side of the quasar spectrum, 3) the signal-to-noise ratio S/N per 100 km/s velocity stretch, and 4) spectral resolution on the blue side of the spectrum, quantified in terms of FWHM of the LSF, keeping the spectral sampling factor fixed to $\text{FWHM}/dv_{\text{blue}} = 2$. When varying one of these parameters, we keep the other parameters fixed to the fiducial values of $n_{\text{latent}} = 5$, $\lambda_{\text{cut}} = 2000 \text{ \AA}$, S/N = 10, and FWHM = 100 km/s (with a velocity spacing of $dv_{\text{blue}} = 50 \text{ km/s}$).

2.3.3.1 Latent Dimension

An important component of our pipeline is the quasar continuum dimensionality reduction (DR) model introduced in Section 2.2.1. It reduces the number of pixels $n_{\lambda} = O(100 - 1000)$ covering the full spectral range to a low n_{latent} -dimensional parametric (PCA) representation based on 14 781 low-redshift SDSS-autofit continua. The PCA vectors are rank-ordered by the amount of variation in the sample that they account for. By only keeping the first n_{latent} vectors (where $n_{\text{latent}} \ll n_{\lambda}$), we can reduce the dimensionality of our continuum model to an arbitrary extent, typically $n_{\text{latent}} = O(1 - 10)$. The precise number of vectors n_{latent} is a free hyper-parameter of the model. It should be large enough to allow for sufficient model flexibility, yet too high of a number of latent factors comes with the danger of overfitting. In

this section we determine the optimal dimensionality of our PCA continuum model for the inference task at hand, which is to say the smallest latent dimension that minimizes the errors on the inferred astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} .

Figure 2.4 shows a representative example of our inference pipeline applied to the same quasar spectrum with different choices for the latent dimension of the continuum model. To analyze the overall contribution of the continuum model to the error budget, we also generated a continuum-normalized version of the spectrum. Normalizing the mock spectrum by the (in reality unknown) continuum completely removes the need for the continuum DR model. In this way, we eliminate all continuum nuisance parameters from the analysis (i.e., PCA coefficients ξ as well as the normalization factor s_{norm} of the continuum), leaving the astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} as the only two parameters to infer. This setting can be seen as an optimal bound as to how constraining the resulting posteriors can theoretically get assuming perfect knowledge of the continuum. In Figure 2.4, we compare such a run to runs with a low-dimensional continuum model with $n_{\text{latent}} = 5$ and a higher-dimensional one with $n_{\text{latent}} = 15$. The continuum inferred with $n_{\text{latent}} = 15$ shows an intense oscillation blueward of the Lyman- α line that is not present in the underlying mock continuum, suggesting that this model is too flexible and that the additional PCA vectors do not capture additional physics.

In the left panel of Figure 2.5 we show the marginalized posterior distributions of $\langle x_{\text{HI}} \rangle$ and t_{Q} for the three inference runs depicted in Figure 2.4. Note that these and all subsequent runs have been coverage corrected according to the procedure described in Section 2.3.1 to assure exact marginal coverage in the 2d-astrophysical parameter space. This assures that we do not misinterpret narrow but overconfident posteriors as overly precise. We show all original inference test results in Figure 2.12 in Appendix 2.B.

It is clearly apparent that the posterior PDF of the $n_{\text{latent}} = 15$ model is significantly less informative and extends further along the axis of degeneracy between the two astrophysical parameters than the PDF of the lower-dimensional model. Independently of the choice of latent dimension, the posteriors are significantly wider than those obtained with the continuum-normalized model where both parameters are very tightly constrained, indicating that continuum reconstruction contributes significantly to the overall error budget.

These individual trends are confirmed when considering the median inference precision $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ and $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$ as introduced in Section 2.3.2 with respect to 100 mock samples drawn from the full prior volume. We depict $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ and $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$ with respect to both the full sample (100 objects) and the fiducial subset (16 objects) as a function of latent dimension in the upper left panels of Figure 2.6. The errorbars are estimated in terms of the 16% and 84%-percentiles with respect




Figure 2.5: Marginal $(\langle x_{\text{HI}} \rangle, t_Q)$ -posterior distributions inferred from one of our 100 mock quasars with $\langle x_{\text{HI}} \rangle = 0.74$ and $\log_{10} t_Q/\text{yr} = 6.11$ with different hyper-parameter settings. IGM transmission skewers are generated by the analytical model introduced in Section 2.2.2.2. All HMC samples are reweighted, assuring that each ensemble passes a marginal coverage test with respect to the two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_Q . Whenever applicable, we marginalized over all nuisance parameters related to the shape of the continuum. We vary latent dimension n_{latent} , red-side wavelength coverage λ_{cut} , signal-to-noise ratio S/N , and spectral resolution FWHM, in the four panels from left to right. Each panel contains three representative choices of the respective parameter. The corresponding inferred models are shown in Figures 2.4 and 2.8. Our reference run with $n_{\text{latent}} = 5$, $\lambda_{\text{cut}} = 2000 \text{ \AA}$, $S/N = 10$ and $\text{FWHM} = 100 \text{ km/s}$ is depicted in black in each panel. The latent dimension plot also shows the theoretically optimal posteriors as obtained from the run with continuum-normalized spectra labelled as ‘astro-only’.




Figure 2.6: Inference precision with respect to IGM neutral fraction (x_{HI}) (upper half-panels) and quasar lifetime t_{Q} (lower half-panels) as a function of different hyperparameters (i.e., n_{latent} , λ_{cut} , S/N, and FWHM). IGM transmission skewers are generated by the analytical model introduced in Section 2.2.2.2. Median precision values with respect to each ensemble of 100 mock quasars are depicted in yellow (light green), and with respect to the subset of 16 fiducial spectra in red (green), shifted slightly in x -direction with respect to each other for visibility. Error bars correspond to the 16% and 84% percentiles of the respective mock ensemble. In the case of spectral resolution, the fiducial samples are taken from a separate set of 100 objects with fixed neutral fraction ($x_{\text{HI}} = 0.5$) and fixed lifetime $t_{\text{Q}} = 10^6$ yr to reduce stochasticity arising from varying noise draws between mocks at different resolutions. Labelled as ‘astro-only’ on the left of the latent dimension plot, we show the optimal precision bounds as obtained from the run with continuum-normalized spectra. All precision values are based on reweighted marginal posterior PDFs, assuring that each mock ensemble passes a marginal coverage test with respect to the two astrophysical parameters (x_{HI}) and t_{Q} .

to the respective set of objects. We observe that regardless of the choice of n_{latent} , the posteriors are 2 – 3 times wider than the theoretically optimal ones obtained in the continuum-normalized case. This indicates significant contributions of continuum reconstruction to the total error budget which we quantify in further detail at the end of this section. Secondly, there is surprisingly little variation among different choices of latent dimension between $2 \lesssim n_{\text{latent}} \lesssim 10$. If any, there is a slight trend of deteriorating inference precision with higher latent dimension, becoming more evident at $n_{\text{latent}} \gtrsim 10$. We observe the same qualitative behavior for the fiducial subset of mock spectra.

In our companion paper (Hennawi et al. 2025), we tested at which latent dimension n_{latent} the relative continuum reconstruction error flattens out when representing continua with a n_{latent} -dimensional PCA continuum model. We found that the reconstruction error continued to decrease with dimension at least out to $n_{\text{latent}} \sim 10$ (cf. Figures 3 and 4 in Hennawi et al. 2025). Note that here we are asking a different question, that is to say at which latent dimension the precision on the inferred astrophysical parameters flattens out.

We compare the two scenarios by computing the continuum reconstruction error δ as defined in Eq. (2.3) with respect to the pure dimensionality-reduced continuum s_{DR} as well as the inferred continuum s_{inf} . Note that we define the inferred continuum as the median of all model continua corresponding to the $\eta = (s_{\text{norm}}, \xi)$ -samples of the full HMC chain. In order to distinguish the constraining power in the Lyman- α forest and damping wing region from that in the smooth emission line region of the spectrum, we separately compute the mean $\langle \delta \rangle$ and the standard deviation $\sigma(\delta)$ of the reconstruction error with respect to the spectral pixels in the blue ($1185 \text{ \AA} \leq \lambda \leq 1260 \text{ \AA}$) and the red part ($1260 \text{ \AA} < \lambda \leq 2000 \text{ \AA}$) of the spectrum. We chose these ranges in agreement with those in Figures 3 and 4 in Hennawi et al. (2025).

We show the median values of these quantities with respect to the ensemble of 100 mock objects as a function of latent dimension n_{latent} in Figure 2.7. The behavior of the inferred and the dimensionality-reduced continua excellently agrees on the red side of the spectrum. Both the bias $\langle \delta \rangle$ and the standard deviation $\sigma(\delta)$ decrease with increasing latent dimension due to the fact that more flexible models can always provide a closer fit to the true shape of the continuum (see also Figures 3 and 4 in Hennawi et al. 2025). As also observed in Hennawi et al. (2025), this curve does not flatten out before $n_{\text{latent}} = 10$, and by naive expectation, the same should hold true on the blue side of the spectrum. Indeed, this is the case for the purely dimensionality-reduced spectrum. We observe the opposite trend, however, for the reconstruction errors of the inferred continuum whose standard deviation reaches a shallow minimum around $n_{\text{latent}} \sim 5$ while the bias $\langle \delta \rangle$ oscillates between -0.02 and 0.0 without a clear dependence on the value of n_{latent} . Most particularly, $\sigma(\delta)$ diverges from the curve of the




Figure 2.7: Moments of the continuum reconstruction error δ as a function of latent dimension n_{latent} of the PCA continuum model. The panels on the left show the mean $\langle \delta \rangle$ and the standard deviation $\sigma(\delta)$ of the continuum reconstruction error with respect to the blue pixel range ($1185 \text{ \AA} \leq \lambda \leq 1260 \text{ \AA}$). The panels on the right depict the same moments with respect to the red pixels of the spectrum ($1260 \text{ \AA} < \lambda \leq 2000 \text{ \AA}$). Each individual data point represents the median of the corresponding moment with respect to the 100 object mock ensemble with $\lambda_{\text{cut}} = 2000 \text{ \AA}$, $S/N = 10$ and $\text{FWHM} = 100 \text{ km/s}$. The moments related to the full inference task δ_{inf} are depicted as solid blue (red) lines, the ones arising solely from the dimensionality reduction task δ_{DR} as dashed purple (yellow) lines. The standard deviation $\sigma(\delta_{\text{inf}})$ of the blue-side reconstruction error of the inferred continuum has a shallow minimum around $n_{\text{latent}} \simeq 4$, whereas $\sigma(\delta_{\text{DR}})$ keeps decreasing monotonically with latent dimension like in the smooth red part of the spectrum.

dimensionality-reduced continuum at around $n_{\text{latent}} \gtrsim 5$ and starts to grow at an increasing rate.

We thus conclude that additional PCA vectors do keep improving the formal fit to the continuum, but they do not yield additional constraining power because they do not encode additional physical information about the blue side of the spectrum that is absorbed by the Lyman- α forest or the IGM damping wing. As a result, the variation of the reconstruction error of the inferred continuum starts to increase instead of continuing to decrease for higher-dimensional PCA models. Choosing a model that is too flexible is therefore not desirable for reconstructing the continuum accurately from a spectrum with a Lyman- α damping wing imprint. This insight, along with the weak dependence of inference precision $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ and $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$ on latent dimension, lead us to the conclusion that $n_{\text{latent}} = 5$ is an adequate choice for our PCA continuum model. Note, however, that this value depends on the spectral coverage which was here analyzed up to 2000 \AA . Spectra covering additional wavelength ranges

redward of 2000 Å would likely require PCA models of a higher dimensionality, but the general behavior found in this section would hold.

2.3.3.2 Wavelength coverage

The smooth emission lines on the red side of a quasar spectrum are correlated with the shape of the continuum blueward of the Lyman- α line (see Figure 7 in Hennawi et al. 2025). How well we can reconstruct the continuum in this region is going to be degenerate with the astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} governing the IGM transmission field. Gaining as much information as possible about the red-side shape of the continuum by covering more emission lines blueward of the Lyman- α line can thus enhance overall inference precision on the astrophysical parameters of interest.

Our fiducial wavelength grid covers the rest-frame wavelength range between 1185 Å and 2000 Å. In this section, we investigate how much information is contained in the outer regions on the red side of the spectrum, informing us about the importance of covering those regions when collecting new spectra for damping wing analysis.

Covering the inevitably required Lyman- α region with a given spectrograph will automatically provide us with coverage out to some wavelength redward of the Lyman- α line. Getting redder coverage out to 2000 Å in the rest-frame would likely incur more observational cost to observe multiple setups (depending on the spectrograph) and as such we aim to quantify in this section the precision gains from higher red-side wavelength coverage.

The two upper panels of Figure 2.8 show the example spectrum already considered in Figure 2.4, cut off at $\lambda_{\text{cut}} = 1300$ Å, as well as the full spectrum with $\lambda_{\text{cut}} = 2000$ Å. The uncertainties of the inferred continuum are notably higher when cutting off the spectrum at $\lambda_{\text{cut}} = 1300$ Å. The same is the case for the inferred astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} whose marginal posterior distributions are depicted in the second panel of Figure 2.5: inference precision significantly deteriorates for the shorter wavelength coverage of $\lambda_{\text{cut}} = 1300$ Å which excludes the S IV emission line around 1397 Å and the O IV] line at 1402 Å, as well as the C IV line around 1549 Å. Particularly the C IV emission line strength and shape correlates strongly with the shape of the continuum in the Lyman- α region where we need to reconstruct the IGM damping wing imprint (again see Figure 7 in Hennawi et al. 2025). An intermediate cut-off at $\lambda_{\text{cut}} = 1600$ Å which includes all these lines produces posteriors whose shape is almost identical to those of the full spectrum (i.e., $\lambda_{\text{cut}} = 2000$ Å).

The same trends hold for the full set of 100 mock quasar spectra. Between covering the full wavelength grid out to $\lambda_{\text{cut}} = 2000$ Å and a cut-off wavelength of $\lambda_{\text{cut}} = 1600$ Å, the posterior widths largely remain constant and gradually start increasing for shorter cut-off wavelengths, as seen in the upper right panels of Figure 2.6. The same is the case for the subset of spectra from the




Figure 2.8: Model components of the true and the inferred spectrum of one of our 100 mock quasars with $\langle x_{\text{HI}} \rangle = 0.74$ and $\log_{10} t_{\text{Q}}/\text{yr} = 6.11$, considered in different observational setups, and compared to our reference setting of $\lambda_{\text{cut}} = 2000 \text{ \AA}$, $S/N = 10$ and $\text{FWHM} = 100 \text{ km/s}$. IGM transmission skewers are generated by the analytical model introduced in Section 2.2.2.2. For better visibility of the damping wing region, the rest-frame wavelength axis is stretched on the left side of the figure between 1185 \AA and 1300 \AA . The full mock spectrum of the quasar is depicted in black and consists of the true continuum (green) combined with IGM transmission (purple) and mock spectral noise (yellow). The inferred model spectrum is depicted in red, consisting of the inferred continuum (blue) combined with the inferred IGM transmission field. Solid lines represent the median inferred models, shaded regions denote the 16% and the 84% percentile variations reflecting parameter uncertainty, continuum reconstruction errors, as well as spectral noise. All inferred model spectra are simulated based on reweighted HMC samples, assuring that each ensemble passes a marginal coverage test with respect to the two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} . Lower red-side wavelength coverage or a lower signal-to-noise ratio increases the uncertainties of the inferred model, whereas the dependence on resolution is relatively weak.

fiducial region of parameter space. Therefore, major emission lines (especially the C IV line) contain important information that significantly enhances the overall inference precision. The correlation with the weaker lines in the outer part of the spectrum, on the other hand, is much less pronounced. We can therefore conclude that inference quality does not suffer notably if wavelengths in the range between $1600 \text{ \AA} \leq \lambda \leq 2000 \text{ \AA}$ are missing due to observational constraints or cost (or if they are excluded to break down the computational complexity). Our analysis does not exclude, however, that important emission lines beyond $\lambda_{\text{cut}} = 2000 \text{ \AA}$, such as Mg II, or even H- β or H- α that are within the reach of JWST instruments, could further enhance inference precision. On the contrary, given the strong impact of the C IV emission line, it is not unlikely that covering these lines would further improve the overall constraints. The challenge in this regard is constructing an adequately large PCA training set for our continuum DR model. We postpone this task to future work.

2.3.3.3 Signal to Noise Ratio

Required exposure times and thus the cost of future observations crucially depends on the desired signal-to-noise ratio S/N of the spectrum. Potential gains in inference precision from obtaining higher- S/N spectra are thus highly informative about the most efficient setting for upcoming observations. Throughout this work, we define S/N as the signal-to-noise ratio per 100 km/s velocity interval, computed over the telluric absorption free observed frame wavelength range $11750 - 13300 \text{ \AA}$ (i.e., $1376 - 1557 \text{ \AA}$ in the rest-frame at $z_{\text{QSO}} = 7.54$). Figure 2.8 shows examples of inferred spectra for a very noisy ($S/N = 2$), a very high signal-to-noise ratio ($S/N = 20$), and our reference spectrum with $S/N = 10$. The random seed is the same for each realization and only the amplitude of the noise is adjusted. While the high signal-to-noise spectrum does not lead to significant improvements over the intermediate one, the low signal-to-noise spectrum has somewhat larger error bars in the inferred continuum. The shape of its marginalized posterior distribution is wider as well, as the contours in the third panel of Figure 2.5 suggest.

The 100 object ensemble considered in the bottom left panels in Figure 2.6 confirms these trends: both neutral fraction and lifetime precision ($\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ and $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$) do not change significantly between $10 \lesssim S/N \lesssim 20$ and start increasing notably at smaller signal-to-noise ratios. The same trends are apparent for the subset of fiducial samples. We therefore conclude that spectra with signal-to-noise ratios of $S/N \sim 10$ are sufficient for a precise parameter inference.

2.3.3.4 Spectral resolution

Spectral resolution is one of the most important factors to consider for both observational and computational reasons. Modern-day instruments such as X-Shooter are able to collect spectra with resolutions of $R \sim 10\,000$ even for the faintest quasars that can be of interest for damping wing analysis.

Resolving the Lyman- α forest with such high resolutions might be beneficial since within the noise, each transmission spike constitutes a lower limit for the continuum at the corresponding wavelength pixel. In other words, the continuum should never fall below the observed data as the IGM transmission values can never exceed unity. Furthermore, resolving the proximity zone with additional resolution elements may further increase constraining power. On the other hand, higher resolutions hold little to no gains in the region of the spectrum redward of the Lyman- α line where the flux is a smooth function of wavelength. Hence, we divide the wavelength grid at $\lambda \simeq 1218 \text{ \AA}$ into a low-resolution red part and a blue part where we vary resolution between low values comparable to the red-side resolution all the way up to X-Shooter-like high-resolution values.

Yet, matrix inversion and therefore the cost of a likelihood evaluation scale cubically with the number of spectral pixels, and even after introducing the aforementioned hybrid wavelength grid, we end up with up to $O(10^4)$ spectral pixels, making the repeated inversion of the corresponding covariance matrices for the likelihood computations in the inference process notably more expensive. Besides the observational cost of collecting such high-resolution spectra, it is hence also relevant from a computational point of view to quantify the gains of analyzing spectra at high spectral resolution.

We remind the reader that we quantify spectral resolution in terms of the FWHM of the LSF. When varying FWHM, we proportionally adjust the velocity pixel scale, keeping the spectral sampling factor fixed to 2. Note also that we always work with a constant signal-to-noise ratio of $S/N = 10$ *per* 100 km/s velocity interval. This assures that we do not erroneously attribute signal-to-noise trends to spectral resolution.

Our technical implementation comes with the minor caveat that each mock realization of the spectrum at a different resolution has its unique noise draw. Though this acts as an additional source of stochasticity, we find that after averaging over all 100 mock spectra the effects are small enough to deduce an overall trend. The sample size of our 16 fiducial spectra, however, is not sufficient to average out the statistical fluctuations due to the different noise draws. To gain a statistically meaningful picture in the fiducial region of $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ -parameter space, we instead fix the astrophysical parameters to the fiducial values of $\langle x_{\text{HI}} \rangle = 0.5$ and $t_{\text{Q}} = 10^6$ yr and simulate a new set of 100 mock spectra which only differ by the draws of the continuum and of the lognormal Lyman- α forest fluctuations.

Strikingly, the effects of changing spectral resolution are smaller than one might expect. Figure 2.8 depicts the inferred continuum of our example spectrum for the reference resolution of FWHM = 100 km/s, a fine resolution of FWHM = 30 km/s, and a coarse value of FWHM = 1000 km/s whose pixel scale of 500 km/s coincides with that on the red side of the spectrum. The inferred continua look remarkably similar; even for the coarsest resolution, the general shape of the continuum is well captured. The corresponding posteriors in Figure 2.5 are somewhat less informative but still largely retain their original shape even for FWHM = 1000 km/s.

We find similarly weak trends for the full ensemble of 100 objects. The width of the posteriors in Figure 2.6 only changes very slightly with resolution. In particular, the difference between the median widths of the $\langle x_{\text{HI}} \rangle$ posteriors from the highest- and lowest-resolution runs is only $\sim 4\%$, and ~ 0.03 dex for $\log_{10} t_{\text{Q}}/\text{yr}$. This shows that spectral resolution can be lowered drastically almost without any notable impact on inference precision. The reason for this appears to be that resolving proximity zone absorption lines is secondary for a precise inference of the astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} . The main constraining power seems to come from the smooth IGM damping wing signature along with the size and shape of the quasar proximity zone which are still determined to within reasonable accuracy even for the lowest-resolution spectra (cf. Figure 2.8). Merely resolutions that resolve all individual Ly- α absorption lines could possibly lead to improvements but might not be preferable from a computational point of view, given the computational cost of analyzing higher resolution spectra and the limited benefit with respect to constraining power.

Note that in the case of the full-simulation IGM transmission model, higher resolution *does* lead to an improvement in inference precision. The inferred posteriors, however, are significantly overconfident. This is a result of our Gaussian approximation for the IGM transmission likelihood (Eq. (2.5)) being less adequate at increasingly high resolutions (see also Figure 11 in Hennawi et al. 2025). We elaborate in detail on coverage tests and inference precision of the full-simulation model in Appendices 2.B and 2.C. After correcting for the overconfidence, the dependence of inference precision on spectral resolution is similarly weak as in the case of the simple analytic IGM transmission model. Informed by these trends, we adopt the coarse red-side resolution (FWHM = 1000 km/s) across the entire wavelength grid, pending further improvements in our likelihood approximation that we leave for future work. Given that the need to apply big coverage corrections is currently obviating the gains from higher resolution spectra, this choice of resolution minimizes the computational and observational cost of the analysis. It is also advantageous from a computational point of view as it reduces the total number of spectral pixels to $n_{\lambda} = 313$. When applied to real-world data from a higher-resolution spectrograph, we opt to achieve a comparable setting by rebinning the spectrum to the corresponding velocity

pixel scale of 500 km/s. This is also the strategy we adopt in the following section to characterize the overall error budget of our inference pipeline as realistically as possible.

2.3.3.5 The optimal hyper-parameter setting

In summary, based on the preceding results, we adopt the following choices as the optimal hyper-parameter setting for the subsequent quantification of the full error budget of our inference pipeline: we choose a latent dimension of $n_{\text{latent}} = 5$ for the PCA dimensionality reduction model (along with the amplitude degree-of-freedom resulting in a six-parameter continuum model), a rest-frame wavelength coverage out to $\lambda_{\text{cut}} = 2000 \text{ \AA}$, an intermediate signal-to-noise ratio of $S/N = 10$ per 100 km/s velocity interval, and a spectral resolution of $\text{FWHM} = 100 \text{ km/s}$ rebinned to a velocity pixel scale of 500 km/s. We emphasize that even though the analysis above was based on the analytical IGM transmission model, we arrive at the same hyper-parameter choices when modelling IGM transmission with our full-simulation model. We discuss these results in Appendix 2.C. As pointed out above, the only notable difference between the two IGM transmission models with regard to hyper-parameter variations is the deteriorating coverage behavior in the full-simulation model which is obviating the gains from high-resolution spectra and led us to conclude that we can adopt the coarse red-side resolution across the entire spectrum. To eliminate any remaining impact from non-optimal coverage behavior, we apply the marginal coverage correction procedure described in Section 2.3.1 also to the mock ensemble considered in the following section.

2.3.4 The Total Error Budget

All inference runs in the previous sections were obtained with the analytical IGM transmission model described in Section 2.2.2.2. Its deterministic relationship between the astrophysical parameters $\theta = (\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ and damping wing strength allowed us to gain a maximally clear picture of how inference precision depends on the hyper-parameters considered. In doing so, we neglected the fact that in realistic models of reionization such as the 21cmFAST topologies of our full-simulation model, the size of reionization bubbles does not depend deterministically on the global IGM neutral fraction. In addition, by performing radiative transfer calculations, we account for the fact that fluctuations in the HI density field will impact the size of the quasar’s ionization front which we previously approximated deterministically according to Eq. (2.7).

In this section, we account for these two sources of stochasticity by modelling IGM transmission according to our full-simulation model introduced in Section 2.2.2.1. Re-running an identical set of mock objects, combined




Figure 2.9: Inference precision with respect to IGM neutral fraction $\langle x_{\text{HI}} \rangle$ as a function of $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ -parameter space, broken up into the individual model components. The two rows compare the two IGM transmission models, the columns show the effect of the continuum model. The dashed boxes enclose the fiducial region of parameter space, centered around $\langle x_{\text{HI}} \rangle = 0.5$ and $t_{\text{Q}} = 10^6$ yr. All precision values are computed from our 1000 object ensemble, ran with the optimal hyper-parameter setting ($n_{\text{latent}} = 5$, $\lambda_{\text{cut}} = 2000$ Å, $S/N = 10$, and $\text{FWHM} = 100$ km/s rebinned to $dv_{\text{blue}} = 500$ km/s). Each pixel shows the average precision $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ with respect to the same 10 continua. All precision values are based on reweighted marginal posterior PDFs, assuring that each mock ensemble passes a marginal coverage test with respect to the two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} . Inference precision increases for objects with a stronger damping wing imprint. Both the continuum model and the stochasticity of reionization encoded in the full-simulation model add a similar amount of uncertainty, as quantified in Table 2.1.




Figure 2.10: Like Figure 2.9, but for lifetime inference precision $\Delta_{68\%}(\log_{10} t_Q/\text{yr})$. For the analytical IGM transmission model, a stronger damping wing imprint implies higher lifetime precision, as seen before in the case of $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$. The stochasticity of the full-simulation model erodes these trends. The longest-lived objects show extraordinarily high precision due to the thermal proximity effect.

with analytical IGM transmission skewers, allows us to precisely quantify the impact of the aforementioned sources of stochasticity on the total error budget of our inference pipeline. By repeating the inference after dividing each spectrum by its underlying continuum, we are further able to quantify the contribution of the continuum reconstruction error on the final inference precision. In other words, we break down the total error budget of our inference scheme into the separate contributions from the continuum model and the two aforementioned additional sources of stochasticity encoded in our full-simulation model.

We further investigate how these contributions change depending on the exact location in $(\langle x_{\text{HI}} \rangle, t_Q)$ -parameter space. To that end, we define a linear (10×10) -grid covering our full $(\langle x_{\text{HI}} \rangle, \log_{10} t_Q/\text{yr})$ -prior range, i.e., $0 \leq \langle x_{\text{HI}} \rangle \leq 1$ and $3 \leq \log_{10} t_Q/\text{yr} \leq 8$. In addition, we draw 10 distinct autofit continua and combine them with 10 distinct IGM transmission skewers at each location of the $(\langle x_{\text{HI}} \rangle, \log_{10} t_Q/\text{yr})$ -grid. By considering the same

10 continua at each point in astrophysical parameter space, we suppress additional stochasticity that would otherwise arise due to varying continuum draws. This results in an overall ensemble of $(10 \times 10) \times 10 = 1000$ mock spectra.

Figure 2.9 summarizes all aforementioned dependencies for $\langle x_{\text{HI}} \rangle$ -inference precision $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$. The four different panels show inference precision in $(\langle x_{\text{HI}} \rangle, \log_{10} t_{\text{Q}}/\text{yr})$ -parameter space for the same 1000 mock objects considered with/without continuum model and with analytic/full-simulation IGM transmission model, respectively. Lifetime inference precision $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$ is shown in Figure 2.10. Table 2.1 lists the respective median precision values with respect to the full prior volume and the fiducial region of parameter space as defined in Section 2.3.2 (enclosed by the dashed boxes in Figures 2.9 and 2.10).

Independently of the model configuration, we see significant variation as a function of the astrophysical parameter values. As an underlying general trend, we find an overall improvement in constraining power for short quasar lifetimes t_{Q} , and, to a somewhat lesser extent, for higher neutral fractions $\langle x_{\text{HI}} \rangle$. This behavior can be directly attributed to the enhanced damping wing strength of quasars in a neutral environment and short-lived quasars that have had less time to carve out a large ionized bubble in their immediate proximity. The trend is particularly pronounced for the analytic IGM transmission model where this dependence on $\langle x_{\text{HI}} \rangle$ and t_{Q} is incorporated explicitly through Eqs. (2.7) and (2.13).

Secondly, we point out that inference precision can vary notably between the same continua combined with analytical or full-simulation IGM transmission skewers. We observe a significantly increased amount of scatter for the full-simulation model and overall deteriorating inference precision in most regions of parameter space. This is an immediate effect of the additional sources of stochasticity due to the non-deterministic bubble-size distribution from reionization, and the size of the ionized bubble sourced by the quasar. As a result, in our full model, $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ can be as low as 3.0% and as high as 42.2% for individual objects. Likewise, $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$ varies between 0.07 dex and 1.12 dex. The full-parameter space medians of $\Delta_{68\%}(\langle x_{\text{HI}} \rangle) = 28.0_{-8.8}^{+8.2}\%$ and $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr}) = 0.80_{-0.55}^{+0.22}$ dex, respectively, are dominated to a similar extent by the errors related to continuum reconstruction and the stochasticity of the ionized bubble size distribution: for the most basic task of exclusively inferring the two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} from a continuum-normalized spectrum with deterministic IGM transmission model, we are able to recover $\langle x_{\text{HI}} \rangle$ with only $2.2_{-0.8}^{+2.5}\%$ uncertainty as quantified in terms of $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$. Modelling the quasar continuum and switching to the full-simulation IGM transmission model encoding the stochasticity of reionization and the quasar ionized bubble sizes raises $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ to a total of $28.0_{-8.8}^{+8.2}\%$. We can see that the addition of continuum reconstruction error and the full-simulation IGM transmission

Table 2.1: Median $\langle x_{\text{HI}} \rangle$ - and t_{Q} -inference precision $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ and $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$, broken up into the individual model components as seen in Figures 2.9 and 2.10. The values listed are medians of all samples in the full (or fiducial) region of parameter space, given the corresponding model configuration. The parameter limits of the fiducial region are enclosed by dashed boxes in Figures 2.9 and 2.10. Uncertainties are estimated according to the respective 16% and 84% percentile values.

	$\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$				$\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$	
	Continuum-normalized	Full continuum model	Continuum-normalized	Full continuum model	Continuum-normalized	Full continuum model
All samples	Analytical Model	$2.2^{+2.5}_{-0.8}\%$	$15.2^{+12.8}_{-6.9}\%$	$0.12^{+0.12}_{-0.06}\text{ dex}$	$0.36^{+0.19}_{-0.18}\text{ dex}$	
	Full-Simulation Model	$15.2^{+3.0}_{-4.2}\%$	$28.0^{+8.2}_{-8.8}\%$	$0.55^{+0.35}_{-0.34}\text{ dex}$	$0.80^{+0.22}_{-0.55}\text{ dex}$	
Fiducial samples	Analytical Model	$2.4^{+1.1}_{-0.7}\%$	$17.4^{+9.3}_{-5.9}\%$	$0.08^{+0.06}_{-0.03}\text{ dex}$	$0.28^{+0.20}_{-0.11}\text{ dex}$	
	Full-Simulation Model	$15.8^{+2.4}_{-2.4}\%$	$28.5^{+7.1}_{-6.5}\%$	$0.72^{+0.19}_{-0.32}\text{ dex}$	$0.95^{+0.09}_{-0.21}\text{ dex}$	

model contribute a roughly equal amount to the total error budget by turning them on separately: the continuum model raises $\Delta_{68\%}(\langle x_{\text{HI}} \rangle)$ to $15.2^{+12.8}_{-6.9}\%$, and the full-simulation IGM transmission model to the almost identical value of $15.2^{+3.0}_{-4.2}\%$ (see Table 2.1). Most notably, this implies that the intrinsic stochasticity of reionization and quasar ionized bubble sizes accounts for almost half of the total error budget. Eliminating this stochasticity from the inference process could lead to significant precision gains in future inference pipelines.

The overall picture for lifetime precision $\Delta_{68\%}(\log_{10} t_{\text{Q}}/\text{yr})$ is similar. We can infer t_{Q} from a continuum-normalized spectrum with analytical IGM transmission model with an uncertainty of $0.12^{+0.12}_{-0.06}$ dex. This increases to $0.36^{+0.19}_{-0.18}$ dex when modelling the quasar continuum and keeping the IGM transmission model deterministic, and to $0.55^{+0.35}_{-0.34}$ dex when switching to the full-simulation IGM transmission model but not performing continuum reconstruction, adding up to a total of $0.80^{+0.22}_{-0.55}$ dex when both models are turned on simultaneously. Thus, the intrinsic stochasticity of reionization and the quasar ionized bubble sizes even dominates over continuum reconstruction uncertainties in the total lifetime error budget. When comparing the full-simulation to the analytical IGM transmission model, an additional feature in $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ -parameter space is worthy of note: lifetime uncertainties of the oldest objects ($t_{\text{Q}} \gtrsim 10^7$ yr) are significantly smaller than those of intermediate-lifetime ones. This opposes the general trend that we pointed out in Section 2.2.4 (i.e., the longer the lifetime, the larger the ionized region around the quasar and hence the weaker the damping wing signature, exacerbating the parameter inference). Notably, the lifetime uncertainties of parts of these objects are even smaller than those of their analytical counterparts. We attribute this improvement in precision to the thermal proximity effect of He II reionization (Meiksin et al. 2010; Bolton et al. 2010, 2012; Khrykin et al. 2016): the photoelectric heating resulting from the He II \rightarrow He III photoionizations gives rise to a thermal proximity zone around the quasar with a significantly increased IGM temperature, causing a change in shape of the IGM transmission profile via the Gunn-Peterson optical depth formula. Eq. (2.14) of our analytical IGM transmission model, on the other hand, does not account for this effect. The size of the thermal proximity zone increases with quasar lifetime, facilitating the inference for the longest-lifetime objects. As a result, objects in the fiducial region of parameter space with weaker He II proximity effect suffer the most under the more realistic IGM transmission model based on cosmological simulations, as evident in a fiducial parameter space median lifetime precision of $0.95^{+0.09}_{-0.21}$ dex, with $0.72^{+0.19}_{-0.32}$ dex caused by the stochasticity of reionization alone.

2.4 Conclusions

In this work we explored a pipeline to infer the global volume-averaged IGM neutral fraction and the quasar lifetime using the damping wing signature of high-redshift quasars. We determined the optimal hyper-parameter setting and rigorously quantified the associated error budget. The inference pipeline has been introduced in Hennawi et al. (2025) and consists of a parametric (PCA) model for the quasar continuum, as well as a realistic forward model for IGM transmission based on skewers extracted from hydrodynamical simulations and a semi-numerical reionization topology combined with 1d ionizing radiative transfer. To isolate the impact of the stochasticity of reionization as well as the quasar ionized bubble sizes, we also introduced a deterministic analytical IGM transmission model. We tested the pipeline on a fixed set of mock observational spectra with heteroscedastic observational noise, varying the latent dimension of our PCA continuum model, the wavelength coverage in the red part of the spectrum, the signal-to-noise ratio as well as spectral resolution. We found that:

1. Latent dimensions of $n_{\text{latent}} \sim 4 - 5$ are sufficient to capture all information that is intrinsically related to the IGM damping wing imprint. This conclusion differs from the pure dimensionality reduction task where Hennawi et al. (2025) found the continuum reconstruction error to decrease monotonically as a function of n_{latent} . Our results suggest that PCA vectors beyond $n_{\text{latent}} > 5$ do not encode information about the damping wing signature and thus do not yield additional constraining power.
2. Longer wavelength coverage in the red part of the spectrum improves inference precision, especially including major emission lines such as the Si IV and the O IV] line around 1400 Å, as well as the C IV line around 1549 Å. However, if this information is not available, even the Lyman- α damping wing signature alone can yield non-trivial constraints. The trends suggest potential additional gains from covering major emission lines beyond our maximum wavelength coverage of 2000 Å, such as Mg II as well as H- β or H- α .
3. A higher signal-to-noise ratio leads to higher precision, with $S/N \gtrsim 10$ per 100 km/s velocity stretch being desirable for sufficiently precise results. We note for the full-simulation IGM transmission model that higher S/N comes at the price of decreasing coverage probability as our likelihood approximation fails to capture non-Gaussian information that is otherwise gaussianized by the noise.
4. Given the current implementation of the inference pipeline, spectral resolution does not have a notable impact on inference precision. All

gains that are achieved in the case of the full-simulation IGM transmission model are obviated by the coverage corrections we have to apply to the increasingly overconfident posteriors. We conclude that, pending improvements in our pipeline, we can reduce computational as well as observational cost by using the coarse red-side resolution of $dv = 500$ km/s for the entire spectrum.

Given the optimal hyper-parameter setting, we quantified the individual uncertainty contributions of the continuum model as well as the stochasticity of the ionized bubble size distribution as well as the quasar ionized bubble sizes as a function of astrophysical parameter space. We concluded that:

1. Inference precision varies strongly across $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ -parameter space, with the highest precision achieved for objects with a strong damping wing imprint (i.e., young quasars in a largely neutral environment). The observationally most relevant parameter values lie in the intermediate precision range. Inferring very long quasar lifetimes ($t_{\text{Q}} \gtrsim 10^7$ yr) is facilitated by the thermal proximity effect.
2. Using our most sophisticated model setting (continuum model + full-simulation IGM transmission model), we can infer $\langle x_{\text{HI}} \rangle$ to a median uncertainty of $28.0^{+8.2}_{-8.8}$ % across the full parameter space covered by our simulation grid (i.e., $0 \leq \langle x_{\text{HI}} \rangle \leq 1$ and $3 \leq \log_{10} t_{\text{Q}}/\text{yr} \leq 8$), and to $28.5^{+7.1}_{-6.5}$ % in the fiducial region between $0.25 \leq \langle x_{\text{HI}} \rangle \leq 0.75$ and $5.25 \leq \log_{10} t_{\text{Q}}/\text{yr} \leq 6.75$. We can constrain quasar lifetimes to $0.80^{+0.22}_{-0.55}$ dex and $0.95^{+0.09}_{-0.21}$ dex in the full and fiducial region of parameter space, respectively.
3. Across full parameter space, the (PCA-based) quasar continuum model contributes $15.2^{+12.8}_{-6.9}$ % to the total error budget of $\langle x_{\text{HI}} \rangle$, and a roughly equal amount of $15.2^{+3.0}_{-4.2}$ % is due to the stochasticity of reionization and quasar ionized bubble sizes. Across the fiducial region of parameter space, we find contributions of $17.4^{+9.3}_{-5.9}$ % and $15.8^{+2.4}_{-2.4}$ %, respectively. Concerning lifetimes, the contributions of continuum model and stochasticity of reionization are $0.36^{+0.19}_{-0.18}$ dex and $0.55^{+0.35}_{-0.34}$ dex, respectively, across the full parameter space, and $0.28^{+0.20}_{-0.11}$ dex and $0.72^{+0.19}_{-0.32}$ dex in the fiducial region, showing that the total error budget is largely dominated by the stochasticity of reionization and quasar ionized bubble sizes.

Additional modelling uncertainties whose consideration exceeds the scope of this work might further contribute to a minor extent to the final inference precision. Specifically, one could consider the effects of varying reionization topologies at fixed global IGM neutral fraction, resulting in different distributions of ionized bubble sizes. We further assumed that *all* halos with masses above $M_{\text{halo}} \gtrsim 2 - 3 \times 10^{11} M_{\odot}$ host quasars without considering non-trivial

quasar-host mass functions. In addition, flickering lightcurves might lead to small deviations from the simple ‘light bulb’ model which we assumed here.

Our results pave the way for the application of our inference machinery to real observational spectra to obtain robust constraints on the timing of reionization as well as the growth of supermassive black holes. Combining the individual constraints of $O(10 - 100)$ independent objects at a similar redshift (as are soon to be detected by the Euclid wide-field survey, [Euclid Collaboration et al. 2019](#)) will scale down the aforementioned uncertainties by a factor of $\sqrt{10} - \sqrt{100}$, allowing us to place $\lesssim 5\%$ -constraints on the evolution of the IGM neutral fraction as a function of redshift.

Acknowledgements

We thank John Tamanas for contributions to an earlier version of the inference code base. We acknowledge helpful conversations with the ENIGMA group at UC Santa Barbara and Leiden University. 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](#)), SkyCalc_ipy ([Leschinski 2021](#)), h5py ([Collette 2013](#)), Matplotlib ([Hunter 2007](#)), corner.py ([Foreman-Mackey 2016](#)), and IPython ([Pérez & Granger 2007](#)). This work was performed using the compute resources from the Academic Leiden Interdisciplinary Cluster Environment (ALICE) provided by Leiden University. 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.

Appendices

2.A Correction to the Analytical Damping Wing Formula

In Section [2.2.2.2](#) we introduced an analytical expression (Eq. [\(2.13\)](#)) governing the damping wing optical depth. Our derivation closely follows that in [Mortlock \(2016\)](#) but differs in a prefactor of 1/8 in the second term of




Figure 2.11: Approximations of the Lyman- α cross section of neutral hydrogen (left) and associated Lyman- α transmission field t in the damping wing region of the spectrum (right). The [Bach & Lee \(2015\)](#) profile is an empirical fit to the full quantum-mechanical cross section; the Lorentzian profile (Eqs. (2.11) - (2.15)) and the two-level model (Eq. (2.19), [Miralda-Escudé 1998](#)) are two simplified analytical approximations thereof. The transmission fields are calculated by integrating Eq. (2.10) directly using the [Bach & Lee \(2015\)](#) fitting formula, or substituting Eqs. (2.15), (2.18) and (2.19), respectively, into Eq. (2.13) assuming $\langle x_{\text{HI}} \rangle = 1$, $z_{\text{QSO}} = 7.54$, $z_{\text{ion}} = 7.5$ (corresponding to a reionization front at $R_{\text{ion}} = 1.5 \text{ pMpc}$) and $z_{\text{end}} = 6.0$. The lower right panel shows the relative error of the approximate transmission profiles with respect to that obtained with the [Bach & Lee \(2015\)](#) fitting formula. The vertical line highlights the Lyman- α resonance $\lambda_{\alpha} = 1216.78 \text{ \AA}$.

Eq. (2.15) instead of $1/2$ in Eq. (40) in [Mortlock \(2016\)](#). This factor arises when solving the dimensionless integral

$$I(x) = \int \frac{x^{1/2}}{4(1-x)^2} dx = \frac{x^{1/2}}{4(1-x)} + \frac{1}{8} \ln \left(\frac{1-x^{1/2}}{1+x^{1/2}} \right) \quad (2.17)$$

$$\neq \frac{x^{1/2}}{4(1-x)} + \frac{1}{2} \ln \left(\frac{1-x^{1/2}}{1+x^{1/2}} \right). \quad (2.18)$$

Parts of our results (i.e., Section 2.3.3 and Appendices 2.B and 2.C), were obtained with the incorrect prefactor of $1/2$. We investigate the effect of the correction in Figure 2.11 where we show the relative deviation of both τ_{DW} prescriptions with respect to that obtained from the full quantum mechanical Kramers-Heisenberg cross section of neutral hydrogen. We treat as the ground truth for this an empirical fitting formula introduced by [Bach & Lee \(2015\)](#) that was designed as a simple but maximally accurate analytical approximation of the quantum mechanical cross section. For completeness, we also compare to the commonly used two-level model introduced by [Miralda-Escudé \(1998\)](#) (originally derived in [Peebles 1993](#)) where

$$I(x) = \int \frac{x^{\frac{9}{2}}}{4(1-x)^2}. \quad (2.19)$$

For each model, we determine the damping wing optical depth according to Eq. (2.13) and compute the relative error with respect to the optical

depth obtained by directly integrating Eq. (2.10) with the [Bach & Lee \(2015\)](#) fitting formula. We find that the Lorentzian assumption for the cross section overestimates damping wing absorption by maximally $\sim 2\%$ over most parts of the relevant wavelength range. Working with the incorrect prefactor instead leads to an error of $\sim 1 - 2\%$, increasing to maximally $\sim 5\%$. The error of the [Miralda-Escudé \(1998\)](#) model is roughly twice as large across the entire wavelength range. In Section 2.3.3 and Appendix 2.C, we were interested in the *scaling* of inference precision with respect to various hyper-parameters rather than exact quantitative constraints. Due to this fact, and given the small error of the Lorentzian model with the incorrect prefactor of $1/2$ compared to the still widely used [Miralda-Escudé \(1998\)](#) model (e.g. [Řurovčiková et al. 2020](#); [Lidz et al. 2021](#); [Curtis-Lake et al. 2023](#); [Heintz et al. 2024](#); [Umeda et al. 2024](#)), we decided not to repeat all computations to save resources by avoiding significant computational overhead. The quantitative breakdown of the total error budget in Section 2.3.4 was obtained with the correct prefactor of $1/8$ and therefore remains unaffected.

2.B Inference Tests

In Section 2.3.1 we introduced the concept of coverage tests to ensure that the inferred distribution of an ensemble of mock objects is a faithful representation of the true distribution. In Figure 2.12, we present the inference test results for the hyper-parameter configurations considered in Section 2.3.3 with the analytical IGM transmission model. For clarity, we show in each panel the expected coverage probability C_α of the marginal $(\langle x_{\text{HI}} \rangle, \log_{10} t_{\text{Q}}/\text{yr})$ -posteriors for three representative choices of a given hyper-parameter. We observe perfect coverage to within 1σ almost independently of the chosen parameter values, where we assigned error bars based on a Binomial distribution $B(N, C_\alpha)$. This owes to the fact that C_α is the fraction of mocks which are fall inside the α -th credibility region, and we estimate C_α by checking this for each of our $N = 100$ mocks. We find that only spectral resolution has a slight impact on coverage probability. Our Gaussian likelihood approximation (Eq. (2.1)) is thus remarkably accurate when IGM transmission is modelled according to the analytical formalism introduced in Section 2.2.2.2. Decreasing resolution changes the coverage probability from slightly overconfident to slightly underconfident for the lowest resolution spectra with FWHM = 1000 km/s, owing to the fact that less resolved spectral features (particularly in the Lyman- α forest) gaussianize the likelihood of the spectrum according to the central limit theorem. Accordingly, Eq. (2.1) becomes increasingly appropriate for lower resolution data. The extent as to which higher resolution models are overconfident is so small that the raw posteriors are almost identical to the reweighted ones that we obtain by applying the coverage correction procedure described




Figure 2.12: Marginal inference test results for the runs considered in Section 2.3.3. In each panel, we show the expected coverage probability C_α as a function of the credibility level α . Uncertainties correspond to the 16-th and 84-th percentile of a Binomial distribution $B(N, C_\alpha)$, where N is the number of objects in our mock sample. Models with perfect coverage follow the black dashed line; the region above it indicates underconfidence, the region below overconfidence. Whenever applicable, we marginalized over all latent parameters related to the shape of the continuum before fitting a Gaussian mixture model (GMM) to the resulting marginal $((x_{\text{HI}}), \log_{10} t_Q/\text{yr})$ -posterior density to compute the coverage probability. We vary latent dimension n_{latent} , red-side wavelength coverage λ_{cut} , signal-to-noise ratio S/N , and spectral resolution FWHM, in the four panels from left to right. All inference test results were obtained using the analytical IGM transmission model. Each panel contains three representative choices of the respective parameter. Our reference run with $n_{\text{latent}} = 5$, $\lambda_{\text{cut}} = 2000 \text{ \AA}$, $S/N = 10$ and $\text{FWHM} = 100 \text{ km/s}$ is depicted in black in each panel. The latent dimension plot also shows the theoretically optimal posteriors as obtained from the run with continuum-normalized spectra labelled as ‘astro-only’.

in Section 2.3.1. Nonetheless, for consistency, we choose to work with the reweighted posteriors for all our results shown in Sections 2.3.3 and 2.3.4 and Appendix 2.C.

Our full-simulation IGM transmission model, on the other hand, treats more complex physical processes which give rise to significant deviations from gaussianity that are not well-captured by our analytic likelihood approximation (cf. Figure 11 in Hennawi et al. 2025). As a result, our reference model is highly overconfident and we have to apply the coverage correction procedure to obtain meaningful posteriors. All inference test results for this IGM transmission model are shown in Figure 2.13. Coverage probability largely remains unchanged with respect to latent dimension and red-side wavelength coverage. The effect of improving coverage for lower resolution spectra becomes more prominent than for the analytic IGM transmission model and also becomes apparent for lower signal-to-noise spectra. For spectra with S/N as low as 2 or $\text{FWHM} \gtrsim 1000 \text{ km/s}$, the inferred posteriors are close to passing the inference test without coverage correction. We note that such low S/N and low resolution spectra naturally come at the price of decreasing inference precision. As shown in Appendix 2.C, however, all precision gains from high S/N or high resolution spectra are neutralized by their deteriorating coverage behavior. Rebinning spectra to coarse $\gtrsim 500 \text{ km/s}$ velocity pixel scales (corresponding to $\text{FWHM} \gtrsim 1000 \text{ km/s}$ in our analysis)




Figure 2.13: Like Figure 2.12, but for the full-simulation IGM transmission model.

therefore provides a simple way of reducing the computational complexity of the inference problem (or the observational cost of obtaining the required spectra) without losing inference precision. An improved prescription of the likelihood (Eq. (2.1), and especially Eq. (2.5)) would be required to make use of the full information encoded in the high-resolution spectra. A promising avenue for this is simulation-based inference, i.e. training a neural network-based estimator to learn the true non-Gaussian likelihood.

2.C Parameter Scan for the Full-Simulation Model

We confirm in this section that the inference precision trends identified in Section 2.3.3 also hold when reionization is modelled realistically based on cosmological simulations. To that end, we re-ran all simulations presented in Section 2.3.3 with the full-simulation IGM transmission model (Section 2.2.2.1) in place of the analytical one (Section 2.2.2.1). Figure 2.14 summarizes all dependencies of $\langle x_{\text{HI}} \rangle$ - and t_{Q} -inference precision on the four hyper-parameters n_{latent} , λ_{cut} , S/N and FWHM (analogous to Figure 2.6 for the analytical IGM transmission model). Apart from the additional stochasticity that dilutes the overall trends, we find the same qualitative behavior for latent dimension and red-side wavelength coverage. Rather than higher precision for less noisy spectra, however, we observe identical $\langle x_{\text{HI}} \rangle$ -precision values independently of the signal-to-noise ratio. The same is the case for lifetime precision of the fiducial subset of samples. Note that all precision values in Figure 2.6 correspond to the widths of the *reweighted* posterior distributions *after* accounting for any degree of overconfidence which becomes increasingly relevant with increasing S/N as seen in Appendix 2.B: all theoretical precision gains are thus neutralized by the coverage corrections applied to the posteriors whose widths are shown in Figure 2.14. Without these corrections, precision would improve monotonically with increasing S/N as intuitively expected. The picture is similar for spectral resolution: after




Figure 2.14: Like Figure 2.6, but for the full-simulation IGM transmission model. Trends are less pronounced due to the additional stochasticity, and the increasingly significant coverage corrections required at high S/N or high resolution undermine all precision gains or even lead to a decrease in precision after accounting for overconfidence.




Figure 2.15: Inference precision with respect to IGM neutral fraction $\langle x_{\text{HI}} \rangle$ (left) and quasar lifetime t_{Q} (right) as a function of lifetime for different choices of spectral resolution. The remaining hyper-parameters are always fixed to their reference values $n_{\text{latent}} = 5$, $\lambda_{\text{cut}} = 2000 \text{ \AA}$, $\text{S/N} = 10$. IGM transmission skewers are generated by the full-simulation model introduced in Section 2.2.2.2. Each curve shows median precision values with respect to an ensemble of 250 mock quasars at fixed neutral fraction $\langle x_{\text{HI}} \rangle = 0.5$, grouped into lifetime bins. All precision values are based on reweighted marginal posterior PDFs, assuring that each mock ensemble passes a marginal coverage test with respect to the two astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} . Neutral fraction precision decreases after accounting for overconfidence, independently of lifetime, whereas lifetime precision keeps improving for short-lived objects.

accounting for overconfidence, inference precision in $\langle x_{\text{HI}} \rangle$ hardly changes with FWHM, and for the fiducial set of objects even *decreases* with higher resolution. Lifetime precision, on the other hand, *does* improve to some extent when considering samples from the full prior range, but we find the opposite trend for the fiducial set of samples. All curves have in common that the scatter increases significantly towards higher resolution, i.e., a certain subset of spectra benefits notably from increasing resolution whereas the opposite is the case for the rest.

To investigate this behavior more carefully, we fix the value of the IGM neutral fraction to $\langle x_{\text{HI}} \rangle = 0.5$ and consider 250 new mock spectra, only varying the lifetime of the quasars. Figure 2.15 depicts inference precision as a function of lifetime for this new set of samples. We can see that $\langle x_{\text{HI}} \rangle$ -precision approximately changes equally with resolution at all lifetime values considered, whereas the bulk of improvements in t_{Q} -precision can be attributed to low-lifetime objects. The intermediate to long-lifetime range stays constant or even decreases with better resolution, in agreement with the fiducial curve in the bottom right panel of Figure 2.14. This shows that our Gaussian likelihood approximation currently limits us to low resolutions because Eq. (2.5) is not able to capture the non-Gaussian information encoded in higher-resolution spectra, an obstacle that could be overcome with simulation-based inference.



3 | A LOCAL, TOPOLOGY-INDEPENDENT PARAMETRIZATION OF QUASAR IGM DAMPING WINGS

Abstract

Lyman- α damping wings towards quasars provide a unique probe of the global volume-averaged neutral hydrogen (HI) fraction $\langle x_{\text{HI}} \rangle$. Cosmic variance in the intergalactic medium (IGM), however, is a major source of stochasticity since the *local* neutral environment around a quasar varies significantly even when $\langle x_{\text{HI}} \rangle$ is fixed. We show that the IGM damping wing carries additional information about this local ionization topology, unexploited by current analysis frameworks. We introduce two new physically motivated summary statistics encoding the local information about the HI distribution in the IGM *before* it is altered by the quasar's ionizing radiation, encompassing 1) the HI column density, weighted by a Lorentzian profile mimicking the frequency dependence of the Lyman- α cross section, and 2) the quasar's distance to the first neutral patch. This description, when combined with the quasar lifetime as a third parameter, reduces the IGM transmission scatter in the damping wing region of the spectrum to $\lesssim 1\%$ across physical parameter space. We introduce a simple procedure for generating synthetic HI sightlines around quasars and demonstrate that the resulting damping wing profiles are statistically indistinguishable from a realistic reionization topology. This opens the door for optimally extracting the damping wing's salient local information in a model-independent fashion. In the context of a specific reionization model, measurements of these local parameters can be translated into constraints on the global timing of reionization, but additionally provide information about the reionization topology, hitherto unused. A marginally modified version of our framework can also be employed for damping wings towards galaxies.

Work published in: **Timo Kist**, Joseph F. Hennawi and Frederick B. Davies, *Monthly Notices of the Royal Astronomical Society*, Volume 544, Issue 2, December 2025, Pages 2316–2339, doi.org/10.1093/mnras/staf1762. Reprinted here in its entirety.

3.1 Introduction

The epoch of reionization was a hallmark event in the evolution of our universe when the until then prevailing dark ages were ended by the light from the first galaxies and quasars. It is these very objects which we can also use as highly sensitive torchlights probing the presence of even the smallest amounts of neutral hydrogen in the intergalactic medium (IGM). This manifests in the form of extended Gunn–Peterson absorption troughs (Gunn & Peterson 1965) appearing at global volume-averaged IGM neutral fractions as low as $\langle x_{\text{HI}} \rangle \gtrsim 10^{-4}$ where the Lyman- α transition saturates, and transitions into a smooth damping wing imprint redward of the Lyman- α line when $\langle x_{\text{HI}} \rangle$ reaches order unity (Miralda-Escudé 1998). As such, every single high-redshift source offers a glimpse at the progress of reionization of the universe at its respective redshift. Existing (D’Odorico et al. 2023; Onorato et al. 2025) and future (Euclid Collaboration et al. 2019) statistical ensembles of their spectra provide us with a unique avenue to constrain the full timeline of the epoch of reionization, a notoriously hard problem due to the model-dependence of all probes going beyond the CMB electron scattering optical depth (Planck Collaboration et al. 2020) and Lyman-series dark pixel fractions (e.g. McGreer et al. 2015; Jin et al. 2023).

In the case of IGM damping wings, this undertaking is exacerbated by the fact that the IGM absorption imprint has to be disentangled from the intrinsic spectra of these sources, and, potentially, local damped Lyman- α (DLA) absorption systems in front of the source which can mimic the cosmological imprint from the IGM (Huberty et al. 2025). While both these tasks remain tractable when considering quasars as background sources, additional complications arise due to the patchy nature of the reionization process. A consequence of this is that at fixed IGM neutral fraction $\langle x_{\text{HI}} \rangle$, the column density of neutral material *along a given sightline* can vary drastically, depending on how many neutral patches it actually contains. The distribution of neutral versus ionized sightlines at given $\langle x_{\text{HI}} \rangle$ depends on the topology of reionization. The relation between damping wing strength and global IGM neutral fraction is thus inevitably *stochastic*. In essence, the Lyman- α damping wing only probes the *local* ionization state of the surrounding IGM as a single, statistical draw from the *global* ionization topology at the corresponding redshift.

To date, existing damping wing analysis frameworks (e.g. Greig et al. 2017b, 2024a; Davies et al. 2018a; Durovčíková et al. 2020, 2024; Hennawi et al. 2025; Umeda et al. 2024, 2025; Mason et al. 2026) have not made a clear methodological distinction between the two. The conventional approach has always consisted of directly inferring (among other parameters) the *global* IGM neutral fraction $\langle x_{\text{HI}} \rangle$, building upon analytical expressions for the damping wing optical depth (Miralda-Escudé 1998), or realistic IGM

transmission profiles based on cosmological simulations. As demonstrated recently in the context of *quasar* damping wings (Kist et al. 2025b), not less than half the total error budget on the inferred IGM neutral fraction $\langle x_{\text{HI}} \rangle$ arises due to density fluctuations and the stochasticity of reionization, comparable to the uncertainty due to the reconstruction of the unknown quasar continuum. An additional major source of uncertainty in the context of *galaxy* damping wings are local proximate Lyman- α absorption systems (Heintz et al. 2024, 2025) which can obfuscate the imprint arising from the IGM (Huberty et al. 2025), and would need to be marginalized out in order to obtain faithful constraints on the ionization state of the IGM (Mason et al. 2026).

In any case, eliminating the uncertainty sourced by the stochasticity of reionization would allow us to extract *all* the information that the damping wing imprint actually encodes—that is, physical line-of-sight information about the neutral hydrogen (HI) density field in front of the quasar which is informative not only about the timing but also the topology of reionization. Instead of directly processing this information into a constraint on the *global* IGM neutral fraction, leaving the latter unconstrained, we here propose a novel set of *physical* summary statistics which quantify the *local* HI content in front of a given object, and hence tightly parameterize the characteristic shape of the IGM damping wing.

Finding a better parameterization of the IGM damping wing has also been the primary motivation of Chen (2024) who proposed to re-center IGM transmission profiles based on the position in the spectrum where the transmission first approaches zero, and Keating et al. (2024b) who showed that the damping wing strength is tightly parameterized by the average HI density weighted with a Lorentzian profile, mimicking the Lorentzian decay of the Lyman- α cross section. Furthermore, Mason et al. (2026) for the first time inferred the distance between the source and the first neutral patch as a local measure in addition to the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$. Here we build upon the labels introduced by Keating et al. (2024b) and Mason et al. (2026), and establish two clearly defined physical summary statistics quantifying the local HI content around the source, readily applicable for astrophysical parameter inference. Leveraging these summary statistics, we show that the damping wing imprint does not only bear information about the *timing* of reionization in form of the evolution of the global IGM neutral fraction as a function of redshift, but even *topological* information about, e.g., the distribution of ionized patch sizes around the most massive halos *before quasars started shining*.

In addition to that, we show that our summary statistics are statistically insensitive to the underlying reionization model. Instead, all stochasticity in this parameterization is sourced by the well-understood distribution of density fluctuations in the IGM. This paves the way for *topology-independent*, local damping wing constraints which carry topological information, and can

be tied *subsequently* to a specific reionization topology, constraining, e.g., the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ *in the context of that given topology*.

We start by introducing our formalism and defining our new local summary statistics in Section 3.2. We quantify the scatter of realistic IGM transmission profiles in the damping wing region within this parameterization in Section 3.3 and proceed in Section 3.4 by introducing a simplistic toy prescription which we then use to demonstrate the topology-independence of our parameterization. We conclude in Section 3.5.

3.2 Towards a three-parameter model of quasar IGM damping wings

Several works have recently been pushing towards a parameterization better capturing the characteristic shape of IGM damping wings (Chen 2024; Keating et al. 2024b). Specifically, Keating et al. (2024b) noted that (in the absence of a strong ionizing source) the Lorentzian-weighted average HI number density calculated along simulated sightlines from the Sherwood-Relics simulation suite (Bolton et al. 2017; Puchwein et al. 2023) coincides remarkably well with predictions from the analytical Miralda-Escudé (1998) damping wing model. Here, we build upon that finding and introduce novel two-parameter summary statistics of the local *pre*-quasar HI field that significantly reduce the scatter of the damping wing optical depth τ_{DW} as compared to the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$.

3.2.1 Introducing a local HI column density label

We start by recalling that the damping wing optical depth τ_{DW} of a strong ionizing source at redshift z_{QSO} as a function of rest-frame wavelength λ_{rest} is obtained by integrating the product of the HI density field $n_{\text{HI}}^{\text{post}}$ and the Lyman- α cross section of neutral hydrogen σ_{α} (evaluated at wavelength $\lambda = \lambda_{\text{rest}}(1 + z_{\text{QSO}})/(1 + z)$) along the (proper) line-of-sight interval dR :

$$\tau_{\text{DW}}(\lambda_{\text{rest}}) = \int_0^{R(z_{\text{QSO}})} n_{\text{HI}}^{\text{post}}(R) \cdot \sigma_{\alpha} \left(\frac{1 + z_{\text{QSO}}}{1 + z(R)} \lambda_{\text{rest}} \right) dR, \quad (3.1)$$

where $R(z_{\text{QSO}})$ is the proper distance from the observer to the source, and $z(R)$ is the inverted distance-redshift relation, and we henceforth denote all associated comoving distances with a small r in place of a capital one.

Our goal is to construct summary statistics of the HI density field that best capture the characteristic shape of the IGM damping wing as given by $\tau_{\text{DW}}(\lambda_{\text{rest}})$. However, note that we are interested in maximizing the information content with respect to the *pre*-quasar HI density field which we denote as $n_{\text{HI}}^{\text{pre}}$ (in contrast to the post-quasar field $n_{\text{HI}}^{\text{post}}$ giving rise to

the actually observed IGM damping wing as per Eq. (3.1)). This is because $n_{\text{HI}}^{\text{pre}}$ directly relates to the *cosmological* reionization topology, removing the impact of the quasar as a local *astrophysical* source of ionizing radiation. In other words, we can decompose the pre-quasar HI field as

$$n_{\text{HI}}^{\text{pre}} = \langle n_{\text{H}} \rangle(z_{\text{QSO}}) \cdot x_{\text{HI}} \cdot \Delta, \quad (3.2)$$

where $\langle n_{\text{H}} \rangle(z_{\text{QSO}})$ is the cosmic mean hydrogen density at the redshift of interest, Δ is the dimensionless matter overdensity field, and x_{HI} is the HI fraction field which informs us about the *global* volume-averaged IGM neutral fraction $\langle x_{\text{HI}} \rangle$.¹

We now aim to condense the information contained in the *field* $n_{\text{HI}}^{\text{pre}}$ into a low-dimensional set of physically motivated summary statistics. The canonical choice for such a summary statistic is the HI column density $N_{\text{HI}} = \int n_{\text{HI}}^{\text{pre}}(R) dR$ along the line of sight. Noting the obvious similarity between this definition and Eq. (3.1), and following the weighting proposed in Keating et al. (2024b), we introduce a weighting function

$$w \equiv (R/R_{\text{T}} + 1)^{-2} \sim (v - v_{\text{T}})^{-2}, \quad (3.3)$$

mimicking the Lorentzian decay of the Lyman- α cross section σ_{α} in Eq. (3.1), evaluated at a fixed reference distance $R = R_{\text{T}}$.² The equivalence with a Lorentzian profile becomes more apparent if we introduce the velocity grid $v = -H(z_{\text{QSO}}) R$, centered at the Lyman- α line (such that $\lambda_{\text{rest}} = \lambda_{\alpha} (1 + \frac{v}{c})$ for the Lyman- α wavelength $\lambda_{\alpha} \simeq 1215.67 \text{ \AA}$), and express R_{T} as a velocity offset $v_{\text{T}} \equiv +H(z_{\text{QSO}}) R_{\text{T}}$.³

Given a field X , we now define the line-of-sight average $\langle\langle \cdot \rangle\rangle_{\text{Lor}}$ of this field with respect to the Lorentzian weighting kernel defined in Eq. (3.3) as

$$\langle\langle X \rangle\rangle_{\text{Lor}} \equiv \frac{1}{\mathcal{N}(u_{\text{min}}, u_{\text{max}})} \int_{u_{\text{min}}}^{u_{\text{max}}} \frac{X(u)}{(u+1)^2} du, \quad (3.4)$$

where we introduced the dimensionless integration variable $u \equiv R/R_{\text{T}}$ with integration limits u_{min} and u_{max} , corresponding to the distances R_{min} and R_{max} whose values we will fix below. Further, $\mathcal{N}(u_{\text{min}}, u_{\text{max}})$ is a normalization factor ensuring $\langle\langle \mathbf{1} \rangle\rangle_{\text{Lor}} = 1$ for the identity field $\mathbf{1}$. This factor can be explicitly determined, yielding

$$\mathcal{N}(u_{\text{min}}, u_{\text{max}}) = \frac{u_{\text{max}} - u_{\text{min}}}{(u_{\text{max}} + 1)(u_{\text{min}} + 1)}. \quad (3.5)$$

¹For brevity, we write $\langle x_{\text{HI}} \rangle \equiv \langle x_{\text{HI}} \rangle(z_{\text{QSO}})$ throughout this work.

²Note that away from the Lyman- α resonance, a Lorentzian profile is an excellent approximation of the full quantum-mechanical Lyman- α cross section σ_{α} .

³Note that our velocity grid follows the standard convention of $v > 0$ redward of the Lyman- α line, and $v < 0$ blueward where $R > 0$. However, to keep all distances positive, we define $R_{\text{T}} > 0$ even though it corresponds to a *positive* velocity offset $v_{\text{T}} > 0$ *redward* of Lyman- α .

Comparing Eq. (3.4) to Eq. (3.1) governing the damping wing optical depth, we now proceed by defining the *Lorentzian-weighted* HI column density as

$$N_{\text{HI}}^{\text{DW}} \equiv \mathcal{N} \left(\frac{R_{\text{min}}}{R_{\text{T}}}, \frac{R_{\text{max}}}{R_{\text{T}}} \right) \cdot R_{\text{T}} \cdot \langle n_{\text{H}} \rangle (z_{\text{QSO}}) \cdot \langle \langle x_{\text{HI}} \cdot \Delta \rangle \rangle_{\text{Lor}}. \quad (3.6)$$

This definition involves three hyperparameters, i.e., the reference distance R_{T} as well as the integration limits R_{min} and R_{max} . An informed choice of these parameters ensures that this summary statistic reflects the characteristic shape of the IGM damping wing not only for the pre-quasar topology (which we would like to maximize the information content about), but also the post-quasar one (which, due to the impact of the quasar ionizing radiation, inevitably carries a somewhat decreased amount of information about the pre-quasar IGM, but on the other hand gives rise to the actually observed damping wing imprint as per Eq. (3.1)).

Firstly, with regards to the *pre*-quasar topology, we show in Appendix 3.A that in the limit where $R_{\text{min}} \rightarrow 0$ and $R_{\text{max}} \rightarrow R(z_{\text{QSO}})$, and where the Lyman- α cross section can be approximated as perfectly Lorentzian, the pre-quasar optical depth $\tau_{\text{DW}}^{\text{pre}}$ (as defined in Eq. (3.11)) evaluated at the velocity offset v_{T} becomes, up to a prefactor, approximately proportional to $N_{\text{HI}}^{\text{DW}}$:

$$\tau_{\text{DW}}^{\text{pre}}(v = v_{\text{T}}) \simeq \frac{e^2}{m_e c^2} f_{\alpha} \gamma_{\alpha} \lambda_{\alpha} (c/v_{\text{T}} - 1)^2 \times N_{\text{HI}}^{\text{DW}}, \quad (3.7)$$

where $f_{\alpha} \simeq 0.416$ is the Lyman- α oscillator strength, and $\gamma_{\alpha} \equiv \Gamma_{\alpha} \lambda_{\alpha} / (4\pi c)$ with the Lyman- α decay constant $\Gamma_{\alpha} = 6.265 \times 10^8 \text{ s}^{-1}$. This motivates our choice of $N_{\text{HI}}^{\text{DW}}$ as the primary summary statistic that minimizes the scatter of the red-side damping wing transmission among different (pre-quasar) sightlines, and hence encodes the characteristic shape of the IGM damping wing in a single number.

Practically, however, we are concerned with profiles originating from the *post*-quasar topology rather than the pre-quasar one. As we will show in the consequent sections, $r_{\text{min}} = 4 \text{ cMpc}$, $r_{\text{max}} = r_{\text{min}} + 100 \text{ cMpc}$, and $r_{\text{T}} = 18 \text{ cMpc}$ (corresponding to $v_{\text{T}} \simeq 2000 \text{ km/s}$ at $z_{\text{QSO}} = 7.54$)⁴ are adequate choices to ensure that Eq. (3.7) remains valid even for the post-quasar optical depth

⁴As our analysis is restricted to models at this fixed redshift, we henceforth interchangeably refer to reference distances r_{T} and reference velocity offsets v_{T} , noting that only a fixed choice of r_{T} would account for redshift evolution in the density field.




Figure 3.1: Lyman- α transmission and physical fields for an example Nyx sightline combined with a semi-numerical x_{HI} skewer extracted from a $\langle x_{\text{HI}} \rangle = 0.65$ topology (left panels/blue lines; see Section 3.3.1) and an analytical x_{HI} skewer generated according to our toy prescription (right panels/red lines; see Section 3.4.1), before (dark) and after (light colors) a quasar has been shining for $t_{\text{Q}} = 10^6$ yr. The panels show from top to bottom the Lyman- α transmission field, the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ as a function of the upper integration limit r_{max} , the HI density field n_{HI} (also shown in units of $x_{\text{HI}} \cdot \Delta$), including its Lorentzian-weighted pre-quasar version (black dash-dotted line), the HI fraction x_{HI} , the overdensity field Δ , and the temperature field T . Our fiducial integration range for the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ is highlighted by the grey-shaded region, and the horizontal dashed line marks the value of $\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2} = 20.1$ corresponding to this integration range. The vertical dotted lines mark the location $r_{\text{patch}} = 8$ cMpc of the first neutral patch in the pre-quasar sightline.

τ_{DW} . Adopting these parameter choices and converting all quantities to comoving units, we can evaluate Eq. (3.6) as

$$N_{\text{HI}}^{\text{DW}} = 5.1 \times 10^{20} \text{ cm}^{-2} \times \left(\frac{\mathcal{N}\left(\frac{r_{\text{min}}}{r_{\text{T}}}, \frac{r_{\text{max}}}{r_{\text{T}}}\right)}{0.67} \right) \left(\frac{r_{\text{T}}}{18 \text{ cMpc}} \right) \quad (3.8)$$

$$\times \left(\frac{1 + z_{\text{QSO}}}{1 + 7.54} \right)^2 \left(\frac{\langle\langle x_{\text{HI}} \cdot \Delta \rangle\rangle_{\text{Lor}}}{1} \right).$$

Here we used $\langle n_{\text{H}} \rangle(z_{\text{QSO}}) = \langle n_{\text{H}} \rangle_0 (1 + z_{\text{QSO}})^3$, where $\langle n_{\text{H}} \rangle_0 = 3H_0^2(1 - Y)\Omega_{\text{b}}/(8\pi Gm_{\text{p}}) \simeq 1.9 \times 10^{-7} \text{ atoms/cm}^3$ with H_0 and Ω_{b} chosen according to a [Planck Collaboration et al. \(2020\)](#) cosmology, and a big bang nucleosynthesis (BBN) hydrogen fraction of $1 - Y = 0.76$.

We show in the top and left panels of Figure 3.1 how we arrive at this summary statistic for a simulated example sightline of a $t_{\text{Q}} = 10^6 \text{ yr}$ quasar in a globally 65% neutral universe. A more detailed description of the underlying simulations will be provided in Section 3.3.1. In each panel, the post-quasar version of the corresponding field is depicted in light blue, the pre-quasar one in dark blue. The top panel shows the IGM transmission profiles,⁵ and the second panel shows the Lorentzian-weighted HI column density (Eq. (3.8)) as a function of the upper integration limit r_{max} , where the $N_{\text{HI}}^{\text{DW}}$ value resulting from our fiducial choice of $r_{\text{max}} = r_{\text{min}} + 100 \text{ cMpc}$ is marked by the horizontal dotted line. The third panel depicts the HI density field n_{HI} (also given in units of $x_{\text{HI}} \cdot \Delta$), and additionally the Lorentzian-weighted version of the pre-quasar field according to Eq. (3.3). The remaining panels show the neutral hydrogen fraction field x_{HI} , the overdensity field Δ , and the temperature field T , respectively. The grey-shaded area in each panel depicts our fiducial integration range of $r_{\text{min}} = 4 \text{ cMpc}$ and $r_{\text{max}} = r_{\text{min}} + 100 \text{ cMpc}$. We can see how distant neutral patches are downweighted by the Lorentzian weighting function (third panel), and therefore contribute less to $N_{\text{HI}}^{\text{DW}}$ (second panel) and hence do not notably affect the IGM damping wing (top panel) either.

3.2.2 Choice of the $N_{\text{HI}}^{\text{DW}}$ integration limits

As Eq. (3.7) shows, our $N_{\text{HI}}^{\text{DW}}$ statistic is an optimal prescription of the damping wing optical depth $\tau_{\text{DW}}^{\text{pre}}(v = v_{\text{T}})$ of the pre-quasar density field $n_{\text{HI}}^{\text{pre}}$ evaluated at the velocity offset v_{T} , provided that we integrate along the *entire* line of sight and assume a perfectly Lorentzian cross section. If we were to consider damping wings around galaxies, it would therefore be immediately clear that we would have to set our lower integration limit to $r_{\text{min}} = 0$.

⁵Note that all profiles which are colored red will be discussed later in Section 3.4.

However, this simple picture changes when considering damping wings towards quasars instead. Even though we are still ultimately interested in the pre-quasar HI field, in practice we are inevitably probing the optical depth arising from the *post*-quasar field $n_{\text{HI}}^{\text{post}}$. Differences arise interior to the quasar ionization front where the pre-quasar neutral material is ionized away and hence the information about this part of the pre-quasar density field is lost; though we can account for this by also constraining the lifetime t_{Q} of the quasar. By tuning our integration limits R_{min} and R_{max} in Eq. (3.6), we can make sure that we still maximize the information content about $n_{\text{HI}}^{\text{pre}}$, and nevertheless obtain a statistic which describes the observed damping wing profile with a minimum amount of scatter. By nature of the required approximations, this scatter will be non-zero, but an informed choice of the integration limits ensures that it stays minimal and certainly remains smaller than the amount of scatter which $\langle x_{\text{HI}} \rangle$ entails.

To do so, we firstly note that the optical depth integral in Eq. (3.1) effectively only starts at the radius of the ionization front R_{ion} that the quasar has carved out, since $n_{\text{HI}}^{\text{post}}$ is drastically reduced within the ionized bubble around the quasar (i.e., at $R < R_{\text{ion}}$). On the other hand, at $R > R_{\text{ion}}$, where the quasar ionization front has not arrived yet, we can safely approximate $n_{\text{HI}}^{\text{post}} \simeq n_{\text{HI}}^{\text{pre}}$. In a simple theoretical picture where every photon emitted by the quasar ionizes exactly one hydrogen atom in a (Strömgren) sphere around the quasar, the radius of the ionization front R_{ion} relates to the lifetime of the quasar t_{Q} and the volume averaged HI density $\langle n_{\text{HI}}^{\text{pre}} \rangle \equiv \langle n_{\text{HI}}^{\text{pre}} \rangle (z_{\text{QSO}})$ as

$$R_{\text{ion}} = \left(\frac{3Q t_{\text{Q}}}{4\pi \langle n_{\text{HI}}^{\text{pre}} \rangle} \right)^{1/3} \quad (3.9)$$

(Cen & Haiman 2000), where Q is the quasar's emission rate of ionizing photons. Due to the approximate nature of this equation, defining $R_{\text{min}} \equiv R_{\text{ion}}$ is not guaranteed to give us the correct integration limit for an individual sightline. However, as we strive to define a sightline-independent integration limit, Eq. (3.9) constitutes a useful starting point for determining a universal value.

Since we aim to define our $N_{\text{HI}}^{\text{DW}}$ label agnostically to the reionization state of the universe, we first have to remove the dependence of the cosmic mean HI density $\langle n_{\text{HI}}^{\text{pre}} \rangle = \langle n_{\text{H}} \rangle_0 (1 + z_{\text{QSO}})^3 \langle x_{\text{HI}} \rangle$ on the IGM neutral fraction $\langle x_{\text{HI}} \rangle$. We can achieve this by assuming $\langle x_{\text{HI}} \rangle = 1$ which gives rise to the shortest R_{ion} possible. By doing so, we make our lower integration limit shorter, and hence conservatively also account for regions closer to the source where neutral material is only still present for sightlines from highly neutral topologies.

Secondly, as $N_{\text{HI}}^{\text{DW}}$ is supposed to be a summary statistic for the *pre*-quasar HI density $n_{\text{HI}}^{\text{pre}}$, the lifetime dependence of Eq. (3.9) is undesirable too, and we can eliminate it by fixing t_{Q} in Eq. (3.9) to a representative value

of $t_Q = 10^{4.5}$ yr. This specific value constitutes an empirical compromise between 1) quasars with longer lifetimes of $t_Q \gtrsim 10^{4.5}$ yr and hence larger R_{ion} , where a fixed lower integration limit of $R_{\text{ion}}(t_Q = 10^{4.5} \text{ yr})$ includes pre-quasar IGM neutral material which has already been ionized away and therefore does *not* contribute to the damping wing, and 2) short-lived quasars with $t_Q \lesssim 10^{4.5}$ yr where this integration limit excludes foreground neutral material that actually *will* impact the damping wing. As we will demonstrate in Section 3.3.3, our choice of $t_Q = 10^{4.5}$ yr provides a reasonably accurate approximation for the full lifetime range expected for observed high-redshift quasars (see e.g. Morey et al. 2021; Khrykin et al. 2021). With these choices, we arrive at the comoving value of

$$r_{\text{ion}} = 4.0 \text{ cMpc} \left(\frac{Q}{10^{57.14} \text{ s}^{-1}} \right)^{1/3} \left(\frac{t_Q}{10^{4.5} \text{ yr}} \right)^{1/3} \left(\frac{\langle x_{\text{HI}} \rangle}{1.0} \right)^{-1/3} \quad (3.10)$$

and fix our lower integration limit accordingly to $r_{\text{min}} = 4.0 \text{ cMpc}$.⁶

The optical depth integral (Eq. (3.1)) formally runs all the way down to redshift zero; however, due to the Lorentzian decline in the outer parts of the HI cross section, we can safely stop the integration in Eq. (3.4) already significantly earlier than this, while still retaining a tight approximation of the full damping wing optical depth as per Eq. (3.7). Specifically, we only have to integrate until the Lorentzian line-of-sight average of neutral fraction and overdensity field $\langle x_{\text{HI}} \cdot \Delta \rangle_{\text{Lor}}$ has converged. We practically find that this is the case after integrating over a range of $\sim 100 \text{ cMpc}$ (as exemplarily seen in the second panel of Figure 3.1), where our weighting kernel (Eq. (3.3)) has decreased to $\sim 3\%$ of its value at r_{min} . Thus, we set the upper integration limit to $r_{\text{max}} = r_{\text{min}} + 100 \text{ cMpc}$.

3.2.3 The damping wing as a one-parameter family

We argued in Section 3.2.1 that we identified a close-to optimal summary statistic $N_{\text{HI}}^{\text{DW}}$ for the damping wing optical depth evaluated at a fixed reference distance r_T (or, equivalently, velocity offset v_T). We now demonstrate that by restricting ourselves to this fixed reference point, we do not give up a significant amount of additional information encoded in the imprint. In other words, we show that the Lyman- α transmission values at different velocity offsets are so correlated with each other that the damping wing profile essentially constitutes a one-parameter family, i.e., the entire spectral shape is determined by a single parameter, namely the transmission $\tau_{\text{DW}}(v = v_T)$ at our reference velocity offset.

To demonstrate this, we depict in Figure 3.2 the correlation between a large number of simulated IGM transmission profiles evaluated at our fixed

⁶Note that due to cosmological expansion manifesting in the density field, we arrived at the same comoving integration limit independently of redshift.




Figure 3.2: IGM transmission values at the fixed reference velocity offset $v_T = 2000$ km/s versus transmission values at a number of velocities offsets $v = v_T + \Delta v$ away from this reference offset. Solid curves and shaded regions depict the median as well as the 68- and 95-percentile scatter of these transmission values, respectively, among $852 \times 21 \times 51$ realistic IGM transmission profiles simulated according to the description in Section 3.3.1. For reference, we are also showing the trivial one-to-one relation for $v = v_T$ as a dashed black line.

reference velocity offset $v_T = 2000$ km/s versus their value at a number of other velocity offsets $v = v_T + \Delta v$. Specifically, we considered $852 \times 21 \times 51$ transmission profiles consisting of 852 distinct sightlines at 21 global IGM neutral fraction values $\langle x_{\text{HI}} \rangle$ covering a linear range between 0 and 1, and 51 lifetime values t_Q covering a logarithmic range between 10^3 and 10^8 yr. The technical details of the simulations will be outlined in Section 3.3.1.

For each velocity offset, we show the median transmission values of all profiles as a solid curve, as well as the 68- and 95-percentile scatter as shaded regions. Our first observation is that in the largest parts of the damping wing region (at all velocity offsets $v \gtrsim 1500$ km/s), the width of the 68- (95-) percentile region never exceeds $\sim 1 - 2\%$ ($\sim 2 - 3\%$), even at the lowest IGM transmission values which are significantly impacted by damping wing absorption, decreasing down to $< 1\%$ ($< 2\%$) at higher transmission values. A number of effects conspire such that we observe a relatively similar degree of scatter for most of these curves: due to continuity reasons, the scatter naturally decreases for velocity offsets closer to our reference offset at $v_T = 2000$ km/s marked by the dashed black line.

However, the scatter does not change symmetrically around v_T , i.e., we find a more significant increase blueward of v_T compared to redward. As demonstrated by Keating et al. (2024b), this is due to the impact of infalling

gas as well as residual neutral gas in ionized patches which increasingly affects the shape of the transmission profiles in the vicinity of the Lyman- α line. These effects are also the reason why we cannot expect the damping wing to remain a one-parameter family at velocity offsets even closer to Lyman- α line center. Specifically, we see that the 68-percentile scatter has already increased to $\sim 6\%$ at $v = 1000$ km/s, and grows to $O(10\%)$ at Lyman- α line center. The scatter also grows in the other direction (i.e, further redward of v_T), albeit significantly more modestly than blueward. In this direction, the scatter is limited by the fact that transmission values are generally higher, approaching their upper bound of unity, and hence allowing for less scatter overall. As a result, the spread of the curves shows a mild peak around $v = 4000$ km/s, and then decreases again for even larger velocity offsets as the transmission values approach unity.

Overall, this demonstrates that the IGM transmission value at the reference velocity of $v_T = 2000$ km/s informs us about the *entire* shape of the IGM damping wing at *any* velocity offset $v \gtrsim 1500$ km/s with remarkable precision, and hence a summary statistic that maximizes the information at this reference location also extracts the bulk of the information encoded in the entire damping wing imprint.

3.2.4 A second label: the distance to the first neutral patch

As discussed in the previous sections, the characteristic shape of the IGM damping wing is encoded to lowest order in the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ as defined in Eq. (3.6). By defining a second summary statistic, we can 1) further minimize the scatter among Lyman- α transmission skewers at a given set of parameter values, and 2) gain additional information about the reionization topology.

An obvious choice to achieve the first objective would be the location r_{ion} of the ionization front of the quasar, as this determines precisely where the optical depth integral in Eq. (3.1) effectively starts (c.f. the discussion in Section 3.2.2). However, at a fixed ionizing photon emission rate Q , and at a fixed column density $N_{\text{HI}}^{\text{DW}}$, the radius of the ionization front r_{ion} is, up to a minor contribution due to density fluctuations, largely degenerate with the quasar lifetime t_Q , as Eq. (3.9) suggests. The lifetime, however, is of direct physical interest and we therefore opt to rather keep t_Q as a post-quasar label.

By instead turning to the *pre*-quasar equivalent of r_{ion} , i.e., the distance r_{patch} between the source and the first neutral patch along the line of sight *in the pre-quasar topology*, we can simultaneously address both aforementioned objectives: while it is obvious that most long-lived quasars have carved out an ionized bubble extending significantly beyond the pre-quasar neutral patch location, the same is not true for objects whose ionization front r_{ion}

does not or only marginally exceed r_{patch} . The latter can be the case primarily for young quasars, but also older ones where the first neutral patch is located at a large distance from the quasar itself. In any case, due to the fixed choice of the integration limit r_{min} for $N_{\text{HI}}^{\text{DW}}$, it is clear that r_{patch} encapsulates additional information about the pre-quasar ionization topology that is not fully captured by $N_{\text{HI}}^{\text{DW}}$. Besides this, r_{patch} is of direct physical interest as a key statistic of the topology of reionization, yielding complementary information to the values estimated from Lyman- α emission from galaxies (Mason & Gronke 2020; Hayes & Scarlata 2023; Umeda et al. 2024; Witstok et al. 2024; Torralba-Torregrosa et al. 2024; Lu et al. 2024a,b; Nikolić et al. 2025), and has already been introduced by Mason et al. (2026) as an additional summary statistic for IGM damping wings around galaxies.

Overall, this provides us with a three-parameter model, parameterizing quasar IGM damping wings as a function of $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}, t_{\text{Q}})$. The two labels $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ are summary statistics of the *pre*-quasar topology, and could hence also be employed in the context of IGM damping wings towards galaxies, provided that the integration range of $N_{\text{HI}}^{\text{DW}}$ is adjusted accordingly. The third parameter t_{Q} is specific to quasars, encapsulating their impact as strong ionizing sources on the IGM transmission field.

3.2.5 Relation to previous definitions

Parameterizations aiming to better capture the characteristic shape of the IGM damping wing have recently been proposed in Chen (2024), Keating et al. (2024b), and Mason et al. (2026). Our $N_{\text{HI}}^{\text{DW}}$ statistic introduced in Section 3.2.1 particularly builds upon the study by Keating et al. (2024b), whereas the distance r_{patch} to the first neutral patch has already been employed by Mason et al. (2026) in the context of galaxy damping wings. In this section we investigate the similarities and differences among these definitions found in the literature and our own ones.

We chose to adopt as our first summary statistic the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ along the line of sight from the source, rather than the average velocity-weighted number density $\langle\langle n_{\text{HI}} \rangle\rangle_{\text{Lor}}$ as introduced in Keating et al. (2024b).⁷ We deem this definition more appropriate for the following reasons: as an inherent line-of-sight quantity, $N_{\text{HI}}^{\text{DW}}$ emphasizes that the damping wing signature really only encodes one-dimensional information about the *local HI content in front of the quasar*. Only by combining the information from several such sightlines do we gain statistical information about the global average of the HI density and hence the global reionization topology. Secondly, the use of a line-of-sight quantity more clearly emphasizes the *physical scale* (i.e., ~ 100 cMpc) impacting the IGM damping wing. In addition, the analogy to the column density of a local proximate damped

⁷Note that the symbol $\langle\langle n_{\text{HI}} \rangle\rangle_{\text{Lor}}$ is chosen in accordance with our own notation.

Lyman- α absorber (DLA) helps building intuition for the physical range of $N_{\text{HI}}^{\text{DW}}$, as well as the degree to which IGM damping wings and proximate DLAs can be disentangled.

We emphasize, however, that in practice, the notions of HI column densities and line-of-sight averaged HI number densities are equivalent, and we could easily convert between the two by realizing that $\langle n_{\text{HI}} \rangle_{\text{Lor}} = \langle n_{\text{H}} \rangle(z_{\text{QSO}}) \cdot \langle x_{\text{HI}} \cdot \Delta \rangle_{\text{Lor}}$ as follows from Eqs. (3.2) and (3.4). According to Eq. (3.6), this implies that the two statistics only differ by the geometrical factor $\mathcal{N}\left(\frac{r_{\text{min}}}{r_{\text{T}}}, \frac{r_{\text{max}}}{r_{\text{T}}}\right)$ as well as a factor of r_{T} . However, this equality demands that the average $\langle n_{\text{HI}} \rangle_{\text{Lor}}$ be computed over the same integration range as $N_{\text{HI}}^{\text{DW}}$. The adequate lower integration limit for galaxy damping wings as considered in Keating et al. (2024b) would be $r_{\text{min}} = 0$, in contrast to our choice of $r_{\text{min}} = 4 \text{ cMpc}$ due to the presence of a strong ionizing source which almost certainly ionizes away all neutral gas within the first few cMpc from the source. Further, with regards to quasar damping wings, Keating et al. (2024b) chose to operate on the post-quasar field, whereas we here opt to use pre-quasar labels with an adjusted integration limit r_{min} . This way, we directly constrain the pre-quasar field—which is of higher physical interest—at the price of a small (and, as we will show, negligible) additional amount of scatter.

In addition, despite the seeming agreement of our weighting function (Eq. (3.3)) with that adopted by Keating et al. (2024b), our summary statistic differs in that we do *not* realign the skewers in the way proposed in that work and in Chen (2024). We always keep the velocity grid centered at the Lyman- α line, whereas Chen (2024) and Keating et al. (2024b) re-center their velocity grid based on the individual realignment point of each transmission profile, defined as the point where the profile approaches zero for the first time. Note further that the meaning of the realignment point differs among the two works: Chen (2024) does not perform radiative transfer but attempts to model quasar lifetime effects through a simplistic model for the quasar ionization front. In return, their realignment point thus corresponds to what is conventionally measured as the *proximity zone size* in the spectrum of a quasar, often used as a summary statistic to gain information about its lifetime. Keating et al. (2024b), on the other hand, models lifetime effects through ionizing radiative transfer and performs the realignment at various but *fixed* lifetime values (as well as a set of pre-quasar skewers). Hence, the pre-quasar realignment point encodes similar information as our r_{patch} statistic, while the post-quasar one largely coincides with the proximity zone size *at a fixed lifetime value*, hence largely encoding the scatter due to density fluctuations in the IGM.

As a result, our three-parameter model makes both types of re-centering obsolete: while Chen (2024)'s realignment point is well characterized by the quasar lifetime t_{Q} which is a separate parameter of our model, Keating et al.

(2024b)’s pre-quasar realignment point is largely represented by our r_{patch} statistic. We find a negligible impact from realigning post-quasar skewers at a fixed lifetime which, while removing some excess scatter due to density fluctuations in the IGM, would be largely degenerate with t_Q . The virtue of t_Q and r_{patch} is that both of these summaries are of direct physical interest, and we thus prefer to use them over abstract realignment points.

Furthermore, the distance r_{patch} between the source and the first neutral patch has been introduced in Mason et al. (2026) as an additional summary statistic in the context of galaxy IGM damping wings. Here we adopt the same *pre*-quasar statistic, and show that it remains a meaningful summary even for *post*-quasar IGM transmission profiles. Also note in this context that one of the earliest damping wing studies (Schroeder et al. 2013) already aimed to constrain the distance between the quasar and the first neutral patch, albeit in the *post*-quasar topology rather than the *pre*-quasar one as targeted in this work. Due to the absence of radiative transfer in their modeling, similar conclusions apply as drawn above regarding Chen (2024)’s realignment point, i.e., at fixed ionizing photon emission rate Q , the post-quasar distance to the first neutral patch encodes very similar information as the quasar lifetime t_Q which is already explicitly included as a parameter in our modeling framework.

3.3 Quantifying the IGM transmission scatter

With the definition of our new set of summary statistics at hand, we now describe how we simulate realistic IGM transmission profiles based on cosmological simulations combined with semi-numerical reionization topologies and 1d radiative transfer. We then use these simulated profiles to demonstrate the advantages of our labels over the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ in parameterizing the characteristic shape of the IGM damping wing by comparing the IGM transmission scatter at fixed parameter values within the global and the local parameterization.

3.3.1 Simulating IGM transmission fields

We generate IGM transmission profiles \mathbf{t} based on the hybrid approach introduced in Davies et al. (2018a), combining hydrodynamical skewers extracted from the Nyx cosmological simulations and x_{HI} skewers from semi-numerical reionization topologies with 1d ionizing radiative transfer.

We start by extracting 6×200 density, velocity, and temperature skewers originating at the 200 most massive halos ($M_{\text{halo}} \geq 2 \times 10^{11} M_{\odot}$) from the $z = 7.0$ snapshot of the Nyx simulations (Almgren et al. 2013; Lukić et al. 2015), containing 4096^3 baryon and another 4096^3 dark matter particles in a $100 \text{ cMpc}/h$ box. Our fiducial reionization topologies are simulated using an

adapted version of **21cmFast** (Mesinger et al. 2011; Davies & Furlanetto 2022) in a larger 400 cMpc box on a 2048^3 initial and a 512^3 output grid, providing us with sufficient statistics to probe the distribution of neutral/ionized regions around the 500 rarest, most massive halos ($M_{\text{halo}} \geq 3 \times 10^{11} M_{\odot}$). We generate such topologies at 21 different global IGM neutral fractions $\langle x_{\text{HI}} \rangle$, where we achieve $\langle x_{\text{HI}} \rangle = 0.05, 0.1, \dots, 0.95$ by tuning the ionizing efficiency ζ , and trivially add a completely ionized ($\langle x_{\text{HI}} \rangle = 0.0$) and a completely neutral ($\langle x_{\text{HI}} \rangle = 1.0$) model.

We subsequently combine each Nyx sightline with a random x_{HI} skewer pointing in a random direction originating at one of the aforementioned 500 most massive halos of the **21cmFast** box. Here we adopt the Nyx temperature field for all initially ionized regions, and assume an initially cold IGM for all neutral regions, fixed to $T = 2000$ K. To model the small-scale impact of the ionizing quasar radiation, we then perform one-dimensional radiative transfer along these sightlines (Davies et al. 2016), assuming a 'light bulb' lightcurve where the quasar has been shining at constant luminosity throughout its lifetime t_{Q} until the redshift z_{QSO} of interest. We here consider a model with $z_{\text{QSO}} = 7.54$ and an ionizing photon emission rate of $Q = 10^{57.14} \text{ s}^{-1}$, resembling the quasar ULAS J1342+0928. We compute the time evolution up to a maximum lifetime of $t_{\text{Q}} = 10^8$ yr, and store 51 intermediate outputs on a logarithmically spaced grid between $t_{\text{Q}} = 10^3$ and 10^8 yr. We then convolve the resulting output fields from the radiative transfer code with a Voigt profile to obtain the final IGM transmission profiles. This provides us with a set of 1200 IGM transmission profiles on a 21×51 grid of $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ values. For the bulk of this work, we focus on the three representative lifetime values of $t_{\text{Q}} = 10^4, 10^6$ and 10^8 yr, representing a young, an intermediate, and a long-lived quasar.

We exclude sightlines that exhibit strong proximate optically thick absorption line systems. This is necessary due to the lack of a subgrid prescription for star formation in the Nyx simulations which implies that the properties and abundance of such systems are not guaranteed to match those found in the real universe. Observationally, such objects can be excluded based on the presence of associated metal absorption lines in the spectrum (Davies et al. 2025).

To identify such sightlines in our simulations, we consider all 1200 skewers from the completely ionized pre-quasar topology, divide each of them into chunks of size 0.1 pMpc, and compute the (unweighted) HI column density within each of these chunks. We then exclude all sightlines where at least one chunk (located within the first 5000 km/s from the source) exceeds a column density threshold of 10^{19} cm^{-2} . After applying this criterion, 852 skewers remain, and we end up with a set of $852 \times 21 \times 51$ transmission profiles covering the full $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ -parameter space.

3.3.2 Global parameterization

We start by investigating the IGM transmission scatter within the global parameterization, depicted in Figure 3.3 at select $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ parameter values. Here, the solid lines mark the median transmission value of the 852 profiles, and the shaded regions denote the 68-percentile scatter (enclosed by the 16-th and the 84-th percentile). Besides the clear emergence of the damping wing imprint in more neutral environments (i.e., at higher global IGM neutral fractions $\langle x_{\text{HI}} \rangle$ and shorter quasar lifetimes t_{Q}), we can see that the bulk of the scatter emerges in the proximity zone due to density fluctuations in the quasar ionized IGM, as well as variation in the location of the quasar ionization front. However, even in the smooth damping wing region redward of the Lyman- α line we can identify a significant amount of scatter. This scatter is a direct consequence of the stochastic nature of reionization. In other words, at a given global value of the IGM neutral fraction $\langle x_{\text{HI}} \rangle$, the *local* HI column density giving rise to the observed damping wing imprint can differ significantly from the global average.

To get a more quantitative idea of this scatter, we focus on the distribution of IGM transmission values at a specific location of the profile, at $v = v_{\text{T}}$.⁸ Figure 3.5 shows a heatmap of the 68-percentile widths $\sigma(t(v = v_{\text{T}}))$ of these distributions as a function of the astrophysical parameters $\langle x_{\text{HI}} \rangle$ and t_{Q} , where $\sigma(t(v = v_{\text{T}}))$ is defined as half the distance between 84-th and 16-th percentile.

For short quasar lifetimes $t_{\text{Q}} \lesssim 10^4$ yr, we find similarly wide distributions as Keating et al. (2024b) do in profiles from the Sherwood-Relics simulations in the complete absence of a quasar. As expected, the scatter decreases with increasing quasar lifetime due to the larger quasar ionized regions which suppress the IGM damping wing such that the transmission is very close to unity redward of line center. Similarly, the scatter is smaller than 1% in a fully ionized universe, even at the shortest quasar lifetimes, and grows with increasing IGM neutral fraction $\langle x_{\text{HI}} \rangle$, peaking around $\langle x_{\text{HI}} \rangle \sim 0.7$, and reaching values of up to $\sim 7\%$ at the shortest lifetimes of $10^3 - 10^4$ yr. Notably, the scatter is not exclusively sourced by the distribution of ionized patches of the underlying reionization topology but also by the mere presence of density fluctuations which are in fact the *only* source of scatter in an entirely neutral universe where $\langle x_{\text{HI}} \rangle = 1$ and thus no topology variations are present.

In realistic settings where the intrinsic continuum of the quasar is unknown and has to be reconstructed too, the amount of scatter among different IGM transmission profiles seen in Figure 3.5 is comparable to the $\sim 5\%$ uncertainties in the reconstructed quasar continuum (c.f. Figure 6 in Hennawi

⁸As demonstrated in Section 3.2.3, due to the strong correlation across the entire IGM damping wing imprint, it suffices to evaluate τ_{DW} at a single pixel to get a sufficient picture of the scatter across the entire IGM transmission profile.




Figure 3.3: Median and 68-percentile scatter of the simulated IGM transmission profiles based on our semi-numerical reionization topology in the global $(\langle x_{\text{HI}} \rangle, t_0)$ parameterization. The profiles are shown on a representative grid of $(\langle x_{\text{HI}} \rangle, t_0)$ parameter values.




Figure 3.4: Median and 68-percentile scatter of the simulated IGM transmission profiles based on our semi-numerical reionization topology in the local ($N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}, f_Q$) parameterization. The profiles are shown on a representative grid of ($N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}, f_Q$) parameter values for the most typical combinations of $N_{\text{HI}}^{\text{DW}}$ and r_{patch} across the full parameter space.




Figure 3.5: Sightline-to-sightline scatter of the IGM damping wing as a function of astrophysical parameter space in the global $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ -parameterization. The scatter is quantified as half the 68-percentile width $\sigma(t(v = v_{\text{T}}))$ of the distribution of IGM transmission values at $v = v_{\text{T}}$ among all sightlines at a given location in $(\langle x_{\text{HI}} \rangle, t_{\text{Q}})$ -parameter space.

et al. 2025). As a result, this excess stochasticity in the IGM transmission model significantly adds to the total error budget on the reconstructed $\langle x_{\text{HI}} \rangle$ and t_{Q} values (Kist et al. 2025b), whereas we could isolate it from the task of reconstructing the transmission profile if we were to work with a parameterization that more tightly captures the characteristic shape of IGM damping wings.

3.3.3 Local parameterization

We now demonstrate that the scatter among individual IGM transmission profiles decreases significantly when labelling the profiles with the local summary statistics introduced in Section 3.2 in place of the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$.

To that end, we compute the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ as defined in Eq. (3.6) and the distance r_{patch} to the first neutral patch for each (pre-quasar) sightline and subsequently aggregate all post-quasar profiles according to these new labels. We fix the velocity offset and the integration range for $N_{\text{HI}}^{\text{DW}}$ as discussed in Sections 3.2.1 and 3.2.2. To compute r_{patch} , we smooth the x_{HI} field with a box-car filter of size 0.5 cMpc and determine r_{patch} as the distance from the source where the smoothed x_{HI} field first becomes more than 50% neutral. This smoothing procedure makes sure that our r_{patch} measurement identifies the first *spatially extended*




Figure 3.6: Distribution of our local summary statistics $N_{\text{HI}}^{\text{DW}}$ and r_{patch} among all simulated IGM transmission profiles based on semi-numerical x_{HI} skewers originating from massive halos (cyan; see Section 3.3.1) and random locations (dark blue), and analytical x_{HI} skewers generated according to our toy prescription (red; see Section 3.4.1). Dark shaded regions denote the 68% contours, lighter ones the 95% ones. The distributions are smoothed via a kernel density estimation (KDE) with a Gaussian kernel of size (0.05, 1.0) in $(\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}, r_{\text{patch}}/\text{cMpc})$ parameter space. The sharp edges of the toy model contours result from the fact that these profiles are simulated on a fixed grid of parameter values.

neutral patch without being impacted by individual fluctuations in the x_{HI} field which do not have a notable impact on the damping wing shape. We find that we are not overly sensitive to the exact choices of these parameters.

The cyan contours in Figure 3.6 depict the 68 and 95% regions of the distribution of $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ values among all 852×21 sightlines.⁹ As our $\langle x_{\text{HI}} \rangle$ parameter grid uniformly covers the range between zero and one, this is the $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ distribution that follows under the assumption that all neutral fraction values $0 \leq \langle x_{\text{HI}} \rangle \leq 1$ are equally likely, i.e., in the absence of any knowledge about the reionization state of the IGM. We observe a clear degeneracy between $N_{\text{HI}}^{\text{DW}}$ and r_{patch} , where a higher column density $N_{\text{HI}}^{\text{DW}}$ implies a smaller r_{patch} value and vice versa. This degeneracy arises due to the fact that a closer distance to the first neutral patch clearly allows for more neutral material distributed within the range over which $N_{\text{HI}}^{\text{DW}}$ is computed. It is also apparent, however, that this is not a perfect degeneracy,

⁹Recall we are considering 852 different sightlines at 21 distinct global IGM neutral fraction values $\langle x_{\text{HI}} \rangle$. Our 51 different quasar lifetime values do not play a role as $N_{\text{HI}}^{\text{DW}}$ and r_{patch} are defined on the pre-quasar topology.

and therefore the r_{patch} statistic contains additional information not encoded in $N_{\text{HI}}^{\text{DW}}$.

To quantify this information, we define (unequally spaced) bins for the two labels and aggregate the transmission profiles accordingly. The binning is informed by the number of available skewers and the degree to which parameter variations impact the resulting IGM transmission profiles. Specifically, we choose a finer binning towards higher column densities and shorter distances to the first neutral patch which cause more variation in the profiles due to increasingly strong damping wing absorption.

The scatter among the aggregated IGM transmission profiles in select parameter bins representative of the variations across the full parameter space is shown in Figure 3.4. While a one-to-one comparison between individual bins in Figures 3.3 and 3.4 is not possible due to the redefined bin labels, we did compute the effective line-of-sight averaged neutral fractions $\langle\langle x_{\text{HI}} \rangle\rangle_{\text{los}}$ of the profiles in each parameter bin, and we depict in Figure 3.4 selected bins where the median $\langle\langle x_{\text{HI}} \rangle\rangle_{\text{los}}$ matches the cosmic average $\langle x_{\text{HI}} \rangle$ values shown in Figure 3.3.¹⁰ With this in mind, we can certainly note that the scatter redward of the Lyman- α line in Figure 3.4 is consistently lower than in Figure 3.3, implying that we identified a tighter parameterization of the characteristic shape of the IGM damping wing. Note that as expected, the scatter in the Lyman- α forest region is not impacted by the new parameterization.

For a more quantitative picture across the full range of parameter space, we summarize the 68-percentile widths $\sigma(t(v = v_{\text{T}}))$ of the distributions of IGM transmission values at $v = v_{\text{T}}$ in all $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter bins in Figure 3.7.¹¹ In each given row, we are showing from left to right the scatter at the three representative lifetimes of $t_{\text{Q}} = 10^4, 10^6$ and 10^8 yr, respectively. For reference, the three upper panels depict the scatter in the global $\langle x_{\text{HI}} \rangle$ parameterization,¹² whereas the three middle panels depict the scatter of the same profiles aggregated according to the $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameterization. The three lower panels will be discussed in Section 3.4. The transparency of the colorbar is proportional to the logarithm of the number of profiles in each bin and thus represents the distribution of $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ values which was already explicitly depicted in Figure 3.6.

¹⁰Recall that we denote cosmic averages with single angle brackets $\langle \cdot \rangle$, and their $z = 0$ values as $\langle \cdot \rangle_0$ throughout. We always denote line-of-sight averages with double angle brackets, where $\langle\langle \cdot \rangle\rangle_{\text{los}}$ is the conventional line-of-sight average, and $\langle\langle \cdot \rangle\rangle_{\text{Lor}}$ as defined in Eq. (3.4) additionally involves a Lorentzian weighting function.

¹¹Recall that we evaluate the IGM transmission scatter at $v = v_{\text{T}}$, i.e., exactly at the reference velocity offset v_{T} used to determine $N_{\text{HI}}^{\text{DW}}$. Due to the strong pixel-by-pixel correlation across the entire IGM damping wing signature, this does *not* mean we are underestimating the scatter, since this scatter can maximally contribute another $\sim 0.5 - 1\%$ to $\sigma(t(v = v_{\text{T}}))$ as demonstrated in Section 3.2.3.

¹²Note that these are exactly the $t_{\text{Q}} = 10^4, 10^6$ and 10^8 yr slices of the heatmap in Figure 3.5.




Figure 3.7: Sightline-to-sightline scatter of the IGM damping wing as a function of astrophysical parameter space for different parameterizations and topologies. The scatter is quantified as half the 68-percentile width $\sigma(t(v = v_T))$ of the distribution of IGM transmission values at $v = v_T$ among all sightlines at a given location in parameter space, and at quasar lifetimes of $t_Q = 10^4$ yr (left column), 10^6 yr (middle column) and 10^8 yr (right column). *Top row:* Profiles based on semi-numerical x_{HI} skewers (see Section 3.3.1) in the global (x_{HI}) parameterization. *Middle row:* Same profiles as above, but in the local $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameterization. *Bottom row:* Profiles based on the analytical toy x_{HI} skewers (see Section 3.4.1) in the local $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameterization. Bins which are occupied by a lower number of sightlines are more transparent, and white if not occupied at all.

Remarkably, the transmission in the vast majority of bins in the local parameterization seen in the three middle panels of Figure 3.7 varies by $\lesssim 1 - 2\%$, whereas it is consistently higher in the global parameterization depicted in the upper panels, reaching up to $\sim 7\%$ at $t_Q = 10^4$ yr. This shows that our local summary statistics capture the characteristic shape of the IGM damping wing much more tightly than the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$.

At all lifetimes, however, we do observe a mild increase in scatter towards the highest column densities where the damping wing imprint gets more pronounced. This increased scatter is a result of the fact that we chose to define the two summaries on the *pre*-quasar x_{HI} field. By fixing the integration limit for $N_{\text{HI}}^{\text{DW}}$ independently of quasar lifetime, we did not fully pay tribute to the variable radius of the ionization front r_{ion} where the integration effectively starts. Hence, by fixing $r_{\text{min}} = 4 \text{ cMpc}$, we are theoretically including too much or too little neutral material in the $N_{\text{HI}}^{\text{DW}}$ statistic, depending on the size of the ionized bubble the quasar has carved out, which in turn depends on the quasar lifetime, as discussed in Section 3.2.2. Nevertheless, the fact that the resulting scatter does not exceed $\sim 2\%$ even at the highest HI column densities shows that our approximation was a valid one, still leading to a tight parameterization of the *post*-quasar IGM transmission field nearly independently of quasar lifetime.

3.4 Towards a topology independent local IGM damping wing parameterization

We introduced in the previous section a new three-parameter model for quasar IGM damping wings. We classified simulated IGM transmission profiles from a given reionization topology according to these labels and showed that the sightline-to-sightline scatter drastically decreases in this new parameterization as compared to the conventional (e.g. Greig et al. 2017b; Davies et al. 2018a; Ďurovčiková et al. 2024; Hennawi et al. 2025) one- or two-parameter model based on the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ (and the quasar lifetime t_Q). Specifically, by replacing $\langle x_{\text{HI}} \rangle$ with the two *local* summary statistics $N_{\text{HI}}^{\text{DW}}$ and r_{patch} , we were able to reduce the scatter in the damping wing region of the profile down to $\lesssim 1\%$ across the entire range of parameter space. In this section, we go beyond these results and demonstrate that our conclusions hold true largely independently of the underlying reionization topology.

This is because once our local summaries are set, the detailed distribution of neutral hydrogen along the line of sight, set by the reionization topology, does not matter anymore for the shape of the profile in the damping wing region of the spectrum. Instead, all relevant information about the damping

wing shape is already contained in $N_{\text{HI}}^{\text{DW}}$ as Eq. (3.7) shows. The only remaining sources of scatter are the required approximations for this equality to hold, i.e., the finite integration limits of $N_{\text{HI}}^{\text{DW}}$, any differences between pre- and post-quasar density field that have not been mitigated by the choice of the lower integration limit and the inclusion of r_{patch} as additional label, as well as any deviations of the damping wing shape from being a one-parameter family determined at a fixed reference distance (or velocity offset). However, all these contributions remain negligible as we showed in the previous section, neither do our conclusions change if the profiles originate from a different reionization topology as we will demonstrate here.

The scatter in the proximity zone region remains unaffected either because here the dominant source of sightline-to-sightline scatter is density fluctuations, while the x_{HI} field—only varying discontinuously between ionized and neutral patches with values close to zero and one—effectively acts as a window function for the fluctuations in the density field (see Eq. (3.2) and the example realization of these fields in Fig. 3.1) which are already accounted for by considering a statistical sample of density skewers (independently of reionization topology). As a result, all effects of the ionization topology on the observed IGM transmission profile for a given sightline are already well captured in the two summary statistics $N_{\text{HI}}^{\text{DW}}$ and r_{patch} .

In Bayesian language, the mere effect of assuming a specific reionization topology is therefore imposing a corresponding prior on the parameters $N_{\text{HI}}^{\text{DW}}$ and r_{patch} according to their distribution within that topology. For future damping wing analyses, this will allow one to perform a clear methodological distinction between 1) the topology-insensitive task of inferring the local damping wing statistics ($N_{\text{HI}}^{\text{DW}}, r_{\text{patch}}, t_{\text{Q}}$), and 2) tying these constraints to a particular reionization topology, constraining not only the timing of reionization via the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$, but also its topology which remains unconstrained by $\langle x_{\text{HI}} \rangle$ alone. It is immediately clear that measurements of $N_{\text{HI}}^{\text{DW}}$ (as a statistical moment of the x_{HI} field) and r_{patch} for a statistical ensemble of objects carry additional topological information that we can extract and leverage to compare reionization models whose topologies are different.

Here we demonstrate the topology independence of our labels¹³ in a simplified fashion by introducing a highly simplistic toy model to generate x_{HI} skewers analytically rather than extracting them from a realistic semi-numerical reionization topology. Specifically, we show that within our local parameter framework both models are statistically equivalent in terms of the resulting median IGM transmission profiles and scatter.

The specifics of the toy prescription which we introduce in the following section are set keeping in mind the practical task of performing inference with

¹³Again, the topology information is contained in the *distribution* of these summaries while the summaries themselves can be defined and measured topology-independently.

respect to these new local summary statistics. Specifically, we introduced in Hennawi et al. (2025) a fully Bayesian framework for inferring astrophysical model parameters from high-redshift quasar spectra. Although originally introduced to directly infer (among other parameters) the *global* IGM neutral fraction, the framework is equally applicable to other parameters governing the IGM transmission field such as our local summary statistics. All which is needed is an estimate of the IGM transmission likelihood given the model parameters. Under the assumption of Gaussianity, this practically means we need to be able to estimate means and covariances of the IGM transmission field smoothly as a function of the model parameters $N_{\text{HI}}^{\text{DW}}$ and r_{patch} .

As $N_{\text{HI}}^{\text{DW}}$ and r_{patch} are derived quantities for IGM transmission profiles from the semi-numerical reionization simulation, the number of available sightlines can vary significantly from parameter bin to parameter bin. Especially for $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ values that are rare in a given reionization topology, our covariance estimates can easily become so noisy that the likelihood function is not guaranteed to vary smoothly as a function of $N_{\text{HI}}^{\text{DW}}$ and r_{patch} . While one could certainly beat down this noise with a massively increased number of simulations, building an analytical toy prescription where the continuity of the IGM transmission field with respect to $N_{\text{HI}}^{\text{DW}}$ and r_{patch} is intrinsically built in drastically reduces the number of required sightlines and facilitates the inference task for future applications, even though establishing the full inference pipeline would exceed the scope of this work.

As far as this work is concerned, the toy prescription provides us with a simple way of testing the sensitivity of our summary statistics to the underlying reionization topology by comparing the scatter of the IGM transmission profiles within these two very different topologies—the latter topology not even physically motivated. The fact that this scatter agrees remarkably well between the two demonstrates the topology-insensitivity of our labels and legitimizes the future use of the analytical toy prescription for the purpose of astrophysical parameter inference.

3.4.1 An analytic toy prescription for generating neutral fraction profiles

To generate IGM transmission profiles according to the toy prescription which we will introduce now, we follow the same hybrid approach outlined in Section 3.3.1, only replacing the semi-numerical x_{HI} skewers with analytically generated ones. This implies in particular that we use the same hydrodynamical skewers as described in that section and the same radiative transfer code to model the impact of the quasar ionizing radiation.

3.4.1.1 Description of the prescription

We start by realizing that for a given density skewer, there exists a minimal and a maximal HI column density value $N_{\text{HI},\text{min}}^{\text{DW}}$ and $N_{\text{HI},\text{max}}^{\text{DW}}$ as computed according to Eq. (3.6). These values can only be achieved if the IGM is completely ionized or completely neutral (within the integration range), respectively. To determine $N_{\text{HI},\text{min}}^{\text{DW}}$, we first have to compute the equilibrium ionization state for a fully ionized sightline based on the input UV background and temperature field, providing us with an x_{HI} field with values on the order of $O(10^{-3} - 10^{-4})$ which we can then use to compute $N_{\text{HI},\text{min}}^{\text{DW}}$. To determine $N_{\text{HI},\text{max}}^{\text{DW}}$, we subsequently set $x_{\text{HI}} = 1$ at all distances greater than r_{patch} . Note that $N_{\text{HI},\text{min}}^{\text{DW}}$ and $N_{\text{HI},\text{max}}^{\text{DW}}$ differ from sightline to sightline, based on the specific realization of density fluctuations along the line of sight.

To achieve any column density value $\hat{N}_{\text{HI}}^{\text{DW}}$ between $N_{\text{HI},\text{min}}^{\text{DW}}$ and $\leq N_{\text{HI},\text{max}}^{\text{DW}}$, we start from the completely ionized sightline and add in neutral material until the desired $\hat{N}_{\text{HI}}^{\text{DW}}$ value is reached. We do so by subsequently adding in neutral patches i of a minimum size Δr_{min} and a maximum size Δr_{max} originating at locations r_i and growing them in the direction towards the observer.¹⁴ Note that this procedure is by no means meant to resemble any physical processes occurring during reionization. It is a purely numerical procedure to ensure that small changes in the model parameters $N_{\text{HI}}^{\text{DW}}$ and r_{patch} result in small changes in the x_{HI} field and hence the IGM transmission profiles, allowing us to obtain smoothly varying means and covariances for future inference applications. To that end, we always follow the same sequence of locations $(r_i)_{i \in \mathbb{N}}$ when inserting new patches into a given sightline. In particular, we obtain the x_{HI} field for a given sightline and a desired parameter tuple $(\hat{N}_{\text{HI}}^{\text{DW}}, \hat{r}_{\text{patch}})$ according to the following procedure:

1. Insert a neutral patch of size Δr_{min} at the location $r_0 = \hat{r}_{\text{patch}}$, extending out to $\hat{r}_{\text{patch}} + \Delta r_{\text{min}}$.
2. Grow the current (i -th) patch (originating at r_i) in the direction of the observer by subsequently setting x_{HI} pixel values adjacent to the outer edge of the patch to unity. In case a given pixel is already neutral because it overlaps with one of the previous neutral patches, this step has no effect.
3. If the i -th patch has reached a size of Δr_{max} , insert a new patch of size Δr_{min} originating at the next location r_{i+1} .
4. If $r_{i+1} < \hat{r}_{\text{patch}}$, do not add the $(i+1)$ -th patch and instead skip to the $(i+2)$ -th one because \hat{r}_{patch} is by definition the distance to the

¹⁴We refer to r_i as the 'origin' of this patch and stress that it can only 'grow' in the direction away from the source.

neutral patch *closest* to the source. If necessary, repeat this step until a location is found where a patch can be added.

5. While $N_{\text{HI}}^{\text{DW}} < \hat{N}_{\text{HI}}^{\text{DW}}$, repeat steps (ii) - (iv).

In this way, we are able to generate x_{HI} skewers on any desired $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ grid, varying smoothly with both parameters. Note, however, that the $N_{\text{HI}}^{\text{DW}}$ ranges are physically restricted by $N_{\text{HI},\text{min}}^{\text{DW}}$ and $N_{\text{HI},\text{max}}^{\text{DW}}$ for each individual skewer, so it is not necessarily the case that there exists a realization of every sightline at each point in $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space.

Moreover, for any given sightline, the maximal $N_{\text{HI}}^{\text{DW}}$ value changes as a function of r_{patch} . This is because a skewer with a certain neutral patch distance r_{patch} is by definition guaranteed to be ionized at $r < r_{\text{patch}}$, so the highest possible column density $N_{\text{HI},\text{max}}^{\text{DW}}$ for a given sightline decreases with increasing r_{patch} . A subspace of the full $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space is therefore *physically inaccessible*, and this is determined *exclusively* by the distribution of density fluctuations in the IGM, while entirely independent of any further assumptions about the reionization topology. Note in particular that our synthetic sightlines do not even necessarily respect the correlation between the locations of ionized patches and overdensity dictated by any inside-out reionization scenario.

The prescription introduced in this section comes with two hyperparameters Δr_{min} and Δr_{max} determining the minimum and the maximum size of neutral patches that we insert into the sightlines. We choose these as $\Delta r_{\text{min}} = 0.5 \text{ cMpc}$ and $\Delta r_{\text{max}} = 5.0 \text{ cMpc}$, loosely informed by the sizes we empirically find for skewers in our semi-numerical reionization topology (compare also [Xu et al. 2017](#)), and noting that we are not sensitive to these exact choices.

3.4.1.2 Generating IGM transmission profiles

Using this prescription, we now generate a grid of IGM transmission profiles based on the same 852 DLA-excluded density sightlines as in Section 3.3.2. We construct the grid in accordance with the 21 and 18 parameter bins used in Section 3.3.3 to classify the realistic profiles with respect to $N_{\text{HI}}^{\text{DW}}$ and r_{patch} , respectively. Accordingly, we perform ionizing radiative transfer for the same lifetime values as in Section 3.3.1, again focusing on the outputs after $t_{\text{Q}} = 10^4$, 10^6 and 10^8 yr .

An example sightline can be found on the right-hand side of Figure 3.1. The resulting IGM transmission profile shown in the top panel is overlaid in red with the one we obtained using a semi-numerical x_{HI} skewer (blue; left-hand panels) as considered in Section 3.3.1. For optimal comparability to the realistic reionization topology, we used the same underlying density skewer and matched the parameter values, i.e., both the semi-numerical sightline (blue) and the analytical toy-sightline (red) have $\log_{10} N_{\text{HI}}^{\text{DW}} / \text{cm}^{-2} =$

20.1, $r_{\text{patch}} = 8 \text{ cMpc}$, and $t_{\text{Q}} = 10^6 \text{ yr}$, where the semi-numerical sightline originates from a topology with a global IGM neutral fraction of $\langle x_{\text{HI}} \rangle = 0.65$. The pre-quasar versions of all physical fields are shown in dark, and those of the IGM transmission field in lighter colors.

Despite the locally very different x_{HI} fields, the resulting IGM transmission profiles look almost identical in the smooth damping wing region redward of the Lyman- α line, and even their proximity zones agree relatively well due to the matching density profiles giving rise to similar Lyman- α absorption signatures. The actual location of the quasar ionization front, however, differs notably among the two examples due to the larger amounts of neutral gas near the quasar for the semi-numerical sightline. Its proximity zone therefore does not extend quite as far out as that of the toy sightline. However, as we will show statistically in the next section, such differences in the proximity zone region due to the distribution of neutral material are entirely outweighed by the differences caused by density fluctuations alone. The damping wing region remains completely unaffected, as our $N_{\text{HI}}^{\text{DW}}$ label by construction accounts for neutral material along the *entire* line of sight with the correct weighting. This suggests that our parameterization well captures the characteristic shape of the IGM damping wing, even in a case where the sightlines originate from highly different reionization topologies. We proceed in the following section with a statistical confirmation of these sightline-based observations.

3.4.2 Comparison to a realistic reionization topology

We now investigate the scatter in the IGM transmission profiles based on our newly introduced toy prescription, and compare it to that found in Section 3.3.3 in the context of a more realistic semi-numerical reionization topology. We depict the median and the 68-percentile scatter of the IGM transmission profiles at representative locations in $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space in Figures 3.8, 3.9 and 3.10. The number of sightlines based on which the percentiles are computed are denoted in each panel. The three different figures show the same sets of profiles at quasar lifetimes of $t_{\text{Q}} = 10^4$, 10^6 and 10^8 yr , respectively. In addition to the profiles based on our toy topology (red), we also show the ones from the realistic reionization topology (blue; see also Figure 3.4), aggregated into the same parameter bins as introduced in Section 3.3.3.¹⁵ For clarity, we omit profiles in noise-dominated parameter bins with no more than 30 sightlines.

¹⁵Note that all toy profiles are simulated at a unique set of parameter values for each given bin, whereas the semi-numerical ones are aggregated from the entire parameter range covered by this bin.




Figure 3.8: Median and 68-percentile scatter of simulated IGM transmission profiles in the local $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameterization at a quasar lifetime of $t_Q = 10^4$ yr, comparing profiles based on semi-numerical x_{HI} skewers (blue; see Section 3.3.1), and analytical ones generated according to our toy prescription (red; see Section 3.4.1). The profiles are shown on a representative grid of $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter values. For clarity, bins which are occupied by no more than 30 sightlines are omitted.

3.4.2.1 Physically accessible regions of parameter space

The first thing of note is that our toy profiles extends over a significantly larger region in $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space as compared to the realistic ones, as we can immediately see by comparing the red and cyan contours depicted in Figure 3.6. Both these models show a non-trivial dependence on the parameter values due to the fact that not all regions of parameter space are accessible given the distribution of density fluctuations in the IGM. However, the toy prescription explores significantly larger regions in parameter space than the more realistic semi-numerical model. This shows that there are parameter configurations which are theoretically possible




Figure 3.9: Like Figure 3.8, but for a quasar lifetime of $t_Q = 10^6$ yr.

for a given sightline but which do not necessarily appear in our simulated reionization topology.

This mainly concerns combinations of small HI column densities $N_{\text{HI}}^{\text{DW}}$ with short neutral patch distances r_{patch} . Such configurations are increasingly hard to achieve the smaller r_{patch} , as neutral material closer to the source gets upweighted according to the Lorentzian weighting function. A very nearby neutral patch therefore has to be extremely small in extent as to not exceed a given HI column density $N_{\text{HI}}^{\text{DW}}$ —a configuration that hardly occurs in the realistic reionization topology. The relatively tight correlation between $N_{\text{HI}}^{\text{DW}}$ and r_{patch} seen for the cyan contours in Figure 3.6 is therefore specific to this topology, and not necessarily a generic feature of other ones.

We demonstrate this fact by also depicting the parameter distribution that would follow if quasars were not to be placed in the most massive halos but in random environments (dark blue contours). For each of our 21

3.4. TOWARDS A TOPOLOGY INDEPENDENT LOCAL IGM DAMPING WING PARAMETERIZATION

116




Figure 3.10: Like Figure 3.8, but for a quasar lifetime of $t_Q = 10^8$ yr.

global topologies, we generate 1200 such sightlines¹⁶ (770 of which remain after excluding sightlines with strong proximate absorbers according to the same criterion as adopted previously), and we analogously compute the values of our local summary statistics for each sightline. The resulting distribution shows clear differences to the one obtained under the assumption that quasars only populate massive halos (cyan). In particular, at short neutral patch distance r_{patch} , the distribution shows an extended tail towards low column densities $N_{\text{HI}}^{\text{DW}}$. Such configurations are physically admitted (see red contours obtained with our toy prescription) but unlikely if quasars were to reside in the most massive halos which, assuming an inside-out reionization scenario, constitute the first regions to be ionized. As a result, largely ionized sightlines (i.e., sightlines where $N_{\text{HI}}^{\text{DW}}$ is low) preferentially go

¹⁶In practice, we take the same 1200 sightlines as before and instead place the source in the center rather than at the origin of the sightline. This suffices for a first qualitative impression as an overdensity ~ 70 cMpc away from the quasar does not significantly impact the damping wing signature anymore.

together with a large distance r_{patch} to the first neutral patch, whereas such a correlation is not necessarily required when quasars are placed in random locations. While this already qualitatively demonstrates the added value of our parameterization for distinguishing between different reionization scenarios and topologies, we leave a more thorough exploration of these application of our parameterization to future work.

In any case, as seen in the three lower panels of Figure 3.7, the IGM transmission scatter as quantified by the 68-percentile widths $\sigma(t(v = v_T))$ remains similarly low even in those regions of parameter space which remain unexplored by the realistic topology with quasars living in the most massive halos. Our parameterization therefore tightly captures the shape of the IGM damping wing in *all* regions of parameter space which are physically allowed.

The unpopulated regions of parameter space at high column densities with large neutral patch distances, on the other hand, are a hard physical constraint, independent of reionization topology. As argued in Section 3.4.1, the maximum column density $N_{\text{HI}}^{\text{DW}}$ for a given sightline at a given neutral patch distance r_{patch} can only be achieved if the skewer is fully neutral starting at $r \geq r_{\text{patch}}$ (and, by definition, fully ionized at $r < r_{\text{patch}}$). This maximum value therefore decreases with increasing r_{patch} for a given sightline. Parameter combinations $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ where we cannot find *any* sightline that reaches the corresponding column density are therefore *physically excluded* since even the strongest density fluctuations do not suffice anymore to achieve this HI column density $N_{\text{HI}}^{\text{DW}}$. Again, this statement is solely based on the underlying distribution of density fluctuations and independent of the topology of reionization.

3.4.2.2 Agreement between realistic reionization model and toy prescription

Moving on to the comparison of the IGM transmission profiles between the realistic and the toy topology, we observe a remarkable agreement throughout nearly all bins that are populated by both models, across the entire lifetime range, as can be seen in the panels of Figures 3.8, 3.9 and 3.10 where profiles from both models are available. The only bins where differences become visible are the bottom right ones in each figure where $\log_{10} N_{\text{HI}}^{\text{DW}} / \text{cm}^{-2} = 20.6$ and $r_{\text{patch}} = 0.3 \text{ cMpc}$.¹⁷ While both models still perfectly agree at the lowest lifetime of $t_Q = 10^4 \text{ yr}$ (see Figure 3.8), the profiles increasingly differ at longer lifetimes of $t_Q = 10^6$ and 10^8 yr as seen in Figures 3.9 and 3.10.

¹⁷Due to the lack of a star formation prescription, the Nyx simulations do not correctly capture the circumgalactic medium of massive halos, and therefore we exclude the first 0.3 cMpc of all sightlines. Hence $r_{\text{patch}} = 0.3 \text{ cMpc}$ is the shortest neutral patch distance possible in our simulations. Note that this does not imply our simulated transmission profiles are unphysical, as material this close to the quasar is certainly guaranteed to be ionized away, even for the shortest quasar lifetimes.

More specifically, we see that while the profiles still agree reasonably well in the damping wing region of the spectrum, the proximity zone transmission of the realistic profiles is systematically higher than that of the toy ones. The reason for this is related to the fact that in the realistic semi-numerical reionization topology, r_{patch} values of 0.3 cMpc are in practice only found for completely neutral sightlines, since a small r_{patch} value always implies a large column density $N_{\text{HI}}^{\text{DW}}$ as can be seen in Figure 3.6. On the other hand, our toy model also allows for many sightlines with a first neutral patch at $r_{\text{patch}} = 0.3$ cMpc and a number of ionized patches farther away from the source, leading to lower column densities, and hence the contours extending out to much smaller $N_{\text{HI}}^{\text{DW}}$. These configurations are extremely rare in realistic reionization topologies as the regions around the most massive halos where our quasars are placed are among the first ones to be reionized, whereas no such constraint applies to our toy prescription where we only place the first neutral patch at the desired location and subsequently add more neutral patches at *any* location behind it until a desired column density is reached.

Having available this broader range of x_{HI} skewers at $r_{\text{patch}} = 0.3$ cMpc means that we can achieve a given HI column density $N_{\text{HI}}^{\text{DW}}$ for many additional density sightlines. These additional sightlines are generically not as underdense as the ones that achieve a given $N_{\text{HI}}^{\text{DW}}$ value in combination with a completely neutral sightline. This is because *both* x_{HI} and Δ contribute to $N_{\text{HI}}^{\text{DW}}$ (see Eq. (3.6)), and if x_{HI} is equal to unity, Δ must be comparatively low. Vice versa, the additional, globally more ionized, skewers from our toy topology can be combined with sightlines less underdense than those required for completely neutral skewers to achieve the same column density $N_{\text{HI}}^{\text{DW}}$. The more typical densities present in these additional sightlines in return lead to less pronounced transmission spikes in the proximity zone region which become increasingly apparent the larger the proximity zone sizes, i.e., the longer the quasar lifetimes, as the progression of the bottom right panels of Figures 3.8, 3.9 and 3.10 shows. The damping wing region on the other hand remains largely unaffected; if any, we observe a slight increase in damping wing absorption strength for the realistic sightlines as they are globally the most neutral ones. However, our $N_{\text{HI}}^{\text{DW}}$ statistic *by construction* already accounts for the competing effects between x_{HI} and Δ with the correct Lorentzian weighting, and hence the damping wing strength still agrees well among both models. The remaining small differences among the two models can be attributed to differences in the neutral hydrogen content in front of our fixed lower integration limit of $r_{\text{min}} = 4$ cMpc.

While Figures 3.8, 3.9 and 3.10 only show the scatter of the IGM transmission profiles in select parameter bins, we verified that the effect described above only becomes relevant in parameter bins with $r_{\text{patch}} = 0.3$ cMpc, while both models perfectly agree across the entire rest of $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space (modulo a somewhat increased statistical noise level for the realistic

topology in bins which are more scarcely populated). The prior volume where our toy prescription introduces a small bias therefore remains minimal, and this bias certainly remains subdominant to other major sources of uncertainty in realistic inference settings where e.g. the intrinsic continuum of the quasar has to be reconstructed as well (Kist et al. 2025b). Correspondingly, the 68-percentile widths $\sigma(t(v = v_T))$ of the IGM transmission distributions coincide remarkably well between the two topologies in all regions of parameter space, as can be seen by comparing the middle to the lower panels of Figure 3.7.

Overall, this leads us to the striking conclusion that the statistical properties (median and 68-percentile scatter) of the IGM transmission profiles are largely insensitive to the peculiarities of *any* given reionization topology. Again, we compared transmission profiles from two entirely different models—the latter of which is not even physically motivated—and yet, no differences are recognizable across the largest parts of parameter space. Even the remaining differences at $r_{\text{patch}} = 0.3 \text{ cMpc}$ are a mere relic of the unphysical fact that our toy prescription does not account for regions around the most massive halos reionizing early. In future applications, we could alleviate this effect by modifying our toy model to respect other spatial correlations between x_{HI} and Δ such as their average radial profile in addition to $N_{\text{HI}}^{\text{DW}}$ and r_{patch} instead of randomly laying down the neutral patches. Our statement about topology independence of our new set summary statistics therefore remains unconditionally true when comparing profiles from any two realistic reionization topologies where this property is accounted for.

3.4.2.3 Parameter dependence of the IGM transmission profiles at short quasar lifetimes

After these general considerations, we now investigate how the shape of the IGM transmission profiles changes as a function of $N_{\text{HI}}^{\text{DW}}$ and r_{patch} . We start by considering the short lifetime ($t_Q = 10^4 \text{ yr}$) profiles depicted in Figure 3.8. As expected, higher HI column densities generally cause stronger IGM damping wing imprints. Also our second summary statistic r_{patch} has a clear impact on the shape of the IGM transmission profiles, even though it is defined as the distance to the first neutral patch in the *pre*-quasar topology. The first obvious reason for this is the fact that $N_{\text{HI}}^{\text{DW}}$ is defined to be insensitive to any kind of structure in front of our lower integration limit $r_{\text{min}} = 4 \text{ cMpc}$. For quasars whose first neutral patch is located at $r_{\text{patch}} < r_{\text{min}}$, and whose lifetime is not long enough to ionize away this neutral patch, r_{patch} is the only label that can carry information about the remaining nearby neutral material which $N_{\text{HI}}^{\text{DW}}$ by definition cannot capture.

However, we see in Figure 3.8 that the IGM transmission profiles still carry information about r_{patch} even if $r_{\text{patch}} > r_{\text{min}}$. This is because on average, a short-lived quasar which has only been shining for a few thousands of years

has not yet had sufficient time to carve out an ionized region extending (significantly) beyond the start of the first neutral patch (regardless how its original location r_{patch} relates to r_{min}), and so r_{patch} constitutes a hard upper limit for the location of its ionization front r_{ion} (marked with a vertical dash-dotted line in each panel).¹⁸ The transmission therefore certainly drops to zero at or prior to r_{patch} (dotted vertical lines), and hence r_{patch} carries additional information about the shape of the profile in this region, clearly supplementing the information about the red-side damping wing shape which is primarily encoded in the $N_{\text{HI}}^{\text{DW}}$ statistic. This can be seen by comparing the panels within a given row (i.e., at fixed $N_{\text{HI}}^{\text{DW}}$) in Figure 3.8. In these regions of parameter space where the transmission drops to zero close to r_{patch} , this label therefore encodes similar information as the velocity realignment point proposed in Chen (2024) and Keating et al. (2024b) (when considered at a fixed quasar lifetime t_{Q}).

On the other hand, no differences are recognizable between the profiles in the two upper rows of Figure 3.8 where all transmission is extinguished at a parameter-independent location significantly closer to the source than the position r_{patch} of the first neutral patch. As the IGM is fully ionized up to the latter point, the Gunn-Peterson optical depth τ_{GP} in this region is inversely proportional to the photoionization rate Γ_{QSO} of the quasar, i.e., $\tau_{\text{GP}} \sim 1/\Gamma_{\text{QSO}}$. Since the quasar’s photoionization rate decreases as the inverse square of the distance r from the quasar, $\Gamma_{\text{QSO}} \sim 1/r^2$, all transmission eventually gets suppressed beyond a certain distance, even if the quasar ionization front has already passed through. As a result, profiles with r_{patch} values exceeding this distance are entirely degenerate. This degeneracy would be ameliorated if we were to consider a brighter quasar whose Γ_{QSO} would be higher and hence this distance would be located farther outwards from the source.

3.4.2.4 Parameter dependence of the IGM transmission profiles at intermediate to long quasar lifetimes

The sensitivity to r_{patch} decreases as we move to older objects with lifetimes of $t_{\text{Q}} = 10^6$ or 10^8 yr such as depicted in Figures 3.9 and 3.10. Such quasars have had a sufficient amount of time to carve out a large ionized region around themselves, in most cases exceeding the first pre-quasar neutral patch location and hence, the transmission does not sharply drop to zero at r_{patch} . Instead, r_{patch} is located somewhere within the proximity zone of most such quasars, and as such is hardly reconstructable from a given IGM transmission

¹⁸The location of the ionization front r_{ion} is measured empirically based on all toy sightlines which are available in a given parameter bin. To that end, we take the same algorithm used to determine r_{patch} for the pre-quasar topology, and apply it to the post-quasar topology instead. The dash-dotted line marks the median r_{ion} of all toy profiles in a given panel, and the grey-shaded region the 68-percentile scatter.

profile. Note that in the two panels of Figure 3.9 where $r_{\text{patch}} = 20 \text{ cMpc}$ (and $t_{\text{Q}} = 10^6 \text{ yr}$), we still do observe a similar drop-off close to r_{patch} as we did for shorter neutral patch distances at $t_{\text{Q}} = 10^4 \text{ yr}$. This shows that at a lifetime of $t_{\text{Q}} = 10^6 \text{ yr}$, a neutral patch located at $r_{\text{patch}} = 20 \text{ cMpc}$ still constitutes a comparably hard limit for the quasar ionization front which comes to a halt shortly thereafter, whereas closer neutral patches easily get ionized away such that r_{ion} significantly exceeds r_{patch} .

Remarkably, however, even in the case where the first pre-quasar neutral patch is ionized away, we statistically still do observe a remnant signature of r_{patch} in form of an excess transmission bump starting at exactly this location. This feature can most clearly be identified in the central panels of Figures 3.9 and 3.10 where both median and scatter are elevated starting precisely at r_{patch} . Most notably, a noisier version of this bump can even be identified for the semi-numerical IGM transmission profiles in the $(\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}, r_{\text{patch}}/\text{cMpc}) = (20.2, 20)$ panel of Figure 3.10. We attribute this feature to photoelectric heating of the (neutral) IGM by the hard quasar spectrum.

The mechanism can be understood by returning to the example sightline depicted in Figure 3.1: comparing the temperature fields before and after the quasar has turned on, we see that the initially neutral regions get heated significantly (bottom row) after getting ionized away by the quasar. This is because at large optical depth, the soft photons all get absorbed, and the optical depth of the harder photons is also high so they get absorbed as well, causing photoelectric heating up to 40 000 K (c.f. also Figure 6 of Davies & Hennawi (2023)). As a result, the transmission is enhanced in these regions. Since this significant amount of heating is restricted to initially neutral patches which recently got ionized by the quasar, the pre-quasar x_{HI} topology is still encoded in the post-quasar temperature field, and hence also leaves an imprint on the IGM transmission field. For an individual sightline, however, this signature is hard to disentangle from the stochasticity of the Lyman- α absorption features in the proximity zone of the quasar due to density fluctuations, even more so when the continuum is unknown, and in the presence of observational noise. Yet, our $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameterization allows us to statistically account for this feature when performing astrophysical parameter inference. We will explore this possibility in future work.

3.5 Conclusions

We introduced in this work a novel three-parameter model that tightly captures the characteristic shape of quasar IGM damping wings. As an alternative to the common parameterization based on the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$, we defined two new *local* summary statistics quantifying the

neutral hydrogen content along the sightline from the quasar *before* it started shining: 1) the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$, and 2) the distance r_{patch} between the quasar and the first neutral patch. We supplemented these two local measures of the *pre*-quasar neutral topology with the quasar lifetime t_{Q} as a third parameter encapsulating the effects of the ionizing quasar radiation.

Since quasars are typically going to ionize away all surrounding neutral gas within the first few cMpc, we found that the damping wing is most sensitive to the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ starting at 4 cMpc from the source, and integrated over a range of 100 cMpc. By adding a Lorentzian weighting function (associated to a fixed reference distance $r_{\text{T}} = 18$ cMpc) to the column density integral, we account for the effect of the Lyman- α cross section σ_{α} in the optical depth integral, assuring a maximally tight parameterization of the IGM damping wing at the spectral pixel corresponding to r_{T} . By demonstrating that the damping wing essentially constitutes a one-parameter family, we showed that $N_{\text{HI}}^{\text{DW}}$ tightly parameterizes the damping wing imprint not only at this specific location but across the *entire* spectral range.

We introduced as a second summary statistic the distance r_{patch} from the source to the first neutral patch in the *pre*-quasar topology, and showed that it still encodes information about *post*-quasar IGM transmission profiles, and that this information is complementary to $N_{\text{HI}}^{\text{DW}}$. We related the definitions of our summary statistics back to the labels recently proposed in [Chen \(2024\)](#), [Keating et al. \(2024b\)](#), and [Mason et al. \(2026\)](#), and argued that by starting the $N_{\text{HI}}^{\text{DW}}$ integration directly at the location of the source, our parameterization is also applicable in the context of IGM damping wings towards galaxies.

We simulated realistic IGM transmission profiles and compared their scatter in the damping wing region of the spectrum at fixed parameter values within both the global and the local parameterization. We found that due to the stochastic distribution of neutral patches during reionization and density fluctuations in the IGM, the 68-percentile scatter of the IGM transmission values at $v_{\text{T}} = 2000$ km/s in the global parameterization can be as large as 7% (at a quasar lifetime of $t_{\text{Q}} = 10^4$ yr), and we demonstrated that this scatter decreases down to $\lesssim 1\%$ across the entire range of physical parameter space when aggregating the same sightlines according to our local parameterization.

We introduced a simple numerical prescription to generate synthetic HI density profiles at any desired location in $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameter space, varying smoothly as a function of both these parameters. Even though the procedure does not attempt to describe any physical processes, it allowed us to demonstrate the robustness of our parameterization against the choice of reionization model since it results in a significantly different neutral topology. We observed an exceptional agreement between the realistic and

the toy profiles at almost every location in physical parameter space. Small differences between the models only became apparent for profiles with a neutral patch directly next to the quasar, which we could trace back to the fact that our toy prescription does not pay tribute to the fact that reionization takes place inside out. As this would be accounted for by any realistic reionization model, the remarkable overall agreement between the two models demonstrates the topology-insensitivity of our parameterization.

In addition, this agreement legitimizes the future use of our toy prescription for the purpose of astrophysical parameter inference, decoupling 1) the topology-insensitive task of inferring the local damping wing statistics introduced in this work, and 2) tying these constraints to a specific reionization model, resulting in near-optimal constraints not only on the global timing of reionization, but also the reionization topology, hitherto unconstrained with quasar IGM damping wings. In this context, all assumptions about the reionization model can be encoded in a prior on $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$, determined by the distribution of these parameters within the model of interest. This clear separation will facilitate future inference endeavors, and even the comparison of different reionization models, harnessing *all* the information encapsulated in quasar IGM damping wings.

Acknowledgements

We acknowledge helpful conversations with the ENIGMA group at UC Santa Barbara and Leiden University and would especially like to thank Shane Bechtel and Benjamin Snyder for comments on an early version of this manuscript. This work made use of NumPy (Harris et al. 2020), SciPy (Virtanen et al. 2020), Astropy (Astropy Collaboration et al. 2013, 2018, 2022), h5py (Collette 2013), Matplotlib (Hunter 2007), 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.

Appendices

3.A On the relation between τ_{DW} and $N_{\text{HI}}^{\text{DW}}$

We introduced in Section 3.2 the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ as a summary statistic of the local pre-quasar neutral hydrogen content along the line of sight from the source. By comparing $N_{\text{HI}}^{\text{DW}}$ as defined in Eq. (3.6) to the pre-quasar damping wing optical depth

$$\tau_{\text{DW}}^{\text{pre}}(\lambda_{\text{rest}}) = \int_0^{R(z_{\text{QSO}})} n_{\text{HI}}^{\text{pre}}(R) \cdot \sigma_{\alpha} \left(\frac{1+z_{\text{QSO}}}{1+z(R)} \lambda_{\text{rest}} \right) dR, \quad (3.11)$$

we motivate in this section why this summary statistic is a near-optimal means to reduce the IGM transmission scatter redward of the Lyman- α line. Specifically, we show that in the limit where the Lorentzian-weighted average $\langle x_{\text{HI}} \cdot \Delta \rangle_{\text{Lor}}$ as defined in Eq. (3.4) is computed along the *entire* line of sight, and where the Lyman- α cross section is approximated as perfectly Lorentzian, Eq. (3.7) holds, or, in other words, $N_{\text{HI}}^{\text{DW}}$ encodes the same information as the pre-quasar damping wing optical depth $\tau_{\text{DW}}^{\text{pre}}(v = v_{\text{T}})$ evaluated at the reference velocity offset v_{T} .

We start by writing the integrand in Eq. (3.11) explicitly in terms of the integration variable R . This is required since we need to evaluate the Lyman- α cross section σ_{α} at wavelength

$$\lambda = \frac{1+z_{\text{QSO}}}{1+z(R)} \lambda_{\text{rest}}. \quad (3.12)$$

The relation $z(R)$ between redshift z and proper distance R from the quasar can be obtained by inverting the light-travel distance relation

$$R(z) = \int_z^{z_{\text{QSO}}} \frac{c}{(1+z')H(z')} dz'. \quad (3.13)$$

Certainly, the endpoint of reionization z_{end} puts an end to all contributions to the damping wing optical depth as per Eq. (3.11).¹⁹ Since the integrand $c/(1+z')H(z')$ does not change significantly from z_{QSO} to z_{end} , we can expand Eq. (3.13) to linear order around $z = z_{\text{QSO}}$, finding

$$R(z) \simeq \frac{c}{(1+z_{\text{QSO}})H(z_{\text{QSO}})} (z_{\text{QSO}} - z). \quad (3.14)$$

This relation can now easily be inverted, and we obtain

$$1+z(R) = (1+z_{\text{QSO}}) \left(1 - \frac{H(z_{\text{QSO}})}{c} R \right) \quad (3.15)$$

¹⁹In fact, Section 3.2.2 led us to the even stronger conclusion that the damping wing imprint is only sensitive to the structure within the first ~ 100 cMpc from the source.

If we now evaluate the Lyman- α rest-frame wavelength $\lambda_{\text{rest}} = \lambda_{\alpha} (1 + \frac{v}{c})$ at the velocity offset $v = v_{\text{T}}$, recalling that $R_{\text{T}} \equiv +v_{\text{T}}/H(z_{\text{QSO}})$, and using Eq. (3.15), we immediately find that Eq. (3.12) turns into

$$\lambda = \lambda_{\alpha} \frac{c + H(z_{\text{QSO}}) R_{\text{T}}}{c - H(z_{\text{QSO}}) R}. \quad (3.16)$$

With an expression at hand for $\lambda(R)$, we now proceed by substituting this relation into the following expression for the scattering cross section σ_{α} of the Lyman- α line:

$$\sigma_{\alpha}(\lambda) = \frac{\pi e^2}{m_e c} f_{\alpha} \phi_{\alpha}(\lambda). \quad (3.17)$$

Here, $f_{\alpha} \simeq 0.416$ is the Lyman- α oscillator strength, and $\phi_{\alpha}(\lambda)$ is the line profile function, commonly assumed as a Voigt profile. Away from the Lyman- α resonance itself, we can assume to good approximation that $\phi_{\alpha}(\lambda)$ is of Lorentzian shape (Draine 2011; Bach & Lee 2015), i.e.,

$$\phi_{\alpha}(\lambda) = \frac{\gamma_{\alpha} \lambda_{\alpha}/(\pi c)}{(\lambda_{\alpha}/\lambda - 1)^2 + \gamma_{\alpha}^2}, \quad (3.18)$$

where $\gamma_{\alpha} \equiv \Gamma_{\alpha} \lambda_{\alpha}/(4\pi c)$ with the decay constant $\Gamma_{\alpha} = 6.265 \times 10^8 \text{ s}^{-1}$ of the Lyman- α transition. Evaluating λ_{rest} at the velocity offset v_{T} assures that λ (as given by Eq. (3.12)) is sufficiently far away from the Lyman- α resonance λ_{α} such that $\phi_{\alpha}(\lambda)$ is not only to good approximation of Lorentzian shape, but we can also safely neglect the second term in the denominator of Eq. (3.18). Utilizing Eq. (3.16), we can explicitly rewrite Eq. (3.17) for the Lyman- α scattering cross section as a function of R :

$$\sigma_{\alpha}(\lambda(R)) \simeq \frac{e^2}{m_e c^2} f_{\alpha} \gamma_{\alpha} \lambda_{\alpha} \frac{(c/H(z_{\text{QSO}}) - R_{\text{T}})^2}{(R + R_{\text{T}})^2}. \quad (3.19)$$

Substituting this back into Eq. (3.11) shows that our weighting function (Eq. (3.3)) exactly mimics the scaling in the integrand of the optical depth integral. If we now let $R_{\text{min}} \rightarrow 0$ and $R_{\text{max}} \rightarrow R(z_{\text{QSO}})$, we formally arrive at Eq. (3.7), showing that in these limits, $N_{\text{HI}}^{\text{DW}}$ is an optimal summary of the pre-quasar damping wing optical depth $\tau_{\text{DW}}^{\text{pre}}$. By instead choosing the integration limits in the way discussed in Section 3.2.2, we can ensure that an analogous version of Eq. (3.7) even holds for the *post*-quasar optical depth τ_{DW} .

3.B On the hyperparameter dependence of $N_{\text{HI}}^{\text{DW}}$

The Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ defined in Eq. (3.6) comes with three hyperparameters, the integration limits r_{min} and r_{max} as well as




Figure 3.11: Distribution of the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$ for different hyperparameter choices among all 21×852 simulated IGM transmission profiles based on semi-numerical x_{HI} skewers. Variations with respect to the lower integration limit r_{min} are shown in the left panel, the upper integration limit r_{max} in the middle, and the reference velocity offset v_{T} in the right panel. The distribution resulting from our reference choices of $r_{\text{min}} = 4.0 \text{ cMpc}$, $r_{\text{max}} = r_{\text{min}} + 100.0 \text{ cMpc}$ and $v_{\text{T}} = 2000 \text{ km/s}$ is shown as black curve in each panel. All distributions are smoothed via a kernel density estimation (KDE) with a Gaussian kernel of size $\Delta(\log_{10} N_{\text{HI}}^{\text{DW}} / \text{cm}^{-2}) = 0.1$.

the reference distance r_{T} (or, equivalently, velocity offset v_{T}) whose values we fix as discussed in Sections 3.2.2 and 3.2.3. We investigate in this appendix the effects of varying these three parameters.

Figure 3.11 depicts how the distribution of our $N_{\text{HI}}^{\text{DW}}$ statistic among all 21×852 simulated sightlines changes as a function of these three parameters. We start by studying the effects of the lower integration limit r_{min} depicted in the left-hand panel. As discussed in Section 3.2.2, this is the most sensitive hyperparameter as it determines how much nearby pre-quasar structure is included in the column density integral. This nearby structure carries the highest Lorentzian weight but also tends to have been ionized away by a quasar with a typical lifetime, and as such, our choice of $r_{\text{min}} = 4 \text{ cMpc}$ constitutes a compromise between maximizing the information about the pre-quasar topology that we can still gain while not including too much material that does not actually contribute to the observed damping wing signature anymore. In line with these considerations, we see in Figure 3.11 that the $N_{\text{HI}}^{\text{DW}}$ -distribution gets more skewed towards higher values the shorter the integration limit is. In particular, we observe that the high- $N_{\text{HI}}^{\text{DW}}$ peak decreases in size and the distribution instead obtains a more extended tail towards the highest $N_{\text{HI}}^{\text{DW}}$ values.

Figure 3.12 depicts— analogously to Figure 3.7— half the 68-percentile scatter $\sigma(t(v = v_{\text{T}}))$ among all 21×852 profiles when aggregated according to our $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ parameterization for three different choices of the lower integration limit $r_{\text{min}} = 0.3 \text{ cMpc}$, 4.0 cMpc and 8.0 cMpc , and three different lifetimes of $t_{\text{Q}} = 10^4 \text{ yr}$, 10^6 yr and 10^8 yr . Although the differences appear relatively weak, we see that the choice of $r_{\text{min}} = 0.3 \text{ cMpc}$ leads to a slight improvement at the shortest lifetime of $t_{\text{Q}} = 10^4 \text{ yr}$ where a shorter




Figure 3.12: Sightline-to-sightline scatter of the IGM damping wing as a function of astrophysical parameter space for different lower integration limits r_{\min} for the HI column density $N_{\text{HI}}^{\text{D,W}}$. Analogously to the middle row of Figure 3.7, each panel shows half the 68-percentile width $\sigma(t(v = v_T))$ of the distribution of IGM transmission values at $v = v_T$ among all sightlines based on semi-numerical x_{HI} skewers as a function of $(N_{\text{HI}}^{\text{D,W}}, r_{\text{patch}})$ -parameter space, for quasar lifetimes of $t_Q = 10^4 \text{ yr}$ (left column), 10^6 yr (middle column) and 10^8 yr (right column), and for integration limits of $r_{\min} = 0.3 \text{ cMpc}$ (upper row), 4.0 cMpc (middle row) and 8.0 cMpc (lower row). Bins which are occupied by a lower number of sightlines are more transparent, and white if not occupied at all.




Figure 3.13: Like Figure 3.12 but for upper integration limits of $r_{\text{max}} = r_{\text{min}} + 80.0 \text{ cMpc}$ (upper row), $r_{\text{min}} + 100.0 \text{ cMpc}$ (middle row) and $r_{\text{min}} + 140.0 \text{ cMpc}$ (lower row) for the HI column density $N_{\text{HI}}^{\text{DW}}$.




Figure 3.14: Like Figure 3.12 but for reference velocity offsets of $v_T = 1000 \text{ km/s}$, 2000 km/s and 3000 km/s for the HI column density $N_{\text{HI}}^{\text{DW}}$. Regardless of the reference offset, the scatter is consistently evaluated at $v = 2000 \text{ km/s}$.

integration limit would in fact be more appropriate, while we observe an opposite trend at longer quasar lifetimes. Vice versa, a large integration limit of $r_{\text{min}} = 8.0$ cMpc notably increases the scatter at the shorter lifetimes where this means we are clearly excluding too much neutral material from the integration range. These trends empirically confirm the aptness of our lower integration limit at 4.0 cMpc.

Going forward, we see that impact of the upper integration limit r_{max} is significantly weaker since all contributions from distant neutral material are anyways suppressed by the Lorentzian decline of our weighting function. As a result, the column density distribution in the middle panel of Figure 3.11 remains largely unaffected between $r_{\text{max}} = r_{\text{min}} + 80.0$ cMpc and $r_{\text{min}} + 140.0$ cMpc corresponding to the maximum length of our sightlines. Note that we are not sensitive to the differences in the low- $N_{\text{HI}}^{\text{DW}}$ peak as these are the column densities that do not cause any observable damping wing imprint and therefore their exact value has little effect. Correspondingly, Figure 3.13 shows that the scatter $\sigma(t(v = v_{\text{T}}))$ stays remarkably constant regardless of our choice of r_{max} and only starts to increase slightly towards our lowest value of $r_{\text{max}} = 80.0$ cMpc.

Finally, the right-hand panel of Figure 3.11 suggests that a change of our reference velocity-offset v_{T} essentially just causes a linear shift in the column density distribution. This is in line with our considerations in Section 3.2.3 where we demonstrated that the damping wing imprint largely constitutes a one-parameter family whose value at any given velocity-offset is determined by the value at a fixed reference offset v_{T} . Correspondingly, the $N_{\text{HI}}^{\text{DW}}$ -distribution only gets shifted by the corresponding proportionality constant. When considering the sightline-to-sightline scatter $\sigma(t(v = 2000 \text{ km/s}))$ shown in Figure 3.14 for the three different reference offsets of $v_{\text{T}} = 1000 \text{ km/s}$, 2000 km/s and 3000 km/s ,²⁰ we see the scatter increases exactly in line with the amount one would deduce from Figure 3.2 based on the Lyman- α transmission itself.

²⁰Note that for this analysis we always evaluate the scatter at $v = 2000 \text{ km/s}$, regardless of the value of v_{T} .

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:

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

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 $\text{H}\beta$ and $[\text{O III}]$ 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.

5

FIRST CONSTRAINTS ON THE
LOCAL IONIZATION TOPOLOGY
IN FRONT OF TWO QUASARS
AT $z \sim 7.5$

Abstract

Thus far, Lyman- α damping wings towards quasars have been used to probe the *global* ionization state of the foreground intergalactic medium (IGM). A new parameterization has demonstrated that the damping wing signature also carries *local* information about the distribution of neutral hydrogen (HI) in front of the quasar before it started shining. Leveraging a recently introduced Bayesian JAX-based Hamiltonian Monte Carlo (HMC) inference framework, we derive constraints on the Lorentzian-weighted HI column density $N_{\text{HI}}^{\text{DW}}$, the quasar's distance r_{patch} to the first neutral patch and its lifetime t_{Q} based on JWST/NIRSpec spectra of the two $z \sim 7.5$ quasars J1007+2115 and J1342+0928. After folding in model-dependent topology information, we find that J1007+2115 (and J1342+0928) is most likely to reside in a $\langle x_{\text{HI}} \rangle = 0.32_{-0.20}^{+0.22}$ ($0.58_{-0.23}^{+0.23}$) neutral IGM while shining for a remarkably short lifetime of $\log_{10} t_{\text{Q}}/\text{yr} = 4.14_{-0.18}^{+0.74}$ (an intermediate lifetime of $5.64_{-0.43}^{+0.25}$) along a sightline with $\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2} = 19.70_{-0.86}^{+0.35}$ ($20.24_{-0.22}^{+0.25}$) and $r_{\text{patch}} = 28.9_{-14.4}^{+54.0}$ cMpc ($10.9_{-5.9}^{+5.6}$ cMpc). In light of the potential presence of local absorbers in the foreground of J1342+0928 as has been recently suggested, we also demonstrate how the Lorentzian-weighted column density $N_{\text{HI}}^{\text{DW}}$ provides a natural means for quantifying their contribution to the observed damping wing signal.

Work published in: **Timo Kist**, Joseph F. Hennawi, Frederick B. Davies, Eduardo Bañados, Sarah E. I. Bosman, Zheng Cai, Anna-Christina Eilers, Xiaohui Fan, Zoltán Haiman, Hyunsung D. Jun, Yichen Liu, Jinyi Yang and Feige Wang, *Monthly Notices of the Royal Astronomical Society*, Volume 545, Issue 3, January 2026, staf2219, doi.org/10.1093/mnras/staf2219. Reprinted here in its entirety.

5.1 Introduction

The highly sensitive Lyman- α transition observed in the spectra of high-redshift sources carries a wealth of information about the epoch of reionization, most particularly when the absorption imprint from the foreground neutral intergalactic medium (IGM) saturates, manifesting in a characteristic *damping wing* signature redward of Lyman- α line center (Miralda-Escudé 1998). Quasars (Mortlock et al. 2011; Bolton et al. 2011; Greig et al. 2017b, 2019, 2022, 2024b; Bañados et al. 2018; Davies et al. 2018a; Wang et al. 2020; Yang et al. 2020a; Reiman et al. 2020; Āurovčíková et al. 2020, 2024) as well as galaxies (Curtis-Lake et al. 2023; Hsiao et al. 2024; Keating et al. 2024b; Umeda et al. 2024, 2025; Park et al. 2025; Mason et al. 2026) have been used as background sources to infer constraints on the global volume-averaged fraction $\langle x_{\text{HI}} \rangle$ of neutral hydrogen (HI) at the redshift of the source, either by stacking spectra at a similar redshift (e.g. Āurovčíková et al. 2024; Umeda et al. 2024, 2025), or statistically combining the inferred $\langle x_{\text{HI}} \rangle$ constraints (e.g. Greig et al. 2024b; Mason et al. 2026). Recently, this has been complemented by claimed detections of damping wings next to individual Gunn-Peterson troughs due to neutral islands persisting until $5.5 \lesssim z \lesssim 6$ (Becker et al. 2024; Spina et al. 2024; Zhu et al. 2024; Sawyer et al. 2025). All these analyses have established IGM damping wings as a natural probe of the timing of reionization. But recent studies have constructed parameterizations that capture the shape of the IGM damping wing significantly more tightly than the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$, and even carry information about the *local* ionization topology in front of a given source (Chen 2024; Keating et al. 2024a; Mason et al. 2026; Kist et al. 2025c).

Specifically, Kist et al. (2025c) put forward a three-parameter model applicable in the context of quasars, consisting of two summary statistics of the x_{HI} field *before* the quasar started shining (which we dub the *pre-quasar* x_{HI} field), and the quasar’s lifetime t_{Q} encoding the effects of its ionizing radiation. The former two statistics of the pre-quasar ionization topology are 1) the column density $N_{\text{HI}}^{\text{DW}}$ of neutral material in front of the quasar, weighted by a Lorentzian profile mimicking the frequency dependence of the Lyman- α cross section σ_{α} appearing in the optical depth integral; and 2) the distance r_{patch} from the quasar to the first neutral patch. In a companion paper, Kist et al. (2025a) introduced a fully Bayesian pipeline to infer these two parameters along with the quasar lifetime t_{Q} in a reionization model-independent fashion based on observed high-redshift quasar spectra, accounting for all relevant sources of uncertainty such as the unknown intrinsic continuum of the quasar, IGM transmission fluctuations and spectral noise. Statistical tests on hundreds of mock spectra demonstrated that the pipeline allows them to constrain t_{Q} to $0.58_{-0.13}^{+0.13}$ dex, $N_{\text{HI}}^{\text{DW}}$ to $0.72_{-0.25}^{+0.08}$ dex

and r_{patch} to $35.7^{+7.8}_{-31.9}$ cMpc in case a noticeable damping wing is present in the spectrum. Here we take advantage of this framework to infer the first *local* IGM damping wing constraints from JWST/NIRSpec spectra of two of the highest-redshift quasars known to date, that is, J1007+2115 at $z = 7.51$ (Yang et al. 2020a) and J1342+0928 at $z = 7.54$ (Bañados et al. 2018).

We briefly summarize in Section 5.2 the underlying theory and our local damping wing analysis framework, and proceed in Section 5.3 by introducing the data and presenting our analysis results for the two aforementioned objects. We put these measurements into context with existing literature constraints and conclude in Section 5.4.

5.2 Methods

We conduct our analysis in the context of the local damping wing parameterization put forward by Kist et al. (2025c), harnessing the inference pipeline originally introduced in Hennawi et al. (2025) and adapted by Kist et al. (2025a) for the use with this parameterization. Kist et al. (2025a) also conceptualized a theoretical framework for tying the inferred local parameter constraints which are agnostic to the underlying reionization topology to a specific reionization model to obtain a global constraint on the timing of reionization. We restrict this section to a brief but self-contained summary of each of these analysis components and refer the reader to the works above for more comprehensive descriptions of the respective parts.

The first of the two local summary statistics of the pre-quasar ionization topology defined by Kist et al. (2025c) that we aim to constrain in this work is the Lorentzian-weighted HI column density

$$N_{\text{HI}}^{\text{DW}} \equiv 5.1 \times 10^{20} \text{ cm}^{-2} \times \left(\frac{r_{\text{T}}}{18 \text{ cMpc}} \right) \times \left(\frac{1 + z_{\text{QSO}}}{1 + 7.54} \right)^2 \int_{r_{\text{min}}/r_{\text{T}}}^{r_{\text{max}}/r_{\text{T}}} \frac{x_{\text{HI}}(r) \cdot \Delta(r)}{(r/r_{\text{T}} + 1)^2} \text{d}(r/r_{\text{T}}), \quad (5.1)$$

where r denotes the comoving distance, and the denominator in the integrand is the Lorentzian weighting kernel which accounts for the frequency dependence of the Lyman- α cross section that the neutral hydrogen field $x_{\text{HI}}(r)$ and the overdensity field $\Delta(r)$ are integrated against in the optical depth integral. The three distances r_{min} , r_{max} and r_{T} are hyperparameters which we fix to $r_{\text{min}} = 4$ cMpc, $r_{\text{max}} = r_{\text{min}} + 100$ cMpc and $r_{\text{T}} = 18$ cMpc throughout (for details see Kist et al. 2025c). Our second statistic r_{patch} can be straightforwardly defined as the (comoving) distance between the source and the first neutral patch. Recall that both $N_{\text{HI}}^{\text{DW}}$ and r_{patch} are summaries of the *pre*-quasar ionization topology x_{HI} as this is the cosmological field we are aiming to constrain. Our third parameter, the lifetime t_{Q} of the quasar, summarizes the effects of the quasar’s ionizing radiation.

We simulate IGM transmission profiles \mathbf{t} based on a hybrid approach following [Davies et al. \(2018a\)](#) where we combine density, velocity and temperature skewers from the Nyx hydrodynamical simulations ([Almgren et al. 2013](#); [Lukić et al. 2015](#)) and synthetic x_{HI} skewers ([Kist et al. 2025c](#)) generated at 21 column density values between $17.48 \leq \log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2} \leq 21.08$ and 18 different neutral patch distances between $0.3 \text{ cMpc} \leq r_{\text{patch}} \leq 143.0 \text{ cMpc}$. We then perform one-dimensional radiative transfer ([Davies et al. 2016](#)) assuming 51 quasar lifetimes between $10^3 \text{ yr} \leq t_{\text{Q}} \leq 10^8 \text{ yr}$, and forward-model realistic instrumental effects ([Hennawi et al. 2025](#)). All our models are generated assuming $z_{\text{QSO}} = 7.54$ with an ionizing photon emission rate of $Q = 10^{57.14} \text{ s}^{-1}$, explicitly resembling J1342+0928 but also appropriate for J1007+2115. Further, to tie our local constraints to the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ in the way described in [Kist et al. \(2025a\)](#), we determine the stochastic mapping $P_{\text{top}}(\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$,¹ by combining the same hydrodynamical Nyx sightlines with x_{HI} skewers extracted from semi-numerical reionization topologies simulated with a modified version of 21cmFAST ([Mesinger et al. 2011](#); [Davies & Furlanetto 2022](#)) at 21 different global IGM neutral fractions between $0 \leq \langle x_{\text{HI}} \rangle \leq 1$.

In combination with a principal component analysis (PCA) model for the intrinsic quasar continuum \mathbf{s} constructed based on a dataset of 44,587 quasar spectra at $1.878 < z < 3.427$ from the SDSS-III Baryon Oscillation Spectroscopic Survey (BOSS) and SDSS-IV Extended BOSS (eBOSS) ([Hennawi et al. 2026](#); [Kist et al. 2025a](#)), this allows us to write down the likelihood of an observed spectrum \mathbf{f} given the astrophysical parameters $\boldsymbol{\theta}$, the PCA coefficients $\boldsymbol{\eta}$ and its observational noise vector $\boldsymbol{\sigma}$ as the following multivariate Gaussian distribution ([Hennawi et al. 2025](#)):

$$L(\mathbf{f} | \boldsymbol{\sigma}, \boldsymbol{\theta}, \boldsymbol{\eta}) = \mathcal{N}(\mathbf{f}; \langle \mathbf{t} \rangle \circ \langle \mathbf{s} \rangle, \boldsymbol{\Sigma} + \langle \mathbf{S} \rangle \mathbf{C}_{\mathbf{t}} \langle \mathbf{S} \rangle + \langle \mathbf{T} \rangle \mathbf{C}_{\mathbf{s}} \langle \mathbf{T} \rangle), \quad (5.2)$$

where $\mathbf{t} \circ \mathbf{s}$ is the element-wise (Hadamard) product of the two mean vectors $\langle \mathbf{t} \rangle$ and $\langle \mathbf{s} \rangle$, and $\mathbf{C}_{\mathbf{t}}$ and $\mathbf{C}_{\mathbf{s}}$ are the covariance matrices of \mathbf{t} and \mathbf{s} , and we write $\boldsymbol{\Sigma} \equiv \text{diag}(\boldsymbol{\sigma})$, $\mathbf{T} \equiv \text{diag}(\mathbf{t})$ and $\mathbf{S} \equiv \text{diag}(\mathbf{s})$.

We follow [Kist et al. \(2025a\)](#) and initially infer the local parameters $\boldsymbol{\theta} \equiv (\log_{10} t_{\text{Q}}/\text{yr}, \log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}, r_{\text{patch}})$ under the assumption of a logarithmically flat prior on the lifetime between $3 \leq \log_{10} t_{\text{Q}}/\text{yr} \leq 8$, and a constant two-dimensional topology-agnostic prior on $(\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}, r_{\text{patch}})$ with a non-trivial boundary enclosing all physically permitted regions of parameter space (for details see [Kist et al. 2025a](#)). After marginalizing out the nuisance parameters $\boldsymbol{\eta}$, we can probabilistically tie these local constraints on $\boldsymbol{\theta}$ to the aforementioned semi-numerical reionization topology via the conditional dis-

¹Throughout, we use a ‘top’ subscript to denote probability distributions that are defined in the context of a specific reionization model.

tribution $P_{\text{top}}(\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle)$ and obtain a joint constraint on the global IGM neutral fraction:

$$P_{\text{top}}(\langle x_{\text{HI}} \rangle, \boldsymbol{\theta} | \mathbf{f}, \boldsymbol{\sigma}) = L(\mathbf{f} | \boldsymbol{\sigma}, \boldsymbol{\theta}) \times P(\langle x_{\text{HI}} \rangle) \times P(\log_{10} t_{\text{Q}}/\text{yr}) \\ \times P_{\text{top}}(\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}, r_{\text{patch}} | \langle x_{\text{HI}} \rangle) / P(\mathbf{f} | \boldsymbol{\sigma}). \quad (5.3)$$

For comparison, we also perform the inference following the conventional approach without local summary statistics where $\boldsymbol{\theta} \equiv (\langle x_{\text{HI}} \rangle, \log_{10} t_{\text{Q}}/\text{yr})$ such that no conversion according to Eq. (5.3) is necessary. We henceforth refer to this as the *global* parameterization (as opposed to our *local* one). Here we assume a flat neutral fraction prior between $0 \leq \langle x_{\text{HI}} \rangle \leq 1$.

Practically, we sample from the respective posterior distribution via the NumPyro Hamiltonian Monte-Carlo (HMC) implementation with No U-Turn Sampler (NUTS). Each inference run consists of 8 HMC chains with 1000 warm-up and 2000 sampling steps per chain (for details, see Hennawi et al. 2025; Kist et al. 2025a). We reweight these samples based on the coverage tests performed in Kist et al. (2025a) to ensure that we are quoting statistically faithful constraints.

5.3 Local IGM damping wing constraints at $z \sim 7.5$

We now proceed by leveraging the framework summarized in the previous section to infer the first *local* IGM damping wing constraints for two of the highest-redshift quasars known to date. That is, J1007+2115 at $z = 7.51$ (Yang et al. 2020a) and J1342+0928 at $z = 7.54$ (Bañados et al. 2018) with absolute magnitudes of $M_{1450} = -26.82$ and $M_{1450} = -26.34$ at 1450 \AA in the rest frame, respectively, determined based on Euclid photometry (Euclid Collaboration et al. 2026).

This work presents the first damping wing analysis of JWST/NIRSpec spectra of these two objects (JWST GO 1764, PI: Fan, and JWST GTO 1219, PI: Luetzgendorf; see also Christensen et al. 2023; Hennawi et al. 2026). The spectra were taken with the NIRSpec Fixed Slits together with the G140H/F070LP and G235H/F170LP grating and filter combinations at a resolution of $R \simeq 2700$ and have signal-to-noise ratios of $S/N \simeq 10 - 25$ and $50 - 100$, respectively, in the smooth Lyman- α region. The data were reduced via a combination of the JWST Science calibration pipeline CALWEBB (version 1.13.4) and the python-based semi-automated reduction pipeline PypeIt (Prochaska et al. 2020). The full details of the reduction procedure are discussed in Hennawi et al. (2026).

The spectral coverage of these observations allows us to operate on a rest-frame wavelength range of $1175 \text{ \AA} \leq \lambda \leq 3000 \text{ \AA}$. After removing narrow absorbers, we rebin both spectra to a velocity pixel scale of $dv = 500 \text{ km/s}$

(and likewise are our forward-modeled IGM transmission profiles) as our current likelihood prescription does not allow us to extract all information from the full resolution data at its native pixel scale in a statistically faithful manner (Kist et al. 2025b,a) which is pending improvements in our likelihood prescription, for example through simulation-based inference (Chen 2024).

5.3.1 J1007+2115

We start by analyzing the quasar J1007+2115 at $z = 7.51$. Previous studies have inferred IGM neutral fractions of $\langle x_{\text{HI}} \rangle = 0.39^{+0.22}_{-0.13}$ (Yang et al. 2020a) and $\langle x_{\text{HI}} \rangle = 0.27^{+0.21}_{-0.17}$ (Greig et al. 2022) from ground-based spectra of this object. Note that the former analysis also obtained a constraint of $\log_{10} t_{\text{Q}}/\text{yr} = 4.89^{+1.15}_{-0.76}$ on its lifetime (Eilers et al. 2021), and used, upon some minor differences, the same simulation models as we do in this work, however with a significantly different analysis pipeline. Our re-analysis is based on the new JWST/NIRSpec spectrum of this object, and we perform the analysis both in the context of the conventional global as well as our local IGM damping wing parameterization. We depict the final rebinned spectrum that forms the input to our pipeline as the black line in Figure 5.1 with associated noise vector in yellow. The upper panel depicts our reconstruction in the context of the local parameterization, and the lower panel in the context of the global one. The median reconstructed continua are shown in blue, and the full models including IGM absorption in red. The shading around these lines highlights the associated 16% and 84% percentile variations, reflecting parameter uncertainty, continuum reconstruction errors, as well as spectral noise (for further details see Hennawi et al. 2025).

Overall, the models show a remarkable agreement, both providing a good fit to the spectrum across the entire spectral range, not only matching the shape of the proximity zone and the Lyman- α damping wing, but also the smooth emission lines redward of Lyman- α . The median models slightly undershoot the observed spectrum around $\lambda \approx 1225 \text{ \AA}$ (somewhat more pronounced in the global parameterization), and slightly overshoot at $\lambda \approx 1245 \text{ \AA}$. However, in both regions the observed value is still within the 68% region. Somewhat more redward, at $\lambda \approx 1265 \text{ \AA}$, three pixels got rejected by the sigma-clipping procedure we apply prior to the inference (see Kist et al. 2025a). These pixels are possibly affected by a mild broad absorption line (BAL) system but this does not appear to impact our conclusions as we obtain the same results when no clipping is applied.

The only main difference between the global and the local model curves is a reduced degree of scatter in the local parameterization as this parameterization excludes the scatter due to the stochastic distribution of neutral patches along the line-of-sight from the inference task. Both models only show a very mild damping wing. The black contours in Figure 5.2 depict the associated local parameter constraints. The lifetime posterior peaks relatively sharply




Figure 5.1: Inferred model for the JWST/NIRSpec spectrum of the quasar J1007+2115, fitted in the context of the local IGM damping wing parameterization (upper row) and the global one (lower row). The observed (and rebinned) spectrum of the quasar is depicted in black, with the noise vector shown in yellow. The inferred model spectrum is depicted in red and the unabsorbed inferred continuum in blue, where solid lines represent the median inferred models, and shaded regions the 16% and the 84% percentile variations reflecting parameter uncertainty, continuum reconstruction errors, as well as spectral noise.




Figure 5.2: Posterior distributions inferred from the JWST/NIRSpec spectrum of the quasar J1007+2115 depicted in Figure 5.1 in the context of the local IGM damping wing parameterization (black) and the global one (yellow). Both distributions are marginalized over 7 nuisance parameters describing the shape of the quasar continuum. Additionally depicted in green is the topology-informed version of the local constraints, entailing the non-trivial prior $P_{\text{top}}(\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2}, r_{\text{patch}})$ (explicitly depicted in purple in the extra panel), and also providing a constraint on the global IGM neutral fraction $\langle x_{\text{HI}} \rangle$ in good agreement with the directly inferred one.

at $\log_{10} t_Q/\text{yr} = 4.09_{-0.15}^{+0.36}$, implying that we are concerned with a comparably young object, as already hinted at by its small proximity zone (see e.g. also the objects in Eilers et al. 2020, 2021). The data are not overly constraining with respect to the two local parameters $N_{\text{HI}}^{\text{DW}}$ and r_{patch} . In essence, only the highest HI column densities (and correspondingly short neutral bubble distances that are still located in the physically permitted region of parameter space enclosed by the dashed line in the $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel) are excluded as these would lead to a noticeable damping wing imprint.

We obtain a global $\langle x_{\text{HI}} \rangle$ constraint on the timing of reionization based on these local constraints by following the procedure introduced in Kist et al. (2025a) and summarized in Section 5.2, folding in the distribution of our local parameters within the realistic semi-numerical reionization topology as a prior. The green contours in Figure 5.2 show the resulting posterior distribution. We see that the prior on $N_{\text{HI}}^{\text{DW}}$ and r_{patch} induced by this topology (shown explicitly in purple in the extra panel of the plot) disfavors the (low- $N_{\text{HI}}^{\text{DW}}$, low- r_{patch}) region of the posterior and therefore gives more weight to the (low- $N_{\text{HI}}^{\text{DW}}$, high- r_{patch}) and (high- $N_{\text{HI}}^{\text{DW}}$, low- r_{patch}) peaks.² The resulting $\langle x_{\text{HI}} \rangle$ posterior suggests a comparably low to intermediate neutral fraction of $\langle x_{\text{HI}} \rangle = 0.32_{-0.20}^{+0.22}$ at this redshift. The lifetime constraint largely remains unaffected by folding in the topology information, still peaking at $\log_{10} t_Q/\text{yr} = 4.14_{-0.18}^{+0.74}$, on the lower end of the literature value of $\log_{10} t_Q/\text{yr} = 4.89_{-0.76}^{+1.15}$ (Yang et al. 2020a; Eilers et al. 2021). Our local parameter constraints, on the other hand, are fully prior-dominated, as a comparison of the green and purple contours in the extra $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel in the same figure suggests.

We test the robustness of the $\langle x_{\text{HI}} \rangle$ constraint by also comparing it to the one we obtain when inferring $\langle x_{\text{HI}} \rangle$ and t_Q directly in the context of the global parameterization. The corresponding model curves are shown in the lower panel of Figure 5.1. The resulting posterior, depicted by the yellow contours in Figure 5.2, is in good agreement with the one obtained by converting our local constraints (green). The only differences between the two are that the converted posterior exhibits a somewhat more extended tail towards longer lifetimes. Since these longer lifetimes go hand in hand with a higher IGM neutral fraction, the marginal $\langle x_{\text{HI}} \rangle$ posterior obtained via the local parameterization (green) shows a very mild peak at around $\langle x_{\text{HI}} \rangle \simeq 0.3$, in contrast to the directly inferred one which peaks very close to $\langle x_{\text{HI}} \rangle \simeq 0$. Overall, however, both posteriors are in clear statistical agreement,

²Note that the green converted posterior seemingly has support in regions where the black unconverted one does not. In fact, *both* posteriors have support in these regions, but before performing the conversion, these are not included in the 95% contour shown in the plot. Further, due to smoothing effects and the statistical reweighting procedure we apply, the contours marginally extend beyond the physical $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ boundary enclosed by the dashed line.

underlining the robustness of our constraints despite the highly different ways in which they were obtained.

5.3.2 J1342+0928

The second object of our study, the quasar J1342+0928 at $z = 7.54$, has already been extensively analyzed in previous literature (Bañados et al. 2018; Davies et al. 2018a; Greig et al. 2019; Ďurovčíková et al. 2020; Reiman et al. 2020; Greig et al. 2022), resulting in varying constraints on the IGM neutral fraction, with median values ranging from $\langle x_{\text{HI}} \rangle \approx 0.2 - 0.6$. We again apply both versions of our inference framework to the JWST/NIRSpec spectrum of this object and obtain the fits and parameter constraints shown in Figures 5.3 and 5.4, again in good statistical agreement with a reduced degree of scatter in the local parameterization. Unlike in the case of J1007+2115, we infer the clear presence of a damping wing with a correspondingly high HI column density of $\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2} = 19.99^{+0.27}_{-0.40}$ and a rather short distance of $r_{\text{patch}} = 12.0^{+5.0}_{-7.4}$ cMpc to the first neutral patch, increasing/reducing further to $\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2} = 20.24^{+0.25}_{-0.22}$ and $r_{\text{patch}} = 10.9^{+5.6}_{-5.9}$ after folding in the topology dependence (green contours). The lifetime posterior shows a clear degeneracy with $N_{\text{HI}}^{\text{DW}}$ in the local parameterization, peaking at $\log_{10} t_{\text{Q}}/\text{yr} = 5.64^{+0.25}_{-0.43}$, and with $\langle x_{\text{HI}} \rangle$ in the global one, preferring a somewhat lower value of $\log_{10} t_{\text{Q}}/\text{yr} = 4.69^{+1.40}_{-0.33}$. However, the global posterior shows a long axis of degeneracy with a second peak closer to the higher value preferred in the local model. Previous analyses inferred $\log_{10} t_{\text{Q}}/\text{yr} = 5.38^{+0.72}_{-1.30}$ with a similarly wide degeneracy (Davies et al. 2018a, 2019; Eilers et al. 2021). The corresponding $\langle x_{\text{HI}} \rangle$ constraints suggest an intermediate to high IGM neutral fraction of $\langle x_{\text{HI}} \rangle = 0.58^{+0.23}_{-0.23}$ in the context of the local parameterization, and a very high one of $\langle x_{\text{HI}} \rangle = 0.78^{+0.12}_{-0.30}$ in the context of the global model.

These results ought to be treated with caution since Davies et al. (2025) recently pointed out the possibility of contamination of the IGM in front of J1342+0928 by proximate DLA absorbers located at $r_{\text{abs}} \simeq 21$ cMpc, 31 cMpc and/or 56 cMpc based on weak Mg II absorption lines they identified in the JWST/NIRSpec spectrum that is also the subject of this analysis. While the authors concluded that such proximate absorbers would have to be unusually metal-poor, we use this possibility to highlight an additional virtue of our local parameterization. That is, our modeling framework allows us to naturally relate the (Lorentzian-weighted) HI column density $N_{\text{HI}}^{\text{DW}}$ to the column density $N_{\text{HI}}^{\text{abs}}$ of a putative absorber. This offers possibilities of 1) retrospectively correcting an inferred $N_{\text{HI}}^{\text{DW}}$ constraint for a contribution $N_{\text{HI}}^{\text{abs}}$ from a potential local absorber, constrained by other means, or even 2) *jointly* inferring $N_{\text{HI}}^{\text{DW}}$ and $N_{\text{HI}}^{\text{abs}}$ in a fully Bayesian way. We discuss the former approach below and leave an implementation of the latter one to future work.




Figure 5.3: Inferred models like in Figure 5.1 but for the JWST/NIRSpec spectrum of the quasar J1342+0928.




Figure 5.4: Like Figure 5.2 but for the posterior distributions inferred from the JWST/NIRSpec spectrum of the quasar J1342+0928 depicted in Figure 5.3. The $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel also shows the putative local absorber constraints obtained by Davies et al. (2025) assuming two different metallicities which provide lower limits on $N_{\text{HI}}^{\text{DW}}$.

To that end, note that the total column density $N_{\text{HI}}^{\text{DW}}$ giving rise to the practically observed damping wing signal comprises both the contribution from the IGM and that from any potential local absorber. Under the assumption of a cosmological HI density field $n_{\text{HI}}^{\text{IGM}}(r) = x_{\text{HI}}(r) \cdot \Delta(r)$ sourced by the IGM, superposed by the absorber's HI density $n_{\text{HI}}^{\text{abs}}(r)$, it directly follows from Eq. (5.1) that we can break up $N_{\text{HI}}^{\text{DW}}$ into their respective contributions:

$$N_{\text{HI}}^{\text{DW}} = N_{\text{HI}}^{\text{DW,IGM}} + N_{\text{HI}}^{\text{DW,abs}}. \quad (5.4)$$

By approximating the localized absorber's density field $n_{\text{HI}}^{\text{abs}}(r)$ as a Dirac delta peak at position r_{abs} , we can straightforwardly use Eq. (5.1) to determine $N_{\text{HI}}^{\text{DW,abs}}$ based on its classical column density $N_{\text{HI}}^{\text{abs}}$:

$$\log N_{\text{HI}}^{\text{DW,abs}} = \log N_{\text{HI}}^{\text{abs}} - 2 \log(r_{\text{abs}}/r_{\text{T}} + 1), \quad (5.5)$$

provided the absorber is located within the integration range for $N_{\text{HI}}^{\text{DW}}$ enclosed by r_{min} and r_{max} .

As per Eq. (5.4), any inferred value of $N_{\text{HI}}^{\text{DW}}$ only constitutes an upper limit on the column density $N_{\text{HI}}^{\text{DW,IGM}}$ sourced by the IGM (and thus relevant to reionization), from which the contributions from any potential absorber would have to be subtracted. Vice versa, the column density $N_{\text{HI}}^{\text{DW,abs}}$ of a given absorber places a lower limit on the full $N_{\text{HI}}^{\text{DW}}$. We plot these limits based on the values determined by Davies et al. (2025) assuming oxygen abundances relative to solar of $[\text{O}/\text{H}] = -2.5$ (-3.5) as light (dark) blue arrows in the $(N_{\text{HI}}^{\text{DW}}, r_{\text{patch}})$ panel of Figure 5.4. Note that the length of the arrows carries no information and we place these constraints at $r_{\text{patch}} = r_{\text{abs}}$ for each absorber even though the two distances are not necessarily related given that we define r_{patch} as the first *extended* neutral patch in front of the quasar.

Note that the $\log N_{\text{HI}}^{\text{DW,abs}}$ constraint from the closest absorber in the $[\text{O}/\text{H}] = -3.5$ case aligns well with the peak of the posterior. This is in excellent agreement with the fact that a such metal-poor absorber at this distance would solely account for the observed damping wing imprint, leaving no room for any additional IGM contribution when subtracted from the total column density $N_{\text{HI}}^{\text{DW}}$. The Lorentzian-weighted HI column density of an $[\text{O}/\text{H}] = -2.5$ absorber, on the other hand, would largely be negligible compared to our inferred $N_{\text{HI}}^{\text{DW}}$, and hence not bias our conclusions about the ionization state of the IGM. These simple considerations also pave the way for a principled, *joint* inference of the damping wing signal due to the IGM and potential local absorbers. To that end, one could follow Davies et al. (2025) and infer an absorber's column density $N_{\text{HI}}^{\text{abs}}$ based on O I and Mg II absorption line fitting (with $[\text{O}/\text{H}]$ as an additional parameter). Eqs. (5.4) and (5.5) then establish the connection between the absorber's column density $N_{\text{HI}}^{\text{abs}}$ and the full damping wing column density $N_{\text{HI}}^{\text{DW}}$, forming the




Figure 5.5: Reionization history constraints obtained from the two quasars J1007+2115 at $z = 7.51$ and J1342+0928 at $z = 7.54$ in the context of the local IGM damping wing parameterization (green stars) and the global one (yellow pentagons). Other symbols show existing literature constraints for the same two objects, slightly offset with respect to each other for clarity, and gray swatches are 68 % and 95 % contours of the [Planck Collaboration et al. \(2020\)](#) CMB constraints.

basis of a full Bayesian framework that would jointly fit for the strength of the IGM damping wing and potential local absorption systems.

5.4 Summary and Conclusions

We summarize in Figure 5.5 our constraints on the reionization history obtained by analyzing the JWST/NIRSpec spectra of the two $z \sim 7.5$ quasars J1007+2115 and J1342+0928 in the context of the local damping wing parameterization (green stars) put forward in [Kist et al. \(2025c,a\)](#) and the conventional global one (yellow pentagons). We compare these constraints to the CMB constraint on the reionization optical depth inferred by the [Planck Collaboration et al. \(2020\)](#) whose 68 %- and 95 %-contours are marked as gray swatches, as well as previous IGM damping wing measurements of the two objects studied in this work ([Bañados et al. 2018](#); [Davies et al. 2018a](#); [Yang et al. 2020a](#); [Reiman et al. 2020](#); [Āurovčíková et al. 2020](#); [Greig et al. 2022](#)). Our constraints are in clear statistical agreement with all these literature constraints, also owing to the comparably large statistical uncertainties that are unavoidable for individual objects due to the stochastic nature of reionization ([Kist et al. 2025b,c,a](#)). Note that despite the methodological improvements in our pipeline as compared to previous approaches ([Hennawi et al. 2025](#); [Kist et al. 2025a](#)), our uncertainties are not necessarily smaller

than those quoted in the literature. This is explained by the fact that we rigorously folded in all relevant uncertainties due to stochastic reionization, the unknown quasar lifetime, continuum reconstruction, and spectral noise.

We conclude that J1007+2115 prefers a somewhat lower neutral fraction of $\langle x_{\text{HI}} \rangle = 0.32^{+0.22}_{-0.20}$, compared to $\langle x_{\text{HI}} \rangle = 0.58^{+0.23}_{-0.23}$ based on J1342+0928 (both inferred in the framework of our local parameterization). The directly inferred global $\langle x_{\text{HI}} \rangle$ posterior of the latter object ($\langle x_{\text{HI}} \rangle = 0.78^{+0.12}_{-0.30}$) peaks at the high end compared to our local and most literature damping wing constraints (though closer to [Planck Collaboration et al. \(2020\)](#)), but note here the extended axis of degeneracy of the full posterior of this object (c.f. Figure 5.4). In addition, *all* damping wing-based $\langle x_{\text{HI}} \rangle$ constraints would be biased high in case an unusually metal-poor absorber was indeed present in the foreground of this object ([Davies et al. 2025](#)) which would pull the actual corrected IGM neutral fraction closer to the value inferred from J1007+2115. Our local parameterization provides a natural framework for studying this possibility more carefully in future work, particularly relevant also to JWST observations of IGM damping wings towards galaxies ([Mason et al. 2026](#); [Huberty et al. 2025](#)). More conclusive statements about the ionization state of the IGM at $z \sim 7.5$ will be enabled by applying our robust inference approach to the spectra of additional objects identified at these redshifts by the Euclid wide field survey ([Euclid Collaboration et al. 2019](#); [Bañados et al. 2025](#); [Euclid Collaboration et al. 2026](#)).

In addition, we constrained the lifetimes of the two objects of this study and found J1007+2115 to be remarkably young with $\log_{10} t_{\text{Q}}/\text{yr} = 4.14^{+0.74}_{-0.18}$, whereas J1342+0928 (with $\log_{10} t_{\text{Q}}/\text{yr} = 5.64^{+0.25}_{-0.43}$) is closer to the average expected lifetimes of $t_{\text{Q}} \sim 10^6$ yr ([Khrykin et al. 2021](#); [Morey et al. 2021](#)). Most importantly, we also obtained the first constraints on the local ionization topology in front of these two sources. With a Lorentzian-weighted HI column density of $\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2} = 19.70^{+0.35}_{-0.86}$ and a first neutral patch at $r_{\text{patch}} = 28.9^{+54.0}_{-14.4}$ cMpc, we remained in the prior-dominated regime for the pre-quasar sightline originating from J1007+2115, but we measured $\log_{10} N_{\text{HI}}^{\text{DW}}/\text{cm}^{-2} = 20.24^{+0.25}_{-0.22}$ and $r_{\text{patch}} = 10.9^{+5.6}_{-5.9}$ cMpc for J1342+0928. While this work was focused on the conversion of these measurements to constraints on the global timing on reionization, constraining $N_{\text{HI}}^{\text{DW}}$ and r_{patch} for larger statistical samples of objects will also open the doors to direct constraints on the topology of reionization with quasar IGM damping wings ([Sharma et al. 2025](#)).

Acknowledgements

We acknowledge helpful conversations with the ENIGMA group at UC Santa Barbara and Leiden University. This work is based on observations made with the NASA/ESA/CSA James Webb Space Telescope. The data were ob-

tained from the Mikulski Archive for Space Telescopes at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127 for JWST. These observations are associated with programs #1219 and #1764. 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. SEIB is supported by the Deutsche Forschungsgemeinschaft (DFG) under Emmy Noether grant number BO 5771/1-1. FW acknowledges support from NSF award AST-2513040.

Data Availability

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

BIBLIOGRAPHY

- Abazajian K. N., et al., 2016, [arXiv e-prints](#), p. arXiv:1610.02743
- Abdalla E., et al., 2022, [Journal of High Energy Astrophysics](#), 34, 49
- Abdul Karim M., et al., 2025, [Phys. Rev. D](#), 112, 083515
- Abdurashidova Z., et al., 2022, [ApJ](#), 925, 221
- Adame A. G., et al., 2025, [J. Cosmology Astropart. Phys.](#), 2025, 021
- Adams N. J., et al., 2024, [ApJ](#), 965, 169
- Ade P., et al., 2019, [J. Cosmology Astropart. Phys.](#), 2019, 056
- Alam S., et al., 2021, [Phys. Rev. D](#), 103, 083533
- Albrecht A., Steinhardt P. J., 1982, [Phys. Rev. Lett.](#), 48, 1220
- Almgren A. S., Bell J. B., Lijewski M. J., Lukić Z., Van Andel E., 2013, [ApJ](#), 765, 39
- Alpher R. A., Bethe H., Gamow G., 1948, [Physical Review](#), 73, 803
- Alsing J., Charnock T., Feeney S., Wandelt B., 2019, [MNRAS](#), 488, 4440
- Asthana S., Haehnelt M. G., Kulkarni G., Aubert D., Bolton J. S., Keating L. C., 2024, [MNRAS](#), 533, 2843
- Astropy Collaboration et al., 2013, [A&A](#), 558, A33
- Astropy Collaboration et al., 2018, [AJ](#), 156, 123
- Astropy Collaboration et al., 2022, [ApJ](#), 935, 167
- Atek H., et al., 2024, [Nature](#), 626, 975
- Aubert D., Teyssier R., 2010, [ApJ](#), 724, 244
- Bach K., Lee H.-W., 2015, [MNRAS](#), 446, 264
- Bayes T., Price R., 1763, *Philosophical Transactions of the Royal Society of London Series I*, 53, 370
- Bañados E., et al., 2018, [Nature](#), 553, 473
- Bañados E., et al., 2025, [MNRAS](#), 542, 1088
- Becker G. D., Bolton J. S., Haehnelt M. G., Sargent W. L. W., 2011, [MNRAS](#), 410, 1096
- Becker G. D., Bolton J. S., Madau P., Pettini M., Ryan-Weber E. V., Venemans B. P., 2015, [MNRAS](#), 447, 3402
- Becker G. D., D'Aloisio A., Christenson H. M., Zhu Y., Worseck G., Bolton J. S., 2021, [MNRAS](#), 508, 1853
- Becker G. D., Bolton J. S., Zhu Y., Hashemi S., 2024, [MNRAS](#), 533, 1525
- Berestetskii V. B., Lifshitz E. M., Pitaevskii V. B., 1971, *Relativistic quantum theory. Pt.1*
- Beringue B., et al., 2025, [arXiv e-prints](#), p. arXiv:2506.06274
- Bethe H. A., Salpeter E. E., 1957, *Quantum Mechanics of One- and Two-Electron Atoms*
- Bingham E., et al., 2018, [arXiv e-prints](#), p. arXiv:1810.09538
- Boera E., Becker G. D., Bolton J. S., Nasir F., 2019, [ApJ](#), 872, 101
- Bogdán Á., et al., 2024, [Nature Astronomy](#), 8, 126

- Bolan P., et al., 2022, *MNRAS*, 517, 3263
- Bolton J. S., Haehnelt M. G., 2007a, *MNRAS*, 374, 493
- Bolton J. S., Haehnelt M. G., 2007b, *MNRAS*, 382, 325
- Bolton J. S., Becker G. D., Wyithe J. S. B., Haehnelt M. G., Sargent W. L. W., 2010, *MNRAS*, 406, 612
- Bolton J. S., Haehnelt M. G., Warren S. J., Hewett P. C., Mortlock D. J., Venemans B. P., McMahon R. G., Simpson C., 2011, *MNRAS*, 416, L70
- Bolton J. S., Becker G. D., Raskutti S., Wyithe J. S. B., Haehnelt M. G., Sargent W. L. W., 2012, *MNRAS*, 419, 2880
- Bolton J. S., Puchwein E., Sijacki D., Haehnelt M. G., Kim T.-S., Meiksin A., Regan J. A., Viel M., 2017, *MNRAS*, 464, 897
- Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, *ApJ*, 379, 440
- Bondi H., Gold T., 1948, *MNRAS*, 108, 252
- Born M., Jordan P., 1925, *Zeitschrift fur Physik*, 34, 858
- Bosman S. E. I., 2021, *arXiv e-prints*, p. arXiv:2108.12446
- Bosman S. E. I., Fan X., Jiang L., Reed S., Matsuoka Y., Becker G., Haehnelt M., 2018, *MNRAS*, 479, 1055
- Bosman S. E. I., et al., 2022, *MNRAS*, 514, 55
- Bouwens R. J., et al., 2015, *ApJ*, 803, 34
- Bouwens R. J., et al., 2021, *AJ*, 162, 47
- Bradbury J., et al., 2018, JAX: composable transformations of Python+NumPy programs, <http://github.com/google/jax>
- Bromm V., 2013, *Reports on Progress in Physics*, 76, 112901
- Bruton S., Lin Y.-H., Scarlata C., Hayes M. J., 2023, *ApJ*, 949, L40
- Bunker A. J., et al., 2023, *A&A*, 677, A88
- Busca N. G., et al., 2013, *A&A*, 552, A96
- Carniani S., et al., 2024, *Nature*, 633, 318
- Cen R., Haiman Z., 2000, *ApJ*, 542, L75
- Chabanier S., et al., 2019, *J. Cosmology Astropart. Phys.*, 2019, 017
- Chen H., 2024, *MNRAS*, 528, L33
- Chen H., Gnedin N. Y., 2021, *ApJ*, 911, 60
- Cheng C., et al., 2018, *ApJ*, 868, 26
- Chiu H.-Y., 1964, *Physics Today*, 17, 21
- Chornock R., Berger E., Fox D. B., Lunnan R., Drout M. R., Fong W.-f., Laskar T., Roth K. C., 2013, *ApJ*, 774, 26
- Choudhury T. R., Paranjape A., Bosman S. E. I., 2021, *MNRAS*, 501, 5782
- Christensen L., et al., 2023, *A&A*, 680, A82
- Collette A., 2013, Python and HDF5. O'Reilly
- Cook W. G., Glushchenko I. A., Ijjas A., Pretorius F., Steinhardt P. J., 2020, *Physics Letters B*, 808, 135690
- Cranmer K., Brehmer J., Louppe G., 2020, *Proceedings of the National Academy of Science*, 117, 30055
- Croft R. A. C., Weinberg D. H., Katz N., Hernquist L., 1998, *ApJ*, 495, 44

- Croft R. A. C., Weinberg D. H., Bolte M., Burles S., Hernquist L., Katz N., Kirkman D., Tytler D., 2002, *ApJ*, 581, 20
- Curtis-Lake E., et al., 2023, *Nature Astronomy*, 7, 622
- D’Aloisio A., McQuinn M., Davies F. B., Furlanetto S. R., 2018, *MNRAS*, 473, 560
- D’Odorico V., et al., 2023, *MNRAS*, 523, 1399
- Davies F. B., Furlanetto S. R., 2016, *MNRAS*, 460, 1328
- Davies F. B., Furlanetto S. R., 2022, *MNRAS*, 514, 1302
- Davies F. B., Hennawi J. F., 2023, *arXiv e-prints*, p. arXiv:2312.06763
- Davies F. B., Furlanetto S. R., McQuinn M., 2016, *MNRAS*, 457, 3006
- Davies F. B., et al., 2018a, *ApJ*, 864, 142
- Davies F. B., et al., 2018b, *ApJ*, 864, 143
- Davies F. B., Hennawi J. F., Eilers A.-C., 2019, *ApJ*, 884, L19
- Davies F. B., Hennawi J. F., Eilers A.-C., 2020, *MNRAS*, 493, 1330
- Davies F. B., Bañados E., Hennawi J. F., Bosman S. E. I., 2025, *ApJ*, 989, L27
- Davies F. B., et al., 2026, *MNRAS*, 545, staf1862
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
- Delubac T., et al., 2015, *A&A*, 574, A59
- Donnan C. T., et al., 2024, *MNRAS*, 533, 3222
- Draine B. T., 2011, *Physics of the Interstellar and Intergalactic Medium*
- Duane S., Kennedy A. D., Pendleton B. J., Roweth D., 1987, *Physics Letters B*, 195, 216
- Ďurovčíková D., Katz H., Bosman S. E. I., Davies F. B., Devriendt J., Slyz A., 2020, *MNRAS*, 493, 4256
- Ďurovčíková D., et al., 2024, *ApJ*, 969, 162
- Dyson F. J., 1949, *Physical Review*, 75, 486
- Edge D. O., Shakeshaft J. R., McAdam W. B., Baldwin J. E., Archer S., 1959, *Mem. RAS*, 68, 37
- Eilers A.-C., Davies F. B., Hennawi J. F., Prochaska J. X., Lukić Z., Mazzucchelli C., 2017, *ApJ*, 840, 24
- Eilers A.-C., Davies F. B., Hennawi J. F., 2018, *ApJ*, 864, 53
- Eilers A.-C., et al., 2020, *ApJ*, 900, 37
- Eilers A.-C., Hennawi J. F., Davies F. B., Simcoe R. A., 2021, *ApJ*, 917, 38
- Einasto J., Kaasik A., Saar E., 1974, *Nature*, 250, 309
- Einstein A., 1916, *Annalen der Physik*, 354, 769
- Einstein A., 1917, *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pp 142–152
- Elbers W., 2025, *arXiv e-prints*, p. arXiv:2508.21069

- Elias J. H., Joyce R. R., Liang M., Muller G. P., Hileman E. A., George J. R., 2006, in McLean I. S., Iye M., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 6269, Ground-based and Airborne Instrumentation for Astronomy. p. 62694C, [doi:10.1117/12.671817](https://doi.org/10.1117/12.671817)
- Endsley R., Stark D. P., Whittler L., Topping M. W., Chen Z., Plat A., Chisholm J., Charlot S., 2023, *MNRAS*, 524, 2312
- Englert F., Brout R., 1964, *Phys. Rev. Lett.*, 13, 321
- Euclid Collaboration et al., 2019, *A&A*, 631, A85
- Euclid Collaboration et al., 2026, in prep.
- Event Horizon Telescope Collaboration et al., 2019, *ApJ*, 875, L1
- Event Horizon Telescope Collaboration et al., 2022, *ApJ*, 930, L12
- Fan X., et al., 2000, *AJ*, 120, 1167
- Fan X., Narayanan V. K., Strauss M. A., White R. L., Becker R. H., Pentericci L., Rix H.-W., 2002, *AJ*, 123, 1247
- Fan X., et al., 2006, *AJ*, 132, 117
- Fan X., Bañados E., Simcoe R. A., 2023, *ARA&A*, 61, 373
- Faucher-Giguère C.-A., Prochaska J. X., Lidz A., Hernquist L., Zaldarriaga M., 2008, *ApJ*, 681, 831
- Fausey H. M., et al., 2025a, *MNRAS*, 536, 2839
- Fausey H. M., van der Horst A. J., Tanvir N. R., Wiersema K., Fynbo J. P. U., Hartmann D., de Ugarte Postigo A., 2025b, *ApJ*, 985, 28
- Feynman R. P., 1948, *Reviews of Modern Physics*, 20, 367
- Finkelstein S. L., et al., 2019, *ApJ*, 879, 36
- Font-Ribera A., et al., 2014, *J. Cosmology Astropart. Phys.*, 2014, 027
- Foreman-Mackey D., 2016, *The Journal of Open Source Software*, 1, 24
- Friedmann A., 1922, *Zeitschrift für Physik*, 10, 377
- Fumagalli M., O'Meara J. M., Prochaska J. X., Worseck G., 2013, *ApJ*, 775, 78
- Furlanetto S. R., Oh S. P., 2005, *MNRAS*, 363, 1031
- Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004, *ApJ*, 613, 1
- Furlanetto S. R., Oh S. P., Briggs F. H., 2006, *Phys. Rep.*, 433, 181
- Gaikwad P., et al., 2020, *MNRAS*, 494, 5091
- Gaikwad P., Srianand R., Haehnelt M. G., Choudhury T. R., 2021, *MNRAS*, 506, 4389
- Gaikwad P., et al., 2023, *MNRAS*, 525, 4093
- Garaldi E., Kannan R., Smith A., Springel V., Pakmor R., Vogelsberger M., Hernquist L., 2022, *MNRAS*, 512, 4909
- Garcia-Gallego O., Iršič V., Haehnelt M. G., Bolton J. S., 2025a, *arXiv e-prints*, p. arXiv:2510.00107
- Garcia-Gallego O., Iršič V., Haehnelt M. G., Viel M., Bolton J. S., 2025b, *Phys. Rev. D*, 112, 043502
- Garzilli A., Boyarsky A., Ruchayskiy O., 2017, *Physics Letters B*, 773, 258
- Gnedin N. Y., 2014, *ApJ*, 793, 29

- Gnedin N. Y., Madau P., 2022, *Living Reviews in Computational Astrophysics*, 8, 3
- Gnedin N. Y., Ostriker J. P., 1997, *ApJ*, 486, 581
- Gnedin N., Zhu H., 2025, *The Open Journal of Astrophysics*, 8, 111
- Goto H., et al., 2021, *ApJ*, 923, 229
- Greig B., Mesinger A., McGreer I. D., Gallerani S., Haiman Z., 2017a, *MNRAS*, 466, 1814
- Greig B., Mesinger A., Haiman Z., Simcoe R. A., 2017b, *MNRAS*, 466, 4239
- Greig B., Mesinger A., Bañados E., 2019, *MNRAS*, 484, 5094
- Greig B., Mesinger A., Davies F. B., Wang F., Yang J., Hennawi J. F., 2022, *MNRAS*, 512, 5390
- Greig B., et al., 2024a, *MNRAS*, 530, 3208
- Greig B., et al., 2024b, *MNRAS*, 533, 3312
- Gruzinov A., Hu W., 1998, *ApJ*, 508, 435
- Gunn J. E., Peterson B. A., 1965, *ApJ*, 142, 1633
- Guralnik G. S., Hagen C. R., Kibble T. W., 1964, *Phys. Rev. Lett.*, 13, 585
- Guth A. H., 1981, *Phys. Rev. D*, 23, 347
- Haardt F., Madau P., 1996, *ApJ*, 461, 20
- Haardt F., Madau P., 2012, *ApJ*, 746, 125
- Hamilton W. R., 1834, *Philosophical Transactions of the Royal Society of London*, pp 247–308
- Hamilton W. R., 1835, *Philosophical Transactions of the Royal Society of London*, pp 95–144
- Harris C. R., et al., 2020, *Nature*, 585, 357
- Hartoog O. E., et al., 2015, *A&A*, 580, A139
- Hastings W. K., 1970, *Biometrika*, 57, 97
- Hayes M. J., Scarlata C., 2023, *ApJ*, 954, L14
- Hazard C., Mackey M. B., Shimmins A. J., 1963, *Nature*, 197, 1037
- Heintz K. E., et al., 2024, *Science*, 384, 890
- Heintz K. E., et al., 2025, *A&A*, 693, A60
- Heisenberg W., 1925, *Zeitschrift fur Physik*, 33, 879
- Heisenberg W., 1927, *Zeitschrift fur Physik*, 43, 172
- Heitler W., 1954, Quantum theory of radiation
- Hennawi J. F., Kist T., Davies F. B., Tamanas J., 2025, *MNRAS*, 539, 2621
- Hennawi J. F., et al., 2026, in prep.
- Higgs P. W., 1964, *Phys. Rev. Lett.*, 13, 508
- Hinshaw G., et al., 2013, *ApJS*, 208, 19
- Hoag A., et al., 2019, *ApJ*, 878, 12
- Hoffman M. D., Gelman A., 2011, *arXiv e-prints*, p. arXiv:1111.4246
- Hoyle F., 1948, *MNRAS*, 108, 372
- Hsiao T. Y.-Y., et al., 2024, *ApJ*, 973, 8
- Hu W., et al., 2021, *Nature Astronomy*, 5, 485
- Hubble E., 1929, *Proceedings of the National Academy of Science*, 15, 168
- Huberty M., Scarlata C., Hayes M. J., Gazagnes S., 2025, *ApJ*, 987, 82

- Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
- Hunter J. D., 2007, *Computing in Science and Engineering*, 9, 90
- Ijjas A., Steinhardt P. J., 2016, *Phys. Rev. Lett.*, 117, 121304
- Ijjas A., Steinhardt P. J., Loeb A., 2013, *Physics Letters B*, 723, 261
- Iliev I. T., Mellema G., Pen U.-L., Merz H., Shapiro P. R., Alvarez M. A., 2006, *MNRAS*, 369, 1625
- Inoue A. K., et al., 2018, *PASJ*, 70, 55
- Iršič V., et al., 2017a, *Phys. Rev. D*, 96, 023522
- Iršič V., Viel M., Haehnelt M. G., Bolton J. S., Becker G. D., 2017b, *Phys. Rev. Lett.*, 119, 031302
- Ishimoto R., et al., 2020, *ApJ*, 903, 60
- Ivanov M. M., Toomey M. W., Karaçaylı N. G., 2025, *Phys. Rev. Lett.*, 134, 091001
- Jhaveri T., Karwal T., Hu W., 2025, *Phys. Rev. D*, 112, 043541
- Jin X., et al., 2023, *ApJ*, 942, 59
- Jones G. C., et al., 2025, *MNRAS*, 536, 2355
- Jung I., et al., 2020, *ApJ*, 904, 144
- Kageura Y., et al., 2025, *ApJS*, 278, 33
- Kageura Y., Ouchi M., Naokawa F., Umeda H., Matsumoto A., Harikane Y., Nakane M., Thai T. T., 2026, *arXiv e-prints*, p. arXiv:2601.09644
- Kannan R., Garaldi E., Smith A., Pakmor R., Springel V., Vogelsberger M., Hernquist L., 2022, *MNRAS*, 511, 4005
- Kant I., 1784, *Berlinische Monatsschrift*, Dezember-Heft, 481
- Kapteyn J. C., 1922, *ApJ*, 55, 302
- Karaçaylı N. G., Font-Ribera A., Padmanabhan N., 2020, *MNRAS*, 497, 4742
- Karaçaylı N. G., et al., 2025, *J. Cosmology Astropart. Phys.*, 2025, 004
- Karzas W. J., Latter R., 1961, *ApJS*, 6, 167
- Kaurov A. A., 2016, *ApJ*, 831, 198
- Kaurov A. A., Gnedin N. Y., 2014, *ApJ*, 787, 146
- Keating L. C., Puchwein E., Bolton J. S., Haehnelt M. G., Kulkarni G., 2024a, *MNRAS*, 531, L34
- Keating L. C., Bolton J. S., Cullen F., Haehnelt M. G., Puchwein E., Kulkarni G., 2024b, *MNRAS*, 532, 1646
- Khrykin I. S., Hennawi J. F., McQuinn M., Worseck G., 2016, *ApJ*, 824, 133
- Khrykin I. S., Hennawi J. F., Worseck G., Davies F. B., 2021, *MNRAS*, 505, 649
- Kist T., Ijjas A., 2022, *J. Cosmology Astropart. Phys.*, 2022, 046
- Kist T., Hennawi J. F., Davies F. B., 2025a, *arXiv e-prints*, p. arXiv:2508.21812
- Kist T., Hennawi J. F., Davies F. B., 2025b, *MNRAS*, 538, 2704
- Kist T., Hennawi J. F., Davies F. B., 2025c, *MNRAS*, 544, 2316
- Kist T., et al., 2026, *MNRAS*, 545, staf2219
- Knox L., Scoccimarro R., Dodelson S., 1998, *Phys. Rev. Lett.*, 81, 2004

- Kogut A., et al., 2003, *ApJS*, 148, 161
- Kokorev V., et al., 2023, *ApJ*, 957, L7
- Konno A., et al., 2014, *ApJ*, 797, 16
- Konno A., et al., 2018, *PASJ*, 70, S16
- Koopmans L., et al., 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*. p. 1, doi:10.22323/1.215.0001
- Lamb D. Q., Reichart D. E., 2000, *ApJ*, 536, 1
- Laplace P. S., 1812, *Théorie analytique des probabilités*. Courcier, Paris, <http://eudml.org/doc/203064>
- Laplace P.-S., 1814, *Essai philosophique sur les probabilités*. Courcier, Paris, <http://eudml.org/doc/203193>
- Laplace P. S., 1820, *Théorie analytique des probabilités*. Courcier, Paris, <http://eudml.org/doc/203444>
- Larson R. L., et al., 2023, *ApJ*, 953, L29
- Lee K.-G., et al., 2015, *ApJ*, 799, 196
- Lemaître G., 1927, *Annales de la Société Scientifique de Bruxelles*, 47, 49
- Leschinski K., 2021, *SkyCalc_ipy*: SkyCalc wrapper for interactive Python
- Lewis J. S. W., et al., 2022, *MNRAS*, 516, 3389
- Lidz A., Faucher-Giguère C.-A., Dall’Aglio A., McQuinn M., Fechner C., Zaldarriaga M., Hernquist L., Dutta S., 2010, *ApJ*, 718, 199
- Lidz A., Chang T.-C., Mas-Ribas L., Sun G., 2021, *ApJ*, 917, 58
- Linde A. D., 1982, *Physics Letters B*, 108, 389
- Liu A., Shaw J. R., 2020, *PASP*, 132, 062001
- Lu T.-Y., et al., 2024a, *arXiv e-prints*, p. arXiv:2411.04176
- Lu T.-Y., Mason C. A., Hutter A., Mesinger A., Qin Y., Stark D. P., Endsley R., 2024b, *MNRAS*, 528, 4872
- Lukić Z., Stark C. W., Nugent P., White M., Meiksin A. A., Almgren A., 2015, *MNRAS*, 446, 3697
- Madau P., Dickinson M., 2014, *ARA&A*, 52, 415
- Madau P., Haardt F., Rees M. J., 1999, *ApJ*, 514, 648
- Madau P., Giallongo E., Grazian A., Haardt F., 2024, *ApJ*, 971, 75
- Maiolino R., et al., 2024a, *Nature*, 627, 59
- Maiolino R., et al., 2024b, *A&A*, 691, A145
- Maitra S., Kulkarni G., Arora V., Viel M., Asthana S., Bolton J. S., Haehnelt M. G., Keating L., 2026, *arXiv e-prints*, p. arXiv:2601.16263
- Malloy M., Lidz A., 2015, *ApJ*, 799, 179
- Marshak N. X., Simotas K., Lukić Z., Park H., Ahrens J., Johnson C. R., 2025, *arXiv e-prints*, p. arXiv:2512.12466
- du Mas des Bourboux H., et al., 2017, *A&A*, 608, A130
- du Mas des Bourboux H., et al., 2020, *ApJ*, 901, 153
- Mason C. A., Gronke M., 2020, *MNRAS*, 499, 1395
- Mason C. A., Treu T., Dijkstra M., Mesinger A., Trenti M., Pentericci L., de Barros S., Vanzella E., 2018, *ApJ*, 856, 2
- Mason C. A., et al., 2019, *MNRAS*, 485, 3947

- Mason C. A., Chen Z., Stark D. P., Yi Lu T., Topping M., Tang M., 2026, *A&A*, 705, A114
- Matsumura T., et al., 2014, *Journal of Low Temperature Physics*, 176, 733
- Matthews B. M., et al., 2021, *ApJS*, 252, 15
- McDonald P., 2003, *ApJ*, 585, 34
- McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., Ostriker J. P., 2000, *ApJ*, 543, 1
- McDonald P., et al., 2006, *ApJS*, 163, 80
- McGreer I. D., Mesinger A., Fan X., 2011, *MNRAS*, 415, 3237
- McGreer I. D., Mesinger A., D’Odorico V., 2015, *MNRAS*, 447, 499
- Meiksin A., Madau P., 1993, *ApJ*, 412, 34
- Meiksin A., Tittley E. R., Brown C. K., 2010, *MNRAS*, 401, 77
- Melandri A., et al., 2015, *A&A*, 581, A86
- Mellema G., Iliev I. T., Alvarez M. A., Shapiro P. R., 2006, *New Astronomy*, 11, 374
- Mertens F. G., et al., 2020, *MNRAS*, 493, 1662
- Mesinger A., 2010, *MNRAS*, 407, 1328
- Mesinger A., Furlanetto S., Cen R., 2011, *MNRAS*, 411, 955
- Mesinger A., Aykutalp A., Vanzella E., Pentericci L., Ferrara A., Dijkstra M., 2015, *MNRAS*, 446, 566
- Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, *J. Chem. Phys.*, 21, 1087
- Miralda-Escudé J., Ostriker J. P., 1990, *ApJ*, 350, 1
- Miralda-Escudé J., 1998, *ApJ*, 501, 15
- Miralda-Escudé J., Haehnelt M., Rees M. J., 2000, *ApJ*, 530, 1
- Morales A. M., Mason C. A., Bruton S., Gronke M., Haardt F., Scarlata C., 2021, *ApJ*, 919, 120
- Morey K. A., Eilers A.-C., Davies F. B., Hennawi J. F., Simcoe R. A., 2021, *ApJ*, 921, 88
- Mortlock D., 2016, in Mesinger A., ed., *Astrophysics and Space Science Library Vol. 423, Understanding the Epoch of Cosmic Reionization: Challenges and Progress*. p. 187, doi:10.1007/978-3-319-21957-8_7
- Mortlock D. J., et al., 2011, *Nature*, 474, 616
- Muñoz J. B., Mirocha J., Chisholm J., Furlanetto S. R., Mason C., 2024, *MNRAS*, 535, L37
- Nakane M., et al., 2024, *ApJ*, 967, 28
- Neal R. M., 1996, in , *Bayesian Learning for Neural Networks*. Springer New York, New York, NY, pp 55–98, doi:10.1007/978-1-4612-0745-0_3, https://doi.org/10.1007/978-1-4612-0745-0_3
- Nikolić I., Mesinger A., Mason C. A., Lu T.-Y., Tang M., Prelogović D., Gagnon-Hartman S., Stark D. P., 2025, *A&A*, 699, A323
- Ning Y., Jiang L., Zheng Z.-Y., Wu J., 2022, *ApJ*, 926, 230
- O’Meara J. M., Prochaska J. X., Worseck G., Chen H.-W., Madau P., 2013, *ApJ*, 765, 137

- Ocvirk P., et al., 2016, *MNRAS*, 463, 1462
- Ocvirk P., et al., 2020, *MNRAS*, 496, 4087
- Ono Y., et al., 2012, *ApJ*, 744, 83
- Onorato S., et al., 2025, *MNRAS*, 540, 1308
- Oort J. H., 1932, *BAIN*, 6, 249
- Ostriker J. P., Peebles P. J. E., 1973, *ApJ*, 186, 467
- Ouchi M., et al., 2010, *ApJ*, 723, 869
- Ouchi M., et al., 2018, *PASJ*, 70, S13
- Padovani P., et al., 2017, *A&A Rev.*, 25, 2
- Pakmor R., et al., 2023, *MNRAS*, 524, 2539
- Palanque-Delabrouille N., et al., 2013, *A&A*, 559, A85
- Palanque-Delabrouille N., et al., 2015, *J. Cosmology Astropart. Phys.*, 2015, 011
- Park H., et al., 2025, *ApJ*, 983, 91
- Patel M., Warren S. J., Mortlock D. J., Fynbo J. P. U., 2010, *A&A*, 512, L3
- Pawlik A. H., Schaye J., van Scherpenzeel E., 2009, *MNRAS*, 394, 1812
- Pedregosa F., et al., 2011, *Journal of Machine Learning Research*, 12, 2825
- Peebles P. J. E., 1993, *Principles of Physical Cosmology*, doi:10.1515/9780691206721.
- Penzias A. A., Wilson R. W., 1965, *ApJ*, 142, 419
- Perlmutter S., et al., 1999, *ApJ*, 517, 565
- Phan D., Pradhan N., Jankowiak M., 2019, *arXiv e-prints*, p. arXiv:1912.11554
- Pillepich A., et al., 2018, *MNRAS*, 473, 4077
- Planck Collaboration et al., 2014, *A&A*, 571, A16
- Planck Collaboration et al., 2016, *A&A*, 594, A13
- Planck Collaboration et al., 2020, *A&A*, 641, A6
- Poincare H., 1906, *Popular Astronomy*, 14, 475
- Press W. H., Schechter P., 1974, *ApJ*, 187, 425
- Pritchard J. R., Loeb A., 2012, *Reports on Progress in Physics*, 75, 086901
- Prochaska J. X., Worseck G., O'Meara J. M., 2009, *ApJ*, 705, L113
- Prochaska J., et al., 2020, *The Journal of Open Source Software*, 5, 2308
- Puchwein E., et al., 2023, *MNRAS*, 519, 6162
- Pâris I., et al., 2011, *A&A*, 530, A50
- Pérez F., Granger B. E., 2007, *Computing in Science and Engineering*, 9, 21
- Ravoux C., et al., 2025, *J. Cosmology Astropart. Phys.*, 2025, 079
- Reichardt C. L., 2016, in Mesinger A., ed., *Astrophysics and Space Science Library Vol. 423, Understanding the Epoch of Cosmic Reionization: Challenges and Progress*. p. 227, doi:10.1007/978-3-319-21957-8_8
- Reichardt C. L., et al., 2021, *ApJ*, 908, 199
- Reiman D. M., Tamanas J., Prochaska J. X., Ďurovčiková D., 2020, *arXiv e-prints*, p. arXiv:2006.00615
- Riess A. G., et al., 1998, *AJ*, 116, 1009

- Robertson B. E., Ellis R. S., Furlanetto S. R., Dunlop J. S., 2015, *ApJ*, 802, L19
- Robertson B., et al., 2024, *ApJ*, 970, 31
- Romano M., Grazian A., Giallongo E., Cristiani S., Fontanot F., Boutsia K., Fiore F., Menci N., 2019, *A&A*, 632, A45
- Rubin V. C., Ford Jr. W. K., 1970, *ApJ*, 159, 379
- Rudie G. C., Steidel C. C., Shapley A. E., Pettini M., 2013, *ApJ*, 769, 146
- Sailer N., Farren G. S., Ferraro S., White M., 2026, *Phys. Rev. Lett.*, 136, 081002
- Salpeter E. E., 1964, *ApJ*, 140, 796
- Satyavolu S., Kulkarni G., Keating L. C., Haehnelt M. G., 2023a, *MNRAS*, 521, 3108
- Satyavolu S., et al., 2023b, *MNRAS*, 522, 4918
- Sawyer F., Bolton J. S., Becker G. D., Conaboy L., Haehnelt M. G., Keating L., Kulkarni G., Puchwein E., 2025, *MNRAS*
- Saxena A., et al., 2024, *A&A*, 684, A84
- Schaye J., et al., 2015, *MNRAS*, 446, 521
- Schaye J., et al., 2023, *MNRAS*, 526, 4978
- Schenker M. A., Stark D. P., Ellis R. S., Robertson B. E., Dunlop J. S., McLure R. J., Kneib J.-P., Richard J., 2012, *ApJ*, 744, 179
- Schmidt M., 1963, *Nature*, 197, 1040
- Schroeder J., Mesinger A., Haiman Z., 2013, *MNRAS*, 428, 3058
- Schrödinger E., 1926a, *Annalen der Physik*, 384, 361
- Schrödinger E., 1926b, *Annalen der Physik*, 384, 489
- Schrödinger E., 1926c, *Annalen der Physik*, 385, 437
- Schrödinger E., 1926d, *Annalen der Physik*, 386, 109
- Schwinger J., 1948, *Physical Review*, 73, 416
- Seljak U., et al., 2005, *Phys. Rev. D*, 71, 103515
- Seljak U., Slosar A., McDonald P., 2006, *J. Cosmology Astropart. Phys.*, 2006, 014
- Sellentin E., Starck J.-L., 2019, *J. Cosmology Astropart. Phys.*, 2019, 021
- Sharma Y. M., Davies F. B., Gaikwad P., Nasir F., Bosman S. E. I., 2025, *ApJ*, 983, 118
- Shen Y., et al., 2007, *AJ*, 133, 2222
- Simmonds C., et al., 2024, *MNRAS*, 535, 2998
- Slipher V. M., 1913, *Lowell Observatory Bulletin*, 2, 56
- Slipher V. M., 1917, *Proceedings of the American Philosophical Society*, 56, 403
- Slosar A., et al., 2013, *J. Cosmology Astropart. Phys.*, 2013, 026
- So G. C., Norman M. L., Reynolds D. R., Wise J. H., 2014, *ApJ*, 789, 149
- Sobacchi E., Mesinger A., 2015, *MNRAS*, 453, 1843
- Songaila A., Cowie L. L., 2010, *ApJ*, 721, 1448
- Spina B., Bosman S. E. I., Davies F. B., Gaikwad P., Zhu Y., 2024, *A&A*, 688, L26

- Springel V., et al., 2005, *Nature*, 435, 629
- Stark D. P., Ellis R. S., Chiu K., Ouchi M., Bunker A., 2010, *MNRAS*, 408, 1628
- Stark D. P., Ellis R. S., Ouchi M., 2011, *ApJ*, 728, L2
- Sun Z., Ting Y.-S., Cai Z., 2023, *ApJS*, 269, 4
- Sunyaev R. A., Zeldovich I. B., 1980, *ARA&A*, 18, 537
- Suzuki N., 2006, *ApJS*, 163, 110
- Suzuki N., Tytler D., Kirkman D., O'Meara J. M., Lubin D., 2005, *ApJ*, 618, 592
- Tang M., Stark D. P., Topping M. W., Mason C., Ellis R. S., 2024, *ApJ*, 975, 208
- Tepper-García T., 2006, *MNRAS*, 369, 2025
- Thomson Baron Kelvin W., 1904, *Baltimore Lectures on Molecular Dynamics and the Wave Theory of Light*. Cambridge Library Collection - Physical Sciences, Cambridge University Press
- Tomonaga S., 1946, *Progress of Theoretical Physics*, 1, 27
- Torralba-Torregrosa A., et al., 2024, *A&A*, 689, A44
- Totani T., Kawai N., Kosugi G., Aoki K., Yamada T., Iye M., Ohta K., Hattori T., 2006, *PASJ*, 58, 485
- Totani T., et al., 2014, *PASJ*, 66, 63
- Totani T., Aoki K., Hattori T., Kawai N., 2016, *PASJ*, 68, 15
- Trac H., Chen N., Holst I., Alvarez M. A., Cen R., 2022, *ApJ*, 927, 186
- Trapp A. C., Furlanetto S. R., Davies F. B., 2023, *MNRAS*, 524, 5891
- Trott C. M., et al., 2020, *MNRAS*, 493, 4711
- Übler H., et al., 2024, *MNRAS*, 531, 355
- Umeda H., Ouchi M., Nakajima K., Harikane Y., Ono Y., Xu Y., Isobe Y., Zhang Y., 2024, *ApJ*, 971, 124
- Umeda H., et al., 2025, *ApJS*, 277, 37
- Umeda H., Ouchi M., Kageura Y., Harikane Y., Nakane M., Thai T. T., Nakajima K., 2026, *ApJ*, 997, 86
- Viel M., Haehnelt M. G., Springel V., 2004, *MNRAS*, 354, 684
- Viel M., Lesgourgues J., Haehnelt M. G., Matarrese S., Riotto A., 2005, *Phys. Rev. D*, 71, 063534
- Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, *Phys. Rev. D*, 88, 043502
- Villaseñor B., Robertson B., Madau P., Schneider E., 2022, *ApJ*, 933, 59
- Villaseñor B., Robertson B., Madau P., Schneider E., 2023, *Phys. Rev. D*, 108, 023502
- Virtanen P., et al., 2020, *Nature Methods*, 17, 261
- Vogelsberger M., et al., 2014, *MNRAS*, 444, 1518
- Voigt W., 1912, *Über das Gesetz der Intensitätsverteilung innerhalb der Linien eines Gasspektrums*. Sitzungsberichte Vol. 1912,25, München, <https://publikationen.badw.de/de/003395768>
- Walther M., Oñorbe J., Hennawi J. F., Lukić Z., 2019, *ApJ*, 872, 13

- Wang F., et al., 2020, *ApJ*, 896, 23
Weinberg S., 1989, *Reviews of Modern Physics*, 61, 1
Witstok J., et al., 2024, *A&A*, 682, A40
Witstok J., et al., 2025, *Nature*, 639, 897
Wold I. G. B., et al., 2022, *ApJ*, 927, 36
Wolfson M., Hennawi J. F., Davies F. B., Oñorbe J., 2023, *MNRAS*, 521, 4056
Worseck G., et al., 2014, *MNRAS*, 445, 1745
Xu Y., Yue B., Chen X., 2017, *ApJ*, 844, 117
Yang J., et al., 2020a, *ApJ*, 897, L14
Yang J., et al., 2020b, *ApJ*, 904, 26
Zel'dovich Y. B., 1970, *A&A*, 5, 84
Zheng Z.-Y., et al., 2017, *ApJ*, 842, L22
Zhou Y., Chen H., Matteo T. D., Ni Y., Croft R. A. C., Bird S., 2024, *MNRAS*, 528, 3730
Zhu Y., et al., 2022, *ApJ*, 932, 76
Zhu Y., et al., 2024, *MNRAS*, 533, L49
Zwicky F., 1933, *Helvetica Physica Acta*, 6, 110

ENGLISH SUMMARY

If I had to summarise the contents of this thesis in just two words, it would be its title: *Les Lumières*. Light is essential to us as astrophysicists, as we inevitably have to infer all information about the nature of our Universe from light³ that we manage to detect with our telescopes on the ground or in space. Particularly suitable for our purposes is a specific type of astrophysical lights that we call *quasars*. These astrophysical objects are shining so extraordinarily bright that they came to be known as *quasi-stellar objects* or *quasars*, even though their distances are anything but comparable to those of stars in our own galaxy. On the contrary, powered by supermassive black holes in the centres of distant galaxies, quasars are so exceptionally powerful and bright that we are still able to detect them at *cosmological* distances, out to redshifts $z \sim 7.5$. Given the time it took until their light that we are currently observing had reached Earth, looking at such distant sources means nothing else but peering into our Universe's first billion years.

A landmark event has shaped the evolution of our Universe at these early times: the *Epoch of Reionisation*. Initiated by the first stars and galaxies that started forming in these years, this era marked the end of the cosmic *dark ages* where the Universe remained devoid of any visible light. However, once these first astrophysical sources had started forming, their light quite literally illuminated our Universe: *Les Lumières* – a cosmic *Age of Enlightenment*. This light even had an impact on the vast voids in between these galaxies that we call the *intergalactic medium* (IGM) – the regions where most cosmic matter actually resides. Being exposed to these new sources of light, the neutral hydrogen atoms in the intergalactic medium successively got *re-ionised* by their radiation. As such, bubbles of ionised hydrogen formed around these ionising sources, and grew and merged until the entire intergalactic medium was reionised.

The progression of this process is characterised in terms of the fraction of neutral (as opposed to ionised) hydrogen atoms in the intergalactic medium. Precisely measuring the evolution of this fraction as a function of cosmic time is a highly challenging task – and also the main subject of this thesis. Probing this progression brings us back to quasars, our astrophysical lights that are particularly useful exactly for these purposes. This is because substantial amounts of neutral hydrogen present in the foreground of a quasar imprint a characteristic absorption signature upon its spectrum, the so-called Lyman- α *damping wing*. The fact that its strength is highly sensitive to the exact amount of neutral hydrogen makes it an exquisite probe of reionisation.

³as well as particles such as cosmic rays and neutrinos, and, most recently, gravitational waves

At the same time, the quasar itself is an exceptionally powerful source of ionising radiation which carves out an ionised bubble around itself whose size grows the longer the quasar has been shining. This reduction of neutral material due to the quasar’s own ionising radiation weakens the intergalactic damping wing signature but also imprints a second characteristic signature upon its spectrum, the so-called *proximity zone*, whose size reflects the extent of the ionised region around the quasar. This, in return, informs us about its lifetime, a crucial quantity for understanding the growth and evolution of the supermassive black holes powering these quasars.

Breaking the degeneracy between the contribution from the global, cosmological reionisation process and the quasar’s own impact as an astrophysical source is one of the fundamental tasks related to this problem and requires both adequate simulations and a sophisticated inference pipeline to relate these simulations to observational data. This thesis is concerned with one such framework that has recently been established for these purposes, trying to push its boundaries to extract even more information from the intergalactic damping wing signature than has been used thus far.

Chapter 2 presents an inventory of the current state of the art of quasar damping wing analysis and quantifies the precision that such measurements are able to achieve when accounting for all relevant sources of uncertainty. We are demonstrating that, besides the uncertain lifetime of the quasar, two contributions are of comparable importance, both with a substantial impact on the overall error budget. First, the damping wing signature is not observable individually but imprinted upon the intrinsic spectrum of the quasar. Reconstructing this intrinsic spectrum and disentangling it from the IGM absorption imprint naturally affects the overall constraining power. Second, reionisation did not proceed uniformly but rather according to a “swiss-cheese” topology where individual ionised bubbles grew and merged until the Universe was reionised in its entirety. As such, the actual distribution of neutral and ionised patches along a given line of sight can vary substantially, even when the *average* fraction of neutral hydrogen is fixed. Sightline-to-sightline variations therefore further add to the error budget on this global average quantity.

This latter source of uncertainty is also the motivation for **Chapter 3**, a theoretical study where we establish two new summary statistics that are defined to be robust against cosmic variance in the IGM. Instead of constraining a global average, these parameters are defined locally, quantifying the neutral hydrogen content directly in front of a given source. As such, they open up the possibility of using the IGM damping wing imprint not only to constrain the global timing but also the local topology of reionisation. In addition, we show that these summary statistics do not require any assumptions about the underlying reionisation model and can thus be constrained in a model-independent fashion.

In **Chapter 4**, we proceed by incorporating our new parameterisation of IGM damping wings into our existing inference codebase. We also put forward a framework that allows us to tie these model-independent, local constraints to a specific reionisation model. This amounts to imposing a non-trivial prior on our local summary statistics and then also results in a model-dependent constraint on the global IGM neutral fraction. By testing our inference pipeline on a large set of mock spectra, we also quantify the precision with which we are able to constrain these new statistics.

The thesis ends with **Chapter 5** which presents the first practical application of our newly developed parameterisation. By analysing JWST/NIRSpec spectra of two of the most distant quasars known to date, J1007+2115 at $z = 7.51$ and J1342+0928 at $z = 7.54$, we obtain, for the first time, constraints on the local ionisation topology in front of these two objects. While the first sightline appears to be largely ionised, the second one shows clear evidence of a damping wing imprint, and thus a non-trivial amount of neutral material along the line of sight. With this analysis, we also take the first step towards the application of our framework to a statistical sample of high-redshift sources that will become available in the coming years, shedding bright new light on our understanding of the Epoch of Reionisation.

NEDERLANDSE SAMENVATTING

Als ik de inhoud van dit proefschrift zou moeten samenvatten in twee woorden, zou het de titel van dit werk zijn: *Les Lumières*. Licht is onmisbaar voor ons als sterrenkundigen omdat we noodzakelijk alle informatie over de natuur van het Universum moeten afleiden van het licht⁴ dat we weten te detecteren met onze telescopen op aarde of in het heelal. Bijzonder geschikt voor deze doelen is een bepaald type van astrofysische bronnen die we quasars noemen. Deze astrofysische objecten schijnen zo uitzonderlijk helder dat ze onder de namen *quasi-stellar objects* of *quasars* bekend werden, hoewel hun afstanden allermindst vergelijkbaar zijn met die van sterren in ons eigen sterrenstelsel. Integendeel, aangedreven door supermassieve zwarte gaten in de centra van verre sterrenstelsels zijn quasars zo uitzonderlijk krachtig en helder dat we ze ook nog vanaf *kosmologische* afstanden kunnen detecteren, tot roodverschuivingen van $z \sim 7.5$. Gezien de tijd die het licht nodig had om de aarde te bereiken, betekent naar deze verre lichtbronnen te kijken niets minder dan naar de eerste miljard jaar van het universum te gluren.

Een bepalend tijdperk heeft de ontwikkeling van het universum tijdens deze vroege tijden gevormd: het *tijdperk van reïonisatie*. Ingeleid door de eerste sterren en sterrenstelsels die in deze jaren werden gevormd, markeerde dit tijdperk het eind van de kosmische *dark ages* of *middeleeuwen* waarin het universum verstoken was van elke vorm van zichtbaar licht. Toen echter deze eerste astrofysische bronnen begonnen te vormen, werd het universum letterlijk door hen verlicht: *Les Lumières* – een kosmisch tijdperk van *Verlichting*. Dit licht had zelfs invloed op de enorme leegtes tussen sterrenstelsels waar zich het meest van de kosmische materie daadwerkelijk bevindt. Blootgesteld aan deze nieuwe lichtbronnen, werd de neutrale waterstof geleidelijk gereïoniseerd door hun straling. Dit leidde tot de vorming van bellen van geïoniseerd waterstof rondom deze ioniserende bronnen die groeiden en zich verbonden tot het hele intergalactische medium gereïoniseerd was.

De vooruitgang van dit proces wordt gekenmerkt door de fractie van neutrale (in tegenstelling tot geïoniseerde) waterstofatomen in het intergalactische medium. Een nauwkeurige meting van de ontwikkeling van deze fractie als functie van kosmische tijd is een uitermate moeilijke taak – en ook het onderwerp van dit proefschrift. Het in kaart brengen van deze ontwikkeling brengt ons terug tot quasars, onze astrofysische bronnen die bijzonder geschikt zijn voor deze doelen. Dat komt doordat significante hoeveelheden van neutraal waterstof in de voorgrond van een quasar een karakteristieke absorptiesignatuur in zijn spectrum achterlaten, de zogenaamde Lyman- α *damping wing*. Het feit de sterkte ervan uitermate afhankelijk is

⁴evenals deeltjes zoals kosmische straling en neutrinos – en sinds kort ook zwaartekrachtgolven

van de exacte hoeveelheid neutraal waterstof maakt het een voortreffelijk meetinstrument voor het tijdperk van reïonisatie.

Tegelijkertijd is een quasar zelf een uitzonderlijk krachtige bron van ioniserende straling die een geïoniseerde bel rondom de quasar beïtelt, die groeit naarmate de quasar langer schijnt. De reductie van neutraal materiaal door de ioniserende straling van de quasar zelf verzwakt de intergalactische damping wing signatuur maar laat ook een tweede karakteristieke signatuur achter in zijn spectrum, de zogenaamde *proximity zone* waarvan de grootte overeenkomt met de omvang van de geïoniseerde bel rond de quasar. Dit informeert ons dan over zijn leeftijd, een essentiële grootheid om de groei en de ontwikkeling van supermassieve zwarte gaten te begrijpen die deze quasars aandrijven.

Deze ontarding tussen de bijdrage van het globale, kosmologische proces van reïonisatie en de bijdrage van de quasar als een astrofysische lichtbron te breken is een van de fundamentele taken in verbinding met dit probleem en vereist zowel adequate simulaties als een doordacht gevolgtrekkingskader om deze simulaties te vergelijken met waargenomen data. Dit proefschrift behandelt een framework dat recentelijk voor deze doelen ontwikkeld werd en is bedoeld om de capaciteiten van dit framework uit te breiden om nog meer informatie van deze intergalactische damping wing signatuur af te leiden.

Hoofdstuk 2 biedt een inventaris van de actuele stand van quasar damping wing analyse en kwantificeert de precisie waarmee dit soort metingen gedaan kunnen worden als rekening gehouden wordt met alle relevante bronnen van onzekerheid. We demonstreren dat (behalve de onzekere leeftijd van de quasar) twee bijdragen vergelijkbaar belangrijk zijn, beide met een significante impact op het totale onzekerheidsbudget. Ten eerste is de damping wing signatuur niet individueel te zien maar afgedrukt op het intrinsieke spectrum van de quasar. Dit intrinsieke spectrum reconstrueren en het ontwarren van de intergalactische signatuur van absorptie beïnvloedt de uiteindelijke precisie van de metingen. Ten tweede verliep het proces van reïonisatie niet overal gelijkmatig maar volgde een “Zwitserse kaas” topologie waar individuele geïoniseerde bellen groeiden en versmolten tot het universum helemaal gereïoniseerd was. De daadwerkelijke verdeling van neutraal waterstof langs een bepaalde zichtlijn kan dus aanzienlijk variëren, ook bij een vaste *gemiddelde* fractie van neutraal waterstof. Variaties tussen zichtlijnen dragen dus verder bij aan het totale onzekerheidsbudget.

De laatste bron van onzekerheid is ook de motivatie voor **Hoofdstuk 3**, een theoretische studie waar we twee nieuwe parameters introduceren die robuust zijn tegen kosmische variantie in het intergalactische medium. In plaats van als globaal gemiddelde zijn deze parameters lokaal gedefinieerd en geven de hoeveelheid neutraal waterstof aan langs de zichtlijn van een bepaalde bron. Dus maken ze het mogelijk de damping wing signatuur niet alleen maar te gebruiken om het globale verloop maar ook de lokale topologie

van reïonisatie te reconstrueren. Ook demonstreren we dat deze parameters geen aannamen over het onderliggende model van reïonisatie vereisen en dus op modelonafhankelijke wijze gemeten kunnen worden.

In **Hoofdstuk 4** gaan we door met het integreren van onze nieuwe parameterisering van quasar damping wings in onze bestaande gevolgtrekkingskader. We presenteren ook een framework dat het mogelijk maakt deze model-onafhankelijke lokale metingen aan een bepaald model van reïonisatie te koppelen. Dit komt neer op het opleggen van een niet-triviale prior op onze lokale parameters en levert ook een modelafhankelijke meting op van de globale fractie van neutraal waterstof met zich mee. Door het testen van ons framework op een groot aantal nepspectra meten we ook de precisie waarmee we in staat zijn deze nieuwe parameters te meten.

Het proefschrift eindigt met **Hoofdstuk 5** dat de eerste praktische toepassing van onze nieuwe parameterisering presenteert. Door het analyseren van JWST/NIRSpec spectra van twee van de meest verre quasars die op dit moment bekend zijn, J1007+2115 met $z = 7.51$ en J1342+0928 met $z = 7.54$, presenteren we voor het eerst metingen van de lokale topologie van reïonisatie voor deze twee objecten. Terwijl de eerste zichtlijn grotendeels geïoniseerd blijkt, vertoont de tweede duidelijk bewijs van een damping wing signatuur en dus een niet-triviale hoeveelheid van neutraal waterstof langs die zichtlijn. Door deze analyse doen we ook een eerste stap in richting van de toepassing van ons framework voor een statistisch steekproef van verre quasars dat beschikbaar zal worden in de komende jaren en helder nieuw licht zal schijnen op ons beeld van het tijdperk van reïonisatie.

PUBLICATIONS

First Author

- *First constraints on the local ionization topology in front of two quasars at $z \sim 7.5$*
Timo Kist, Joseph F. Hennawi, Frederick B. Davies, Eduardo Bañados, Sarah E. I. Bosman, Zheng Cai, Anna-Christina Eilers, Xiaohui Fan, Zoltán Haiman, Hyunsung D. Jun, Yichen Liu, Jinyi Yang and Feige Wang
Monthly Notices of the Royal Astronomical Society, Volume 545, Issue 3, January 2026, staf2219.
doi.org/10.1093/mnras/staf2219
- *Inferring local quasar IGM damping wing constraints*
Timo Kist, Joseph F. Hennawi and Frederick B. Davies
submitted to Monthly Notices of the Royal Astronomical Society.
doi.org/10.48550/arXiv.2508.21812
- *A local, topology-independent parameterization of quasar IGM damping wings*
Timo Kist, Joseph F. Hennawi and Frederick B. Davies
Monthly Notices of the Royal Astronomical Society, Volume 544, Issue 2, December 2025, Pages 2316-2339.
doi.org/10.1093/mnras/staf1762
- *Quantifying the precision of IGM damping wing measurements towards quasars*
Timo Kist, Joseph F. Hennawi and Frederick B. Davies
Monthly Notices of the Royal Astronomical Society, Volume 538, Issue 4, April 2025, Pages 2704-2728.
doi.org/10.1093/mnras/staf460

Unrelated to the topics of this thesis:

- *The robustness of slow contraction and the shape of the scalar field potential*
Timo Kist and Anna Ijjas
Journal of Cosmology and Astroparticle Physics, Volume 2022, Issue 08, August 2022, 046.
doi.org/10.1088/1475-7516/2022/08/046
- *Lee-Yang theory of criticality in interacting quantum many-body systems*
Timo Kist, Jose L. Lado and Christian Flindt
Physical Review Research, Volume 3, Issue 3, September 2021, 033206.
doi.org/10.1103/PhysRevResearch.3.033206

Contributing Author

- *Precisely measuring the cosmic reionization history from IGM damping wings towards quasars*
Joseph F. Hennawi, **Timo Kist**, Frederick B. Davies and John Tamanas
Monthly Notices of the Royal Astronomical Society, Volume 539, Issue 3, May 2025, Pages 2621-2653.
doi.org/10.1093/mnras/staf643

CURRICULUM VITAE

Born on the 15th of December 1997, I spent the first 18 years of my life in the small town of Zeven in Northern Germany. I managed to do well in school at the Grundschule Klostergang and later the St.-Viti-Gymnasium because I could compensate my shyness with an insatiable thirst for knowledge, although I cannot claim that the skies have always been my one and only passion. Especially mathematics came naturally to me early on and I enjoyed solving all kinds of mathematical tasks and riddles, but nonetheless, I was similarly excited about the natural sciences – physics in particular – but also history or Latin once they would finally be added to my timetable.

Probably one of the most important events that brought me to the point where I am today is the Mathematical Olympiad. It was in fifth grade that I was finally old enough to attend the central state-level finals at the venerable Mathematical Institute of the University of Göttingen. The event, hosted in an actual city, three hours away from my small home town, was an entirely new experience for me, and also the first time I set foot into a university. But most importantly, I was suddenly surrounded by other students that enjoyed mathematics just as much as I did. Not only was I lucky enough to get back home with a silver medal, but also to return to Göttingen (almost) every year until the end of school, and being invited to several training seminars, a foretaste of future summer schools during my PhD.

As I approached my final years of school, it was rather clear to me that I would want to pursue studies at a university, but I still had to figure out the subject and the place. In essence, I decided to tread the predetermined path and settle on the city that I had already come to know over the past eight years. Subject-wise, it naturally came down to the choice between physics and mathematics. What ultimately turned the balance was the fact that physics would still heavily involve mathematics as a tool to describe the fundamental laws of nature – and maybe also two great teachers, Mr. Dietz and Ms. Wagner, that I had the luck to learn from in secondary school.

With Göttingen's rich history in mathematics and theoretical physics, the boundary between the two disciplines was fluid, and being enrolled as a physics student, I could still attend the 'real' calculus and linear algebra lectures in my first semester – as well as many more advanced ones in the coming years, again blurring my decision towards physics. Attending so many lectures and learning so many new things was exciting but nonetheless intense, and even the semester breaks were often filled with additional programming courses that did not fit into the regular semester schedule.

Two and a half years later – when I had to choose a topic for my Bachelor's thesis – I knew I would want to be a theorist, but I did not know yet in which field. In the end, I had the luck that Fabian Heidrich-Meisner pitched

a project to me where I could try out different numerical methods for solving the Gross-Pitaevskii equation to model Bose-Einstein condensates. In spite of its immediate roots in quantum many body physics, the project also had an astrophysical component, with the same equation being known as the Schrödinger–Poisson equation in the context of fuzzy dark matter.

With my Bachelor’s thesis just submitted, a new chapter of my life followed right away with my first Master’s year that I spent, via the Erasmus programme, in Finland at the University of Helsinki. Living abroad for the first time in my life was certainly a whole new and exciting experience, and it was here that I attended my first lectures on cosmology. Certainly one of the highlights were student-run observing nights at the old observatory domes on the premises of the Metsähovi Radio Observatory.

Despite these first allusions towards astrophysics, I also kept my passion for quantum physics and got offered a summer research internship on a new method to detect criticality in quantum phase transitions with Christian Flindt and Jose Lado at the nearby Aalto University. Even though the COVID-19 pandemic hit in the middle of my second semester, I decided to stay in Finland and go through with the project – the correct decision, as it turned out, since I could not only spend a rather safe pandemic summer in sparsely populated Finland, but we also obtained interesting results that we decided to publish in what turned out to be my first scientific publication.

Upon my return to Göttingen, given the amount of mathematical credits I had accumulated over time, I was essentially only a thesis away from a second Bachelor’s degree in mathematics which I then obtained with a thesis on random Schrödinger operators. Nonetheless, I was rather certain by that time that I would ultimately want to stay in physics. Although I could have continued with my summer project, I remained very curious about astrophysics and cosmology. In the end, Anna Ijjas at the Max Planck Institute for Gravitational Physics had an exciting theoretical cosmology project for me where I could explore ways of addressing the cosmic initial conditions problem with a slow contraction phase in the early Universe. Once my thesis was complete, I even had the luck to get to visit Paul Steinhardt at Princeton University and turn my thesis into a full publication.

Subsequently, the roulette of PhD positions brought me to the Netherlands where Joe Hennawi offered me an exciting PhD project at Leiden University where I could use statistical inference techniques to place better constraints on the Epoch of Reionisation. Four years later, I can certainly say that I learned a great deal the Observatory and that I got much closer from being a theoretical cosmologist to an astrophysicist. Ultimately, in my final year, I was lucky enough to win a Rubicon grant from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) that will allow me to spend two years at the Institute of Astronomy at the University of Cambridge, and I can’t wait to see what this new chapter of my life will entail.

ACKNOWLEDGEMENTS

I have always been wondering what it tells us that the following two pages will most likely be the most-read ones of this entire thesis. Much more than negating the relevance of our research, I hope it rather is a statement about the importance of others in the academic world – after all, four years of astrophysical research are certainly not possible alone.

As such, of course I have to start by thanking Joe for being such a great supervisor over these past four years. Thanks for believing in my potential and abilities rather than my previous research experience – after all, you hired a theoretical cosmology and/or quantum physics student for a PhD project in astrophysics and statistics! Even though I would still not go far enough to call myself a true astronomer, these past four years have certainly brought me much closer to what is considered an astrophysicist. Fred, it was a pleasure having you as a collaborator on this project and to meet and discuss at various different conferences over the past years!

Leiden University has been an outstanding research environment, and it was amazing to work and study in this place after my first two pandemic-shaped research projects that I largely had to carry out in my own room. Especially, I would like to thank Joop and Matthieu for the weekly preprint meetings, Gijs, and later Lukas and Vysakh for running the machine learning journal club, Zephyr and Mathilde for keeping the Green Committee alive, and the Lorentz Institute for hosting the Colloquia Ehrenfestii – sadly a highly under-appreciated event at the Observatory.

Elena, Marijn and Sjoert, it was a pleasure to be a teaching assistant in your courses where I think I learned at least as much as all the other students. Benjamin and Vinay, thanks for all the hard work on helium reionisation, I hope could pass on to you at least a little bit of what I learned in these past years! I certainly don't want to forget to thank the secretaries, the computerworkers, the caretakers, the cleaning staff and all the other less visible professions at the University that kept everything running so smoothly every day, although I have to leave the University while still eagerly awaiting the Universitair Facilitair Bedrijf (UFB) to establish a warm, healthy and diverse lunch menu.

Thanking my group in fact already brings me beyond the bounds of Leiden University. Despite its practical challenges, being part of a group that extends across two continents truly was an enriching experience. Ben, Diego, Jiamu, Kathlynn, Koki, Linda, Molly, Shane, Suk Sien, Teng, Vikram and Yi, after seeing you week by week on the screen in group meeting, it was great to finally meet most of you in person on Corsica, in Boston, Heidelberg, Helsingør, and, most importantly, even at the University of Santa Barbara!

And of course the Leiden branch: Ben, Caitlin, Da-Ming, Elia, JT, Lars, Riccardo and Silvia, it was great to see at least half the group on a daily basis and to even share offices with some of you – which of course also holds for Beth, Elina, Kira, Lukas, Mulan, Pengyu, Piyush, Richelle, Tram and Yannick! Most especially Elia and Silvia, not only were you two amazing office mates (only in combination with Beth and Piyush of course!), but also great friends outside of the university in conference airbnbs, ski chalets, the climbing gym or even on the mountain! The latter of course is also true for all the other climbers, Darío, Erik, Joey, Manuel, Thijs, Yuze and everyone else who occasionally joined for the countless bouldering and climbing sessions!

Thanks to Amy, Christiaan, Joey, Roi, Chloe, Lucy, Ivana, Verónica, Darío, Louis, Thomas, Zeynab, Casper, Celine, Esther, Manuel and Yuze for organising four amazing ski trips, and of course to everyone who joined and made each trip an unforgettable experience! I still can't believe I got introduced to this activity by moving to a country whose spacetime structure is the textbook example of a $(2+1)$ Minkowski metric.

And of course Aaron, Alessia, Alfred, Andrew, Andrés, Aniruddh, Anniek, Anselmo, Arielle, Beth, Bianca, Billy, Brigitte, Christian, Cecilie, Cheng-Lin, Ciarán, Colin, Dag, Dennis, Elena, Emily, Evgenii, Femke, Fraser, Frits, Gimena, Greta, Hylke, Jay, Jeger, Jelle, Jeroen, Jessica, Jinyi, Josh, Joshiwa, Jurjen, Karin, Keqin, Kevin, Kutay, Leoni, Logan, Luna, Mantas, Mariia, Markus, Marta, Martje, Matus, Michiel, Naadiyah, Niccolò, Nicole, Orestis, Osmar, Paola, Pato, Pranjali, Rico, Rob, Roland, Ryan, Sai, Sam, Scott, Sid, Sill, Tobias, Tommy, Vanesa, Victor, Victorine, Will, Willeke, Wout, Zhenlin, Zorry and anyone I might have missed, thanks for all the countless coffee breaks, lunches, borrels, defence parties and so many other things!

Thanks to all my friends from Göttingen for staying in touch even though I did not just leave behind the city but even the entire country for my PhD. Most especially, thanks to Sarah and Sonja for all the visits and the climbing trips, to the entire APAH crew, especially to Annika, Julien and Heiko for the yearly refuge from the Dutch new year celebrations, and most of all to Annika for all the cows!

Nico, Evelyn en Saskia, jullie lessen waren geweldig! Niet alleen om Nederlands te leren maar ook om kennis te maken met zo veel andere leuke mensen. Teresa, dit is natuurlijk des te meer waar in het geval van jou. Ik ben zo blij dat ik je anderhalf jaar geleden in de cursus en bij de demonstratie tegen de onderwijsbezuinigingen heb ontmoet. Dankjewel voor de ongelofelijk mooie tijd die daarop volgde! En hoewel ik nog steeds niet weet wat verder zal gebeuren, hofte ich, dass du weißt, wie viel ich von dir halte.

Zu guter Letzt fehlt natürlich noch meine Familie, ohne die ich ganz sicher nicht da wäre wo ich heute bin. Mama, Papa, danke dass ihr mich immer unterstützt und mir immer den Raum gegeben habt, meinen eigenen Interessen zu folgen, und vor allem danke für ein Zuhause, an das ich seit jeher jederzeit zurückkehren kann und konnte!