UNIVERSITÀ DEGLI STUDI DI TRIESTE

Dipartimento di Fisica

Università degli studi di Trieste

XXXVIII Ciclo del dottorato di ricerca in Fisica

The Lyman-α Forest as a Probe of Fundamental Cosmology:

Detecting Dark Photon Dark Matter and Measuring the Expansion of the Universe

---

Settore scientifico disciplinare: FIS/05 Astronomia e Astrofisica

Dottorando:
Andrea Trost

Coordinatore:
Prof. Francesco Longo

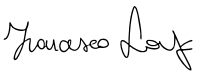

Supervisore:
Prof. Stefano Cristiani

Dr. Gabriella De Lucia

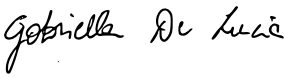

Anno Accademico 2024/2025

Abstract

The absorption line patterns imprinted in the spectra of bright distant quasars are a powerful probe of cosmology and fundamental physics. These spectral features arise from the interaction of quasar light with the diffuse intergalactic medium (IGM), where neutral hydrogen and heavy elements absorb photons at specific rest-frame wavelengths and create absorption lines that retain the physical and gas dynamical information about their environment. In this thesis, the Lyman-α forest of neutral hydrogen has been exploited as a cosmic laboratory for fundamental physics and precision cosmology, focusing in particular on two topics.

The first part of the thesis constrains exotic dark matter scenarios, in particular Dark Photon Dark Matter. These hypothetical particles belong to the dark sector and can kinetically mix with photons in the IGM, leading to oscillations that inject thermal energy into under-dense gas. Such heating leaves a detectable imprint in the Lyman-α forest. Using state-of-the-art hydrodynamical simulations enriched with dark photon phenomenology, I explored the subtle effects of this process. A novel statistical method was developed to compare simulated sightlines with the ultra-high signal-to-noise UVES spectrum of quasar HE0940-1050 ($z_{\rm em} = 3.07$). This approach enabled sensitivity to faint Lyman-α lines arising from under-dense gas. By generating tens of thousands of mock spectra across a wide dark photon parameter space, I placed the most stringent constraints to date on these models. The results show that while much of the parameter space is now excluded, the dark photon remains a viable candidate for resolving the long-standing tension between simulated and observed IGM temperatures at low redshift.

The second part of the thesis contributes to the cosmological redshift drift experiment, a direct and model-independent probe of cosmic expansion. In an accelerating universe, the redshift of a distant source at fixed comoving distance evolves with time. The Lyman-α forest provides an ideal laboratory for this measurement, as it is stable over centuries and each absorption line acts as a comoving marker. I have carried out the first three epochs of this experiment by obtaining 22 hours of ESPRESSO observations of the bright quasar J052915.80-435152.0 ($z_{\rm em} = 3.96$) and 9 hours of J212540.97-171951.4 ($z_{\rm em} = 3.89$). Detailed modelling of hydrogen and metal absorption lines was performed to assess systematics in the reduction pipeline and to validate the achievable precision. The data are consistent with no measurable redshift drift over the short

observational baseline, as expected, and the analysis confirms that the measurement uncertainties scale according to theoretical predictions. The current data provides the first epochs of the long-term measurement of the redshift drift that is expected to produce a detection by ~ 2070 with the next generation of extremely large telescopes.

Together, these results highlight the potential of quasar spectroscopy as precision tools for both fundamental physics and cosmology. The constraints on dark photon dark matter advance our understanding of viable dark sector models and particle dark matter candidates, while the groundwork for the redshift drift experiment establishes the foundations for a decades-long international effort to directly measure the expansion of the Universe with the ELT.

Contents

List of Figures

List of Tables

Acronyms

ΛCDM	Λ Cold Dark Matter.
AGN	Active Galactic Nucleus.
ANDES	ArmazoNes high Dispersion Echelle Spectrograph.
BAO	Baryonic Acoustic Oscillations.
BERV	Barycentric Earth Radial Velocity.
BLR	Broad Line Region.
BSM	Beyond Standard Model.
CCD	Charge-Coupled Device.
CCF	Cross Correlation Function.
CDDF	Column Density Distribution Function.
CGM	Circumgalactic Medium.
CMB	Cosmic Microwave Background.
COS	Cosmic Origins Spectrograph.
CUBES	Cassegrain U-Band Efficient Spectrograph.
DLA	Damped Lyman-α System.
DM	Dark Matter.
DRS	Data Reduction Software.
EFOSC2	ESO Faint Object Spectrograph and Camera 2.
ELT	Extremely Large Telescope.
ESO	European Southern Observatory.
ESPRESSO	Echelle SPectrograph for Rocky Exoplanets and Stable Spectroscopic Observations.
EW	Equivalent Width.
FDM	Fuzzy Dark Matter.
FGPA	Fluctuating Gunn-Peterson Approximation.
FP	Fabry-Pérot etalon.
FRW	Friedmann–Robertson–Walker.

FWHM Full Width Half Maximum.

GBT Green Bank Telescope.
GRB Gamma Ray Burst.
GS Golden Sample.

HIRES High Resolution Echelle Spectrometer.
HST Hubble Space Telescope.

ICM Intra Cluster Medium.
IGM Intergalactic Medium.
ISM Interstellar Medium.

LAE Lyman-α Emitter.
LBG Lyman-Break Galaxy.
LFC Laser Frequency Comb.
LG Local Group.
LLS Lyman Limit System.
LSF Line Spread Function.

MCMC Monte Carlo Markov Chain.
MUSE Multi Unit Spectroscopic Explorer.
MW Milky Way.

NFW Navarro–Frenk–White.
NLR Narrow Line Region.
NTT New Technology Telescope.

OB Observing Block.

PDF Probability Distribution Function.

QE Quantum Efficiency.
QED Quantum Electrodynamics.
QSO Quasi-Stellar Object, or Quasar.
QUBRICS QUasars as BRIght beacons for Cosmology in the Southern hemisphere.

RV Radial Velocity.

S/N Signal-to-Noise ratio.
SB1 Super Bright 1 – J212540.97-171951.4.
SB2 Super Bright 2 – J052915.80-435152.0.
SED Spectral Energy Distribution.

SL Sandage-Loeb test.
SM Standard Model.
SMBH Supermassive Black Hole.
SPH Smoothed Particle Hydrodynamics.

TPCF Two-Point Correlation Function.
tPDF Regualted and Transformed Probability Distribution Function.

UT Unit Telescope.
UVB Ultraviolet Background.
UVES Ultraviolet and Visual Echelle Spectrograph.

VEV Vacuum Expectation Value.
VLT Very Large Telescope.

WDM Warm Dark Matter.
WHIM Warm-Hot Intergalactic Medium.

Part I

The Intergalactic Medium and the Lyman-α forest

1

The Intergalactic Medium

1.1 Cosmology overview

Modern cosmology is grounded in General Relativity, which describes gravity as the curvature of space-time induced by matter and energy. Under the assumption that the Universe is statistically homogeneous and isotropic on large scales, its geometry can be described by the Friedmann–Robertson–Walker (FRW) metric:

$$ds^2 = c^2 dt^2 - a^2(t)\left[\frac{dr^2}{1-kr^2} + r^2(d\theta^2 + \sin^2\theta\, d\phi^2)\right], \tag{1.1}$$

where $a(t)$ is the cosmic scale factor, k denotes the spatial curvature ($k = -1, 0, +1$ corresponding to open, flat, and closed geometries), and c is the speed of light. This metric captures the uniform expansion or contraction of space itself.

Inserting this form of the metric into Einstein's field equations yields the Friedmann equations, which govern the dynamics of the cosmic expansion:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} + \frac{\Lambda c^2}{3}, \tag{1.2}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\rho + \frac{3P}{c^2}\right) + \frac{\Lambda c^2}{3}, \tag{1.3}$$

where ρ and P are the total energy density and pressure of all components of our Universe, G is the gravitational constant, and Λ represents the cosmological constant.

1.1.1 The Standard Cosmological Model

The standard model of cosmology, known as the Λ Cold Dark Matter (ΛCDM) model, provides a remarkably successful description of the Universe's origin, evolution, and large-scale properties. However, its foundation rests on two enigmatic components that, together, constitute over 95% of the cosmic energy density budget: dark matter and dark energy.

The first component is Dark Matter (DM), a massive, non-baryonic, non-collisional, and non-luminous particle that interacts only through gravity with the baryonic matter. The existence of DM was first inferred to support the observation of flat rotation curves of spiral galaxies (Rubin et al., 1980), and later on reiterated by observations of gravitational lensing by massive galaxy clusters (Tyson et al., 1990), anisotropies in the temperature of the Cosmic Microwave Background (CMB, (Smoot et al., 1992)), and the clustering properties of galaxies on large scales (Efstathiou et al., 1984). In the ΛCDM paradigm, DM is categorised as cold, meaning its constituent particles were non-relativistic in the early Universe. This low-velocity dispersion is essential for the hierarchical growth of structure observed today.

The second component, Dark Energy, is responsible for the observed accelerated expansion of the Universe, a discovery made through observations of distant Type Ia supernovae (Riess et al., 1998; Perlmutter et al., 1999), and supported by other indirect evidence based on galaxy clustering and CMB anisotropies (Efstathiou et al., 1990). In general relativity, such an acceleration requires the Universe to be dominated by a substance with a strong negative pressure. While the gravity from ordinary matter and energy density, ρ, slows cosmic expansion, a sufficiently negative pressure ($P < -\frac{1}{3}\rho c^2$) creates a repulsive gravitational force that pushes spacetime apart. The simplest and most successful model for dark energy is Einstein's cosmological constant, Λ, which represents the energy of the vacuum. This vacuum energy has a constant density and an equation of state parameter of $w = P/(\rho c^2) = -1$, perfectly matching the conditions for acceleration. Despite this phenomenological success, the physical origin of Λ remains a profound puzzle, as theoretical calculations of the vacuum energy density exceed the observed value by many orders of magnitude (Weinberg, 1989).

The relative contribution of each energy component to the total energy density of the Universe is conveniently expressed through the dimensionless density parameters,

$$\Omega_i = \frac{\rho_i}{\rho_{\rm crit}}, \quad \text{with} \quad \rho_{\rm crit} = \frac{3H_0^2}{8\pi G}, \tag{1.4}$$

where ρ_i is the energy density of component i and $\rho_{\rm crit}$ is the critical density required for a spatially flat Universe. The present-day value of the Hubble parameter, $H_0 = \dot{a}(t_0)/a(t_0)$, sets the overall expansion scale. In terms of these parameters, the Friedmann equation can be recast as an expression for the Hubble parameter at any redshift z:

$$H^2(z) = H_0^2 \left[\Omega_{\rm r}(1+z)^4 + \Omega_{\rm m}(1+z)^3 + \Omega_k(1+z)^2 + \Omega_{\rm DE}(1+z)^{3(1+w)}\right], \tag{1.5}$$

where $\Omega_{\rm r}$, $\Omega_{\rm m}$, Ω_k, and $\Omega_{\rm DE}$ denote the present-day density parameters for radiation, matter (baryonic and dark), spatial curvature[1], and dark energy, respectively. The dark energy term is presented in generalized form, where w is the dark energy equation of state, being $w = -1$ for a cosmological constant model. This form of the Friedmann

[1] $\Omega_k = 1 - \Omega_{\rm m} - \Omega_{\rm r} - \Omega_{\rm DE}$

equation highlights how each component dominates at different cosmic epochs: radiation at early times, matter during structure formation, and dark energy at late times.

Decades of increasingly precise observations have constrained the ΛCDM key parameters. The latest results from a combination of CMB (from Planck and the Atacama Cosmology Telescope; Planck Collaboration et al., 2020; Louis et al., 2025) and galaxy clustering (DESI; DESI Collaboration et al., 2025b) measurements show that the Universe is spatially flat ($\Omega_k \approx 0$) and composed by:

$$\begin{aligned} \Omega_{\rm m} &= 0.3003 \pm 0.0035, \\ \Omega_\Lambda &= 0.6997 \pm 0.0035, \\ H_0 &= 68.43 \pm 0.27\ {\rm km\,s^{-1}\,Mpc^{-1}}, \end{aligned} \tag{1.6}$$

1.2 The Intergalactic Medium

Within the ΛCDM framework, all the structures observed in the present-day Universe, ranging from individual galaxies to massive galaxy clusters, originated from minuscule primordial density fluctuations, likely seeded by quantum processes during the epoch of inflation. The subsequent evolution of these density perturbations is fundamentally governed by gravitational instability. Two regimes can be identified by defining the density contrast, $\delta = (\rho - \bar{\rho})/\bar{\rho}$, i.e. the fractional difference between the density at a given point and time, and the average density of the Universe. Overdense regions ($\delta > 1$) attract surrounding matter, thereby amplifying their density contrast, while underdense regions ($\delta < 1$) become progressively emptier and less dense.

The crucial role in this process is played by Cold Dark Matter. Being collisionless and possessing negligible thermal velocities (i.e., being non-relativistic), DM perturbations could begin their growth immediately after matter-radiation equality. This is distinct from baryonic matter, which remained tightly coupled to the photon field until the epoch of recombination ($z \approx 1100$), when radiation pressure had prevented the collapse of baryonic overdensities. The gravitational "head start" of DM enabled the formation of the first deep gravitational potential wells. Only after recombination, when the baryonic gas became neutral and decoupled from photons, was it free to cool and fall into these pre-existing DM structures, ultimately forming stars and galaxies.

The structure formation process is therefore hierarchical, proceeding in a "bottom-up" fashion (Cole et al., 2000). Small, low-mass dark matter haloes formed first, and through continuous hierarchical merging and accretion, they grew into the larger, virialised haloes that host today's massive galaxies and galaxy clusters (White et al., 1991). This gravitational collapse is generally described by the linear growth of perturbations for density contrasts $\delta \ll 1$ (Mo et al., 2010). However, when the density contrast reaches a non-linear threshold ($\delta \sim 1$), the clumps decouple from the Hubble flow and collapse under self-gravity into distinct, virialised haloes. The formalism describing the evolution and mass function of these haloes is encapsulated by the Press-Schechter model (Press et al., 1974; Kauffmann et al., 1993).

The persistent gravitational infall of matter, coupled with the overall expansion of spacetime, has arranged the cosmic material into a vast, interconnected network known as the Cosmic Web (Davis et al., 1985; Bond et al., 1996). This web is structurally characterised by:

Nodes (or Knots): The densest and most gravitationally bound regions of the Cosmic Web, located at the intersections of multiple filaments. They host the most massive virialised structures in the Universe, such as galaxy clusters and groups, with typical overdensities $\delta \gtrsim 10^3$.

Filaments: Elongated, quasi-cylindrical structures that connect the nodes and channel matter towards them through coherent inflows. They contain a substantial fraction of the total baryonic mass in the form of diffuse, photoionised or shock-heated gas with moderate overdensities ($\delta \sim 10 - 100$).

Sheets (or Walls): Flattened, two-dimensional structures forming the interfaces between neighbouring voids. They are less dense than filaments but still mark regions of mild baryonic and dark-matter compression, with characteristic overdensities of $\delta \sim 5 - 10$. Sheets often fragment into filaments as nonlinear collapse progresses.

Voids: Vast, underdense regions occupying most of the cosmic volume, where gravitational expansion evacuates matter towards the surrounding filaments and sheets. Their baryonic overdensities are typically $\delta \lesssim 1$, and their gas is cool and highly photoionised.

While the densest regions of the Cosmic Web collapse to form galaxies, the overwhelming majority of the Universe's baryonic content does not reside in stars or the interstellar medium of galaxies. Instead, these baryons trace the underlying dark-matter distribution, filling the filaments, sheets, and voids as a diffuse, highly ionised plasma dominated by hydrogen and helium (see Fig. 1.1). This pervasive component is known as the *Intergalactic Medium* (*IGM*), and it represents the dominant reservoir of baryonic matter in the Universe.

By definition, all baryons in the Universe were part of the IGM at sufficiently early times, before the first stars and galaxies formed. Observationally, the total stellar mass density today accounts for only about $\sim 6-10\,\%$ of the cosmic baryon budget (Fukugita et al., 2004; Madau et al., 2014). Including the cold atomic and molecular gas associated with galaxies raises this fraction to roughly $\sim 15\%$. The hot, X-ray–emitting Intra Cluster Medium (ICM) in galaxy clusters contributes an additional $\sim 4 - 5\%$. Even after accounting for these components, more than half of the expected baryons predicted by Big Bang nucleosynthesis and CMB measurements remain outside collapsed structures, residing instead in the diffuse IGM. This missing fraction is thought to be distributed between the cool, photoionised phase traced by Lyman-α absorption and a warmer, shock-heated component, the Warm-Hot Intergalactic Medium (WHIM), with temperatures in the range $10^5 - 10^7$ K.

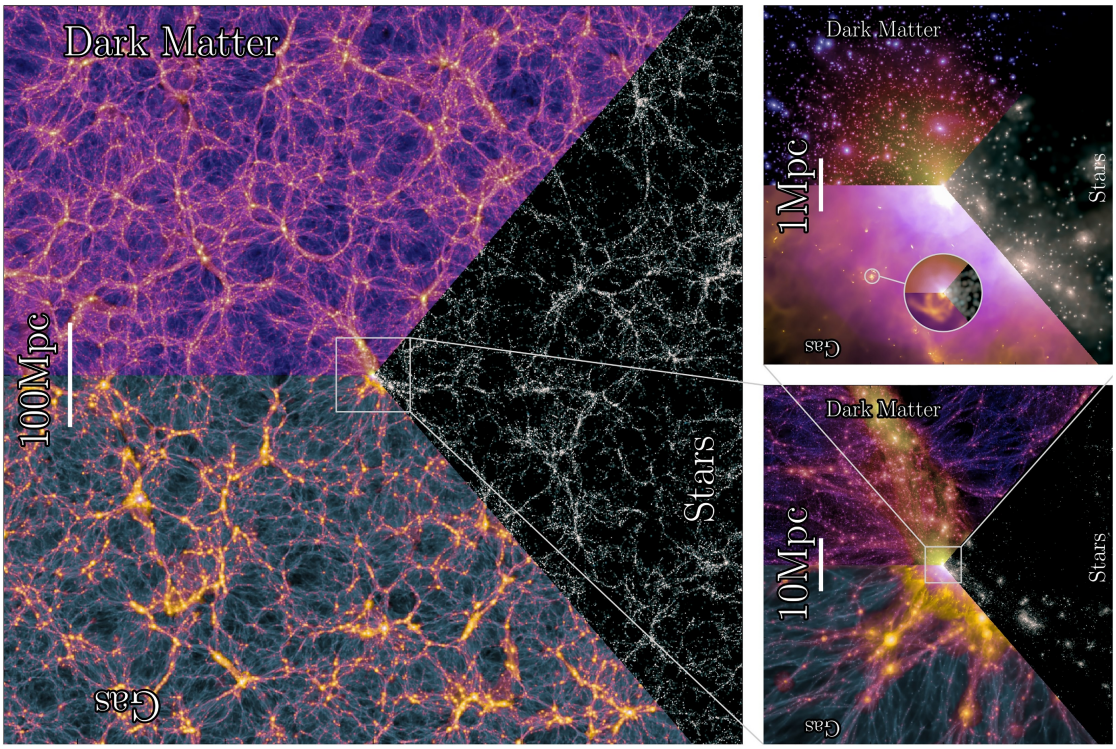

Figure 1.1: *Snapshot of the state of the Universe at $z = 0$ as recreated by the Millennium TNG simulation. The distribution of dark matter is shown in purple. On the largest scales, intergalactic gas dominates the baryon budget, faithfully tracing the underlying DM distribution down to the gas pressure scales, whereas galaxies arise only in the densest environments where DM filaments intertwine. The left side shows a slab of 740 Mpc on the side, and 10 Mpc thick. The zoomed views on the central knot are a factor 10 and 100 smaller, respectively. From Pakmor et al. (2023).*

1.2.1 The Physics of the Intergalactic Medium

As introduced above, the IGM plays a fundamental role in the cosmic baryon cycle: it is both the source of the material that fuels galaxy formation and the sink into which galaxies expel energy, metals, and radiation. Quantifying its distribution, ionisation state, and thermal evolution is essential for understanding the growth of cosmic structures and the evolution the global baryonic mass budget across cosmic time.

Beyond its importance for galaxy evolution, the IGM is also a key probe of cosmology, as we will thoroughly discuss in this thesis. Its diffuse gas traces the large-scale distribution of dark matter with high fidelity on linear and mildly non-linear scales, making it a sensitive tracer of the underlying matter distribution and a powerful source of information on the expansion history of the Universe. The IGM also provides precise constraints on fundamental cosmological parameters such as the matter density Ω_m, the amplitude and slope of primordial fluctuations (σ_8, n_s), and the sum of neutrino masses. Its thermal and ionisation histories further encode the timing and topology of reionisation, the evolution of the ultraviolet background, and the integrated effects of structure formation heating. In this sense, the IGM serves as a unique bridge between cosmology and astrophysics: it connects the physics of the early Universe to the observable structures of the present day, and offers a direct window onto the interplay

between dark matter, baryons, and radiation across cosmic time.

Here we summarise the core concepts required to understand the physics of the IGM and to interpret its observables as cosmological probes.

Density structure and the cosmic web

The intergalactic medium is composed predominantly of hydrogen and helium, the primordial products of Big Bang nucleosynthesis, with mass fractions of roughly 75% and 25%, respectively. Only trace amounts of heavier nuclei, such as lithium and beryllium, were forged in the early Universe, and they contribute negligibly to the overall baryonic mass. At the present epoch, the mean cosmic baryon density is $\rho_{\mathrm{b},0} = \Omega_{\mathrm{b}}\,\rho_{\mathrm{c},0} \simeq 4.2 \times 10^{-31}\,\mathrm{g\,cm^{-3}}$ (Planck Collaboration et al., 2020), corresponding to an average hydrogen number density of

$$\langle n_{\mathrm{H}} \rangle_0 \simeq 2.5 \times 10^{-7}\,\mathrm{cm^{-3}} \approx 0.25\,\mathrm{m^{-3}}. \tag{1.7}$$

This is roughly a million times lower than typical interstellar medium densities, highlighting the extreme rarity of baryons in intergalactic space. Through the uniform expansion of the Universe, the mean density scales as $\langle n_{\mathrm{H}} \rangle(z) = \langle n_{\mathrm{H}} \rangle_0 (1+z)^3$, so that at $z = 3$ the mean hydrogen density was about sixty times higher ($\langle n_{\mathrm{H}} \rangle \simeq 1.6 \times 10^{-5}\,\mathrm{cm^{-3}}$). Helium atoms add an additional $\sim 10\%$ to the particle number density, while the contribution of metals remains dynamically insignificant.

Within the standard ΛCDM framework, this nearly uniform primordial medium has evolved into the intricate network of structures that define the Cosmic Web. Driven by gravitational instability, dark matter collapses into a hierarchy of voids, sheets, filaments, and nodes that span the observable Universe. The baryons, coupled to the gravitational potential of the dark matter, generally trace these same structures, though with less contrast on small scales where gas pressure resists compression. While the Jeans length,

$$\lambda_{\mathrm{J}} = c_{\mathrm{s}} \left(\frac{\pi}{G\rho} \right)^{1/2}, \tag{1.8}$$

with $c_{\mathrm{s}} = \sqrt{k_{\mathrm{B}} T / (\mu m_{\mathrm{p}})}$, characterizes the instantaneous balance between gravity and pressure. In an expanding Universe, however, baryonic perturbations do not respond instantaneously to changes in the thermal state of the gas. As a result, the scale on which pressure smooths baryonic fluctuations is not set by the instantaneous Jeans length, but by the so-called filtering scale (Gnedin et al., 1998). The filtering scale incorporates the full thermal history of the IGM and determines the wavenumber k_{F} above which baryonic density fluctuations are suppressed relative to those of the dark matter. The filtering wavenumber is defined implicitly through

$$\frac{1}{k_{\mathrm{F}}^2}(a) = \frac{3}{a} \int_0^a \mathrm{d}a' \, \frac{c_{\mathrm{s}}^2(a')}{a' H(a')} \left[1 - \left(\frac{a'}{a} \right)^{1/2} \right], \tag{1.9}$$

where a is the scale factor, $H(a)$ is the Hubble parameter, and $c_s(a)$ is the adiabatic sound speed as a function of cosmic time. The corresponding physical filtering scale is

$$\lambda_F = \frac{2\pi}{k_F}, \tag{1.10}$$

which sets the effective pressure-smoothing scale of the baryons, typically $\lambda_F \sim 50 - 150$ kpc in proper units at $z \sim 2 - 4$. On scales smaller than λ_F, baryonic density fluctuations are exponentially damped relative to the underlying dark-matter field.

Thermal and ionisation balance

At such low densities, collisions between particles are rare, and the IGM is nearly optically thin to ultraviolet radiation. Its thermal state is governed by a competition between several heating and cooling processes. The main ones are:

Photoionisation heating: photons from the Ultraviolet Background (UVB) ionise hydrogen and helium atoms. The excess energy of each absorbed photon is converted into kinetic energy of the liberated electrons, which thermalises and heats the gas.

Adiabatic cooling and heating: as the Universe expands, underdense regions cool because the gas performs work on the expanding background, whereas gas collapsing into overdense regions is compressed and heated.

Compton cooling: energetic electrons scatter off CMB photons, transferring energy to them; this process is especially important at high redshift.

Radiative cooling: excited or recombining ions emit photons, removing energy from the gas. The dominant channels are recombination, collisional excitation, and free–free (Bremsstrahlung) emission.

Shock and feedback heating: gravitational infall and galactic outflows drive shocks that locally heat intergalactic gas, particularly in filaments and near galaxies.

The relative importance of these processes changes with redshift (Schaye et al., 2000). After the cosmic recombination epoch ($z \simeq 1100$), the baryons in the Universe were predominantly neutral and cold. As the first stars, galaxies and quasars formed, their ultraviolet and X-ray photons began to ionise the surrounding hydrogen and helium, gradually transforming the neutral intergalactic gas into a fully ionised plasma (Madau et al., 1999; Madau et al., 2015; Grazian et al., 2024). This period, known as the Epoch of Reionisation, likely unfolded between $z \sim 15$ and $z \sim 5.3$ (Kulkarni et al., 2019). After reionisation, photoheating and adiabatic cooling largely control the IGM temperature, maintaining equilibrium at mean temperature of $T \simeq 10^4$ K. The thermal history of the IGM is therefore intimately tied to its ionisation state: an initial sharp rise in temperature accompanies hydrogen reionisation, followed by gradual cooling as the Universe expands. A second, more modest heating event occurs around $z \sim 3$ during the reionisation of singly ionised helium, He II, when hard photons from quasars inject additional energy into the gas (Gaikwad et al., 2021).

Photoionisation equilibrium.

For most of the post-reionisation Universe, the IGM remains in approximate photoionisation equilibrium, in which the rates of ionisations and recombinations balance:

$$n_{\mathrm{H\,I}}\,\Gamma_{\mathrm{H\,I}} = n_{\mathrm{e}} n_{\mathrm{p}} \alpha(T), \tag{1.11}$$

where $\Gamma_{\mathrm{H\,I}}$ is the photoionisation rate and $\alpha(T)$ the recombination coefficient. This balance implies an extremely small neutral fraction, $x_{\mathrm{H\,I}} = n_{\mathrm{H\,I}}/n_{\mathrm{H}} \sim 10^{-6}$, consistent with the observed transparency of the Universe to ionising radiation.

Equation of state of the low-density IGM.

Because heating and cooling operate smoothly across the diffuse gas, the temperature and density are tightly correlated. Cosmological simulations and analytic models show that, after the end of reionization, the IGM follows a power-law relation (Hui et al., 1997a),

$$T = T_0\,\Delta^{\gamma-1}, \tag{1.12}$$

where $\Delta = \rho/\bar{\rho}$, T_0 is the temperature at mean density and γ characterises the slope of the relation. Typical values at redshifts $z \simeq 2-4$ are $T_0 \approx (1-2)\times 10^4\,\mathrm{K}$ and $\gamma \approx 1.3-1.6$. This temperature-density relation encapsulates the balance between photoheating, adiabatic expansion, and recombination cooling. Denser gas cools more efficiently and therefore lies at lower temperatures, while rarefied regions retain heat longer.

1.3 Cosmological simulations

Because the IGM is too diffuse to be directly imaged, numerical simulations are essential to connect its microscopic physics to the observable Universe. Modern simulations follow the coupled evolution of dark matter and baryons within large cosmological volumes, solving for gravity, hydrodynamics, and radiative processes under the framework of the ΛCDM model. Dark matter is evolved with N-body techniques that track the trajectories of massive particles under gravity, while baryons are treated with hydrodynamical solvers that compute the evolution of density, momentum, and thermal energy across a discretised mesh or a set of adaptive particles.

Two broad numerical approaches are commonly employed. The first are Smoothed Particle Hydrodynamics (SPH) codes, such as `GADGET` (Springel, 2005) and its successors (up to `GADGET-4` (Springel et al., 2021)), where the gas is represented by particles whose properties (e.g. temperature, density, peculiar velocity, ionised fraction) are smoothed over a local kernel. The second are grid-based or moving-mesh codes, including `ENZO` (Bryan et al., 2014), `RAMSES` (Teyssier, 2002), and `AREPO` (Weinberger et al., 2020), which solve the Euler equations on a (possibly adaptive) mesh. Each technique has advantages: SPH handles large dynamic ranges efficiently, while grid and moving-mesh codes offer superior shock-capturing and fluid-mixing behaviour. In both cases,

baryonic physics, photoheating, radiative cooling, chemical enrichment, star formation, and feedback from stars and active galactic nuclei is implemented through subgrid prescriptions calibrated to reproduce observed galaxy populations and thermal histories.

For modelling the low-density IGM, the key ingredients are an accurate description of the photoionising ultraviolet background and sufficient spatial resolution to capture the small-scale structure of the Lyman-α forest. State-of-the-art suites such as `IllustrisTNG` (Nelson et al., 2019), `EAGLE` (Schaye et al., 2015), `Sherwood` (Bolton et al., 2017a), and `Nyx` (Almgren et al., 2013) have achieved this balance by combining volumes of tens to hundreds of comoving megaparsecs with kiloparsec-scale resolution. These simulations include spatially uniform UV backgrounds derived from radiative transfer models (e.g. Haardt et al., 2012; Faucher-Giguère, 2020; Puchwein et al., 2019) and self-consistently track the thermal and ionisation evolution of the IGM from reionisation to the present epoch.

When the focus shifts to the reionisation epoch itself, simulations must relax the assumption of a spatially uniform UV background and instead resolve the inhomogeneous, time-dependent propagation of ionisation fronts. This requires coupling radiative transfer to hydrodynamics in so-called radiation-hydrodynamical simulations, which track the growth of ionised bubbles around galaxies and quasars and their eventual overlap. Suites such as the `Sherwood-Relics` (Puchwein et al., 2023) extension of the Sherwood project, or other patchy-reionisation models (e.g. `CoDa` (Ocvirk et al., 2016), `THESAN` (Kannan et al., 2022)), incorporate these effects to capture the complex morphology of the ionising background and the spatial variations in temperature that persist as reionisation relics well after the process has completed.

A striking success of these computational efforts is the quantitative agreement between simulated and observed Lyman-α forest statistics (see next Chapter). When the IGM is evolved under realistic photoheating rates, the resulting optical-depth fluctuations, flux probability distribution functions, and 1D power spectra reproduce the data from high-resolution quasar spectra to within a few per cent (Viel et al., 2006). This agreement confirms that the forest arises naturally from the mildly non-linear density field of the IGM, shaped by the same gravitational physics that governs dark matter. Simulations further predict that most baryons at $z \gtrsim 2$ reside in a diffuse, photoionised phase with $T \sim 10^4$ K, while at lower redshift, a growing fraction is heated to $T \sim 10^5 - 10^7$ K (the WHIM) by structure formation shocks and feedback from galaxies and quasars.

Beyond reproducing existing observations, numerical models provide predictive power. They enable tests of alternative reionisation histories, variations in the UV background, or additional heating mechanisms (such as TeV blazars (Broderick et al., 2012) or exotic energy injection), and they inform the interpretation of new observables such as the Lyman-α forest tomography, fast radio bursts, or metal-line absorption. Modern high-resolution simulations, including those used in this thesis, have become indispensable laboratories for linking the small-scale thermal and ionisation structure of

the IGM to its cosmological context and to properly interpret astronomical observations.

2

The Lyman-α forest

As shown in the previous Chapter, the intergalactic medium is a diffuse ionised gas that permeates the cosmos on the largest scales. The properties of this gas are intimately related to cosmology and fundamental physical processes, as well as cosmic structure formation and galaxy evolution. Despite being a valuable source for our understanding of the Universe, the IGM is tenuous, does not emit visible light, and cannot therefore be imaged directly as a star, galaxy, or quasar can. Its presence is instead observed in *absorption*, meaning that the effect of the IGM appears as an absorption feature on the emission spectra of background sources. Most commonly, the preferred sources used to study the IGM are Quasi-Stellar Objects, or Quasars (QSOs), but also Gamma Ray Bursts (GRBs) and bright star-forming galaxies such as Lyman Break Galaxies (LBGs) and Lyman-α Emitters (LAEs) provide independent probes at different redshifts. For this reason, the study of the IGM is fundamentally connected to the nature and the study of the specific probe used.

In this Thesis, we focus on the study of the IGM in the optical range of QSO spectra. In this chapter, we outline the basics of QSO physics and spectroscopy, with a focus on the Lyman-α forest of Hydrogen lines, our main tracer of the physical state of the IGM. We finally discuss the practical[1] side of astronomical high-resolution spectroscopy with ground-based optical telescopes.

2.1 Quasars

Among the most luminous persistent sources in the Universe are QSOs, a subclass of Active Galactic Nuclei (AGNs) powered by accretion onto Supermassive Black Holes (SMBHs). Their remarkable brightness, often exceeding that of their host galaxies by several orders of magnitude, makes them observable up to the highest known redshifts (Wang et al., 2021). In the context of this work, quasars are primarily relevant as luminous background sources used to trace the intervening intergalactic medium.

[1] and sometimes unpractical...

Nevertheless, an understanding of their physical nature and emission mechanisms is essential for interpreting the spectra they produce.

2.1.1 Active Galactic Nuclei and the Unified Model

AGN are thought to represent a common phase in the evolution of galaxies (Padovani et al., 2017), where a central SMBH, with masses ranging from 10^6 to 10^{10} $M_\odot$, actively accretes the surrounding gas. The gravitational energy released through accretion is converted into radiation with a remarkably high efficiency, up to $\sim 10\%$ of the rest-mass energy, resulting in luminosities that can approach or even exceed $10^{47} - 10^{48}$ erg s^{-1}.

Despite the wide observational diversity of AGN, ranging from relatively faint Seyfert galaxies to extremely luminous quasars, many of their properties can be explained within the framework of a unified model (Antonucci, 1993; Urry et al., 1995). In this paradigm, the primary components of an AGN are similar across the population: a central SMBH, an accretion disk, fast-moving gas clouds that form the Broad Line Region (BLR), more distant clouds forming the Narrow Line Region (NLR), and a dusty torus that obscures part of the emission depending on the line of sight orientation. Relativistic jets may also emerge perpendicular to the plane of the disk in some sources, contributing to radio and high-energy emission.

Variations in observed properties, such as the width of emission lines, the slope of the continuum, or the presence of strong radio emission, are largely attributable to differences in accretion rate, obscuration, and orientation relative to the observer, as schematically described in Fig. 2.1. Within this picture, quasars can be regarded as the most luminous manifestation of AGN activity, typically corresponding to systems with high accretion rates and relatively unobscured lines of sight toward the central engine.

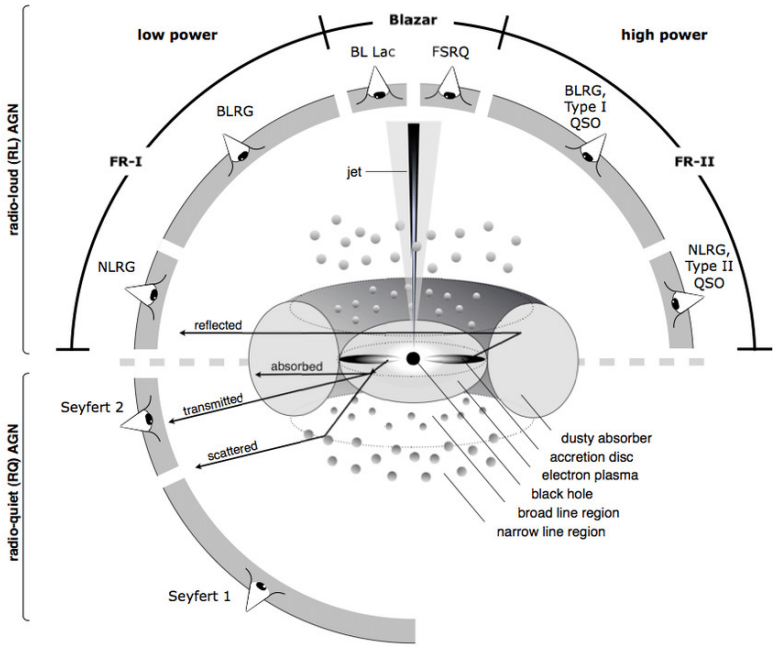

Figure 2.1: *Simplified schematic representation of the unified AGN model and its core components: the SMBH, the accreting disk, the BLR, the NLR, and the dusty torus. Different classes of astronomical sources arise from the model, depending on the viewing angle, as reported in the figure. From Beckmann et al. (2012).*

2.1.2 Structure and Emission Mechanisms

The energy output of a quasar arises from a combination of physical components, each dominating a different region of the Spectral Energy Distribution (SED). The overall continuum spans from radio to X-ray and even γ-ray energies, but in the optical and ultraviolet regimes, the main contributors are the accretion disk, the BLR, and reprocessed emission from surrounding material.

Accretion Disk: Gas accreting onto the SMBH settles into a geometrically thin, optically thick disk, where viscous and turbulent processes transport angular momentum outward and allow matter to spiral inward. As the gas falls deeper into the gravitational potential well, its temperature increases, reaching up to $\sim 10^5$ K in the innermost regions. The disk radiates roughly as a superposition of blackbodies at different temperatures, producing a smooth, thermal continuum that peaks in the ultraviolet. The outer, cooler parts of the disk contribute to the optical and near-infrared emission. This multi-temperature structure gives rise to the characteristic Big Blue Bump seen in the SEDs of type-1 quasars.

Broad-Line Region: Surrounding the accretion disk is a dense ensemble of gas clouds moving at high velocities (typically a few thousand $\mathrm{km\,s^{-1}}$) under the gravitational influence of the central black hole. This region is photoionised by the intense radiation from the accretion disk, producing strong emission lines such as Lyman-α , C IV λ1550, Mg II λ2798, and Hβ. The large peculiar velocities of the clouds result in the wide emission features characteristic of type-1 AGN. The detailed line ratios and profiles depend on the gas density, ionisation parameter, and geometry of the emitting region.

Narrow-Line Region: At larger radii (hundreds to thousands of parsecs) lies the narrow-line region, a more diffuse and less dense medium producing forbidden transitions such as [O III] λ5007 and [N II] λ6583. The lower velocities of these clouds (hundreds rather than thousands of km $\mathrm{s^{-1}}$) yield narrower line widths. Although its contribution to the total luminosity is modest compared to that of the BLR or the accretion disk, the NLR provides important diagnostics of the ionising spectrum and the chemical composition of the circumnuclear gas.

Dusty Torus and Reprocessed Emission: The central regions of the AGN are embedded within a toroidal distribution of dusty material that absorbs a significant fraction of the ultraviolet and optical radiation, re-emitting it in the infrared. The geometry and clumpiness of this torus determine whether the observer has a direct or obscured view of the central engine. When the line of sight lies close to the polar axis, the nucleus is unobscured and broad lines are visible (type-1 objects, including quasars). When viewed edge-on, the torus obscures the inner regions, yielding type-2 AGN dominated by narrow emission lines. This orientation-dependent obscuration naturally accounts for much of the observed diversity within the AGN population. For details, see e.g. Buchner et al. (2015) and Ramos Almeida et al. (2011).

High-Energy Components: Above the optical/UV regime, additional structures contribute to the emission. A hot, optically thin corona located above the accretion disk can upscatter UV photons to X-ray energies through inverse Compton processes. In radio-loud quasars, relativistic jets contribute synchrotron emission that extends from the radio to the X-ray domain (Blandford et al., 2019). These high-energy processes, while not dominant in the optical spectrum, provide important constraints on the physical conditions in the inner accretion flow.

2.1.3 The Optical Spectrum of Quasars

In the optical and near-UV bands, the observed quasar spectrum is a combination of the thermal continuum from the accretion disk and a forest of broad and narrow emission lines originating in the BLR and NLR. The continuum is typically well described by a power law, $f_\nu \propto \nu^{-\alpha}$, with spectral indices in the range $\alpha \approx 0.5 - 0.7$ (Cristiani et al., 2016). Prominent emission features include Mg II λ2798, Hβ, and C IV λ1550, superimposed on the smooth blue continuum.

The relative strength of these components varies significantly between sources, reflecting differences in accretion rate, black hole mass, and viewing geometry. Nonetheless, the general form of the optical spectrum remains remarkably consistent across redshift and luminosity, allowing quasars to serve as reliable and well-characterised background sources in cosmological studies.

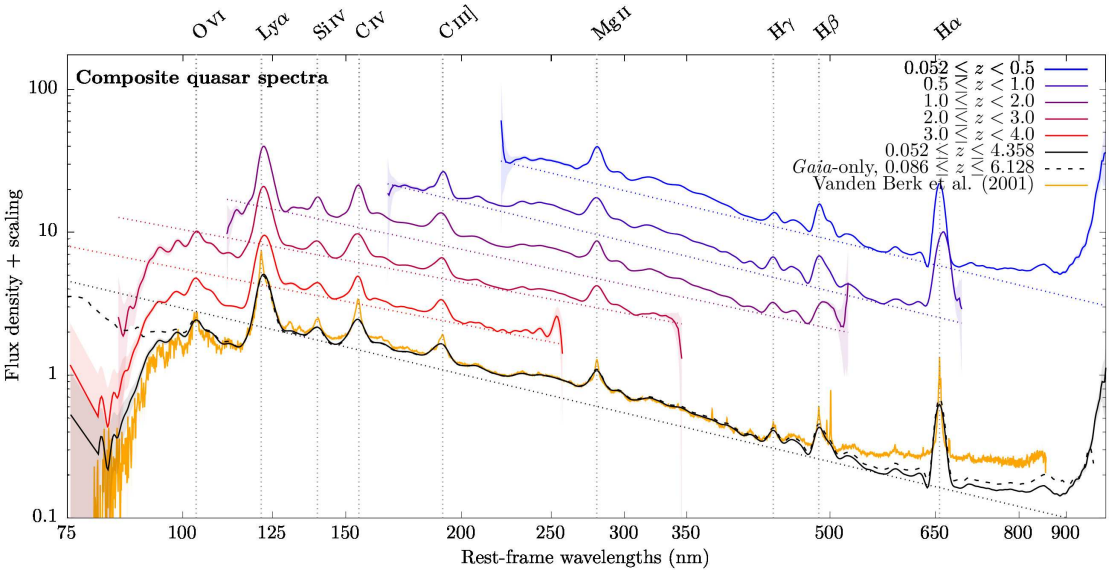

Figure 2.2: *Composite QSO emission spectra in different redshift bins, as a function of rest frame wavelength. Continuum power-law and line emissions are defined. From Gaia Collaboration et al. (2023)*

2.1.4 Intervening Absorption

As the light from a distant quasar travels toward the observer, it encounters various structures in the intergalactic medium. These intervening systems, composed primarily of neutral or partially ionised hydrogen, absorb photons at specific rest-frame

wavelengths corresponding to electronic transitions of the hydrogen atom and of trace metals. The result is a characteristic pattern of absorption lines superimposed on the quasar's intrinsic emission spectrum.

Because the Universe expands, each intervening system imprints its absorption feature at a wavelength that depends on its cosmological redshift. If the quasar is located at an emission redshift $z_{\rm QSO}$ and emits continuum radiation at a wavelength $\lambda_{\rm em}$, this radiation will be redshifted by the cosmological expansion as it travels through the IGM and will be absorbed by an atomic transition of rest-frame wavelength $\lambda_0^{\rm rf} > \lambda_{\rm em}$ at a redshift $z_{\rm abs} < z_{\rm QSO}$, where

$$\lambda_{\rm em} = \lambda_0^{\rm rf} \frac{1 + z_{\rm abs}}{1 + z_{\rm QSO}}. \tag{2.1}$$

This absorption imprints a line in the quasar spectrum at $\lambda_{\rm em}$ in the quasar reference frame, whereas an observer on Earth detects such a line at

$$\lambda_{\rm em}^{\rm obs} = \lambda_{\rm em}(1 + z_{\rm QSO}) = \lambda_0^{\rm rf}(1 + z_{\rm abs}). \tag{2.2}$$

Along the line of sight, the quasar's continuum passes through regions of the intergalactic medium with varying neutral hydrogen density. Each fluctuation in density produces absorption at the corresponding redshifted wavelength. As the radiation accumulates these contributions over cosmological distances, the observed spectrum becomes a dense sequence of absorption features—the Lyman-α forest—arising from the continuous density structure of the low-density IGM (Viel et al., 2006), as shown in Fig. 2.3.

This simple relation, combined with the immense brightness of quasars, makes absorption spectroscopy one of the most powerful tools in observational cosmology. By analysing the pattern, depth, and distribution of absorption lines imprinted on quasar spectra, it becomes possible to reconstruct the physical conditions and spatial distribution of the gas that fills the intergalactic medium. In this way, quasars serve not only as luminous astrophysical laboratories in their own right, but also as backlights that reveal the structure and evolution of the cosmic web across most of the history of the Universe.

2.2 The Lyman-α forest

Among the various species present in the intergalactic medium, neutral hydrogen plays a dominant role in shaping quasar absorption spectra. As the most abundant element in the Universe, it leaves a distinctive imprint on the light of distant quasars. The high absorption cross-section of hydrogen allows transitions from multiple electronic states to contribute to the observed spectrum. Most of the features located blueward of Lyman-α belong to the so-called *Lyman series* (Lyman, 1906), which corresponds to transitions terminating at the ground state of hydrogen. The wavelengths of these lines

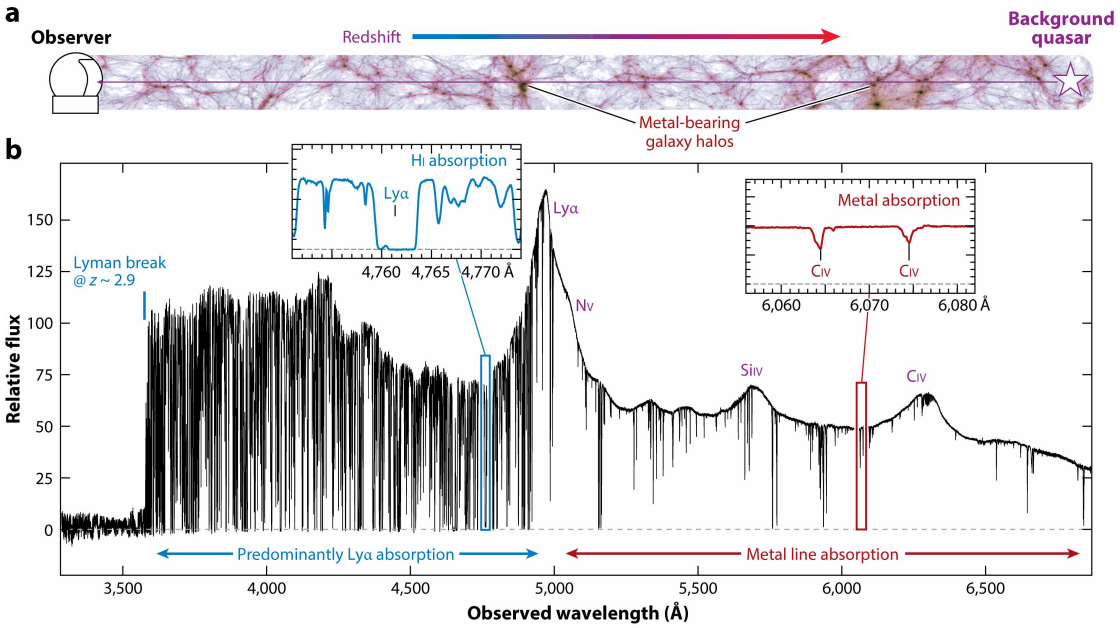

Figure 2.3: *General depiction of how the absorption features in a quasar spectrum arise. a) A ray of light emitted by a quasar (right) travels through the IGM and reaches an observer (left). Along its path, the ray pierces the cosmic web's filaments, where H I gas absorbs light and gives rise to a forest of lines, and galactic haloes, where absorption from metal ions occurs, as well as deep H I absorption. b) High-resolution spectrum of a $z = 3$ quasar (from D'Odorico et al. (2016)). Intrinsic quasar emission features are indicated along the spectrum in purple. Numerous absorption lines, shifted by cosmological redshift, are visible in the spectrum. At wavelengths longer than the quasar's Lyman-α emission, the absorption is dominated by metals, whereas at shorter wavelengths the forest is primarily produced by H I , with some contributions from metal lines. These absorption signatures provide key diagnostics of the physical conditions in the intervening gas. The insets highlight examples of absorption profiles: H I on the left, and representative metal transitions on the right, originating from the $z \sim 2.9$ absorber responsible for the prominent Lyman break near 3600 Å. From Péroux et al. (2020).*

can be derived from the Rydberg relation, expressed as

$$\lambda = \frac{c\, h_P}{E_m - E_n} = \frac{12398.4 \text{ Å eV}}{E_m - E_n}, \tag{2.3}$$

where h_P is Planck's constant, E_m and E_n denote the energies of the initial and final levels of the electron. Since the energy of a given state scales as $1/n^2$ with the integer quantum number n, each allowed transition is uniquely determined by two integers, m and n, with $n > m$. This yields

$$\frac{1}{\lambda} = \frac{1}{911.8 \text{ Å}} \left(\frac{1}{m^2} - \frac{1}{n^2} \right). \tag{2.4}$$

In this framework, the case $m = 1$ defines the Lyman series. The longest-wavelength member of the series is Lyman-α ($n = 2$), located at $\lambda = 1215.67$ Å. The next line, Lyman-β ($n = 3$), appears at a shorter wavelength $\lambda = 1025.72$ Å, while subsequent transitions are conventionally labelled with successive Greek letters (γ, δ, etc.) and have increasingly shorter wavelengths. In the limit of $n \to \infty$, this sequence converges to a limiting wavelength known as the Lyman limit. This threshold occurs at the Ly-

man limit, $\lambda_{\mathrm{LL}} = 911.7\,\text{Å}$, corresponding to the ionization energy of neutral hydrogen (13.6 eV). Photons with wavelengths shorter than this limit no longer drive bound–bound transitions but instead ionise the atom, producing a continuum absorption feature rather than a discrete line. In quasar spectra, intervening hydrogen clouds with neutral-hydrogen columns large enough to yield an optical depth of order unity at 912 Å (typically $N_{\mathrm{HI}} \gtrsim 10^{17.2}\,\mathrm{cm}^{-2}$) significantly suppress the continuum blueward of this limit, generating the characteristic Lyman break.

2.2.1 Absorption Line Profiles

The Lyman-α transition of H I leaves strong traces in the spectra of distant quasars. The exact shape and depth of the absorption lines are tightly related to the physical state of the absorbing gas and can be used to measure its properties. This applies to all absorption lines, regardless of the absorbing ion.

To mathematically define the absorption line profile, we first start by considering that the cross-section of an atomic transition between atomic states m and n is defined, as a function of frequency, as

$$\sigma_\nu = \frac{\pi e^2}{m_e c} f_{mn} \varphi_\nu, \tag{2.5}$$

where φ_ν is the line profile in frequency space[2], and f_{mn} is the oscillator strength of the transition. Its value defines the probability of such a transition to take place. The oscillator strength of the Hydrogen Lyman-α transition is $f_{\mathrm{Ly}\alpha} = 0.4160$.

Unlike what one would expect, observed atomic lines do not have infinitely narrow profiles, i.e. line profiles are not Dirac's Delta functions $\delta_D(\nu - \nu_0)$ centred around the transition's frequency ν_0 (defined by Eq. 2.4), but are widened by three main mechanisms, plus one instrumental effect:

Natural Width: The quantum mechanical uncertainty on the energy E of levels with finite lifetimes is translated into an uncertainty of the transition's frequency, thus widening the profile.

Collisional Broadening: Collisions between atoms reduce the effective lifetimes of a state, leading to broader lines. This effect depends on the pressure of the interacting gas and is usually found in the spectra of stellar atmospheres, whereas it can be neglected in IGM spectroscopy.

Doppler Broadening: The motion of the individual atoms in the gas relative to the observer causes a Doppler shift in the emission/absorption frequency, widening the line. Large-scale turbulent motions of the gas can be treated similarly.

Instrumental Broadening: Spectroscopic instruments that measure emission and absorption lines have a finite resolving power, defined as the minimum Full Width Half Maximum (FWHM) of a resolvable line. As the spectral lines pass through the instrument's optics, they are convolved by the Line Spread Function (LSF) whose FWHM is determined by this resolving power.

[2] Frequency and wavelength are related by $\lambda = c/\nu$

We want to briefly summarise these effects to lead to a comprehensive profile model.

Natural Broadening

The energy levels of an atom are defined by the Schrödinger Equation, with exact values. However, the energy of an atomic level has an intrinsic uncertainty related to its lifetime Δt through Heisenberg's uncertainty principle

$$\Delta E \sim \hbar/\Delta t, \tag{2.6}$$

where $\hbar = h_P/2\pi$. Suppose an atom makes a transition from this state to the ground state. The emitted photon will then have a range of possible frequencies

$$\Delta\nu \sim \frac{\Delta E}{h} \sim \frac{1}{2\pi\Delta t}. \tag{2.7}$$

around the central frequency of the line $\nu_0 = E/h_P$. It can be shown that the exact line profile given by this broadening is defined by a Lorentzian curve

$$\varphi_\nu = \frac{\Gamma/4\pi^2}{(\nu-\nu_0)^2 + (\Gamma/4\pi^2)^2}, \tag{2.8}$$

where the variable Γ is called the radiative damping constant, and the value $\Gamma/2\pi$ represents the full width of the φ_v profile at half - intensity.[3]

Collisional Broadening

In simple atomic models, energy levels are calculated considering only the potential from the nucleus and the bound electrons. In an actual plasma, however, additional particles perturb this potential, which leads to modifications of the atomic energy levels. These perturbations can shift or split the levels, and as a result, spectral lines become broadened. The broadening follows a Lorentzian profile and can be described by including a collisional term in the radiative damping constant of the line $\Gamma = \Gamma_{\text{natural}} + \Gamma_{\text{collision}}$. This effect is usually negligible in the low-density and low-pressure environments of the IGM.

Thermal Broadening

In a gas, atoms have a velocity distribution that depends on the local temperature. The most probable speed u of the atoms is given by

$$u = \sqrt{2k_BT/m}, \tag{2.9}$$

where m is the mass of the atom, T the local temperature of the gas, and k_B the Boltzmann's constant.

[3] Γ is the full width at half-maximum of the profile φ_ω, where $\omega = 2\pi\nu$.

In a gas of temperature T, following the one-component Maxwell-Boltzmann distribution, the probability of finding an atom of mass m with a velocity component between u_x and $u_x + du_x$ is

$$f(u_x)du_x = \sqrt{\frac{m}{2\pi k_B T}}\, e^{-\frac{mu_x^2}{2k_B T}} du_x. \tag{2.10}$$

Assume that the atoms at rest absorb photons at frequency ν_0. The thermal-induced velocity u_x with respect to the observer induces a Doppler shift in the absorbed frequencies, ν, following

$$\frac{\nu - \nu_0}{\nu_0} = \frac{u_x}{c}. \tag{2.11}$$

Combining this Doppler shift with the one-dimensional distribution of the velocities, the profile function becomes

$$\varphi_\nu = \frac{1}{\Delta\nu_D\sqrt{\pi}}\, e^{-\frac{(\nu-\nu_0)^2}{\Delta\nu_D^2}}, \tag{2.12}$$

where the Doppler width $\Delta\nu_D$ is defined as

$$\Delta\nu_D = \frac{\nu_0}{c}\sqrt{\frac{2k_B T}{m}}. \tag{2.13}$$

If the gas has also turbulent motions that can be described by a similar velocity distribution, the effective Doppler width becomes

$$\Delta\nu_D = \frac{\nu_0}{c}\sqrt{\frac{2k_B T}{m} + v_{\text{turb}}^2}, \tag{2.14}$$

where v_{turb} is the root-mean-squared of the turbulent velocities[4]. The thermal line profile is just a Gaussian, which falls off very rapidly away from the line centre. The term

$$b = \sqrt{v_{\text{th}}^2 + v_{\text{turb}}^2} = \sqrt{\frac{2k_B T}{m} + v_{\text{turb}}^2}, \tag{2.15}$$

is called the Doppler Parameter, is measured in km s^{-1}, and takes into account both the thermal and turbulent components of the line broadening effect.

Voigt Profiles

The line profile resulting from the combination of the aforementioned effects is called a *Voigt Profile* and arises from the convolution of a Gaussian and a Lorentzian profile. This profile is defined by the Hjerting function $H(a, x)$

$$H(a, x) = \frac{a}{\pi}\int_{-\infty}^{\infty} \frac{e^{-y^2}}{(x-y)^2 + a^2} dy, \tag{2.16}$$

[4] The turbulence term is usually negligible in the Lyman-α forest.

where x is the dimensionless frequency

$$x = \frac{\nu - \nu_0}{\Delta\nu_D}, \tag{2.17}$$

and the other parameters are defined as

$$y = \frac{\nu_0}{\Delta\nu_D}\frac{u_x}{c}; \tag{2.18}$$

$$a = \frac{\Gamma}{2\pi\Delta\nu_D}. \tag{2.19}$$

Substituting in Eq. 2.5, the absorption cross-section becomes

$$\sigma_\nu = \frac{\sqrt{\pi}e^2}{m_e\, c} f_{ij} \frac{H(a,x)}{\Delta\nu_D}. \tag{2.20}$$

The absorption observed in a spectrum depends on the dimensionless optical depth at the centre of the line τ_{ν_0}, which is defined as

$$\tau_{\nu_0} = \int n\sigma_\nu\, dl, \tag{2.21}$$

where the integral runs over the span of sightline piercing a cloud of density n (in cm^{-3}) and absorption cross section σ_{ν_0}. If the number density is uniform along the integration path, we can write

$$\tau_{\nu_0} = N\frac{\sqrt{\pi}e^2}{m_e\, c}\frac{f_{ij}}{\Delta\nu_D}, \tag{2.22}$$

where N is the *column density*, in units of cm^{-2}, that represents the number of atoms in a column of unit cross section along the sightlines. It is related to the number density of atoms by

$$N = \int n\, dl. \tag{2.23}$$

Note that the optical depth at line centre τ_{ν_0} increases with column density. Moreover, as the temperature of the atoms increases, τ_{ν_0} decreases. If $\tau_{\nu_0} > 1$, the line is said to be saturated and its width changes only slightly with the increase of the column density.

The formalism leading to the optical depth at line centre, τ_{ν_0}, is easily generalized to the optical depth at any frequency, any wavelength, τ_λ. Given the optical depth of an absorber, the effect on the measured flux density of quasar spectra is given by

$$F_\lambda = F_\lambda^{\mathrm{cont}} e^{-\tau_\lambda}, \tag{2.24}$$

where $F_\lambda^{\mathrm{cont}}$ is the continuum emission flux density of the quasar. Note that we moved to units of wavelength instead of frequency to be consistent with the usual optical spectroscopy formalism. Since the physical information on the IGM is contained only in the absorption of the background emission radiation, the observed flux density is usu-

ally renormalised to the continuum emission before proceeding with spectral analysis. This step requires the modelling of the quasar continuum emission $F_\lambda^{\rm cont}$, which is not known a priori and can introduce strong biases depending on the considered methodology and type of analysis that one carries out. Throughout the thesis, we refer to normalised and non-normalised fluxes when this transformation is applied or not, respectively.

Fig. 2.4 shows the normalised line profile of the H I Lyman-α transition for different values of column density and Doppler parameter, as a function of relative velocity[5] from the transition central wavelength $\lambda_{{\rm Ly}\alpha} = 1215.67$ Å.

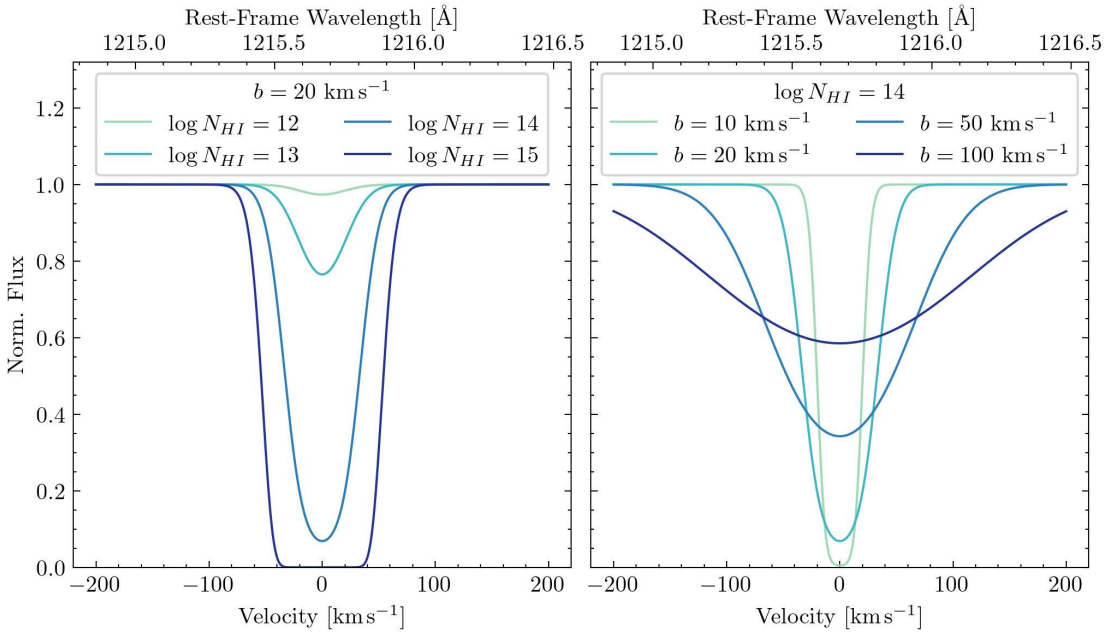

Figure 2.4: *Absorption Voigt profiles of H I Lyman-α lines with different column densities (left panel) and Doppler parameters (right panel).*

Instrumental Line Spread Function

In addition to the intrinsic and physical line broadening effects described above that lead to the definition of the theoretical Voigt profile, an instrumental effect must be included in the analysis and line modelling procedure of real spectral analysis to properly measure physical quantities.

The instrumental broadening, commonly described by the LSF, originates from the finite resolving power of the spectrograph used and is determined by the combined effects of telescope optics, instrument configuration, and the quality of the incoming image, which for ground-based observations also includes atmospheric seeing. The LSF sets a fundamental lower limit on the width of spectral features: if the intrinsic FWHM of a line is narrower than this limit, the line will remain unresolved in the observed spectrum. In practice, the effect of the instrument is included by convolving the

[5] One can easily transform the dispersion axis of a spectrum as $dv = c\, d\lambda/\lambda = -c\, d\nu/\nu$

theoretical Voigt line profile with the LSF before comparing the model to the observed spectrum. Mathematically, the measured flux density can be expressed as

$$F_{\lambda}^{\mathrm{obs}} = \int F_{\lambda'}^{\mathrm{true}} \mathrm{LSF}(\lambda - \lambda') d\lambda', \tag{2.25}$$

where F^{true} is the intrinsic line profile and LSF is the instrumental profile. Neglecting this convolution leads to systematic errors in the retrieved line parameters, such as central wavelength, width, and depth, especially when the intrinsic broadening mechanisms (thermal, turbulent, or collisional) produce features comparable to the instrumental resolution.

In many instruments, the LSF is well approximated by a Gaussian function, which simplifies the convolution and can be treated analytically in some cases. However, more complex spectrographs, aimed at high precision measurements, may exhibit non-Gaussian or asymmetric instrument profiles that lack a simple analytical description. In such cases, the LSF must be characterised empirically and provided as tabulated data to be included in the line-fitting process (e.g. Milaković et al., 2024a).

2.2.2 Metal Absorption

In astrophysics, all elements heavier than helium are collectively referred to as metals. These species are primarily produced through stellar nucleosynthesis (Hoyle, 1946; Burbidge et al., 1957). Broadly, three pathways dominate heavy-element formation: (i) *nuclear fusion*, which synthesises elements up to calcium via processes such as the PP-chain, CNO cycle, and successive burning stages of helium, carbon, oxygen, and silicon (Adelberger et al., 2011); (ii) *photodisintegration*, through which high-energy γ-ray photons eject nucleons or α particles, leading to nuclei up to the iron peak; and (iii) *neutron capture*, which produces most nuclei beyond ^{56}Fe through slow (s-process) or rapid (r-process) neutron captures (Käppeler et al., 2011; Kajino et al., 2019). Once formed, these metals are redistributed into the interstellar medium by stellar winds, supernova explosions, and AGN outflows (Madau et al., 2001; Oppenheimer et al., 2008; Fabjan et al., 2010; Choi et al., 2020), enriching subsequent generations of stars.

The chemical content of a system is commonly quantified by its *metallicity*, expressed in logarithmic units relative to solar values:

$$[X/Y] = \log\left(\frac{X}{Y}\right) - \log\left(\frac{X}{Y}\right)_{\odot}. \tag{2.26}$$

Surveys have shown that a large fraction of intergalactic gas has been enriched to detectable levels: for instance, Simcoe et al. (2004) estimated that up to 70% of absorbers with $N_{\mathrm{H\,I}} > 10^{13.6}$ cm^{-2} exhibit metal lines, corresponding to at least half the cosmic gas mass.

In quasar spectra, metals are detected as absorption features imprinted at both shorter and longer wavelengths than the Lyman-α emission line (McQuinn, 2016).

Lines with rest-frame wavelengths below 1215.67 Å often overlap with the H I forest, complicating their identification, while those at longer wavelengths are more easily recognised in the red part of the spectrum[6]. Reliable identification requires multiple lines at a consistent redshift, since a single observed feature at $\lambda_{\rm obs}$ may correspond to different transitions. Confirmation is aided by two approaches: (i) the presence of *multiplets*, such as the Si IV (1393–1402 Å), C IV (1548–1550 Å) and Mg II (2796-2803 Å) doublets, which share common atomic origins and must exhibit consistent line profiles; or (ii) *co-occurring species*, where different elements (e.g. C, Si, O) trace the same absorbing gas. In the latter case, modest velocity offsets or variations in line width may arise due to species-dependent spatial distributions, but the coincidence in redshift remains the key diagnostic.

Once metal absorption is established at a given redshift, additional transitions can be searched systematically across the spectrum, allowing a detailed reconstruction of the physical and chemical state of the absorbing gas.

2.3 Statistics of the Lyman-α forest

At first glance, spectra of quasars at different redshifts exhibit a large variation in strength and number of Lyman-α absorption lines. Low redshift sightlines appear almost unabsorbed, with few and weak Lyman-α lines, whereas spectra at $z \sim 2-5$ show a crowded forest of deep lines, with heavier blending at the high-redshift end. At the highest redshifts closer to the reionisation epoch, $z \gtrsim 5.7$, the IGM is increasingly more neutral and opaque to Lyman-α radiation, resulting in almost completely absorbed sightlines, where lines are saturated and blended in one large absorption trough blueward of the quasar's Lyman-α emission, leaving only seldom transmission spikes. As shown in the previous section, fitting a Voigt profile to the H I Lyman-α absorption lines yields three parameters: the position (or redshift) of the absorber, its column density and its Doppler parameter, defining either its temperature, its turbulent motion or a combination of both. The evolution across cosmic time and across different cosmic environments of these three quantities is an important tool in the characterisation of the IGM properties. Below, we summarise the evolution and the distributions of the main measurable quantities of the Lyman-α forest.

2.3.1 Lyman-α Opacity

Given the relation between observed flux density and optical depth reported in Eq. 2.24, we can compute the *effective optical depth*, $\tau_{\rm eff}$, of the Lyman-α forest of a spectrum as

$$\tau_{\rm eff} = -\ln\langle e^{-\tau_\lambda}\rangle_\lambda, \tag{2.27}$$

[6] While not blending with H I lines, metal lines with rest-frame wavelength $\lambda > 1215.67$ Å suffer from sky contamination with atmospheric telluric lines.

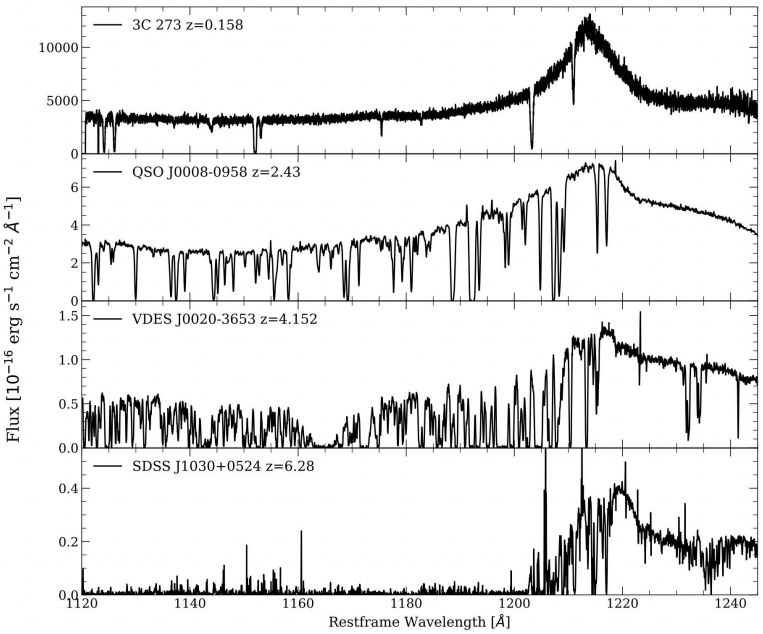

Figure 2.5: *Spectra of quasars spanning a redshift range* $0.1 \lesssim z_{\rm em} \lesssim 6.5$*, reported as a function of rest-frame wavelength, such their Lyman-α emission is located at* 1215.67 Å*. The Lyman-α forest spans the blue side of the spectra, and evolves in number and depth of absorption lines with redshift. At the highest redshifts, the Lyman-α forest is completely absorbed. Regardless of redshift, continuum emission redward of the Lyman-α emission is almost unabsorbed, showing only seldom narrow metal lines. From Tejos (2025).*

so that

$$\left\langle \frac{F_\lambda}{F^{\rm cont}} \right\rangle_\lambda = e^{-\tau_{\rm eff}}, \tag{2.28}$$

where the $\langle\rangle_\lambda$ indicates the averaging over all wavelengths. This definition was proposed to account for the difficulty of measuring τ_λ directly on observed spectra, especially for saturated lines[7]. Note that $\tau_{\rm eff} \neq \langle\tau\rangle$.

Observationally, the effective opacity of the Lyman-α forest evolves as a power law

$$\tau_{\rm eff} = \tau_0 \left(\frac{1+z}{1+z_0} \right)^\beta + C, \tag{2.29}$$

with $\beta = 2.9 \pm 0.11$ at low redshift ($2 < z < 5$, $z_0 = 3.5$, Becker et al., 2007) and $\beta = 13.7 \pm 1.5$ at high redshift ($4.8 < z < 6.2$, $z_0 = 4.8$, Bosman et al., 2022).This steep evolution of $\tau_{\rm eff}$ with redshift reflects the rapid increase in the neutral hydrogen fraction of the IGM, making the Lyman-α forest an important probe of the ionisation state of the Universe across cosmic time.

2.3.2 Number Density Evolution

Similarly to the evolution of the overall average optical depth of the Lyman-α forest, one can study the evolution of the number of absorption lines, $\mathcal{N}_{\rm abs}$, found in a spectrum. This evolution can be encoded in a measurable way as the *cosmic incidence* of absorption lines per unit redshift, $d\mathcal{N}_{\rm abs}/dz$. This quantity can be related to the physical number density of intervening absorption systems, $n_{\rm abs}(z)$, and their cross-section,

[7] For a saturated line $F_\lambda \sim 0$, where $\tau_\lambda \gtrsim 1$ is not univocally defined and a direct $F \rightarrow \tau$ is lost.

$\sigma_{\rm abs}(z)$. Following the standard formulation, this relation can be expressed as

$$\frac{d\mathcal{N}_{\rm abs}}{dz} = \frac{c}{H_0}\, n_{\rm abs}(z)\, \sigma_{\rm abs}(z)\, \frac{dX}{dz}. \tag{2.30}$$

The expansion of the Universe plays a role in the evolution of the line density through the differential absorption distance, dX/dz, i.e. the comoving path length along the sightline per unit redshift. In a flat ΛCDM Universe, dX/dz is given by (Bahcall et al., 1969):

$$\frac{dX}{dz} \equiv \frac{(1+z)^2}{\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}}. \tag{2.31}$$

It is convenient to group the physical terms into a single function, describing the physical evolution of the absorbers

$$\mathcal{N}_{\rm abs}(z) = \frac{c}{H_0}\, n_{\rm abs}(z)\, \sigma_{\rm abs}(z), \tag{2.32}$$

so that Equation 2.30 becomes

$$\frac{d\mathcal{N}_{\rm abs}}{dz} = \mathcal{N}_{\rm abs}(z)\, \frac{dX}{dz}. \tag{2.33}$$

A common parametrization assumes that $d\mathcal{N}_{\rm abs}(z)/dz$ follows a power-law evolution

$$\frac{d\mathcal{N}_{\rm abs}}{dz} = \mathcal{N}_{\rm abs,0}\, (1+z)^{\gamma_n}, \tag{2.34}$$

where $\mathcal{N}_{\rm abs,0}$ is the number density at $z = 0$, and the index γ_n quantifies the redshift evolution. For a non-evolving population with a fixed proper cross-section in an EdS Universe, the power-law evolution is expected to take place with an exponent $\gamma = 0.5$. Measurements (e.g. Janknecht et al. 2006; Kim et al. 2021) suggest that cosmic incidence evolution is better described with a double power-law. At $z \gtrsim 1.5$, absorber populations evolve significantly, with $\gamma_n \sim 1 - 3$[8], while at lower redshift the results are consistent with pure cosmological expansion, defining a stringent constraint on the evolution of the ionising UV background. Note that this approach cannot distinguish between changes in absorber number density and changes in their cross-section, since observations are sensitive only to the product $n_{\rm abs}(z)\, \sigma_{\rm abs}(z)$.

2.3.3 Column Density Distribution

The distribution of neutral hydrogen absorbers can be characterised through the Column Density Distribution Function (CDDF), defined as the number of systems per unit column density and unit absorption path length

$$f(N_{\rm H\,I}, z) \equiv \frac{\partial^2 \mathcal{N}_{\rm abs}}{\partial N_{\rm H\,I}\, \partial X} = \frac{\partial^2 \mathcal{N}_{\rm abs}}{\partial N_{\rm H\,I}\, \partial z} \left(\frac{dX}{dz}\right)^{-1}, \tag{2.35}$$

[8] γ_n depends on the range of column densities considered.

where dX/dz is the absorption distance. This function describes how absorbers of different column densities are distributed along the line of sight at a given redshift, and thus provides a fundamental statistical measure of the physical properties of the IGM.

Over a wide range of column densities, the CDDF is well approximated by a power law,

$$f(N_{\rm H\,I}, z) \propto {N_{\rm H\,I}}^{-\beta}, \tag{2.36}$$

with a slope of $\beta \simeq 1.5$ (Tytler, 1987; Danforth et al., 2016) for the low-density Lyman-α forest ($N_{\rm H\,I} \sim 10^{12} - 10^{17}\,{\rm cm}^{-2}$). At higher column densities, corresponding to Lyman limit systems (LLSs, $N_{\rm H\,I} \gtrsim 10^{17.2}\,{\rm cm}^{-2}$) and damped Lyman-$\alpha$ systems (DLAs, $N_{\rm H\,I} \gtrsim 10^{20.3}\,{\rm cm}^{-2}$), the distribution steepens and deviates from a single power law (Prochaska et al., 2010). Beyond simple power-law fits, the shape of $f(N_{\rm H\,I}, z)$ provides important constraints on both cosmology and astrophysics. Numerical simulations of structure formation reproduce the observed CDDF across more than ten orders of magnitude in column density (Rahmati et al., 2013), once ionisation corrections, self-shielding, and transitions to molecular gas are properly included. The remarkable agreement between observations and simulations suggests that our current understanding of the physical origin of hydrogen absorbers, spanning from tenuous Lyman-α forest clouds to dense damped systems associated with galaxies, is both coherent and robust.

The CDDF also serves as the starting point for deriving integrated quantities of cosmological interest. For instance, the contribution of neutral hydrogen to the cosmic mass budget can be expressed as (Schaye, 2001)

$$\Omega_{\rm H\,I}(z) = \frac{8\pi G m_{\rm H}}{3cH_0} \int N_{\rm H\,I}\; f(N_{\rm H\,I}, z)\, dN_{\rm H\,I}, \tag{2.37}$$

from which the mean gas density in the Universe is

$$\Omega_{\rm g}(z) = \frac{8\pi G m_{\rm H}}{3cH_0(1-Y)} \int \frac{N_{\rm H\,I}}{x_{\rm H\,I}(N_{\rm H\,I}, z)}\, f(N_{\rm H\,I}, z)\, dN_{\rm H\,I}, \tag{2.38}$$

where $Y \approx 0.24$ is the primordial helium mass fraction, and $x_{\rm H\,I} = n_{\rm H\,I}/n_{\rm H}$ is the hydrogen neutral fraction. Note that the neutral fraction does depend on neutral hydrogen column density, having typical values at $z \sim 2-4$ of $x_{\rm H\,I} \sim 10^{-5.5}$ in the low density sparse IGM ($N_{\rm H\,I} \lesssim 10^{16}$), $x_{\rm H\,I} \sim 10^{-3} - 10^{-1}$ in LLSs, and $x_{\rm H\,I} \sim 1$ in DLAs.

2.3.4 Line Width Distribution and the $b - N$ Diagram

As shown in Sect 2.2.1 (Eq. 2.14), the intrinsic width of H I Lyman-α absorption lines is defined by the Doppler parameter, b, which is a direct measure of the velocity dispersion within the absorbing gas cloud and is defined by both thermal and non-thermal (turbulent) motions. For this reason, the distribution of b-parameters is a primary diagnostic for measuring the temperature of the IGM, T_0, at mean cosmic density, providing critical results at low redshifts ($z \lesssim 0.5$) where discrepancies between observational data and hydrodynamical simulations have highlighted the need for refined

models. Simulations (e.g. Sherwood (Bolton et al., 2017b; Bolton et al., 2022b), Illustris TNG, Nelson et al., 2019) predict a narrower distribution of Doppler parameters compared to observations, indicating that the simulated IGM may be too cold or lack sufficient turbulence to match the observed line widths. In Part II of this Dissertation, we tackle this problem and exploit the high-z Lyman-α forest to constrain an exotic Dark Matter model that was put forward to solve this discrepancy.

The temperature of the IGM is also probed by the correlation between measured b-parameter and column density of Lyman-α lines. Observationally, the b–$\log N_{\mathrm{H\,I}}$ distribution exhibits a characteristic lower envelope, with a cutoff at low b values that increases weakly with column density. This cutoff is interpreted as the thermal limit of the IGM, corresponding to the temperature-density relation

$$T = T_0 \Delta^{\gamma-1}, \tag{2.39}$$

where $\Delta = \rho/\bar{\rho}$ is the overdensity, T_0 is the temperature at mean density, and γ parametrises the power-law scaling. The observed lower envelope therefore provides a direct measure of T_0 and γ, enabling reconstruction of the thermal history of the IGM.

Above this cutoff, the distribution shows a broad scatter, reflecting non-thermal contributions such as small-scale turbulence, bulk flows, and line blending. High column density systems often exhibit higher b values due to saturation effects and unresolved substructure. Comparisons with hydrodynamical simulations indicate that reproducing both the shape of the lower envelope and the scatter at higher b remains a stringent test for models of the IGM and feedback processes.

2.3.5 First Order Flux Statistics: Flux Probability Distribution Function

Within the Fluctuating Gunn-Peterson approximation (FGPA, Gunn et al., 1965), the optical depth (and therefore the flux) of a gas in photoionisation equilibrium is proportional to the density multiplied by the recombination rate. At $T \sim 10^4$ K this can be expressed as (Croft et al., 1998; Viel et al., 2005)

$$\tau \propto \rho_b^2 T^{-0.7} = A(z)(\rho_b/\bar{\rho}_b)^\beta, \tag{2.40}$$

where $\beta = 2 - 0.7(\gamma - 1)$ with γ being the power law scaling of the IGM equation of state, and

$$A = 0.433 \left(\frac{1+z}{3.5}\right)^6 \left(\frac{\Omega_b h^2}{0.02}\right)^2 \left(\frac{T_0}{6000\ \mathrm{K}}\right)^{-0.7} \left(\frac{h}{0.65}\right)^{-1} \left(\frac{H(z)/H_0}{3.68}\right)^{-1} \left(\frac{\Gamma_{\mathrm{H\,I}}}{1.5\times10^{-12}\ \mathrm{s}^{-1}}\right)^{-1}, \tag{2.41}$$

where $\Gamma_{\mathrm{H\,I}}$ is the H I photoionisation rate. This shows that the observed Lyman-α flux depends directly on

1. The underlying matter density distribution, ρ_b.
2. The thermal state (T_0 and γ), which affects line widths and the flux-density rela-

tion.

Instead of adopting a Voigt profile decomposition, IGM physical quantities can therefore be directly extracted from the statistical properties of the flux, the simplest of which is the flux Probability Distribution Function (PDF), P(F). High-flux pixels ($F \sim 1$) correspond to underdense, highly ionised gas, while low-flux pixels ($F \ll 1$) trace overdense structures. Thus, the PDF reflects both thermal broadening and the spatial distribution of density fluctuations (McDonald et al., 2000; Bolton et al., 2008; Calura et al., 2012; Cieplak et al., 2017).

Comparisons between observed PDFs and those from hydrodynamical simulations constrain the IGM's thermal and ionisation state. The abundance of high-flux pixels is sensitive to T_0 and γ, while the PDF width informs on small-scale density fluctuations and the hydrogen photoionisation rate (Viel et al., 2004; Peeples et al., 2010). At low redshifts ($z \lesssim 0.5$), the PDF is particularly useful due to the sparse line density, enabling precise measurements of diffuse, low-density gas.

2.3.6 Second-Order Flux Statistics: TPCF and 1D Power Spectrum

While the flux PDF provides a first-order characterisation of the Lyman-α forest by measuring the distribution of flux values, it does not capture the spatial correlations of absorption along the line of sight. Second-order statistics, such as the Two-Point Correlation Function (TPCF) and the 1D flux power spectrum, probe these correlations, offering complementary information about the underlying matter distribution and the thermal and dynamical state of the IGM.

The flux two-point correlation function along the line of sight is defined as

$$\xi_F(v) = \langle \delta_F(v_0)\, \delta_F(v_0 + v) \rangle, \tag{2.42}$$

where $\delta_F = F/\langle F \rangle - 1$ is the fractional fluctuation of the normalized flux, v is the velocity separation between two pixels, and $\langle \cdot \rangle$ denotes an ensemble or line-of-sight average. The TPCF quantifies the degree of clustering of absorption features: a positive $\xi_F(v)$ indicates that high or low flux values tend to cluster over a scale v, reflecting the underlying matter density correlations and thermal broadening effects.

The 1D flux power spectrum is the Fourier transform of the TPCF along the line of sight:

$$P_{\rm 1D}(k) = \int \xi_F(v)\, e^{-ikv}\, dv, \tag{2.43}$$

where k is the wavenumber corresponding to a velocity scale v. The 1D power spectrum provides a scale-dependent measure of fluctuations in the transmitted flux, sensitive to the amplitude and slope of the matter power spectrum on small scales, as well as to the thermal state of the IGM through line broadening. Unlike the PDF, which is a one-point statistic, $P_{\rm 1D}(k)$ encodes the clustering information and correlations between different regions along the line of sight.

The TPCF and the $P_{1D}(k)$ have become a pivotal tool in cosmology, enabling precise measurements of various fundamental parameters (Fernandez et al., 2024; Walther et al., 2025; DESI Collaboration et al., 2025a), including constraints on Warm Dark Matter (WDM, Irsic et al., 2017; Villasenor et al., 2023; Garcia-Gallego et al., 2025), Fuzzy Dark Matter (FDM, Iršič et al., 2017; Armengaud et al., 2017; Kobayashi et al., 2017; Rogers et al., 2021) and neutrino masses (Viel et al., 2005; Palanque-Delabrouille et al., 2015; Palanque-Delabrouille et al., 2020; Yèche et al., 2017; Chabanier et al., 2019; Ivanov et al., 2025).

Beyond single line-of-sight statistics, an even richer characterisation of the Lyman-α forest arises from correlations between neighbouring sightlines. The three-dimensional flux power spectrum, $P_{3D}(k)$, extends the 1D case by capturing clustering both along and transverse to the line of sight. This statistic is constructed from the Fourier transform of the full 3D flux two-point correlation function and encodes anisotropic information about large-scale structure, shaped by peculiar velocities and redshift-space distortions. Measurements of $P_{3D}(k)$ from ensembles of quasar spectra allow one to probe the matter distribution on cosmological scales and to detect features such as baryon acoustic oscillations (BAO, DESI Collaboration et al., 2025a) in the high-redshift universe. Complementary to this, the cross-correlation of flux fluctuations in closely separated quasar pairs or quasar lenses directly constrains the coherence length of absorption structures transverse to the line of sight, providing insights into the geometry, thermal state, and clustering of the IGM, as explored in Chapter 6. Together, these higher-order statistics link the small-scale physics of the IGM to the large-scale distribution of matter, strengthening the role of the Lyman-α forest as a precision cosmological probe.

2.4 Observational Techniques

The study of the Lyman-α forest relies on the acquisition of high Signal-to-Noise ratio (S/N) spectra of distant, luminous quasars. Given the faintness of these sources, such observations necessitate the use of the world's largest telescopes coupled with ultra-stable, high-resolution spectrographs. The following section details the fundamental principles of optical spectroscopy and the specific instrumental landscape that has enabled, and will continue to push, cosmological constraints derived from Lyman-α forest analysis.

2.4.1 Foundations of High-Resolution Optical Spectroscopy

Astronomical spectroscopy is the process of dispersing incident light into its constituent wavelengths to measure the emitted, absorbed, or scattered flux as a function of wavelength. This process is generally achieved using a spectrograph, which consists of four main optical components: an *entrance aperture* (slit or fibre), which isolates the sources of interest in the field, a *collimator* focusing the collected light on the *dispersive*

element, which separates wavelengths. The dispersive element can be a prism, a grism, an etalon, or an echelle grating, such as the ones used in this thesis. A camera then redirects the rays to a *detector*, which digitally registers the spectrum.

For high-resolution applications in cosmology, echelle spectrographs are the standard design. An echelle grating is a specialised diffraction grating characterised by deep, broad grooves and a high blaze angle. This design is leveraged to achieve high resolution while operating at very high diffraction orders. Because these high orders overlap, a second, less dispersive element, known as a cross-disperser (typically a grism or a prism), is required to separate the orders. This two-dimensional separation, high dispersion along one axis (echelle) and cross-dispersion along the orthogonal axis, efficiently packs a broad spectral range onto the rectangular surface of a modern detector, a format essential for capturing broad wavelength ranges in a single exposure. Figure 2.6 reports a detail of the detector image of a star spectrum taken with ESPRESSO, where three spectral orders are shown.

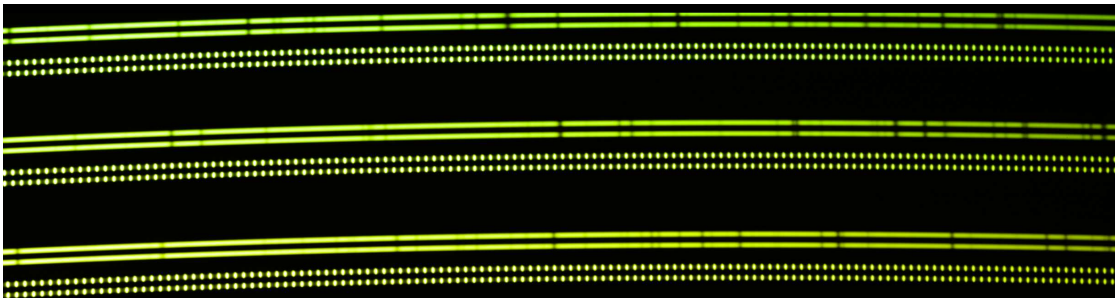

Figure 2.6: *Detail of a detector image of an ESPRESSO spectrum of a star taken during the instrument's commissioning. Solid arched paths across the detector are the spectral orders, with two traces per order due to slicing, where dark shaded regions are due to absorption lines in the stellar spectrum. The series of equispaced bright dots comes from the LFC calibration lamp and acts as an anchor point to provide a precise wavelength calibration of the spectrum. Credit: ESO/ESPRESSO Team.*

Spectral Resolution

The performance of a spectrograph is primarily defined by its spectral resolution, encoded in the resolving power R, which quantifies the instrument's ability to distinguish two closely spaced spectral features. It is formally defined as:

$$R \equiv \frac{\lambda}{\Delta\lambda}, \tag{2.44}$$

where λ is the central wavelength and $\Delta\lambda$ is the smallest resolvable separation between two spectral features, conventionally measured as the FWHM of the instrument's LSF.

For studies of the Lyman-α forest, high resolution is essential to faithfully resolve the line profiles, with typical $b \approx 20 - 30\,\mathrm{km\,s^{-1}}$. To accurately deconvolve the instrumental LSF from these thermal line widths, spectrographs must achieve $R \gtrsim 50,000$. Instruments used for precision cosmology, such as ESPRESSO, push this limit to $R \approx 140,000$ ($\Delta v \approx 2.1\,\mathrm{km\,s^{-1}}$) to resolve even the narrowest features and minimise broadening of the spectral signal.

2.4.2 The Very Large Telescope

The Very Large Telescope (VLT), operated by the European Southern Observatory (ESO) atop Cerro Paranal[9] in the Atacama Desert, has been a foundational facility for cosmological spectroscopy for over two decades. Comprised of four $8.2\ \mathrm{m}$ Unit Telescopes (UTs), the VLT provides the necessary photon-collecting area to target the most distant quasars. Its suite of instrumentation, particularly the high-resolution echelle spectrographs, has been instrumental in mapping the Lyman-α forest, facilitating the detailed analysis of individual absorption systems and two-point statistics. The VLT continues to serve as a high-precision platform, with its capabilities driving the necessary long-term monitoring campaigns for precision cosmology experiments.

2.4.3 High Resolution Spectroscopy: ESPRESSO

The Echelle SPectrograph for Rocky Exoplanets and Stable Spectroscopic Observations (ESPRESSO, Pepe et al., 2014; Pepe et al., 2021a) represents the current state-of-the-art in precision spectroscopy at the VLT. Its mission is driven by two seemingly disparate goals: exoplanet hunting[10] and precision fundamental physics[11]. Both science cases require extreme wavelength calibration accuracy, resolution and stability over decades-long periods of time.

Installed at the Coudé focus of the VLT, ESPRESSO can receive light from up to all four UTs, which are combined incoherently via optical fibres and fed into the spectrograph. This combined light mode (4UT mode) yields an effective collecting area of a 16m-class telescope. The instrument's resolution depends on the mode in which it is operated, varying between $R \sim 140,000$ and $R \sim 70,000$ for 1UT and 4UT modes, respectively. Both modes resolve Lyman-α lines.

Optics and Detectors

ESPRESSO's core optical system is housed in a pressure- and temperature-stabilised vacuum vessel to eliminate wavelength shifts induced by mechanical and thermal instabilities, as shown in Fig. 2.7. The light from the telescope is fed to the instrument by an optical fibre and is then split into two identical spectrograph units (a blue and a red arm) to simultaneously maximise efficiency across the entire operational bandwidth ($380\ \mathrm{nm}$ to $788\ \mathrm{nm}$). The heart of the dispersion system is a large, high-efficiency echelle grating coupled with a grism as the cross-disperser, projecting the spectrum onto the detector. The spectrograph employs low-noise Charge-Coupled Device (CCD) detectors. Specifically, each of the two arms is equipped with a large-format CCD of 92 mm x 92 mm, composed of $9k$ x $9k$ pixels of $10\ \mu\mathrm{m}$ in size. These devices are deeply

[9] 24°37′39″S 70°24′15″W, 2636 m from sea level

[10] Exoplanets: Suárez Mascareño et al. (2020), Tabernero et al. (2021), Doyle et al. (2022), Basant et al. (2025), Sartori et al. (2023), and Passegger et al. (2024)

[11] Fundamental physics: Murphy et al. (2022), Martins et al. (2022), da Fonseca et al. (2022), Guarneri et al. (2024), and Trost et al. (2025)

cooled to minimise thermal noise (dark current) and ensure that the noise in the final spectrum is dominated by unavoidable photon shot noise from the source, critical for faint quasar observations. The exceptional Quantum Efficiency (QE) and low readout noise of these detectors are paramount to achieving the necessary S/N for cosmological studies. Moreover, the optics are enclosed in a vacuum vessel ($P \sim 10^{-5}$ mbar) with controlled pressure and temperature ($\Delta T \lesssim$ mK) to avoid optical index refraction variations, thus ensuring the stability of the instrument over many years.

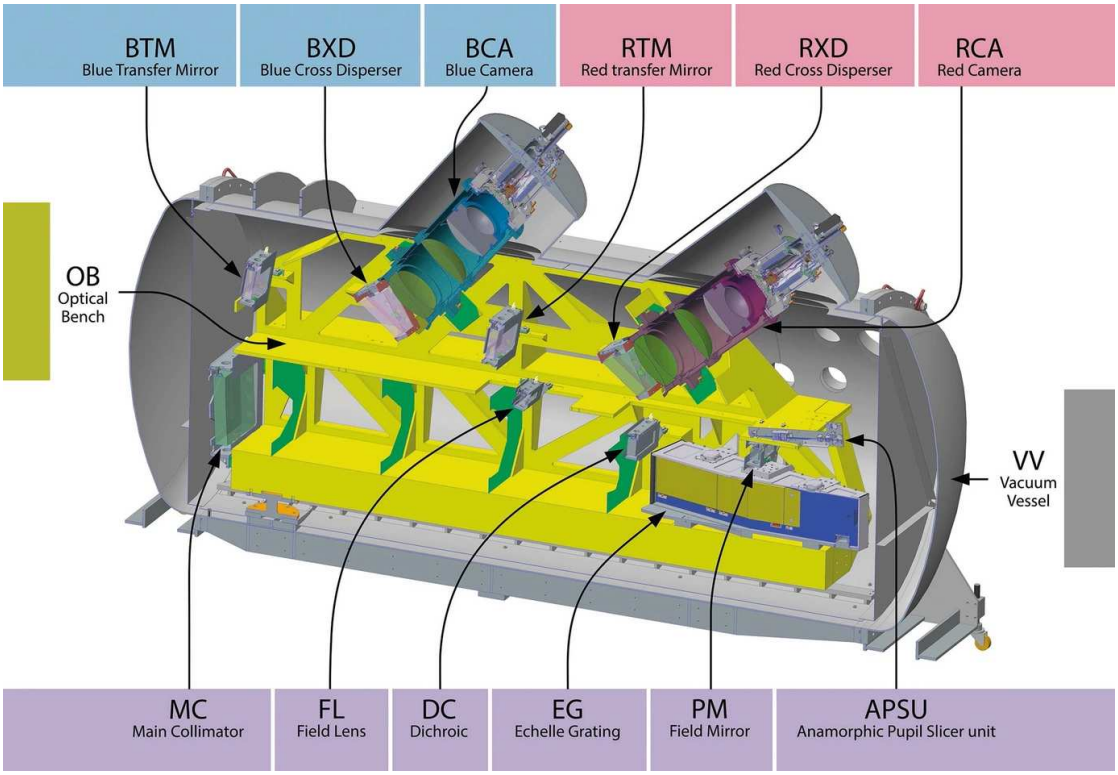

Figure 2.7: *Opto-Mechanical design of the ESPRESSO spectrograph. Credit:ESO/ESPRESSO Team.*

2.4.4 The Future: ELT and ANDES

The next generation of Lyman-α forest cosmology will be defined by the transition to 40-m class telescopes. The ESO Extremely Large Telescope (ELT; Padovani et al., 2023; Cirasuolo et al., 2024), currently under construction on top of Cerro Armazones[12] in the Atacama desert, will possess a 39-meter primary mirror, offering a photon-collecting area roughly 20 times that of a single VLT Unit Telescope. The key instrument for this research on the ELT will be the ArmazoNes high-Dispersion Echelle Spectrograph (ANDES; Marconi et al., 2022; Martins et al., 2024). ANDES is being designed to deliver ultra-high resolving power and stability superior to ESPRESSO, promising to revolutionise time-dependent cosmological measurements. Forecasts, including those presented in this thesis, indicate that the ELT, equipped with ANDES, will reduce the timescale required to achieve significant detection of the cosmological redshift drift

[12] 24°35′21″S 70°11′30″W, 3046 m above sea level.

signal from centuries (with VLT/ESPRESSO alone) to decades. The VLT will remain a crucial platform during this transition, acting both as a testbed for techniques (e.g., wavelength calibration and stability control) and as a complementary monitor in long-term observational campaigns.

2.4.5 Challenges

Achieving the extreme precision required for Lyman-α forest cosmology, particularly for long-baseline experiments like the redshift drift measurement, introduces significant observational and instrumental challenges that must be mitigated or modelled with high fidelity.

Wavelength Calibration

The precision of velocity measurements is fundamentally limited by the accuracy of the spectrograph's wavelength calibration. For Lyman-α forest work, especially the redshift drift, velocities must be known to the $\mathrm{cm\,s^{-1}}$ level across years. Traditional calibration sources, such as Thorium-Argon (ThAr) lamps, suffer from spectral instability and an insufficient density of reference lines in critical wavelength regions. Modern high-precision instruments overcome this by utilising a Laser Frequency Comb (LFC). The LFC provides a dense, stable grid of spectral lines whose frequencies are known to within an absolute accuracy better than $1\mathrm{cm\,s^{-1}}$ (Schmidt, 2024), enabling the required level of fidelity for comparison between observations separated by many years. Moreover, since the LFC lines are not resolved, the analysis of their profiles gives a faithful reconstruction of the instrument's LSF, which can be included in the profile analysis pipeline for more accurate astrophysical measurements and spectral modelling.

Sky Subtraction and Telluric Lines

Ground-based observations are unavoidably contaminated by the night sky background, which consists of a continuum component as well as numerous narrow emission features, mainly arising from hydroxyl (OH) airglow in the upper atmosphere. In ESPRESSO observations, the removal of the sky background is performed using the dedicated fibre B, which is placed on a blank region of the sky to record the night-sky spectrum simultaneously with the target spectrum in fibre A. The effectiveness of this subtraction depends on the relative throughput of the two fibres, the stability of the instrument, and the spatial uniformity of the sky background across the small angular separation between fibres. Residuals may persist due to fibre-to-fibre response differences, scattered light, or temporal variability of the airglow emission during the exposure. These residuals can contaminate faint quasar continua or introduce spurious structures in regions where the signal-to-noise is limited.

Beyond airglow emission and continuum, quasar spectra are also affected by atmospheric absorption lines, known as telluric lines. Telluric absorption lines, primarily from molecular oxygen (O_2) and water vapour (H_2O), complicate the analysis of the

Lyman-α forest, especially for high-z targets for which the Lyman-α transition shifts to longer wavelengths. Their accurate removal requires precise modelling of the atmospheric transmission, which depends on variable conditions such as pressure, temperature, airmass, and humidity throughout the exposure. Residuals from imperfect telluric correction can introduce spurious velocity shifts or artificially broaden spectral lines, thereby biasing constraints on IGM physics and cosmological parameters.

For a detailed model of sky emission and transmission at Cerro Paranal, ESO provides the `SkyCalc` tool, based on the Cerro Paranal Advanced Sky Model (Noll et al., 2012; Jones et al., 2013). Within the ESPRESSO pipeline, the `modelfit` recipe makes use of such synthetic spectra to correct for telluric absorption features (ESO, 2022).

Stability

Instrumental stability is paramount for precision spectroscopy. It is typically divided into short-term (jitter) and long-term (drift) components. Short-term mechanical and thermal instabilities lead to wavelength fluctuations during a single exposure, adding noise to the final co-added spectrum. Long-term stability is the dominant challenge for large Lyman-α cosmology projects. Even small thermal variations can cause refractive index changes in the spectrograph optics, leading to systematic shifts in the measured wavelengths that mimic a spurious cosmological signal. Spectrographs must operate within highly controlled environments (vacuum tanks, temperature stabilisation) to ensure that any observed spectral shift can be confidently attributed to cosmological evolution rather than instrumental drift.

In this Chapter, we have described the properties of the Lyman-α forest and its potential use as a cosmological tool. We showed its evolution with redshift and how this relates to the evolution of the IGM physical properties. We showed that the statistical distributions of the Lyman-α forest's properties are an excellent probe of cosmology, capable of constraining not only the cosmological parameters of ΛCDM , but also putting limits on alternative Dark Matter models and neutrino masses. On the astrophysical side, the Lyman-α is a key probe for the study of the cosmic UV background and the epoch of reionisation, as well as a tool to shed light onto the processes that shape galaxies, their formation and interactions with the cosmic web. The rest of this Thesis is devoted to the exploitation of the Lyman-α forest as a cosmological tool, through the use of both state-of-the-art instrumentation and computational techniques, first in the context of Dark Matter research (Part II, Chap. 3-4), and later as a probe of the accelerated expansion of the Universe (Part III, Chap. 5-7).

Part II

Dark Photon Dark Matter - Constraints from the Lyman-α forest

3

The Standard Model and Beyond

3.1 Introduction

The Standard Model of cosmology, ΛCDM , provides an exceptionally successful framework for describing the evolution and large-scale structure of the Universe. Its agreement with a wide array of observational data, from the CMB to the distribution of galaxies, is a landmark achievement of modern science. However, the model's foundational pillars – Dark Matter and Dark Energy – are placeholders for physical phenomena that remain profoundly misunderstood. This highlights a fundamental disconnect between our understanding of the cosmos on gravitational scales and the quantum theories describing fundamental particles. It is widely anticipated that the resolution of these mysteries will emerge from new physics Beyond the Standard Model (BSM), necessitating a collective approach that bridges cosmology and particle physics.

On the quantum scale, the Standard Model (SM) of Particle Physics represents our most fundamental, experimentally verified theory of matter and its interactions. Its predictive power has been confirmed with unparalleled precision at particle accelerators worldwide. Despite its successes, the SM is manifestly incomplete. The search for BSM physics is thus a primary objective of modern science, driven by both theoretical inconsistencies and observational evidence from cosmology.

This Chapter provides an overview on the foundational principles of particle physics necessary to contextualise searches for new, weakly interacting particles. We will begin by reviewing the architecture of the Standard Model, with an emphasis on its construction from the principle of local gauge symmetry. We will then discuss the theoretical and observational motivations for extending this framework. Finally, this formalism will be employed to introduce a simple yet highly motivated BSM paradigm: a *dark sector* featuring a new $U(1)$ gauge boson, the dark photon, which may serve as a viable dark matter candidate.

3.2 The Standard Model as a Gauge Theory

The Standard Model is a chiral quantum field theory whose structure is dictated by the principle of local gauge invariance under the symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$. Its components and interactions are not arbitrary but are a direct consequence of this underlying symmetry.

3.2.1 Symmetries and the Gauge Principle

The theoretical foundation of the SM is gauge theory, wherein forces arise from the requirement that the Lagrangian of the theory remains invariant under a set of local symmetry transformations. This is a profound physical principle, extending from the idea of symmetry as described by Noether's Theorem, which links every continuous symmetry to a conserved quantity.

A symmetry group is a set of transformations that leave the physical system unchanged. The groups relevant to the SM are Lie groups, specifically the unitary groups:

$U(1)$: the unitary group in one dimension, describes transformations of the form $\psi \to e^{i\alpha}\psi$, where ψ is a complex field. These are phase rotations. It is an Abelian group, meaning its elements commute.

$SU(N)$: the special unitary group in N dimensions, consists of $N \times N$ unitary matrices[1] with determinant +1. These groups are non-Abelian, i.e. their element do not commute. $SU(2)$ transformations act on two-component vectors (doublets), while $SU(3)$ acts on three-component vectors (triplets).

The core idea of a gauge theory is to promote a global symmetry, where the transformation is the same at all spacetime points, to a local one, where the transformation can vary with spacetime coordinates. Consider the Lagrangian for a free Dirac fermion field ψ with mass m:

$$\mathcal{L}_{\text{free}} = \bar{\psi}(i\gamma^\mu \partial_\mu - m)\psi, \tag{3.1}$$

where $\bar{\psi} = \psi^\dagger \gamma^0$ and γ^μ are defined as

$$\gamma^0 = \begin{pmatrix} I & 0 \\ 0 & -I \end{pmatrix}, \qquad \gamma^i = \begin{pmatrix} 0 & \sigma^i \\ -\sigma^i & 0 \end{pmatrix}, \quad (i = 1, 2, 3), \tag{3.2}$$

where I is the 2×2 identity matrix and σ^i are the Pauli matrices. This Lagrangian is invariant under a global $U(1)$ transformation $\psi(x) \to e^{i\alpha}\psi(x)$, where α is a constant. If we demand this symmetry hold only locally, i.e. for $\alpha = \alpha(x)$, the derivative term breaks the invariance. To restore it, one must introduce a new vector field A_μ, the gauge field, and define a covariant derivative, D_μ, as

$$D_\mu \equiv \partial_\mu - iA_\mu(x). \tag{3.3}$$

[1] A unitary matrix is a complex square matrix A that satisfies $A^\dagger A = AA^\dagger = 1$.

This new derivative is constructed to transform covariantly, $D_\mu\psi \rightarrow e^{i\alpha(x)}D_\mu\psi$, provided that the gauge field transforms as $A_\mu \rightarrow A_\mu + \partial_\mu\alpha(x)$. By replacing ∂_μ with D_μ and adding a kinetic term for the gauge field in Eq. 3.1, we obtain the Lagrangian for Quantum Electrodynamics (QED):

$$\mathcal{L}_{\text{QED}} = \bar{\psi}(i\gamma^\mu D_\mu - m)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}, \tag{3.4}$$

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the electromagnetic field strength tensor. The gauge principle has thus mandated the existence of the photon (A_μ) and prescribed the nature of its interaction with the electron (ψ).

3.2.2 The components of the Standard Model

Invoking the gauge principle, the entire zoo of observed particles acquires a natural organisation: fermions appear as matter fields arranged in representations of the gauge symmetries, gauge bosons emerge as the corresponding force carriers, and the Higgs scalar provides the mechanism for mass generation. Here, we briefly highlight the basic properties of each family of particles and their interactions.

The Particles of Matter: Fermions

Fermions are the building blocks of all matter we see. They are characterised by having half-integer spin $(\frac{1}{2}, \frac{3}{2}, \ldots)$ and they obey the Pauli Exclusion Principle, which states that no two identical fermions can occupy the same quantum state. This is essentially why solid matter exists and doesn't collapse on itself.

The SM fermions are organised into three "generations" of increasing mass. Every day matter is made almost exclusively of the first generation. They are further divided into two classes:

Quarks: These particles interact via the strong nuclear force $(SU(3)_C)$ and are always found bound together inside composite particles like protons and neutrons (collectively known as hadrons). There are six "flavours" of quarks:

- *Generation 1:* Up (u), Down (d)
- *Generation 2:* Charm (c), Strange (s)
- *Generation 3:* Top (t), Bottom (b)

A proton, for instance, is composed of two up quarks and one down quark (uud), while a neutron is made of one up and two down quarks (udd).

Leptons: These particles do not interact via the strong nuclear force. The most famous lepton is the electron (e^-). The six leptons are:

- *Generation 1:* Electron (e^-), Electron Neutrino (ν_e)
- *Generation 2:* Muon (μ^-), Muon Neutrino (ν_μ)
- *Generation 3:* Tau (τ^-), Tau Neutrino (ν_τ)

A key feature of the Standard Model is that it is a chiral theory. This means it treats left-handed and right-handed particles (defined by the orientation of their spin relative to their momentum) differently. Specifically, the weak nuclear force ($SU(2)_L$) only interacts with left-handed fermions, a fundamental asymmetry of nature.

The Forces and Their Carriers: Gauge Bosons

Differently from matter particles, forces are described as the exchange of force-carrying particles, known as gauge bosons. These particles have integer spin. The structure of these forces is dictated by the properties of the symmetry groups they arise from.

Electromagnetism: mediated by the photon (γ, or A_μ in field terms as in Eq. 3.4). The photon is massless, giving electromagnetism an infinite range. It interacts with any particle that has an electric charge and arises from the $U(1)$ symmetry, as shown above.

The Strong Nuclear Force: mediated by eight massless gluons (g). This force, arising from the $SU(3)_C$ symmetry of Quantum Chromodynamics, acts on the colour charge of quarks, binding them together into hadrons. Even though, just like photons, gluons are massless, their self-interaction leads to particle confinement and an effective range of the strong nuclear force on the order of 1 femtometre ($\sim 10^{-15}$ m), comparable to the typical size of an atomic nucleus.

The Weak Nuclear Force: mediated by the massive W^+, W^-, and Z bosons. Arising from the $SU(2)_L$ symmetry, this force is responsible for radioactive decay processes, like the beta decay of a neutron ($n \rightarrow p + e^- + \bar{\nu}_e$). The carriers have rest masses approximately $m_W \sim 80.36$ GeV c^{-2} (Chekhovsky et al., 2024) and $m_Z \sim 91.19$ GeV c^{-2} (Aaij et al., 2025), which is about 10^5 times the proton mass. Because of these large masses, the weak force has a very short effective range, roughly 10^{-8} m.

The Origin of Mass: The Higgs Mechanism

The gauge symmetry principle that so elegantly predicts the force carriers initially requires them to be massless. To account for the observed masses of the W and Z bosons, the SM incorporates the Higgs mechanism (Higgs, 1964). The universe is permeated by a Higgs field. In the hot, early universe, this field had a zero value, and the electroweak force was unified. As the universe cooled, the field underwent a phase transition and settled into a non-zero value, known as its Vacuum Expectation Value (VEV). This process, called spontaneous symmetry breaking, breaks the unified electroweak symmetry ($SU(2)_L \times U(1)_Y$) down to the $U(1)$ symmetry of electromagnetism. The W and Z bosons acquire their mass through their interaction with this non-zero VEV, while the photon remains massless. Fermions also acquire their mass by interacting with the Higgs field. The excitation of this field is the famous Higgs boson, H, discovered at CERN in 2012 (Aad et al., 2012; Chatrchyan et al., 2012).

3.3 Limitations of the Standard Model and BSM Portals

Despite its profound success, the SM is incomplete. Several observational and theoretical issues motivate the search for new physics:

Dark Matter: The SM contains no particle with the correct properties to account for the observed cosmological dark matter. A viable candidate must be massive, stable on cosmological timescales, and interact very weakly with SM particles.

Neutrino Masses: In the minimal SM, neutrinos are massless. However, the observation of neutrino oscillations proves they have small but non-zero masses, requiring an extension to the model. The most stringent measurements limit the mass of the neutrinos to be $m_\nu < 0.45\ \mathrm{eV\,c^{-2}}$ at 90% confidence level (KATRIN Collaboration et al., 2025).

The Hierarchy Problem: The mass of the Higgs boson receives enormous quantum corrections that would naturally drive it to the Planck scale ($\sim 10^{19}$ GeV). An extreme fine-tuning of parameters is required to keep it at its observed value (~ 125 GeV), suggesting a missing mechanism or symmetry (Hook, 2023).

Matter-Antimatter Asymmetry: The SM does not contain sufficient sources of CP violation to explain the observed baryon asymmetry of the Universe (Alonso-Álvarez et al., 2024).

Quantum Gravity: The SM does not include a quantum description of gravity, which is necessary for a complete theory of nature.

3.3.1 Portals to a Dark Sector

Many BSM theories postulate the existence of a *dark sector* containing new particles, including dark matter, that are neutral under the SM gauge groups. These particles can interact with the SM through "portals", i.e. operators consistent with all known symmetries that mediate interactions between the two sectors. The simplest such portals involve a new, low-mass mediator:

Vector Portal – The Dark Photon: A new $U(1)$ gauge boson, A', kinetically mixes with the SM photon (Fabbrichesi et al., 2020).

Scalar Portal – The Dark Higgs: A new scalar particle, S, couples to the SM Higgs boson via a term like $|S|^2|H|^2$ in the Lagrangian, allowing the Higgs to act as a bridge (Arcadi et al., 2021).

Neutrino Portal: New heavy neutral leptons (sterile neutrinos) that mix with the SM neutrinos (Fernández-Martínez et al., 2023).

These portals provide the lowest-dimensional, renormalizable interactions between a dark sector and the Standard Model, making them the most natural and well-motivated places to search for new physics. Among these, the vector portal is particularly compelling due to its direct analogy with QED and its rich phenomenological consequences, which span from particle physics experiments to cosmological observations. This thesis will focus exclusively on the dark photon and its potential role as dark matter.

3.4 The Dark Photon as Dark Matter

The construction of the SM from gauge principles provides a clear template for building well-motivated BSM theories. The most essential of such extensions involves positing a new, hidden $U(1)_D$ gauge symmetry. In direct analogy with QED, this symmetry gives rise to a new gauge boson, the dark photon, A', which mediates a dark electromagnetic force. While SM particles are neutral under $U(1)_D$, a portal connecting the two sectors can exist via kinetic mixing. This interaction is described by a gauge-invariant, renormalizable operator that couples the field strength tensors of the SM photon and the dark photon.

The effective Lagrangian describing the photon-dark photon system is (Holdom, 1986):

$$\mathcal{L}_{\gamma A'} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} - \frac{\epsilon}{2}F^{\mu\nu}F'^{\mu\nu} + \frac{1}{2}m^2_{A'}A'_\mu A'^\mu, \tag{3.5}$$

where $F_{\mu\nu}$ and $F'_{\mu\nu}$ are the field strength tensors for the ordinary and dark photons, respectively, ϵ is the dimensionless kinetic mixing parameter, and $m_{A'}$ is the dark photon mass. Two primary scenarios can be envisioned:

Massless Case ($m_{A'} = 0$)**:** A massless dark photon mediates a new long-range Coulomb-like interaction among dark-sector particles, effectively acting as an additional relativistic degree of freedom (*dark radiation*) in the early Universe. If kinetic mixing with the Standard Model photon is present ($\epsilon \neq 0$), it can also induce tiny electromagnetic couplings for hidden-sector states, leading to fifth-force effects that can be contained with precision QED test (Foot et al., 2015; Berlin et al., 2023).

Massive Case ($m_{A'} > 0$)**:** The dark photon can acquire mass, for instance, through a dark Higgs mechanism. If stable, this massive particle is a compelling dark matter candidate. For it to be the cold dark matter, it must be produced non-thermally in the early universe, as its interactions with the SM are too weak for thermal freeze-out to yield the correct relic abundance. Various production mechanisms have been proposed, such as the misalignment mechanism with a nonminimal coupling to the Ricci scalar (Arias et al., 2012; Graham et al., 2016; Alonso-Álvarez et al., 2020), via tachyonic instabilities with misaligned axions (Agrawal et al., 2020; Co et al., 2019; Bastero-Gil et al., 2019), or the decay of topological defects (Long et al., 2019) (see also East et al., 2022; Cyncynates et al., 2023; Cyncynates et al., 2024 for more critical studies on production mechanisms).

Throughout, we will focus only on the latter scenario, presenting models in which massive dark photons constitute the entirety of the dark matter budget of our Universe

Experimental Limits

A vast experimental program has been developed to *directly detect* dark photon DM in the laboratory (Caputo et al., 2021a), using cavities, dielectric disks, dish antennae, plasmas, LC circuits, electric-field radios, and small electric-dipole antennas. How-

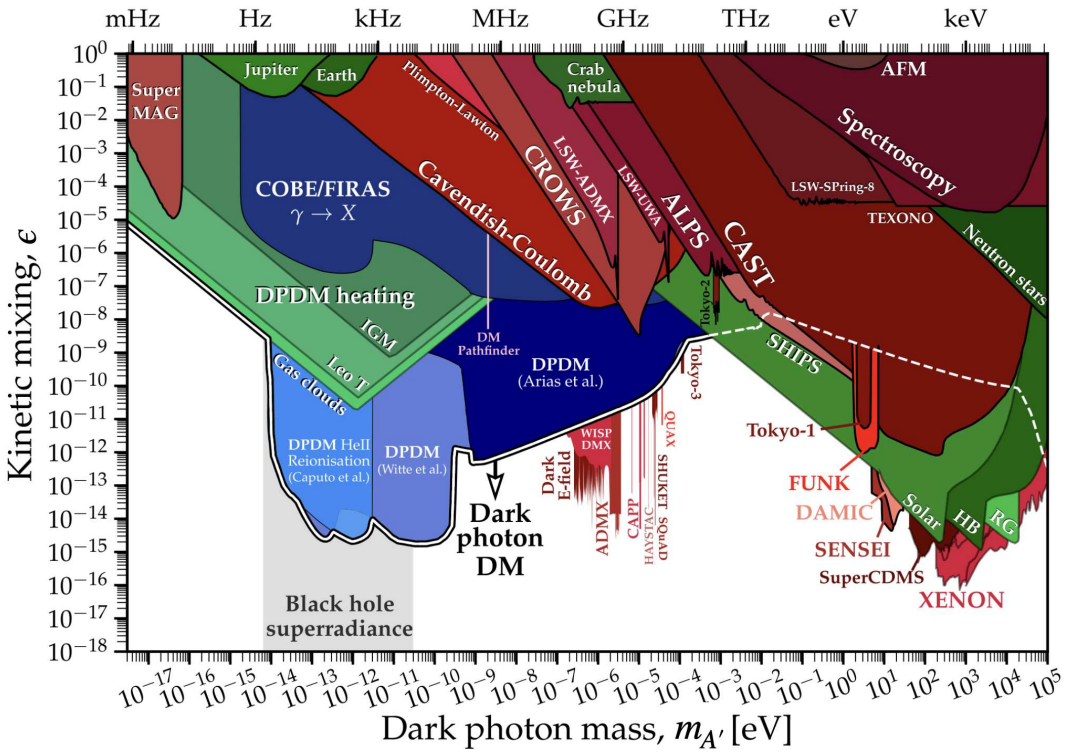

Figure 3.1: *Collection of constraints on dark photon mass, $m_{A'}$, and kinetic mixing parameter, ϵ, with the SM photon. The general colour scheme is cosmological bounds in blue, experimental bounds in red, and astrophysical bounds in green. The thick white line that divides the parameter space in two is the upper limit for which dark photons are a viable candidate for 100% of the DM. Adapted from Caputo et al. (2021a).*

ever, laboratory experiments in the quest for dark matter are typically limited by their size $L_{\rm exp}$, and have difficulty accessing A' masses smaller than $m_{A'} \lesssim \hbar/(L_{\rm exp}c^3)$. Astrophysical and cosmological probes are generally better suited for smaller masses, since they provide longer baselines for the interaction between A' DM and photons. Some observables of this kind include spectral distortions to the CMB (Mirizzi et al., 2009; McDermott et al., 2020; Witte et al., 2020; Caputo et al., 2020a; Caputo et al., 2020b), the solar corona (An et al., 2021; An et al., 2024), and the Earth's ionosphere (Beadle et al., 2024), among others. Figure 3.1 reports a comprehensive set of cosmological, experimental and astrophysical bounds on dark photon mass and mixing.

3.4.1 Cosmological Phenomenology and Resonant Conversion

The coupling between ordinary and dark photons gives rise to a rich phenomenology, and recent years have seen the field bloom with new detection ideas. In a cosmological setting, the oscillations between photon and dark photon enabled by the mixing operator are strongly affected by the properties of the medium in which the particles live. In an ionised plasma, photons acquire an effective mass that depends on the local free

electron density, n_e, defined by

$$m_\gamma^2(\mathbf{x}, z) \approx \frac{4\pi\alpha_{\rm EM}\hbar^2}{m_e c^2} n_e(\mathbf{x}, z) \simeq 3.5 \times 10^{-21}\ {\rm eV}^2\, c^{-4} \left(\frac{n_e(z, \mathbf{x})}{\rm cm^{-3}}\right). \tag{3.6}$$

Whenever and wherever the effective photon mass matches the dark photon mass, $m_{A'}^2 = m_\gamma^2(z, \mathbf{x})$, the conversion between the two species is resonantly enhanced. Assuming this condition is met at time $t_{\rm res}$, the transition probability is given by (Caputo et al., 2020a; Caputo et al., 2020b):

$$P_{A'\to\gamma}(\mathbf{x}, t_{\rm res}) = \frac{\pi\epsilon^2 m_{A'} c^2}{2\hbar} \left|\frac{d \ln m_\gamma^2(\mathbf{x}, t)}{dt}\right|^{-1}_{t=t_{\rm res}}. \tag{3.7}$$

Low-mass dark photons ($m_{A'} \sim 10^{-15} - 10^{-12}$ eV c^{-2}) convert to low-frequency ($\nu = m_{A'}c^2/h_P \sim 0.2 - 240$ Hz)[2] photons that are rapidly thermalised by the IGM plasma, primarily through inverse bremsstrahlung (free-free absorption). The mean free path of these photons (in proper units) is approximated by (Bolton et al., 2022a)

$$\lambda_{ff} \sim \frac{14\ {\rm kpc}}{(1+z)^6} \Delta^{-2} \left(\frac{T}{10^4\ {\rm K}}\right)^{3/2} \left(\frac{m_{A'}}{10^{-13}\ {\rm eV}\, c^{-2}}\right)^2, \tag{3.8}$$

where $\Delta = \rho/\bar{\rho}$ is the gas overdensity relative to the cosmic mean. This mean free path is significantly smaller than the typical scale of Lyman-α absorbers (see Chapter 6), ensuring that their energy is deposited locally, within the environment that prompted the oscillation. This process results in a net energy injection into the IGM, increasing the gas temperature at the density where the resonance occurs. The energy injection per baryon can be approximated as:

$$E_{A'\to\gamma} \sim 2.5\ {\rm eV} \left(\frac{\epsilon_{-14}}{0.5}\right)^2 \left(\frac{m_{-13}}{0.8}\right) \left(\frac{1+z_{\rm res}}{3}\right)^{-3/2}, \tag{3.9}$$

where $z_{\rm res}$ is the redshift of resonance, $\epsilon_{-14} \equiv \epsilon/10^{-14}$, and $m_{-13} \equiv m_{A'}/\left(10^{-13}\ {\rm eV}\, c^{-2}\right)$. This energy injection results in a temperature increase of the hydrogen gas at specific densities, causing Lyman-α lines associated with such environments to broaden, while leaving other lines unmodified.

An approximate expression for $z_{\rm res}$ is:

$$1 + z_{\rm res} \sim 2.75 \left(\frac{m_{-13}}{0.8}\right)^{2/3} \frac{1}{\Delta^{1/3}}. \tag{3.10}$$

At fixed overdensity, lighter models met the resonance condition at later times than heavier models. At fixed mass, resonant conversions occur in underdense regions first (voids), before proceeding to overdense ones (filaments and nodes).

Recent observations of the low-redshift IGM (Viel et al., 2017; Gaikwad et al., 2017;

[2] $h_P = 2\pi\hbar$ is the Planck's constants, not to be confused with h, the dimensionless Hubble parameter.

Burkhart et al., 2022) have revealed a tension between the observed gas temperatures and predictions from standard cosmological hydrodynamical simulations. Specifically, the IGM at $z < 1$ is significantly warmer than can be explained by photo-heating from the UV background and shock heating by AGN feedback alone. This long-standing discrepancy suggests that an additional, yet-unidentified source of energy injection may be at play.

In this context, Bolton et al. (2022a) proposed that this additional heating can be attributed to ultralight A' DM mixing in the IGM. By performing the first hydrodynamical simulations to self-consistently include the thermal energy injection from the resonant conversion of dark photons, the authors demonstrate that a dark photon with a mass $m_{A'} \sim 0.8 \times 10^{-13}\mathrm{eV}\, c^{-2}$ and a kinetic mixing parameter $\epsilon \sim 0.5 \times 10^{-14}$ can successfully reconcile the observed temperatures of the low-redshift IGM with the simulations.

The next Chapter follows the findings of Bolton et al. (2022a) and complements them by searching for the thermal trace of A' DM oscillations in the IGM at $z \sim 2.7$.

4

Dark Photon Dark Matter in the IGM

The previous Chapter introduced the Dark Photon as a simple extension of the standard model of particle physics and a viable candidate to explain the dark matter in our Universe. We summarised the phenomenology of such a DM particle and how it was proposed to solve the low-z IGM temperature tension between observations and simulations. In this Chapter, we want to contribute to the effort of putting even more stringent limits on the particle mass and mixing, exploiting the IGM as a particle detector and the Lyman-α forest as a calorimeter. We aim to complement the low-z findings of Bolton et al. (2022a) with an analysis of the Lyman-α forest in an ultra-high-S/N spectrum at $z \sim 2.7$ and a comparison with state-of-the-art cosmological simulation, via a novel flux summary statistic that is sensitive to the temperature of the underdense gas in which A' DM oscillations occur.

4.1 A' DM heating in cosmological hydrodynamical simulations

To study the impact of A' DM on the Lyman-α forest, we adopted high-resolution cosmological hydrodynamical simulations of the Lyman-α forest performed with a modified version of P-Gadget-3, which is an extended version of the publicly available Gadget-2 code Springel (2005). We used the code originally modified for the Sherwood-Relics simulation project (Puchwein et al., 2023; Bolton et al., 2017a), where it was demonstrated that these models reproduce the observed Lyman-α forest at $2 \leq z \leq 5$. Our fiducial simulation box size is $L = 40h^{-1}\mathrm{Mpc}$ with 2×1024^3 DM and gas particles, giving particle masses of $M_{\mathrm{dm}} = 4.3{\times}10^6 h^{-1}\,\mathrm{M}_\odot$ and $M_{\mathrm{gas}} = 7.97{\times}10^5 h^{-1}\,\mathrm{M}_\odot$. In Sect. 4.3.2 we demonstrate that this choice yields well-converged results.

The simulations were all started at $z = 99$, with initial conditions generated on a regular grid using a ΛCDM transfer function generated by CAMB (Lewis et al., 2000).

The simulations were evolved to $z = 2$, and the gas thermo-chemistry was obtained by solving the full network of non-equilibrium photo-ionisation equations for hydrogen and helium (Puchwein et al., 2015). Ionisation and heating by UV photons was followed using the empirically-calibrated synthesis UV background model of Puchwein et al. (2019), with photo-heating rates adjusted by a factor of 0.8 to match more recent IGM temperature measurements presented by Gaikwad et al. (2021). This adjustment is equivalent to using a slightly softer UV background spectrum relative to the original model used by Puchwein et al. (2019); lowering the IGM temperature by a few thousand degrees Kelvin. The exact spectral shape of the UV background is uncertain, so this approach ensures good agreement with more recent observational constraints on the IGM gas temperature that were unavailable when Puchwein et al. (2019) was published. Note also that the redshift evolution of the UV background model remains unchanged. All gas particles in the simulations with density $\Delta > 10^3$ and temperature $T < 10^5\,\mathrm{K}$ were converted into collisionless star particles Viel et al. (2004). We assumed an underlying ΛCDM cosmology with the following parameters: $\Omega_\mathrm{m} = 0.308$, $\Omega_\Lambda = 0.692$, $h = 0.678$, $\Omega_\mathrm{b} = 0.0482$, $\sigma_8 = 0.829$, $n = 0.961$, and a helium mass fraction $Y_\mathrm{p} = 0.24$.

We modified these simulations to include energy injection into the baryonic gas particles due to $A' \rightarrow \gamma$ following Bolton et al. (2022a). This was included *in addition* to photo-heating from our baseline UV background model, by applying a direct energy injection into the gas particles (following Eq. 3.9) when the resonant condition $m_\gamma^2 = m_{A'}^2$ is met. We performed 25 simulations spanning the A' parameter space, with masses $-13.4 \leq \log(m_{A'}/\mathrm{eV}\,c^{-2}) \leq -11.9$ and kinetic mixing $-15.0 \leq \log\epsilon \leq -13.5$. Table **??** reports the A' parameters considered. We also performed 4 additional simulations with no A' heating and different box sizes and number of particles to test for convergence with mass resolution and box size.

$\log(m_{A'}/\mathrm{eV}\,c^{-2})$	$\log\epsilon$
−13.40	−15.00, −14.50, −14.00
−13.15	−14.75, −14.15, −14.00, −13.50
−12.90	−15.00, −14.50, −14.15, −14.00
−12.65	−14.75, −14.15, −14.00
−12.40	−15.00, −14.50, −14.15, −14.00
−12.15	−14.75, −14.15, −14.00, −13.50
−11.90	−15.00, −14.50, −14.00

Table 4.1: *Summary of the A′ DM parameters assumed in the hydrodynamical simulations used in this work.*

Fig. 4.1 shows slices of a simulated box that assumes the best-fit A' DM model proposed by Bolton et al. (2022a) ($m_{A'} = 8\times10^{-14}\,\mathrm{eV}\,c^{-2}$, $\epsilon = 5\times10^{-15}$) at $z = 3.2$, 2.8 and 2.4 from left to right. Resonant conversion occurs at $z \sim 1.75$ for mean density regions. The middle and bottom panels show the gas temperature and overdensity distributions. In contrast, the top panels show the ratio between the gas temperature with the

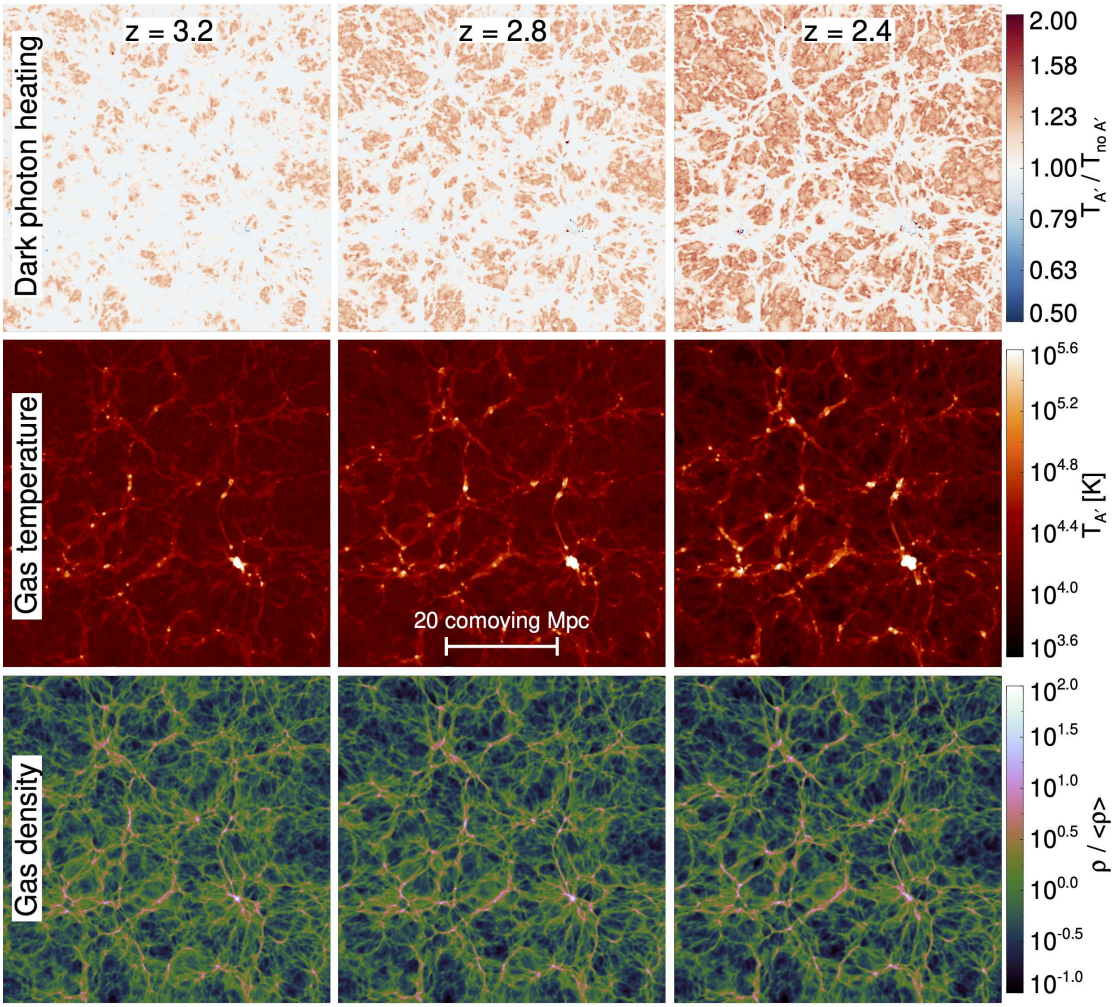

Figure 4.1: *Slices* $(500\ h^{-1}\mathrm{kpc}$ *thick*$)$ *of a simulation box with* A' *DM heating of the best-fit model of Bolton et al. (2022a)* $(m_{A'} = 8\times10^{-14}\ \mathrm{eV}\,c^{-2},\ \epsilon = 5\times10^{-15})$ *at redshift* $z = 3.2, 2.8$ *and* 2.4 *from left to right. Each panel is* $40\ h^{-1}\mathrm{Mpc}$ *on the side. The middle and bottom panels show the time evolution of the baryonic gas temperature and overdensity, respectively. The upper panels show the ratio of the gas temperature in the* A' *DM run relative to the scenario without extra heating. These panels highlight the amount of energy deposited by dark photon resonant conversions and the regions in the cosmic web that are affected by it. Note that the colour bar scales are all logarithmic. Cosmic time evolves from left to right.*

A' DM model and the baseline ΛCDM scenario without exotic heating, highlighting the regions affected by dark photon resonant conversions and the amount of energy injected into the IGM. The plot shows that the A' DM model proposed by Bolton et al. (2022a), in the considered redshift range, modifies the temperature of the baryonic gas up to a factor ~ 2 exclusively in the underdense regions where $\log\Delta \lesssim -0.5$ (e.g. voids) while not affecting the gas in overdense regions (e.g. filaments and nodes). Due to the redshift evolution of the free electron density n_e, given a specific mass, A' DM converts into photons in increasingly dense gas regions as cosmic time passes. Eventually, for the model shown, filaments and knots will be heated around $z \sim 0$.

Fig. 4.2 shows the distribution of the volume weighted temperature as a function of overdensity for the baseline ΛCDM model without extra heating and three A' DM sce-

narios at redshifts $z = 3.4, 2.7, 2.0$. The three models shown have different A' masses but the same kinetic mixing $\log \epsilon = -14.0$.

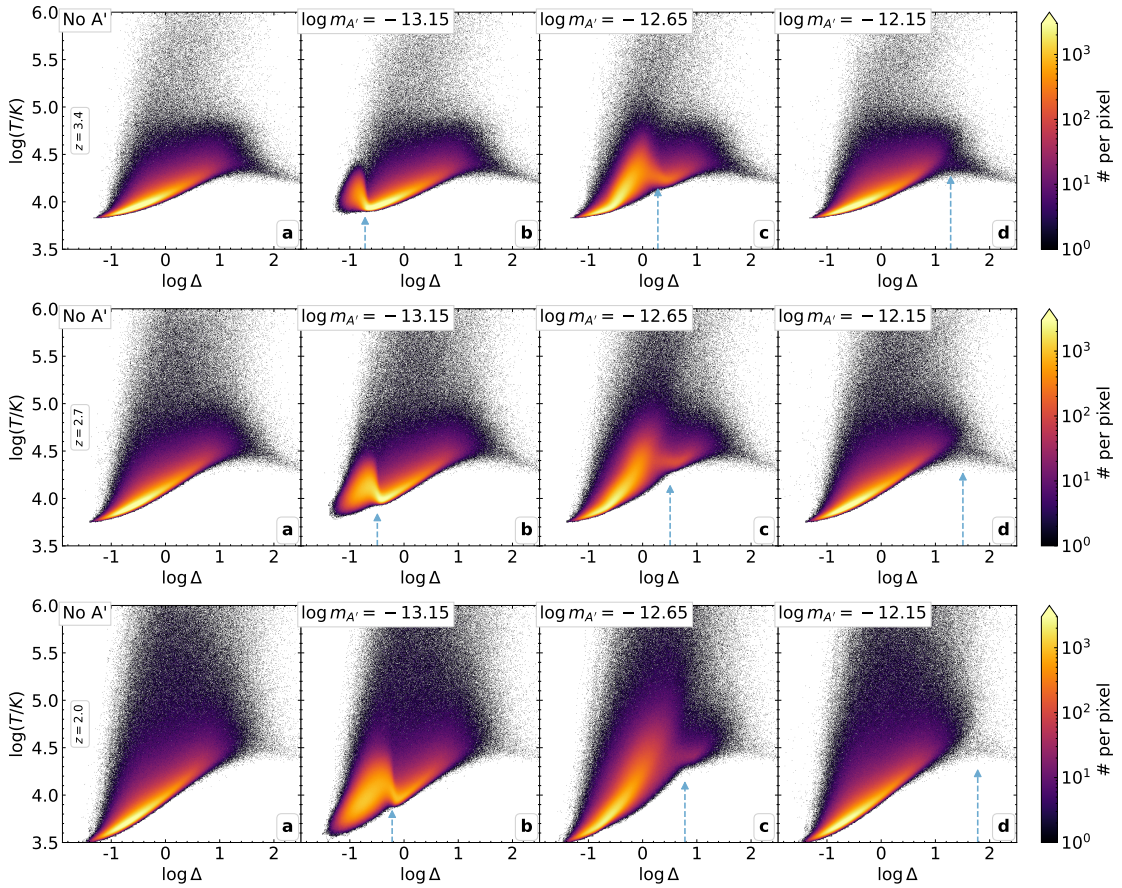

Figure 4.2: *Distribution of the volume weighted gas temperature and overdensities $\Delta = \rho/\bar{\rho}$, in terms of mean density, of different models at redshift $z = 3.4$ (top panels), $z = 2.7$ (middle panels) and $z = 2.0$ (bottom panels).* ***a)*** *Baryonic gas distribution in the baseline ΛCDM model without exotic heating. The core of the distribution at low densities is well described by a power law of the form $T = T_0 \Delta^{\gamma-1}$ with $\gamma = 1.3$.* ***b - d)*** *Temperature-density distribution of gas in models with A' DM heating, where dark photons have masses of $\log(m_{A'}/\mathrm{eV}\, c^{-2}) = -13.15$ (b), -12.65 (c) and -12.15 (d). All three models assume a kinetic mixing $\log \epsilon = -14.0$. A' DM heating induces a sharp increase in the gas temperature at the overdensities for which the resonant condition is met. This overdensity is highlighted with a dashed blue arrow, scaling as $\Delta_{\rm res} \propto m_{A'}^2 (1+z)^{-3}$. The amplitude of the temperature increase depends on the amount of energy injected into the gas due to dark photon oscillations, scaling as $\Delta T_{\rm res} \propto E_{A' \to \gamma} \propto m_{A'} \epsilon^2 (1+z)^{-3/2}$. For a fixed $m_{A'}$, a weaker kinetic mixing produces a less severe temperature feature, maintaining its density position. At the limit, models with $\log \epsilon \ll -14$ are not distinguishable from the model without extra heating, regardless of $m_{A'}$.*

The presence of A' DM heating can be summarised as a sudden increase, relative to the baseline model, of the temperature distribution in a narrow density interval, corresponding to regions where the resonance condition was met at or just before the observed redshift ($m_\gamma^2(z, \mathbf{x}) = m_{A'}^2$). The overdensity position of the resonance feature, $\Delta_{\rm res}$, depends solely on A' mass and redshift, scaling as

$$\Delta_{\rm res} \propto m_{A'}^2 (1+z)^{-3}. \tag{4.1}$$

Increasing the kinetic mixing ϵ increases the amplitude of the temperature feature, scaling as

$$\Delta T \propto E_{A' \to \gamma} \propto m_{A'} \epsilon^2 (1+z)^{-3/2}. \tag{4.2}$$

Therefore, models with weaker mixing than those shown in Fig. 4.2, produce a smaller temperature increase.

The distributions show that the IGM is strongly sensitive to the presence of A' DM, displaying an anomalous and measurable thermal excess proportional to the amount of energy injected in the gas due to dark photon conversions, making it an appealing proxy to probe the presence of A' DM.

Note that if the A' DM mass is such that a resonance never happens within the gas densities probed by the Lyman-α forest, then we do not see any significant effect. Thus, we expect to be sensitive only to a restricted range of A' DM masses; higher or lower masses may be accessed with other probes (see Caputo et al. (2021a) for an overview).

4.2 Discovering A' DM with the Lyman-α forest transmitted flux

The presence of A' DM may impact the Lyman-α forest in the spectra of distant bright quasars, caused by intergalactic neutral hydrogen gas. The shape of the absorption lines is directly related to the thermodynamic state of the absorbing gas, and the presence of A' DM heating will therefore modify the line widths and depths in a unique way, depending on the A' DM mass and kinetic mixing. In this section, we carried out a complete comparison between synthetic spectra extracted from the simulation boxes and a high-resolution, high-signal-to-noise observed spectrum of the Lyman-α forest. We exploited the distribution of flux values across the sightlines, which is sensitive to the presence of A' DM heating, to compare data and simulations.

We focused on the redshift range $z > 2$ for which the Lyman-α forest is accessible by ground-based telescopes. For lower redshifts, the Lyman-α transition falls below the atmospheric UV absorption cut-off ($\lambda < 300$ nm). Moreover, we limited our analysis to $z < 4$, as at higher redshifts the amount of neutral hydrogen in the IGM increases, leading to a tightly crowded and blended Lyman-α forest in the spectra of high redshift quasars. In this instance, a proper assessment of absorption line profiles becomes complex, and the effects of A' DM heating become more difficult to detect.

4.2.1 Simulated Lyman-α spectra

From the simulation boxes, we produced synthetic quasar sightlines by extracting 5000 pencil-beam skewers piercing the simulation volume, at redshift intervals of $\Delta z = 0.1$, each divided into 2048 bins along its length, saving the gas temperature, gas overdensity, local peculiar velocity and neutral hydrogen fraction. The skewers were then used to compute the Lyman-α optical depth, $\tau_{\mathrm{Ly}\alpha}$, as a function of wavelength, using the

interpolation scheme of Theuns et al. (1998) combined with the Voigt profile approximation of Tepper Garcia (2006). The normalised transmitted flux, i.e. the quasar flux fraction not absorbed by the neutral hydrogen, is defined as $F = e^{-\tau_{\mathrm{Ly}\alpha}}$.

Fig. 4.3 shows an example of the simulated normalised transmitted flux at $z = 2.7$, extracted from the same models reported in Fig. 4.2. The baseline model without extra heating due to A' DM is shown as a dashed pink line, whereas the three models with A' DM are shown with blue shaded lines, with darker tones denoting heavier masses. The excess heating due to A' DM resonant conversions translates into the broadening of the Lyman-α lines associated with gas particles having overdensities close to the resonance density Δ_{res}.

For light mass A' DM models (e.g. $\log(m_{A'}/\,\mathrm{eV}\,c^{-2}) = -13.15$), the resonance condition happens in underdense gas at $\Delta_{\mathrm{res}} \ll 1$ at the redshift in consideration, $z = 2.7$. This low-density gas absorbs only a small fraction of the incoming flux, and thus any modification of its thermodynamic state only affects the weak Lyman-α lines close to the continuum, leading to a spectral shape that is not distinguishable by visual inspection from the baseline model without exotic heating.

On the other hand, for intermediate masses (e.g. $\log(m_{A'}/\,\mathrm{eV}\,c^{-2}) = -12.65$) the conversion occurs at $\Delta_{\mathrm{res}} \sim 1$. This gas is sufficiently dense to leave many strong, yet not saturated, lines in the spectrum, and an exotic energy injection leads to a visible difference with respect to the baseline no-A' model.

Lastly, more massive A' DM models (e.g. $\log(m_{A'}/\,\mathrm{eV}\,c^{-2}) = -12.15$) inject energy into gas at $\Delta_{\mathrm{res}} \gg 1$. This gas is mainly associated with very dense structures, such as filaments and knots in the cosmic web, whose absorption lines are deep, saturated and less likely to appear in a spectrum, due to the reduced number and relatively small cross-section of such structures. Moreover, the profile of saturated lines, i.e. lines that absorb 100% of the incoming photons around the transition's wavelength, is not as sensitive to a change in gas temperature as for non-saturated lines, limiting the detectable effects of A' DM heating. The spectrum of such massive A' DM models only marginally differs from the no-A' case.

Note, however, that the spectra shown in Fig. 4.3 are noiseless and with infinite spectral resolution, i.e. are not affected by instrumental defects that will further reduce the differences between A' DM and no-A' models. All three A' DM models shown assume a kinetic mixing $\log \epsilon = -14.0$. Models with weaker mixing show lines with a less severe anomalous broadening, down to the limit of producing the same normalised transmitted flux as the model without A' DM heating, regardless of $m_{A'}$.

4.2.2 Data and mock spectra

The presence of A' DM heating in the IGM, at the redshift in consideration, broadens the weak and non-saturated Lyman-α absorption lines close to the quasar emission continuum. As shown in Fig. 4.3, we are chasing differences in the line profile shape with and without A' that may be fairly small. Therefore, in order to be able to detect (or

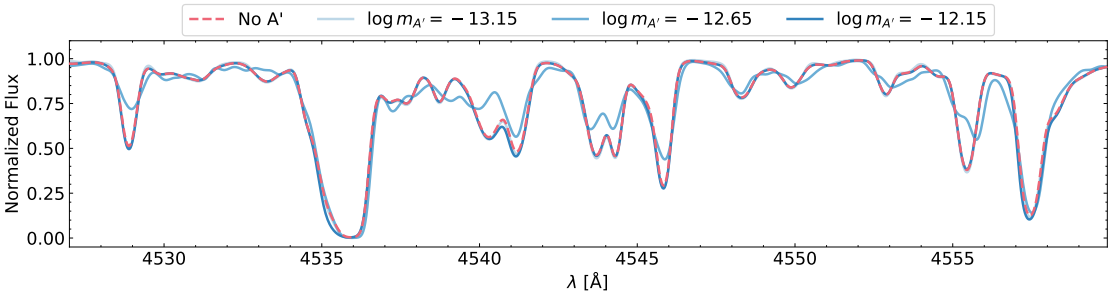

Figure 4.3: *Section of a synthetic noiseless sightline extracted from different simulated boxes with and without A′ DM heating. The dashed pink line is the baseline ΛCDM case without A′. Blue shaded lines show the transmitted fluxes with A′ DM heating, where darker tones represent more massive models. The three A′ DM models reported are the same as in Fig. 4.2 and all assume strong kinetic mixing* ($\log \epsilon = -14.0$). *Only one model visibly differs significantly from the no-heating case* ($\log m_{A'}/\,\mathrm{eV}\,c^{-2} = -12.65$) *and is ruled out by our later analysis.*

constrain) the presence of exotic A' DM heating, we ought to analyse high-resolution quasar spectra with very high S/N levels.

With this in mind, we analysed the high-resolution spectrum of the quasar HE0940-1050 ($z_{\mathrm{em}} = 3.0932$) (Rorai et al., 2016) observed with Ultraviolet and Visual Echelle Spectrograph (UVES) at VLT (Dekker et al., 2000) as the brightest object in the UVES Large Program (Bergeron et al., 2004), having an average S/N in the Lyman-α forest of 280 per resolution element. This is the high-resolution spectrum of a quasar at $2 < z < 3$ with the highest S/N found in the literature and thus is the best candidate for the following analysis.

This spectrum has been used by Rorai et al. (2016) to constrain the thermal properties of the underdense IGM, through a novel manipulation of the normalised flux. By regulating the flux continuum level and rescaling the Lyman-α optical depth before computing the distribution of the flux values of each pixel across the sightline, i.e. computing the Regualted and Transformed Probability Distribution Function (tPDF), Rorai et al. (2016) showed that it is possible to enhance the statistic's sensitivity to the thermal state of the underdense gas, stored in the high flux levels of the transmission spectrum. We carried out a similar analysis to constrain the presence of A' DM heating in the IGM, comparing the synthetic skewers extracted from the simulation outputs to the tPDF of HE0940-1050 (Rorai et al., 2016).

Here we summarise the main steps of the procedure. First of all, we created a mock dataset of Lyman-α sightlines that matches the redshift range of the Lyman-α forest found on the spectrum of HE0940-1050 (about $31\,000\,\mathrm{km\,s^{-1}}$ long, equivalent to $\sim 290\,\mathrm{Mpc}\,h^{-1}$, centred at $z \sim 2.75$). This long spectrum was composed by stitching multiple $40\,\mathrm{Mpc}\,h^{-1}$ long sightlines extracted from simulation outputs at different redshifts, spanning the redshift range $z \sim 2.5 - 3$, to appropriately represent the redshift evolution of A' DM heating in the forest. We reproduced the instrumental effects on the synthetic sightlines by convolving the spectra with a Gaussian kernel with a full-width at half-maximum of $7.2\,\mathrm{km\,s^{-1}}$ to match the UVES resolving power for a 1 arcsec slit aperture and then rebinned the data onto a pixel grid with constant spacing

$\Delta v = 2.5\,\mathrm{km\,s^{-1}}$. The assumed spectral resolution only holds if the seeing at the moment of observation is smaller than the 1″ slit, and could marginally change across the multiple exposures of the spectrum. In Sect. 4.3.1 we investigate whether assuming a different spectral resolution for the mock spectra induces differences in the analysis results.

The noise distribution of the original spectrum was modelled following the recipe of Rorai et al. (2016). Summarising, we

1. Divided the pixels of the original spectrum into 50 bins according to their flux values. The errors on each pixel produced 50 noise distributions for each flux value.
2. To each simulated pixel, we assigned a noise value randomly picked from the distribution associated with its flux level.

Lastly, since the mean transmitted flux value $\bar{F}$ of a Lyman-α forest spectrum can exhibit large fluctuations among different sightlines, and can be rather sensitive to the number of strong absorption systems pierced by the sightline, it is reasonable to consider it as an independent free parameter within the measurement errors. This is particularly important when analysing a single sightline, as in our case, and not a large ensemble of spectra. In this way, we could put a constraint directly on the presence of A' DM , marginalising over cosmic variance effects. Therefore, we scaled the mock flux skewers to a set of different mean flux values spanning the range $\bar{F} = 0.71 - 0.82$, around the observed mean transmitted flux at $z \sim 2.7$ (see Sect. 2.3.1 and Becker et al., 2013). We produced $N_m = 10,000$ mock spectra for each $\{m_{A'}, \epsilon, \bar{F}\}$ parameter point. The same procedure was carried out on the fiducial ΛCDM boxes without A' heating.

Once we had a set of realistic data that matched the properties of the observed spectrum, we moved to the comparison between the two. We first noted that the transmitted flux PDF is particularly sensitive to continuum placement uncertainty, whose erroneous estimation can lead to a shift of the flux levels and therefore to a stretch/-compression of the probability distribution itself.

We reduced the uncertainty due to continuum placement by adopting a standardised flux renormalisation. Taking the original fitted continuum as a starting point, the spectra were divided into regions of $10\,\mathrm{Mpc}\ h^{-1}$. For each region, we found F_{95}, corresponding to the 95th percentile of the flux distribution within that region, and normalised the flux as $F_r = F/F_{95}$. This regulation aligned the spectra from simulations with different A' DM parameters at high fluxes, effectively removing any continuum bias placement in the PDF evaluation.

Then, to focus on fluctuations near the continuum, we enhanced the optical depth $\tau_{\mathrm{Ly}\alpha}$ of the Lyman-α lines by an arbitrary constant factor A as

$$\tau_{\mathrm{A}} = A\,\tau_{\mathrm{Ly}\alpha}, \tag{4.3}$$

thus scaling the normalised (and regulated) flux as

$$F_{\rm A} = \exp(A \ln |F_r|) = |F_r|^A \,. \tag{4.4}$$

We adopted a constant $A = 10$ to compare our results to Rorai et al. (2016). This transformation is equivalent to changing the binning of the standard PDF in a flux-dependent way, such that the sensitivity at high fluxes is enhanced. The optical depth transformation was applied to the continuum-regulated flux F_r described above, on both simulated and real data, before comparison.

Fig. 4.4 shows the regulated and rescaled flux, F_{10}, of the same noiseless sightline reported in Fig. 4.3. The procedure amplifies the differences among models and shows a visible distinction between the no-A' and the light A' DM model ($\log(m_{A'}/\,\mathrm{eV}\,c^{-2}) = -13.15$), which is negligible in the standard flux of Fig. 4.3.

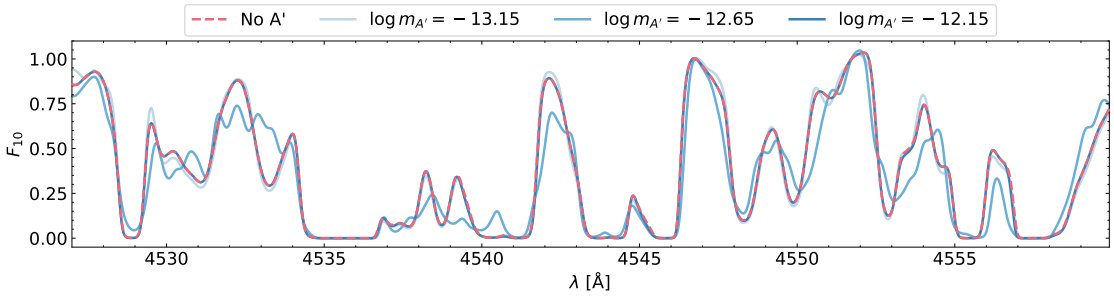

Figure 4.4: *Same section of the synthetic sightlines shown in Fig. 4.3, after continuum regulation and optical depth rescaling have been applied. The dashed pink line is the baseline ΛCDM case without A′. Blue shaded lines show the transmitted fluxes with A′ DM heating, where darker tones represent more massive models. The three A′ DM models reported are the same as in Figs. 4.2 and 4.3, and all assume strong kinetic mixing* ($\log\epsilon = -14.0$). *After the continuum regulation and optical depth transformation, the light A′ DM model* ($\log(m_{A'}/\,\mathrm{eV}\,c^{-2}) = -13.15$) *is visibly distinguishable from the no-A′ case.*

Finally, we computed the probability distribution of F_{10}, called the regulated and transformed PDF, or tPDF for short, throughout the mock datasets.

Fig. 4.5 shows the tPDF of the same A' DM models reported in Figs. 4.2 - 4.4, compared to the tPDF of HE0940-1050 (Rorai et al., 2016) shown as orange diamonds with error bars. The dashed pink line shows the tPDF of the baseline no-A' model, whereas the A' DM models are shown in blue shaded solid lines, with darker tones denoting heavier A' DM masses. Most importantly, we see that the tPDF is able to distinguish between different A' DM prescriptions up to several σ, where a change in $m_{A'}$ modifies the shape of the tPDF. Changes in ϵ and $\bar{F}$ modify the tPDF shape in a non-degenerate way (not shown for conciseness). For low values of ϵ, the energy injection due to A' DM oscillation is effectively negligible and the tPDF is indistinguishable from the no-A' case, regardless of mass. As expected, the massive case ($\log(m_{A'}/\,\mathrm{eV}\,c^{-2}) = -12.15$) shows the same tPDF as the baseline model, since it induces heating only in very dense and rare environments, affecting a relatively small number of lines and thus having small statistical impact on the full spectrum flux distribution.

Notably, the IGM at $z \sim 2.7$ shows good agreement with our no-A' simulation results.

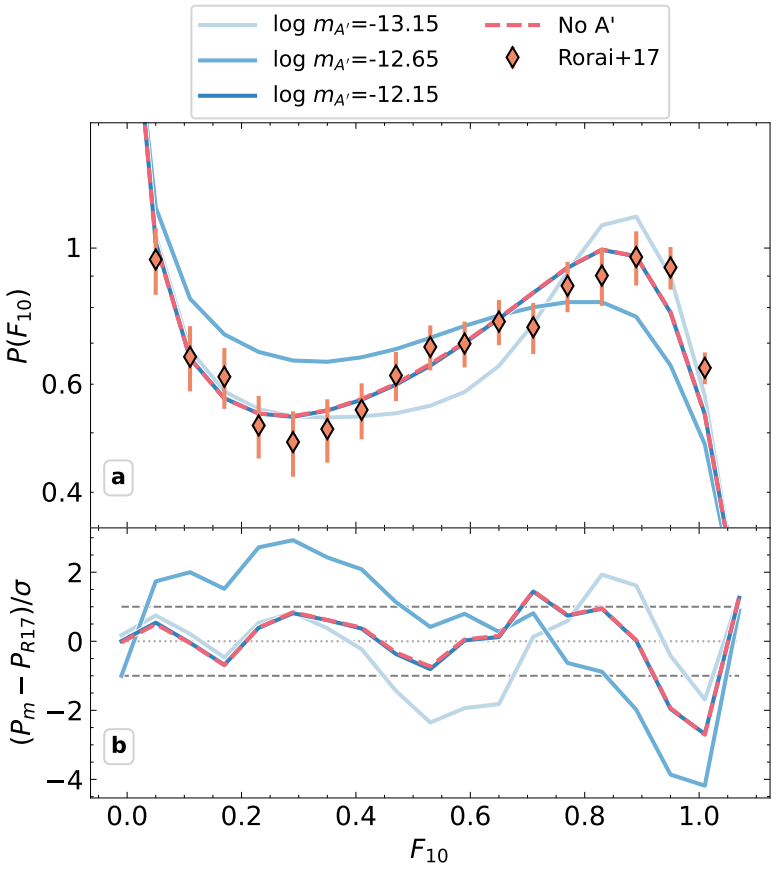

Figure 4.5: ***a***) *Regulated and transformed PDF (or tPDF) of mock spectra emulating the resolution, pixel size and noise distribution of the VLT/UVES spectrum of HE0940-1050, whose tPDF is reported in orange diamonds with error bars, as computed by Rorai et al. (2016). The tPDF of the same models reported in Figs. 4.2 - 4.4 are shown. All models reported have* $\log \epsilon = -14.0$ *and* $\bar{F} = 0.76$. ***b***) *The fractional difference between each model's tPDF and the observed data, in terms of the data error estimate. The horizontal grey dashed lines define the* 1σ *region around the data points reported in Rorai et al. (2016).*

4.2.3 Likelihood

Following the standard assumption that the 19 flux bins in the tPDF are distributed as a multivariate Gaussian, we computed the likelihood between the model and data tPDF for each set of mock sightlines through the function

$$\mathcal{L}(d|m_{A'}, \epsilon; \bar{F}) = \frac{1}{\sqrt{(2\pi)^k\,|\Sigma|}} \exp\left[-\frac{1}{2}\Delta\mathbf{P}^T\Sigma^{-1}\Delta\mathbf{P}\right], \tag{4.5}$$

where $\Delta\mathbf{P} = \mathbf{P}(m'_A, \epsilon; \bar{F}) - \mathbf{P}_d$ is the difference between the average tPDF of the mock spectra, extracted from the A' DM model with mass $m_{A'}$, kinetic mixing ϵ and scaled to mean flux $\bar{F}$, and the measured tPDF $\mathbf{P}_d$ of Rorai et al. (2016). k is the number of tPDF bins considered, $|\Sigma|$ and Σ^{-1} denote the determinant and the inverse of the covariance matrix, respectively, whose ij-th element is computed as

$$\Sigma_{i,j}(m_{A'}, \epsilon; \bar{F}) = \frac{1}{N_m}\sum_{n}^{N_m}\left(P_{i,n} - \bar{P}_i\right)\left(P_{j,n} - \bar{P}_j\right), \tag{4.6}$$

where $P_{i,n}$ and $\bar{P}_i$ are the tPDF values of the ith flux bin computed on the nth sightline and averaged on all mocks, respectively. The sum is performed over all N_m mock spec-

tra of the same model parameter combination. We stress that the covariance matrix is model dependent, especially in the off-diagonal terms, and was recomputed on each $\{m_{A'}, \epsilon, \bar{F}\}$ configuration. Two degrees of freedom were removed by the tPDF normalisation condition ($\sum P(F_i) = 1$) and by the percentile continuum regulation performed on the flux skewers, hence we removed from our analysis the highest and lowest flux bins, following Rorai et al. (2016).

4.3 Convergence of the transformed PDF

4.3.1 Instrumental resolution

The precise estimation of the instrumental resolution of the HE0940-1050 spectrum, especially in the Lyman-α forest, can have a certain degree of uncertainty, as it is defined by the slit width and the sky seeing conditions at the moment of observations. As the spectrum used in our analysis is a composite of several archival observations, the final resolution is not straightforward to evaluate, and misestimation could impact the analysis outcomes.

Specifically, incorrectly assuming the width of the instrumental LSF alters the smoothing scale applied to the mock spectra, which might not reflect the real smoothing occurring in the observational data. This mismatch in smoothing and subsequent line broadening may be mistakenly attributed to a thermal effect from A' DM heating, skewing the analysis results. To investigate this effect, we fitted Gaussian profiles to the telluric lines in the spectrum of HE0940-1050 at ~ 690 nm. Since these lines are unresolved by the instrument, they serve to estimate the width of the spectrograph's line spread function and, by extension, its resolution. On average, the considered lines exhibit a FWHM of $6.6\,\mathrm{km\,s^{-1}}$, which aligns closely with the fiducial FWHM of $7.2\,\mathrm{km\,s^{-1}}$ used in our analysis during the mock creation process, accounting for the instrument's resolution degradation at the blue end of the spectrum.

Additionally, to effectively verify the absence of systematic errors in the A' DM limit's estimation, we created a set of mock spectra derived from the simulations without exotic heating used in the main analysis, while assuming a FWHM of $5\,\mathrm{km\,s^{-1}}$ and $10\,\mathrm{km\,s^{-1}}$, and compute their tPDF. Fig. 4.6 displays the tPDF of the fiducial no-A' mock dataset featuring a LSF with a FWHM of $7.2\,\mathrm{km\,s^{-1}}$ (black), along with the two test sets with FWHM of $5\,\mathrm{km\,s^{-1}}$ (red) and $10\,\mathrm{km\,s^{-1}}$ (blue). The second panel illustrates the differences between the test mocks and the fiducial dataset, showing that the tPDF remains well-converged within uncertainties and is not impacted by the spectral resolution choice. Likewise, the tPDF uncertainty, depicted in the bottom panel as the square root of the diagonal terms of the covariance matrix, converges within approximately 3% when varying the LSF width. Therefore, we expect that the results of the analysis do not suffer from systematics under the choice of the fiducial spectral resolution assumed in the mock spectra, within realistic values of the FWHM, as estimated from the unresolved telluric lines.

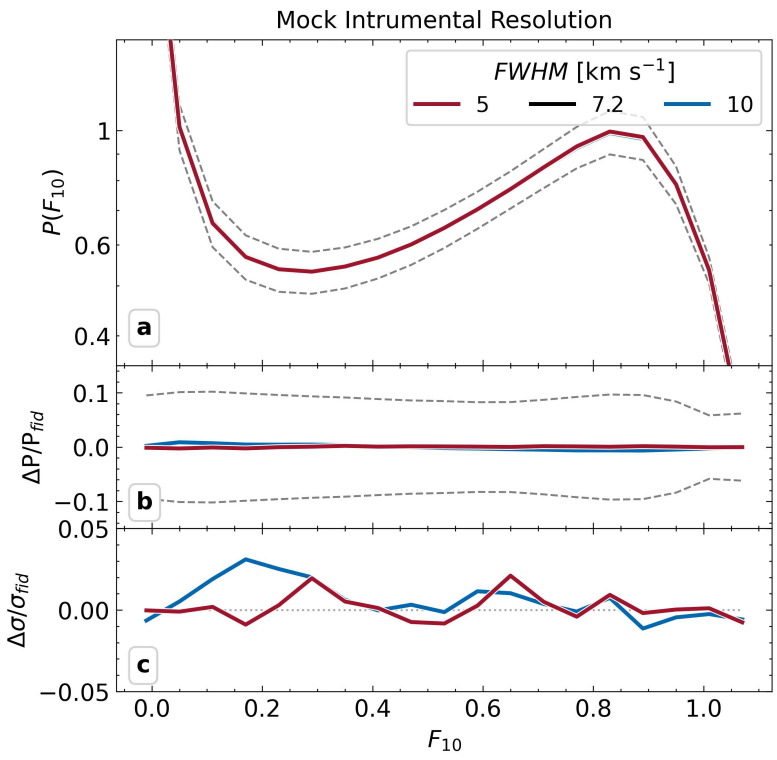

Figure 4.6: *Convergence test of the tPDF with respect to the spectral resolution FWHM assumed in the mock spectra. Top panel: The tPDF of mock spectra from the baseline no-A′ simulations are shown, assuming a spectral resolution FHWM of* $5\,\mathrm{km\,s^{-1}}$ *(red),* $7.2\,\mathrm{km\,s^{-1}}$ *(black, fiducial) and* $10\,\mathrm{km\,s^{-1}}$ *(blue). Dashed grey lines denote the* 1σ *uncertainty on the fiducial tPDF. Middle panel: relative variation of the tPDF under different assumptions of spectral resolution, with respect to the fiducial case. Bottom panel: relative variation of the tPDF error estimates, with respect to the fiducial case.*

4.3.2 Simulation mass resolution and box size convergence

We assessed the numerical convergence of the tPDF by varying the mass resolution and box size of the simulations. The fiducial no A' model used throughout the work is kept as our reference point, with a box size $40\ h^{-1}\mathrm{Mpc}$ and 2×1024^3 gas and dark matter particles. We contrasted the tPDF predicted by this baseline run with those derived from two simulations featuring the same box size but with 2×768^3 and 2×512^3 particles, respectively. This allowed a check of convergence with simulation mass resolution. Similarly, we checked for convergence with simulation volume by comparing our fiducial run to a simulation with box size $20\ h^{-1}\mathrm{Mpc}$ and 2×512^3 particles, thus retaining the same mass resolution between the two runs.

Fig. 4.7 displays the tPDF of the simulations with different mass resolution (left panels) and box size (right panels). The grey dashed curves are the uncertainties obtained from the square root of the diagonal elements of the simulated covariance matrix. In all cases, the tPDF is converged within the expected uncertainties. The diagonal elements of the covariance matrix are also well converged, and are at most 6% larger in the low-resolution runs. The covariance of the tPDF is only slightly influenced by box size, indicating that the baseline simulation adequately captures the cosmic web and does not suffer from large cosmic variance effects. In conclusion, we found that our simulations are well converged. We therefore did not consider any correction to the tPDF or its covariance to address poor convergence with either mass resolution or simulation volume.

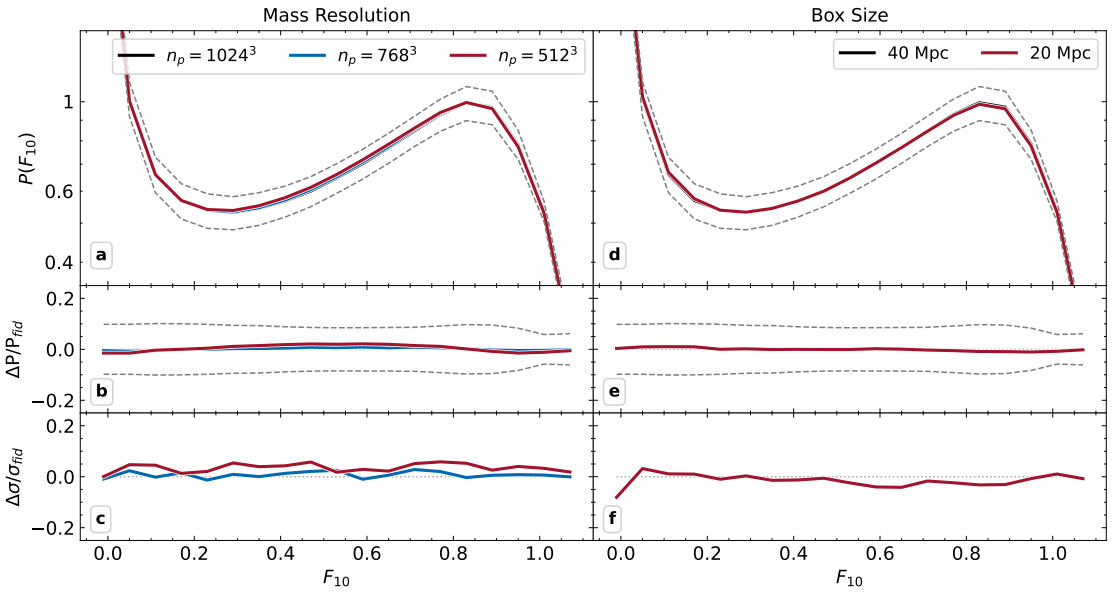

Figure 4.7: *Convergence test of the tPDF with respect to mass resolution (left panels, a-c) and box size (right panels, d-f). The mass resolution test considers three simulations with the same box size (and initial conditions) and different numbers of particles:* 2×1024^3 *(our fiducial run),* 2×768^3 *and* 2×512^3*. The box size test considers two simulations with the same mass resolution but different box sizes:* $40\ h^{-1}$Mpc *(fiducial run) and* $20\ h^{-1}$Mpc*. Top panels (a, d): The tPDF of the Lyman-α forest in the no A′ scenario. The dashed grey line represents the error estimates on the fiducial simulation. Middle panels (b, e): relative variation of the tPDF with respect to the fiducial run. Bottom panels (d, f): relative variation in the tPDF error estimates, which are obtained from the square root of the covariance matrix diagonal terms, with respect to the fiducial run.*

4.4 Dark photon dark matter limits

4.4.1 Posterior Analysis

We explored the posterior probability distribution of the A' DM parameters using a Monte Carlo Markov Chain (MCMC) approach. We linearly interpolated the likelihood evaluated on the simulation grid points and obtained MCMC chains using the `emcee` python package (Foreman-Mackey et al., 2013). We considered two instances, assuming flat priors on $\log m_{A'}$, $\log \epsilon$ and $\bar{F}$ within the boundaries of the simulation grid, as described in Table 4.1 and $0.71 < \bar{F} < 0.82$; or flat priors on $\log m_{A'}$ and $\log \epsilon$, with a Gaussian prior on mean transmitted flux $\mathcal{N}_{\bar{F}}(0.7371, 0.01)$, based on the measurement of Becker et al. (2013). We checked for chain convergence using the Gelman-Rubin test (Gelman et al., 1992).

Fig. 4.8 shows the 1σ and 2σ levels of the posterior distribution in the mass-kinetic mixing plane, marginalised over transmitted mean flux, for the two choices of priors: all flat (dot-dashed green) and Gaussian over mean flux (solid blue). The contours in the two cases are similar, highlighting how the result is not prior dominated. The dashed red line reports the bounds on A' DM mass and mixing computed by Caputo et al. (2020b) based on simple analytical energy distribution assumptions. The best fit model proposed by Bolton et al. (2022a) is shown in orange and lies within the 2σ

region of both posteriors.

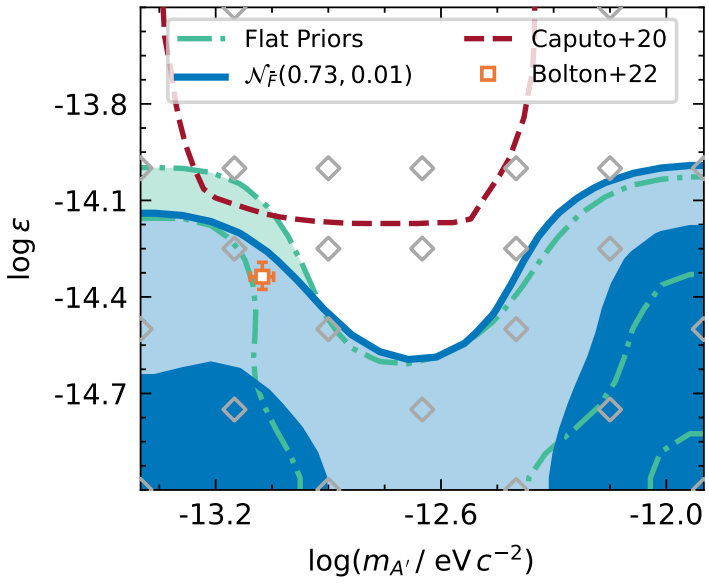

Figure 4.8: *Posterior distribution over the A' DM mass - kinetic mixing plane, marginalized over the transmitted mean flux, assuming flat priors for each parameter (green dot-dashed line), and assuming flat priors on mass and mixing with a Gaussian prior on mean flux $\mathcal{N}_{\bar{F}}(0.7371, 0.01)$ Becker et al. (2013) (blue area). The contours represent 1σ and 2σ levels. Approximate constraints from Caputo et al. (2020b) are also shown (dashed red). The best fit A' DM model of Bolton et al. (2022a) is shown in orange, which is within 2σ of both posterior distributions. Grey squares denote the grid of simulated models used in this work.*

The posterior distribution does not show a preference towards a particular A' DM model. Instead, the most likely models have either low kinetic mixing or very high mass. Both instances would produce a model that does not provide any non-negligible additional heating to the gas in the density range probed by the Lyman-α forest, yielding a tPDF that is indistinguishable from the baseline no-A' case within the error bars. Note moreover that the posterior distribution is limited, for high and low masses, at $\log \epsilon \sim -14.0$ due to the choice of parameter sampling we assumed and, requiring more simulations, would likely extend to higher mixings. Overall, the analysed data did not demonstrate any substantial preference for A' DM heating.

4.4.2 Upper limits

We next turn our attention to setting limits on ϵ as a function of $m_{A'}$. We considered the Lyman-α forest mean flux as a nuisance parameter. We linearly interpolated the likelihood values computed in Eq. 4.5 on a fine grid in the $\epsilon - \bar{F}$ plane, for a given fixed $m_{A'}$, and profiled over the mean flux to obtain $\mathcal{L}(d|m_{A'}, \epsilon)$, the maximum value of $\mathcal{L}(d|m_{A'}, \epsilon; \bar{F})$ over all values of $\bar{F}$ for each pair of $(m_{A'}, \epsilon)$. We then defined the test statistic

$$\lambda(m_{A'}, \epsilon) = 2\left[\ln \mathcal{L}_{\max}(d|m_{A'}) - \ln \mathcal{L}(d|m_{A'}, \epsilon)\right], \tag{4.7}$$

where $\mathcal{L}_{\max}(d|m_{A'})$ is the maximum likelihood over all interpolated $\{\epsilon, \bar{F}\}$ points at a given mass $m_{A'}$. Wilks' theorem (Wilks, 1938) then ensures that λ follows a half-chi-square distribution (Cowan et al., 2011). We therefore obtained the A' DM kinetic mixing 95% confidence limits by finding the value of ϵ at which $\lambda = 2.71$. These lim-

its are reported in Fig. 4.9 for all mass values in our simulation ensemble, compared to the bounds obtained analytically by Caputo et al. (2020b). Note that for low mass ($\log(m_{A'}/\,\mathrm{eV}\,c^{-2}) = -13.4$) and high mass models ($\log(m_{A'}/\,\mathrm{eV}\,c^{-2}) \geq -12.15$) we do not reach the 95% confidence limit within the parameter space explored by our simulations.

We also show the best-fit A' DM model of Bolton et al. (2022a), proposed to reconcile the discrepancy between the IGM temperature $z \sim 0$ found in simulations and the one inferred by spectroscopic observations carried out with the Cosmic Origins Spectrograph (COS, Green et al., 2003; Green et al., 2012) on board of the Hubble Space Telescope (HST). This model is allowed at 95% C.L. by our bounds and therefore still a viable solution to the low-z IGM temperature tension between observations and simulation.

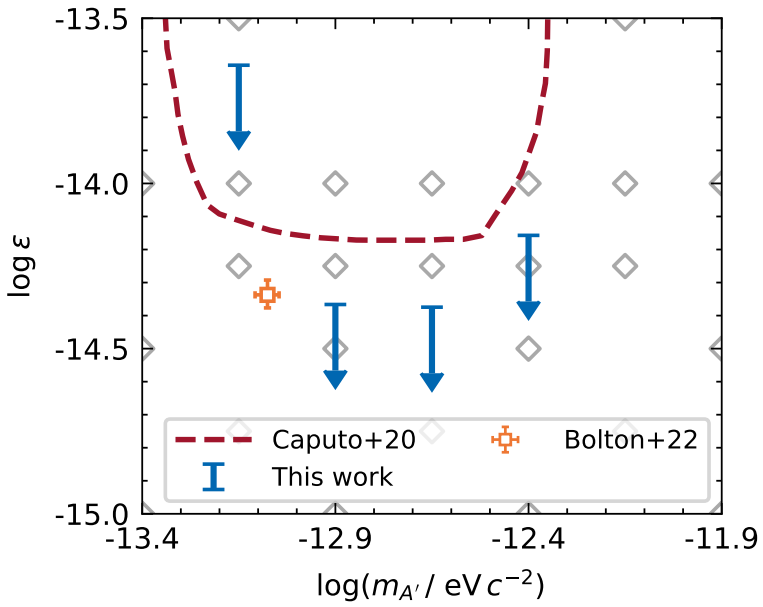

Figure 4.9: *The 95% confidence limits on the A' DM kinetic mixing parameter ϵ at fixed dark photon mass $m_{A'}$ arising from the analysis of the transformed flux tPDF at $z \sim 2.7$. The kinetic mixing at the high and low masses is not constrained by our data at $z \sim 2.75$ within our parameter grid. Approximate constraints from Caputo et al. (2020b) based on requiring that the total heat injected from A' DM not exceed $1\,\mathrm{eV}$ are also shown (dashed red). The best fit A' DM model of Bolton et al. (2022a) is shown in orange. Grey squares denote the grid of simulated models used in this work.*

Part III

The Redshift Drift Experiment

5

The Sandage test

Over the past century, cosmology has undergone a remarkable transformation from a speculative field to a precision science. This progress has been driven by a series of increasingly sophisticated observational probes. The discovery of the Hubble expansion from galaxy redshifts (Hubble, 1926) established the dynamic nature of the Universe, while the detection of the CMB provided a snapshot of the early Universe at $z \sim 1100$, encoding the initial conditions for structure formation. Later, the study of large-scale structure and BAO offered a statistical standard ruler for the expansion history, while Type Ia supernovae, through their role as standardizable candles, revealed the unexpected accelerated expansion of the Universe at late times (Riess et al., 1998; Perlmutter et al., 1999). Together, these probes have established the ΛCDM model as the standard paradigm of cosmology.

Despite their success, all of these methods share an important limitation: they are fundamentally *geometrical* probes. They constrain the expansion history indirectly by measuring distances, angles, or integrated properties of the background geometry, mapping out space, its curvature and its evolution. For example, supernovae provide luminosity distances, BAO constrain angular diameter distances, and the CMB encodes the angular scale of acoustic features at recombination. Interpreting these observables requires assuming General Relativity as the underlying theory of gravity and adopting a specific cosmological model. If either of these assumptions fails, the inferred parameters may not represent the true dynamical state of the Universe.

By contrast, a *kinematical* or *dynamical* probe measures the time evolution of the expansion directly, without relying on integrated distances or geometrical reconstructions. The redshift drift experiment, also known as the Sandage test, provides such a direct probe. Proposed in its modern form by Sandage (1962) and later revisited by Loeb (1998), it is based on the idea that the redshift of distant sources should slowly evolve with the observer's cosmic time as the scale factor changes. This effect, though extremely small, allows a direct, model-independent measurement of the expansion rate as a function of time.

5.1 Theoretical Background

The observable redshift of a photon emitted at comoving distance χ and cosmic time t_e, and received at t_0 in a FRW Universe, i.e. under the only assumptions of isotropy and homogeneity, is given by

$$1 + z(t_0, t_e) = \frac{a(t_0)}{a(t_e)}\,, \tag{5.1}$$

where $a(t)$ is the cosmic scale factor. Assuming fixed χ, if the observer waits for a small interval Δt_0, the redshift of the same source changes because the scale factor evolves. Expanding to first order

$$dz = \frac{\partial z}{\partial t_0} dt_0 + \frac{\partial z}{\partial t_e} dt_e, \tag{5.2}$$

one obtains the time derivative of the redshift – the *redshift drift*:

$$\dot{z} = \frac{dz}{dt_0} = \frac{\partial z}{\partial t_0} + \frac{\partial z}{\partial t_e}\frac{dt_e}{dt_0} = \frac{\dot{a}(t_0)}{a(t_e)} - \frac{\dot{a}(t_e)}{a(t_e)}\frac{a(t_0)}{a(t_e)}\frac{1}{1+z}, \tag{5.3}$$

where, substituting $H = \dot{a}a^{-1}$, using Eq. 5.1 and replacing the unknown emission time, t_e, with its corresponding redshift, yields

$$\dot{z} \equiv \frac{dz}{dt_0} = H_0\big(1+z\big) - H(z)\,, \tag{5.4}$$

where $H_0 = H(t_0)$ is the Hubble constant.

This expression reveals the distinctive power of the redshift drift: it directly measures the difference between the present expansion rate and that at the source's redshift. Unlike luminosity or angular diameter distances, which involve integrals over $H(z)$, the redshift drift gives a direct local derivative, $\dot{a}(z)$, of the expansion history $a(t)$. In other words, a measurement of $\dot{z}(z)$ provides a purely *dynamical* reconstruction of the expansion history of the Universe, detached from any model of gravity and geometrical assumptions.

The sign of $\dot{z}$ is already physically informative:

- In a decelerating Universe, $H(z) > H_0(1+z)$, so $\dot{z} < 0$, at all redshifts.
- In an accelerating Universe, as expected at low redshift in ΛCDM, $H(z) < H_0(1+z)$, so $\dot{z} > 0$.

Thus, simply detecting the sign of the redshift drift would constitute a direct dynamical confirmation of cosmic acceleration, whose cause and interpretation have been at the centre of the cosmological discussion for the last 30 years.

Within the framework of General Relativity, the acceleration can be explained by introducing a new component to the stress–energy budget of the Universe with strongly negative pressure. In its simplest and most successful form, this "dark energy" corresponds to a cosmological constant Λ (Carroll et al., 1992), a smooth and time-independent

contribution with an equation of state $w = -1$. However, a wide range of alternative scenarios remains viable: dynamical scalar fields evolving in a potential (quintessence, Ratra et al., 1988; Caldwell et al., 1998), more exotic fluids such as phantom energy (Caldwell et al., 2003) or Chaplygin gas (Kamenshchik et al., 2001), or phenomenological parametrisations in which w deviates from -1 and evolves with time (DESI Collaboration et al., 2025b).

A conceptually different approach is to abandon the assumption that General Relativity is the correct description of gravity on cosmic scales. In such modified gravity theories, the observed acceleration does not require a dark energy component but arises instead from changes in the law of gravity itself (e.g. braneworld models by Deffayet et al. (2002) or $f(R)$ theories, see Carroll et al. (2005)). By construction, many of these models reproduce the same background expansion history as ΛCDM, but often predict distinct signatures in the growth of cosmic structures or gravitational lensing.

Distinguishing among these possibilities requires probes that are both independent of geometric assumptions and sensitive to the dynamical state of the expansion. The redshift drift is uniquely suited for this role. By directly measuring $\dot{z}(z)$, it provides a kinematical handle on $H(z)$, capable of breaking degeneracies inherent in distance-based observables. When combined with probes of structure growth, it can be used to test the internal consistency of General Relativity: any mismatch between the drift-inferred expansion history and observed clustering would be strong evidence for modified gravity rather than a new energy component. In this sense, while extraordinarily challenging to measure, the redshift drift represents one of the most promising tools for unravelling the true origin of cosmic acceleration.

If such an effect has the potential to revolutionise modern observational cosmology, how come no one has ever achieved a detection? We can predict the magnitude of $\dot{z}$ by adopting a theory of gravity and specifying $H(z)$ through the Friedmann equation. In the case of general relativity, by invoking Eq. 1.5 and plugging it into Eq. 5.4 for different cosmological models, one can compute the expected magnitude of the effect, i.e. how much the redshift of an object at fixed comoving distance is expected to change per year due to the cosmic expansion. In practice, it is usually more common to refer to the effect in terms of radial accelerations, $\dot{v}$, as this is more directly related to spectroscopic observations. In first approximation, $\dot{z}$ translates into $\dot{v}$ as

$$\dot{v} \simeq \frac{c}{1+z}\,\dot{z}. \tag{5.5}$$

Figure 5.1 shows the magnitude of the effect in three cosmological scenarios, with various Ω_m and Ω_Λ values. All models assume flatness and $H_0 = 70\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$. As noted above, only a scenario with $\Omega_\Lambda \neq 0$ defines the presence of a redshift region where $\dot{z} > 0$.

The most important takeaway of the figure, and the reason why no one yet achieved a significant measurement, is the scale of the effect: at $z = 4$ the redshift drift is of the order $10^{-10}\ \mathrm{yr^{-1}}$ or $0.5\ \mathrm{cm\,s^{-1}\,yr^{-1}}$. For comparison, the typical accuracy in extrasolar

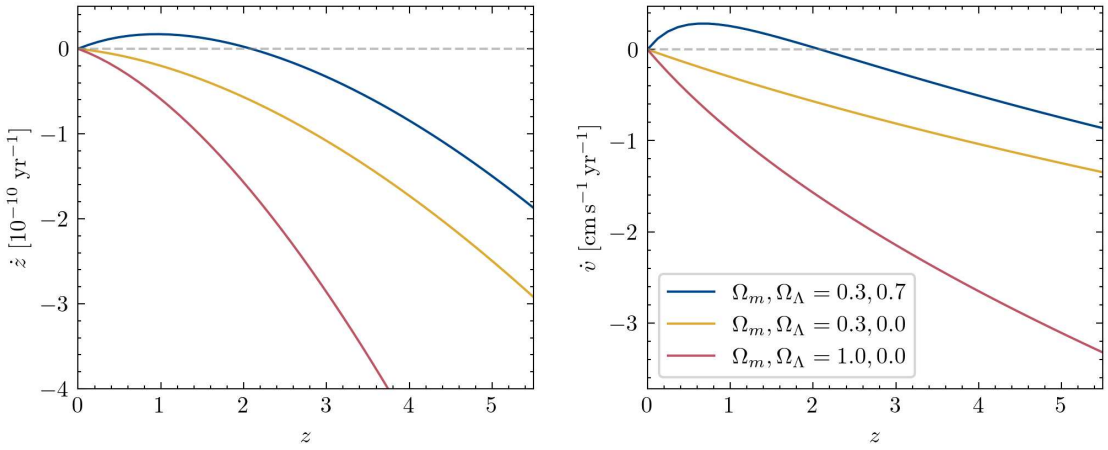

Figure 5.1: *Redshift drift as a function of redshift for various combinations of Ω_m and Ω_Λ, assuming flatness and $H_0 = 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$. Left panel reports the drift in units of redshift ($\dot{z}$), the right panel shows it in units of acceleration ($\dot{v}$).*

planet searches with ESPRESSO at the VLT is of the order of $30\,\mathrm{cm\,s^{-1}}$ (Suárez Mascareño et al., 2020), two orders of magnitude larger than the sought-after cosmological signal.

5.2 The tracers of the drift

The detection of such a tiny effect requires an outstanding measurement precision, while keeping systematic effects and spurious signals to a minimum. This problem can be tackled in two fairly independent ways: by choosing an adequate tracer of the effect, and by observing it with state-of-the-art facilities.

The choice of a tracer is not unique nor straightforward; however, we can define a list of properties that an ideal tracer should have to maximise profitability and detection of the drift: In order to serve as an effective probe of the cosmic redshift drift, a tracer must satisfy a number of physical and observational requirements. These can be summarised as follows:

Minimal peculiar motions: The tracer should follow the Hubble flow as closely as possible, with negligible contributions from local dynamics such as peculiar velocities within galaxies, groups, or clusters. Only in this case can the observed drift be interpreted as a clean measurement of the global expansion.

Well-defined spectral features: The spectral lines associated with the tracer must be sharp, stable, and unambiguously identifiable. Narrow absorption or emission features provide the precision required to detect extremely small shifts in wavelength over decade-long timescales.

Large number of measurable features: Ideally, a tracer should provide many independent spectral lines within a single object. This multiplicity reduces statistical uncertainties and allows robust averaging, maximising the scientific return from limited observational campaigns.

High intrinsic brightness: Since the redshift drift signal is exceedingly faint, tracers must be sufficiently luminous to achieve very high signal-to-noise ratios in a reasonable exposure time, even with next-generation telescopes.

Coverage over a wide (and preferably high) redshift range: The expected signal increases with redshift, and a broad distribution of tracers allows mapping the expansion history across cosmic epochs. Tracers available at $z \gtrsim 3$ are particularly valuable, as the amplitude of the drift is maximised in this regime. On the other hand, low redshift probes ($z \sim 0.5 - 1$), where $\dot{z} > 0$, are relevant for the direct detection of the acceleration.

Note that some of the points above are somewhat in tension with each other; e.g. sharp spectral lines are usually related to cold materials primarily found in dense regions, within strong potential wells and thus heavily affected by local motions. Similarly, high-redshift objects are rarely bright.

Among the possible tracers for this effect, the Lyman-α forest stands out as a particularly promising candidate (Loeb, 1998). The suitability of the Lyman-α forest for redshift drift measurements arises from several key factors:

Low cosmic overdensities: The Lyman-α forest arises from sparse intergalactic neutral hydrogen, which is physically unconnected with the background source, faithfully tracing the Hubble flow, apart from seldom strong absorbers (e.g. DLA) due to dense intervening galactic environments (Rauch, 1998).

Abundance of tracers: The Lyman-α forest provides a dense sampling of absorption features along the line of sight (with number density $dn/dz \approx 100 - 200$ at $z = 2 - 3$, Kim et al. (2002)), enabling the measurement of the redshift drift across numerous independent systems at varying redshifts.

Spectroscopic accessibility: Even though the Lyman-α lines are not particularly narrow, with a typical linewidth of $\sim 30\,\mathrm{km\,s^{-1}}$ (Kim et al., 2002), through the use of high-resolution spectroscopy, their positions can be determined with sufficiently high precision to enable the measurement of the redshift drift. Moreover, the Lyman-α lines at $z > 1.5$ are observable with large ground facilities.

High-redshift sensitivity: the Lyman-α forest can be observed from the ground up to $z \sim 5$. The same tracer and the same observational techniques can be used to cover a large portion of redshift, especially the high-z regime where the (1+z) term of Eq.5.4 amplifies the magnitude of the drift.

Quasar brightness: Quasars are among the most luminous objects at the redshift of interest. High-S/N spectra can be obtained with relatively limited integration time. Moreover, the Lyman-α lines are physically unrelated to the background QSO, and their dynamics are unaffected by the strong energetic phenomena fuelling the quasar emission.

When the Lyman-α forest is used as a tracer, we refer to the experiment as the Sandage-Loeb test (SL).

Another possible probe is the metal absorption lines found in the same QSO spectra observed to measure the Lyman-α forest. These lines are produced by ions heavier than hydrogen and are therefore narrower than Lyman-α lines, thus providing a more precise estimation of their positions. However, the strong metal lines that provide such precise tracers are produced in galactic and circumgalactic environments, where strong gas dynamical processes occur. The stability of such lines over the timescales of the experiment is not guaranteed. The next chapter of this thesis is dedicated to understanding the stability of Lyman-α and metal lines and their use as tracers of cosmological expansion.

5.3 Measurements and forecasts

5.3.1 Earlier attempts

Although visionary, the Sandage test has already been attempted by a few authors adopting different strategies and probes. Here, we summarise two noteworthy publications that put constraints on the drift effect and have complementary value to the analysis we performed.

Darling (2012) analysed multi-epoch H I 21-cm absorption spectra for a sample of 10 absorbers spanning $z = 0.09$–0.69, an important range for constraining dark matter models, with an observational baseline of approximately 13.5 years, relying on Green Bank Telescope (GBT) digital spectra. For the combined ten-object sample, the author reports

$$\dot{z} = (-2.3 \pm 0.8) \times 10^{-8}\,\mathrm{yr}^{-1} \quad (\text{equivalently } \dot{v} = -5.5 \pm 2.2\,\mathrm{m\,s^{-1}\,yr^{-1}}), \tag{5.6}$$

providing the tightest constraint to date. The best single-object limit (3C286; $\langle z \rangle = 0.692$) is

$$\dot{z} = (1.6 \pm 4.7) \times 10^{-8}\,\mathrm{yr}^{-1} \quad (\text{equivalently } \dot{v} = 2.8 \pm 8.4\,\mathrm{m\,s^{-1}\,yr^{-1}}). \tag{5.7}$$

Moreover, the author showed how archival analogue radio data could greatly extend the measurement baseline (up to 30 years); however, the pre-digital data show significant acceleration likely attributable to systematic instrumental and uncorrectable errors and are therefore unsuitable for the drift measurement.

On the other hand, Cooke (2020) presented a novel technique to measure the relative redshift drift occurring between lines at different redshifts that fall close to each other on the quasar spectrum. This technique, called *Ly-α cell*, utilises the centroids of Lyman-α lines as a calibration grid to measure the movement of narrow metal lines. This approach boosts the sought-after signal and minimises wavelength calibration systematics, while also opening up the possibility to exploit old (and so far unreliable) spectra, expanding the measurement temporal baseline (see Martins et al. (2025) for a more detailed study on the potential of the technique). The obvious shortcoming of

this technique is that a Ly-α cell measurement is restricted to a small wavelength range around each cell, reducing the amount of constraining features.

The author applied the technique to four absorption systems along the sightline of QSO HS1700+6416 ($z_{\rm em} = 2.73$), analysing data obtained with the High Resolution Echelle Spectrometer (HIRES) mounted at Keck Observatory (Vogt et al., 1994; O'Meara et al., 2015), taken over a timespan of 10 years. The joint analysis reports a 2σ upper limit on the cosmological acceleration at $z \sim 2$ of $\dot{v} < 65\ \mathrm{m\,s^{-1}\,yr^{-1}}$, posing the best constraint to date based on optical Lyman-α forest data.

5.3.2 Forecasts

The results shown above pose the best limits on the cosmic acceleration to date, all about three orders of magnitude larger than the signal itself, showing how much more work the field requires before a significant detection. In this regard, many other authors have put forward ideas to help speed up the process and properly estimate how much time is still required to conclude the experiment. In the context of the standard Lyman-α forest-based SL test, one work set the stage for the current and next generation of facilities. Liske et al. (2008) has provided the first comprehensive analysis to forecast the measurement precision reached by a large observing programme involving continuous observations of a large sample of quasars.

The authors developed and ran a comprehensive end-to-end simulation of the experiment, based on simulated observational quasar spectra, utilising both Monte Carlo methods to construct synthetic sightlines and line lists extracted from real datasets. A second epoch of each sightline was produced by inducing a small shift in the absorption lines, recreating the effect of the redshift drift in different cosmological scenarios. Finally, each sightline was degraded to realistic noise, pixel size, and resolution, covering a wide grid of parameters. By comparing pairs of mock sightlines, the authors derived a scaling relation for the measurement's precision, based on the concept of total radial velocity information of a spectrum (Bouchy et al., 2001). In the simplest version of the experiment, each target is observed twice, once at the beginning and once at the end of the experiment, so that both epoch spectra have the same S/N. If only the Lyman-α forest is used to measure the drift, the measurement precision in the velocity difference between the spectra scales as

$$\sigma_v = 2 \left(\frac{\mathrm{S/N}}{2370}\right)^{-1} \left(\frac{\mathrm{N_{QSO}}}{30}\right)^{-1/2} \left(\frac{1+\mathrm{z_{QSO}}}{5}\right)^{-\gamma} \mathrm{cm\ s^{-1}}, \tag{5.8}$$

where S/N is the signal-to-noise ratio (per 0.0125 Å pixel) for each epoch spectrum, $N_{\rm QSO}$ is the number of targets considered, and $z_{\rm QSO}$ is their emission redshift. The exponent $\gamma = 1.7$ for $z_{\rm QSO} < 4$ and $\gamma = 0.9$ otherwise. When considering all available absorption lines, including metal lines, over the entire accessible optical wavelength range down to the QSO's Ly-γ emission line, the numerical factor 2 in the equation reduces to 1.35.

In a more realistic context, the observations of a target will not be taken only at the beginning and at the end of the experiment. For a more practical scheduling, it's more realistic to envision a scenario in which multiple epochs are taken during the years, until a detection is carried out. In this context, a new term called *form factor*, g, is added to the equation, where $g = 1$ when only two epochs with equal integration time are considered, and $g = 1.7$ in the case of a uniform distribution of exposure time across a large number of epochs.

Based on these relations, Liske et al. (2008) concluded that a significant detection of the redshift drift would be feasible with the ELT's high-resolution spectrograph, specifically the ANDES instrument. They estimated that a 3σ detection would require approximately 1500 hours of telescope time over a 20-year baseline. This initial forecast, based on a sample of about 30 quasars, established the ELT as the flagship facility for this experiment, transitioning the SL test from a theoretical concept to a viable observational goal.

The successive step in this transition from concept to actual observations lies in the definition of the perfect sample of objects on which the experiment will be carried out. Unfortunately, most wide-area quasar surveys to date, such as SDSS, BOSS, and DESI, have primarily targeted the northern hemisphere, leaving the southern sky comparatively underexplored. This has limited the availability of bright high-redshift quasar catalogues and thus reduced the capabilities of ESO's instruments in carrying out the SL test. The QUasars as BRIght beacons for Cosmology in the Southern hemisphere (QUBRICS) survey (Calderone et al., 2019; Cristiani et al., 2023) was designed to fill this gap and provide a sample of bright high-z quasars complementary to the large northern surveys. By combining multi-wavelength photometric catalogues with machine-learning selection techniques, QUBRICS has assembled a well-defined sample of quasars in the redshift range $z \approx 2 - 5.5$, with emphasis on the most luminous sources accessible to large ESO telescopes. From the large number of newly discovered targets, a sample of 7 bright quasars has been selected as the Golden Sample (GS) to carry out the SL test, as the 7 brightest objects spanning $2.9 \leq z \leq 4.7$. This sample offers a nearly optimal set of targets for ESPRESSO and ANDES monitoring, well distributed across the sky for excellent observational scheduling over a long-term effort, with at least one suitable QSO available for observations from the VLT and ELT sites at any time of the year. Figure 5.2 shows all the QSO known in the literature and those discovered by the QUBRICS survey (with $2.5 \leq z \leq 5$ and $i > 16$), highlighting the 7 bright objects of the GS.

Once the GS is defined, one can easily compute the measurement precision reached in an ANDES-based experiment following Eq. 5.8, where the total S/N per pixel for photon-noise limited observations can be modelled as

$$\mathrm{S/N} = 650 \left[\frac{Z_X}{Z_r} 10^{0.4(16-m_X)} \left(\frac{D}{39\,\mathrm{m}} \right)^2 \frac{t_{\mathrm{exp}}}{10\,\mathrm{h}} \frac{\epsilon}{0.25} \right], \tag{5.9}$$

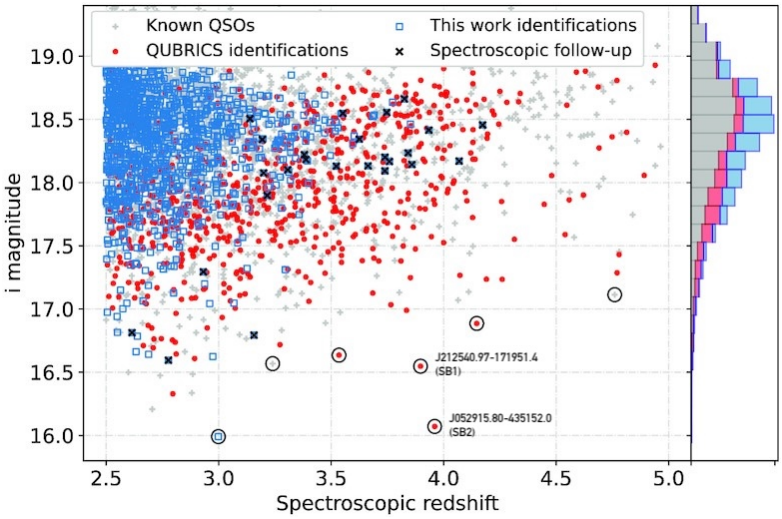

Figure 5.2: *All QSOs in the southern sky known in the literature (grey crosses) and the ones discovered by the QUBRICS survey (red circles, blue squares). The 7 bright quasars of the Golden Sample are highlighted with a black circle mark. Labels define the two objects studied in the next chapters. Adapted from Cristiani et al. (2023).*

where D, $t_{\rm exp}$ and ϵ are the telescope diameter, total integration time per target and total instrumental efficiency. Z_X and m_X are the zeropoint and apparent magnitude of the source in the X-band, respectively, and $Z_r = 8.88 \times 10^{10}\ \mathrm{s^{-1} m^{-2} \mu m^{-1}}$ is the AB zeropoint for an effective wavelength of 6170 Å (corresponding to the SDSS r-band). The normalisation of the above equation assumes a pixel size of 0.0125 Å and a central obscuration of the telescope's primary collecting area of 10 per cent. Assuming a fixed allocation of $1500/7 = 214$ h of observation for each QSO in the sample over 25 years, and assuming to observe all targets twice, at the beginning and at the end of the programme (i.e. $g = 1$ in Eq. 5.8), Cristiani et al. (2023) claim that the redshift drift is expected to be detected with a 99 per cent confidence level, assuming a 25% efficiency of ANDES. Advancements in the instrument's design by the ANDES consortium show that the instrument's efficiency is realistically on the order of $\sim 10\%$ or less, thus lengthening this estimate.

5.4 What's next?

The previous sections highlight how future spectroscopic facilities will bring the SL test to life, successfully reaching significant detection in a reasonable amount of time. However, such forecasts rely on heavy assumptions that will require proper investigation and testing before attempting a conclusive detection.

The most pressing issues that still require a profound comprehension and might have a deep impact on the feasibility and timeline of the SL test are:

The stability of the tracers: The measurement of the cosmic expansion relies on tracers that, ideally, should have fixed comoving coordinates throughout the time of the experiment. As discussed above, IGM absorption lines found in QSO spectra are thought to be sufficiently stable, with limited peculiar velocities, to not induce strong systematic effects in the measurement. However, these absorption lines arise from different cosmic and galactic environments, and a deeper understand-

ing of their dynamics is necessary to trust them as the tracers of our experiment. Specifically, it is of great interest to understand whether metal lines, narrower than Lyman-α lines and thus providing a more precise position measurement, suffer from peculiar velocities of galactic nature. The stability of Lyman-α and metal lines will greatly increase the precision of the measurement, shortening the required temporal baseline, the total integration time and overall schedule of the experiment.

The nature of the quasar sightlines: The presented forecasts of the experiment rely only on the overall properties of the selected sample of quasars, e.g. redshift and magnitude, however, an in-depth high-resolution study of each single sightline is required to understand how much information on the cosmic acceleration is brought by each target. This will depend on the amount of usable tracers, on their density in the spectrum and on the presence of strong interlopers (e.g. low-z strong metal systems, DLAs). Each sightline is unique, and more precise forecasts must be based on thorough studies of the GS.

The measurement procedure: Thus far, the SL test has not been attempted with high-resolution QSO optical spectroscopy, and therefore, the community lacks a consolidated and well-tested procedure to carry out such a measurement. The aforementioned forecasts rely on simple photon-noise assumptions to provide the best possible achieved measurement precision; however, we must develop the tools and the workflow to perform such measurements. The final result will depend on the actual technique adopted and might differ from the theoretical predictions presented so far.

Instrumental systematics: Due to the extremely small nature of the cosmic signal, a profound and precise understanding of the instrumental systematics is required and has not yet been developed. In particular, the most pressing argument is the wavelength calibration precision and stability over the timeline of the experiment, where the LFC is expected to be the only calibration unit able to satisfy the stringent requirements of our experiment. A proper evaluation of the LFC's performances and the estimation of the systematic effect induced by different calibration plans are necessary to properly forecast the detection timeline. Additionally, all ANDES-based forecasts rely on the specifics of a non-finalised design. Any change in the instrument's design and its forecasted performances, especially its total efficiency (ϵ in Eq. 5.9), will have strong effects on the experiment. Effectively, only with the commissioning of ANDES can a complete evaluation of such effects be carried out.

The next chapters of this thesis are devoted to answering these points. In particular, in Chapter 6 we exploit spectroscopic observations of lensed quasars to study the dynamics and dimensions of IGM absorbers, showing how Lyman-α lines are the optimal tracers for the SL test, whereas metal lines are subject to intrinsic peculiar velocities of the same order of magnitude as the cosmological signal itself. Chapter 7 addresses the

other three points above by beginning the measurement of the redshift drift with the two brightest objects in the golden sample with ESPRESSO, to set the first cornerstone of a decade-long effort.

6

Stability of spectral lines as tracers of expansion

The main goal of this chapter is to directly probe the physical sizes of metal and H I absorbers and study their behaviour at sub-kpc scales. With the knowledge of cloud size, we can measure their gas densities from the observed column densities, thereby inferring other physical properties and the dynamical state of the gaseous structure. We are particularly interested in the very small scales, in the kpc to sub-kpc range, where turbulent and complex gas phenomena are expected to arise and, possibly, induce local peculiar velocities and/or modifications to the ionisation states of the absorbing gas. Under such effects, absorption lines caused by these environments might evolve during long temporal baselines, as the one foreseen for the SL test, and thus not be faithful tracers of the expansion. In the following pages, we measured the sizes and dynamics of Lyman-α and metal lines arising in the spectra of a lensed quasar, finding that the former are in fact stable over decades, while the latter arise from clumpy structures bounded by small dark matter haloes, with velocity variations over time at the level of the cosmic signal itself.

6.1 Structure and dynamics at small scales

The crucial problem of the physical scale of the Lyman-α forest clouds has been tackled by a number of studies[1], finding in general large sizes of few hundred kiloparsecs. Observations (e.g. Smette et al., 1992; Smette et al., 1995) and simulations (e.g. Bolton et al., 2017b) have shown that these absorbing structures are part of a Cosmic Web (Bond et al., 1996) and are expected to undergo, on average, the general Hubble expansion (Davé et al., 1999). However, on intermediate scales (of order 100 kpc) the effects of gravitational collapse may become more pronounced, and galactic and subgalactic potential wells may impart kinetic energy to the gas, whereas on the smallest

[1] E.g.. Weymann et al. (1983); Smette et al. (1992); Crotts et al. (1998); D'Odorico et al. (1998); D'Odorico et al. (2002); Rauch et al. (2005).

(subkiloparsec) scales stellar evolution and gasdynamical processes in the interstellar medium (ISM; supernova remnants, winds) must be the dominant sources of kinetic energy and momentum. Such small-scale galaxy-driven effects primarily affect clumpy metal-bearing haloes in the circumgalactic medium and related to DLA systems. Earlier observations of small-scale structure in Lyman-α forest systems (Rauch et al., 2001a; Rauch et al., 2005) have shown that there is also a trend of the motions to increase in strength with increasing density, e.g., the higher density gas appears to be more turbulent than the more typical Lyman-α forest cloud.

The sizes of absorption clouds can be estimated directly by studying the coincidence of lines in the spectra of close quasar pairs or the images of a lensed quasar. In the case of quasar pairs, the separation (in proper units) of the two lines-of-sight at the redshift of absorption z_a is related to the angular separation of the two quasars, ϑ, as (Mo et al., 2010)

$$l_\perp = a(z_a) r(z_a) \vartheta = 145 h^{-1} \mathrm{kpc} \left(\frac{\vartheta}{10''} \right) \frac{H_0 a_0 r(z_a)}{c(1+z_a)}, \tag{6.1}$$

where $r(z_a)$ is the radial coordinate at redshift z_a. For an Einstein-de Sitter Universe, $\vartheta = 10''$ corresponds to a separation $l_\perp \approx 40 h^{-1}$ kpc at $z_a = 2$. If the absorbing cloud at redshift z_a has a typical size $D > l_\perp$, then it can intersect the two lines-of-sight simultaneously, and the spectrum of each quasar will contain an absorption line produced by the same cloud. It is therefore possible to estimate the typical cloud size by analysing the frequency of line coincidences in the spectra of quasar pairs with various angular separations. Applications of this method to real data can be found in Crotts et al. (1998) and D'Odorico et al. (1998). The sizes of Lyman-α forest clouds appear to be large, with typical values of the order of $200 - 500 h^{-1}$ kpc.

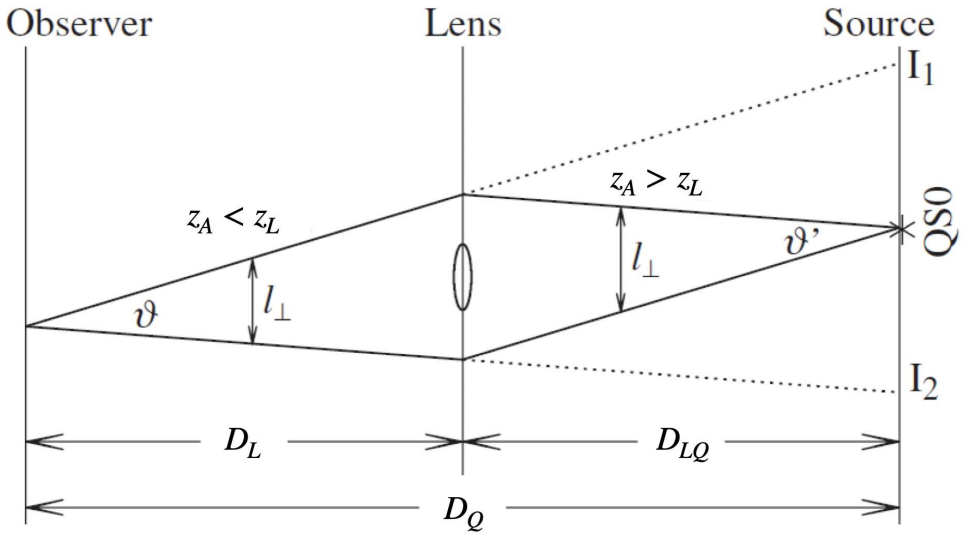

Figure 6.1: *Schematic representation of the geometry of a gravitational lens. Source: Mo et al. (2010).*

Clearly, from Eq. 6.1 one can see that QSO pairs probe sub-kpc scales only at very

low redshift[2], and thus can not be exploited to study the small-scale structure of the IGM, to which we are interested in this work. Luckily, thanks to the geometry of a lensed quasar, as shown in Fig. 6.1, the separation between two sightlines at $z_A > z_L$ (where z_L is the redshift of the lens) is given by

$$l_\perp = d_{A,Q}(z_A)\vartheta', \tag{6.2}$$

where $d_{A,Q}(z_A)$ is the angular-diameter distance from the source to the absorber at z_A, and ϑ' is the angle between the two lines-of-sight at the source position. Since $\vartheta D_L = \vartheta' D_{LQ}$, we have

$$l_\perp = \frac{D_L}{D_{LQ}} d_{A,Q}(z)\vartheta \quad \text{for} \quad z_a > z_L, \tag{6.3}$$

where the sightline separation decreases close to the source, when $d_{A,Q} \to 0$, showing how small scales can be probed at high redshift. This separation model is only an approximation given by geometrical optics, and, in most cases, much more precise models are obtained through the study of the exact mass distribution of the lens and the relativistic modelling of geodesic light propagation.

Lastly, we want to highlight the fact that an absorption feature present in two adjacent sightlines, and unambiguously related to a single absorbing cloud, will show a shift in velocity space proportional to the expansion (or collapse) velocity of the cloud. This phenomenon will be heavily exploited in the following analysis, where the observed shift will be correlated to the separation between the two sightlines to investigate the physical processes that occur on the smallest scales.

6.2 Target Selection

We have selected, as a target to investigate, the small-scale structure of the IGM, the brightest known lensed QSO in the southern sky for which the Lyman-α forest is conveniently observable from the ground ($z > 2.5$) and the separation of at least two images of the lens is above 2″ (in order to avoid flux contamination, see Sect. 6.3). J014516.6-094517 A, B is a lensed QSO with emission redshift $z = 2.719$ and images separated in the sky by 2.24″. The QSO, also known as UM673, was discovered by MacAlpine et al. (1982) and was classified as a double-lensed QSO by Surdej et al. (1987). Surdej et al. (1988) showed the presence of a lensing galaxy at $z = 0.49$, as later confirmed by Eigenbrod et al. (2007). Figure 6.2 shows three images of UM673 in different bands (H, I, V) taken from the CASTLES survey (Muñoz et al., 1998).

The image A has an apparent magnitude around $R \simeq 16.2$, while the image B has $R \simeq 18.3$, with an observed variability of about 0.1 mag (Koptelova et al., 2010). The extinction-corrected flux ratio between components A and B has been estimated to be

[2] In order to probe sub-kpc scales at $z_a \sim 2$ with a QSO pair, the angular separation between the object would need to be small, leading to huge flux contamination between the images.

about 2.14 mag (Koptelova et al., 2014). The Lyman-α forests along the lines of sight towards UM673A,B cover the redshift range from 2.2 to 2.7 and correspond, given the geometry of the lens (see Sect. 6.5), to transverse scales from $\sim 1 h_{69.6}^{-1}$ physical kiloparsecs down to the few hundred pc range.

Previous measurements at intermediate resolution are available for a similar object, RXJ0911.4+0551 (Rauch et al., 2005), and show a remarkable similarity in the Lyman-α forests at these scales. However, the available data on RXJ0911.4+0551 are limited by the precision and stability of the ESI spectrograph at Keck, which implies that differences in velocity between the pairs of absorption features are fully compatible with the measurement error of about $5\,\mathrm{km\,s^{-1}}$.

6.3 Observations of UM 673 A, B

We have taken advantage of the outstanding stability of the ESPRESSO spectrograph (Pepe et al., 2021b) to significantly improve the precision of the measurements with respect to previous studies of the small-scale structure of the IGM (e.g. Rauch et al., 2001b; Rauch et al., 2005), i.e. to reach a sensitivity of about $0.5\,\mathrm{km\,s^{-1}}$ per spectral feature (assumed to be a Lyman-α line with a Doppler parameter of $\sim 20\,\mathrm{km\,s^{-1}}$).

Observations were carried out on Nov 26th 2019 at ESO Paranal, in 4UT mode, i.e. incoherently combining the light of the 4 unit telescopes of the ESO Very Large Telescope. ESPRESSO is a fibre-fed spectrograph with the fibre size corresponding to 1" on the sky. Three exposures of 3600s were obtained on image B and two of 1200s on image A. The different exposure times were chosen in order to (partially) compensate the (~ 9) flux ratio between the two images, as measured in the g band by the DES DR2 Survey (Abbott et al., 2021, $g_A = 16.8965$ and $g_B = 19.2846$), the rough expectation being of a $\sim \sqrt{2}$ ratio of the final S/N between the image A and B combined spectra (see Sect. 6.4).

The seeing conditions were on average around 0.7" and the expected fibre flux contamination in the observation of the image B from the image A turns out to be negligible (less than 10^{-7}, assuming a 2-D Gaussian for the seeing). The CCD detectors were binned by a factor 8×4 along the X and Y (dispersion) directions, respectively. The resolving power of ESPRESSO in the 4UT mode is $R = 70000$, roughly constant in wavelength, as measured by Pepe et al. (2021b).

Other data of UM 673 A, B are also found in the KODIAQ database (O'Meara et al., 2017), obtained with the HIRES spectrograph at Keck, for a total of 14400 and 28600 s for images A and B, respectively. More details about these 2005 and 2008 observations are given in Cooke et al. (2010). These data were not combined with the ESPRESSO spectra we carried out. Still, they were analysed in parallel to compare the two instruments ' performances and pinpoint possible systematic effects in stability and calibration.

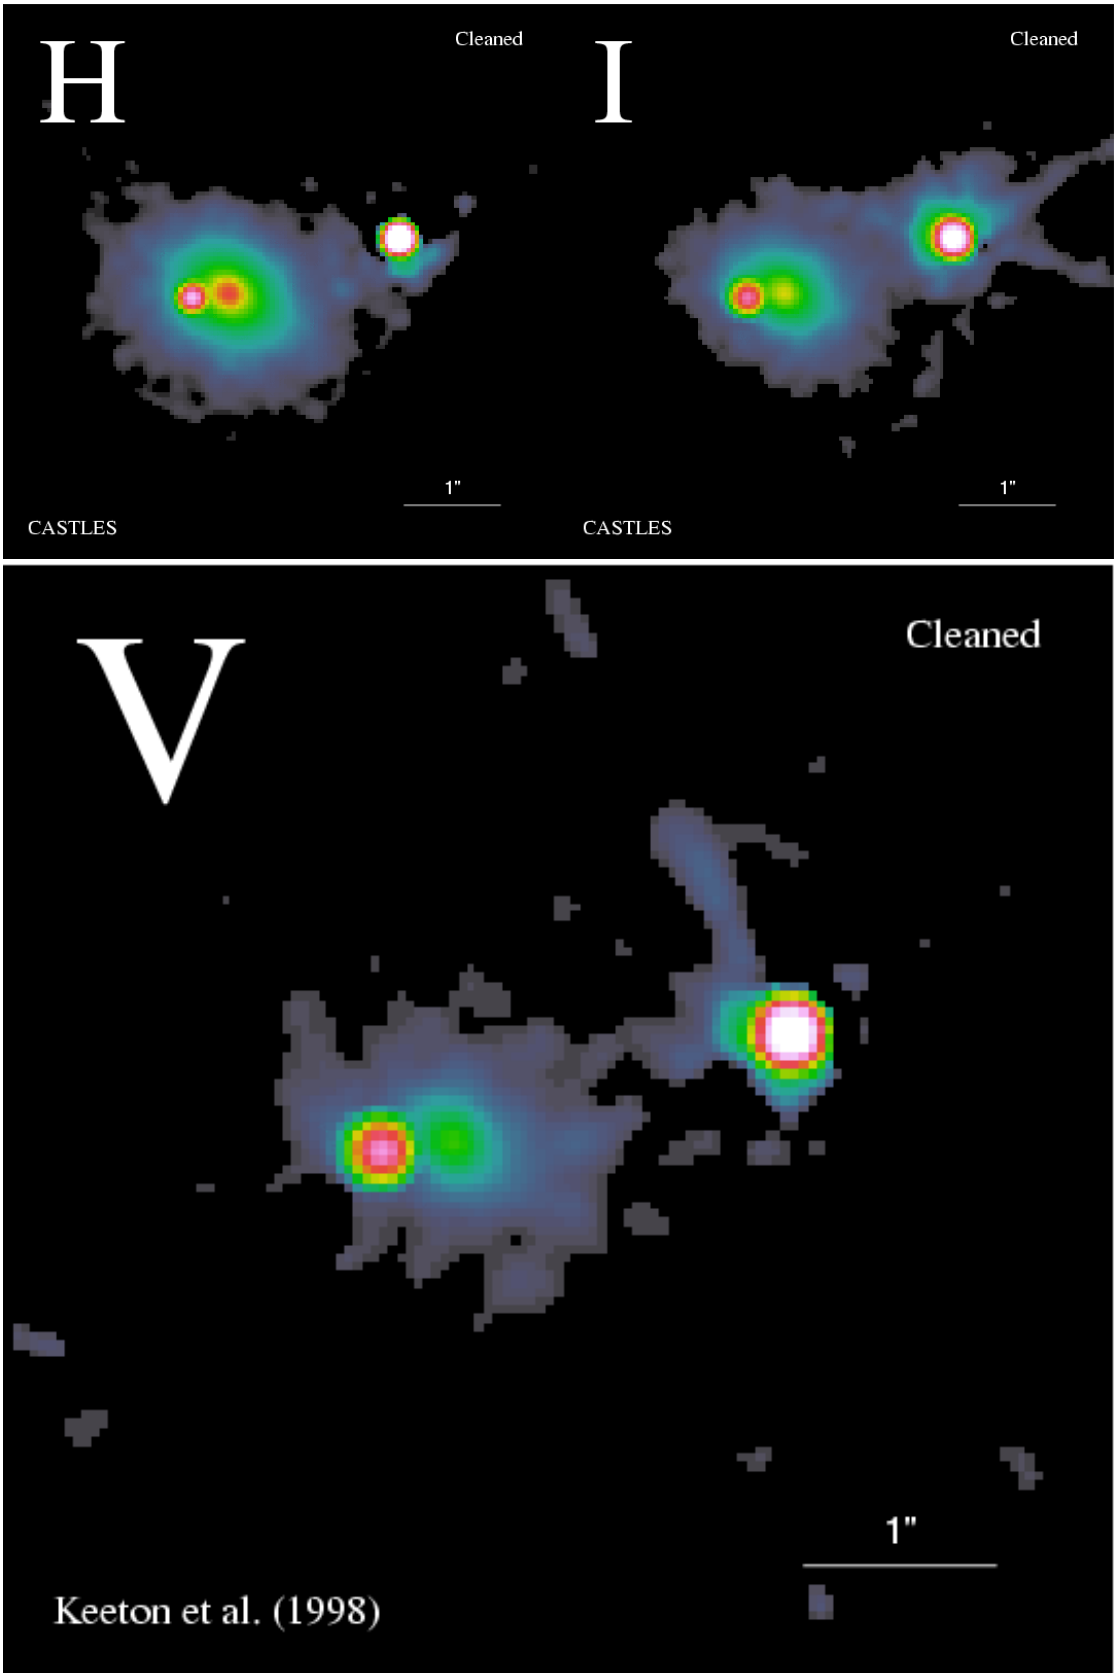

Figure 6.2: *Observations of UM673 in the H band (upper left panel), I band (upper right panel) and V band (lower panel) from the CASTLES survey (Muñoz et al., 1998). In all panels, images A and B are located on the upper right and lower left, respectively. The bright emission between the two QSO images is the lensing galaxy.*

6.4 Data Reduction and Analysis

The observed spectra were reduced with the ESPRESSO Data Reduction Software[3] (DRS, Di Marcantonio et al., 2018). Data reduction included wavelength calibration (with ThAr lamps, converted to the Solar System barycentric frame in vacuum), sky and background subtraction, and optimal extraction of the echelle orders along the cross-dispersion direction (i.e. without any rebinning in the dispersion direction).

The reduced exposures were merged into two spectra, one for image A and one for image B (Fig. 6.3), with the ESPRESSO Data Analysis Software[3] (Cupani et al., 2019) and with the Astrocook package[4] (Cupani et al., 2020). The exposures were first adjusted by the ratio of their median flux densities to account for discrepancies in their relative photon counts; the discrepancy was in all cases within 15% and showed no significant dependence on wavelength. The exposures were then combined by defining a final wavelength grid and computing a weighted average of the contributions to each bin of this grid from the original pixels of the echelle orders. The grid was chosen to be identical for both images, to allow pixel-by-pixel comparison of the two sightlines, with a wavelength range $380\,\mathrm{nm}$ to $780\,\mathrm{nm}$ and a log-wavelength step corresponding to $\delta v = 2.0\,\mathrm{km\,s^{-1}}$ ($\delta\lambda \simeq 3.7\,10^{-3}\mathrm{nm}\times\lambda/550\,\mathrm{nm}$), matching the typical size of the detector pixels for the adopted instrumental setup.

The median S/N at continuum for image A and image B is respectively $\simeq 70$ and $\simeq 50$ per $4.28\,\mathrm{km\,s^{-1}}$resolution element ($\simeq 50$ and $\simeq 35$ per spectral bin, with $\simeq 2$ bins per resolution element), ranging between $\simeq 10$ and $\simeq 80$ per spectral bin in the Lyman-α forest and rapidly decreasing bluewards of $400\,\mathrm{nm}$ in the observed frame. These values are consistent with the expectations, given the observing conditions, as confirmed by the ESPRESSO Exposure Time Calculator[5]. For comparison, the HIRES spectra were binned with a log-wavelength step $\delta v = 2.5\,\mathrm{km\,s^{-1}}$ and achieved a median S/N at continuum of $\simeq 130$ and $\simeq 60$ ($\simeq 75$ and $\simeq 35$ per spectral bin, given $\simeq 3$ bins per resolution element) for image A and image B respectively. A histogram of the S/N per bin in different spectral ranges is given in Fig. 6.5.

The emission continuum in the ESPRESSO spectra was modelled by Astrocook recipes computing a running average of the flux within a $\simeq 400\,\mathrm{km\,s^{-1}}$ window and rejecting outliers at $3\sigma_{\mathrm{clip}}$, with σ_{clip} the formal error on flux. Only outliers below the running average were rejected, and the procedure was iterated until convergence. The continuum model was obtained from the final running average, smoothed by a Gaussian kernel with $\sigma_{\mathrm{smooth}} = 300\,\mathrm{km\,s^{-1}}$. In this way, the observed spectra have been normalised to the respective continua and represent the transmission of the IGM (see Fig. 6.3).

[3] https://www.eso.org/sci/software/pipelines/, version 3.0.0.
[4] https://github.com/DAS-OATs/astrocook
[5] https://www.eso.org/observing/etc/bin/gen/form?INS.NAME=ESPRESSO+INS.MODE=spectro

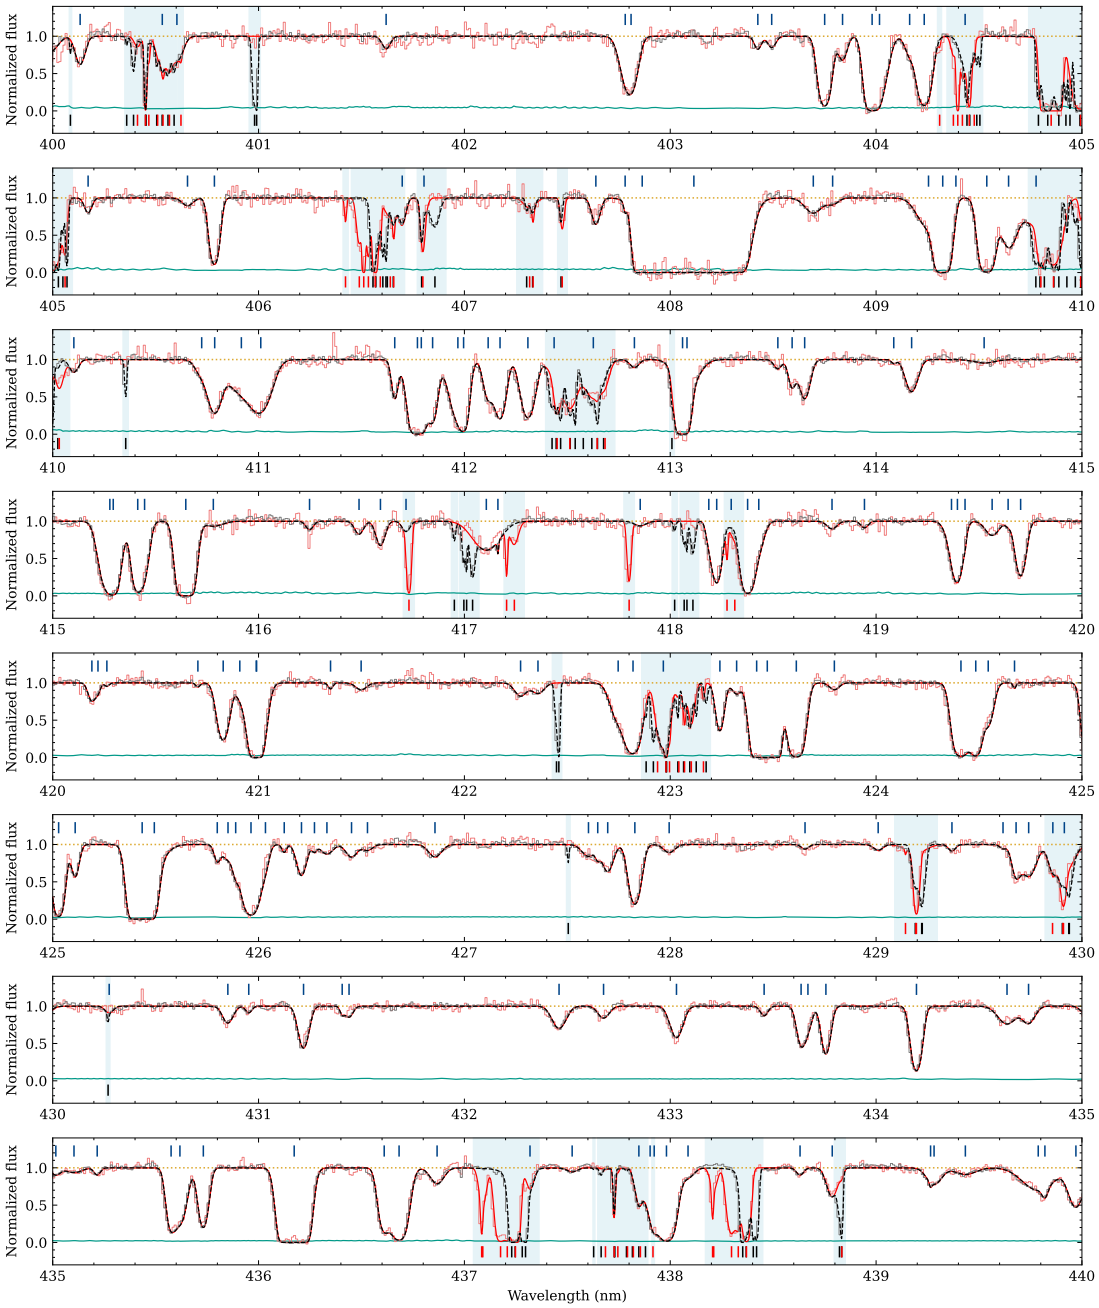

Figure 6.3: *a) A section of the spectrum of UM 673, normalised in transmission. The spectra of the A (grey solid line) and B (light red solid line) images are plotted on top of each other. The absorption models for the two images are also shown: A - black dashed line, B - red solid line. The teal line shows the noise in the spectrum of image B, predominant with respect to component A. The position of the Lyman-α components is shown with blue vertical ticks above the spectrum, while the black and red ticks below it define the positions of the metal components in images A and B, respectively. The vertical blue shaded regions in both segments define areas within the Lyman-α forest where masking has been applied due to metal contamination. The two spectra appear indistinguishable within the noise, except for features corresponding to metal lines.*

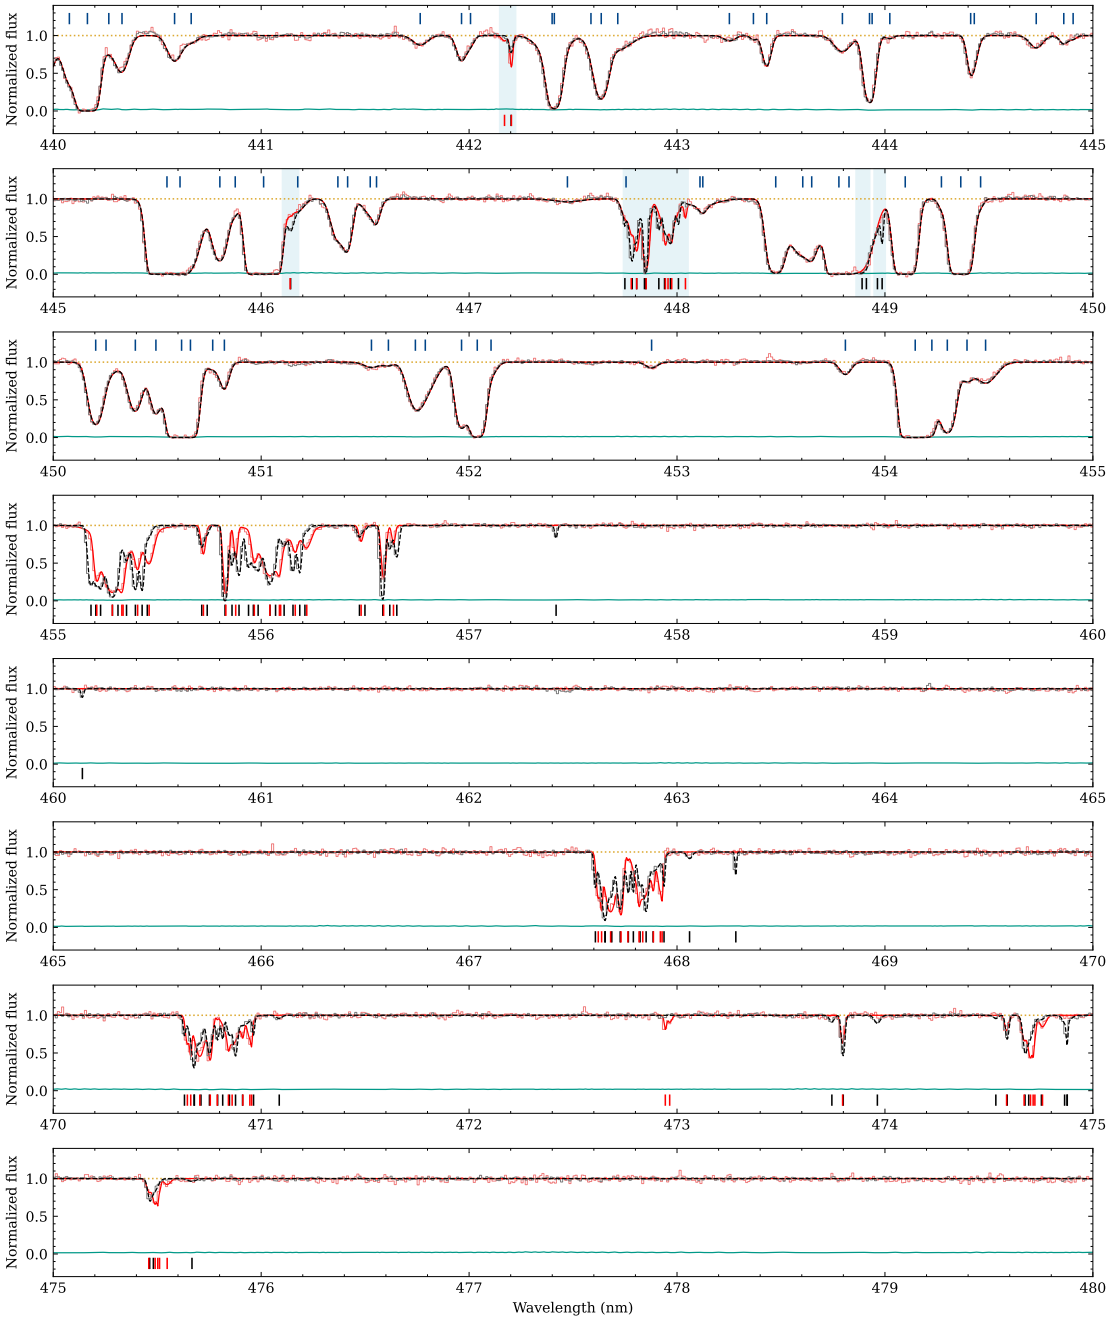

Figure 6.4: *b) Continuation of Fig. 6.3 up to* 480 nm

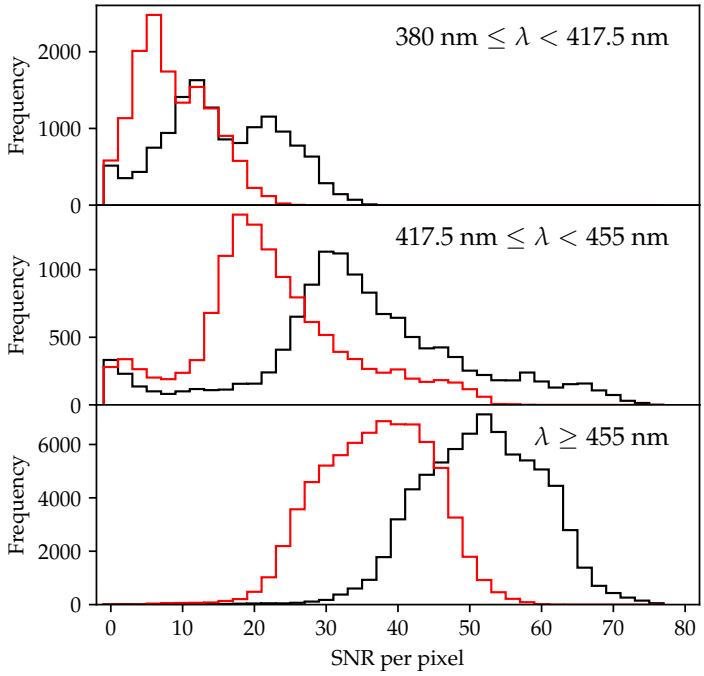

Figure 6.5: *S/N per 2* $\mathrm{km\,s^{-1}}$ *bin in different spectral ranges, for image A and image B (black and red, respectively). The top and centre panel cover the two halves of the Lyman-α forest, while the bottom panel covers the region redwards of it.*

6.4.1 Modelling of the absorption features

We have used Astrocook to model all the absorption features with composite Voigt profiles. The modelling involves three main steps:

1. **Detection and identification of the absorbing species.** The standard approach of Astrocook combines detection and identification of the absorbing species in a single procedure. To this purpose, a reference list of ionic transitions typical of the IGM is used. If there are n transitions in the list, n realizations of the spectrum are obtained by converting the spectral wavelengths into redshifts: $z+1 = \lambda/\lambda_{\mathrm{r},i}$, where $\lambda_{\mathrm{r},i}$ is the rest wavelength of each transition, for $1 \leq i \leq n$. The normalised flux f_λ is converted into a measure of *prominence* for the spectral features, which expresses how much a feature stands out from the local continuum in terms of the local noise: $p_\lambda \propto (1-f_\lambda)/\sigma_\lambda$, where σ_λ is the error on normalised flux. The prominence profiles of the n realisations are then multiplied in redshift space: in this way, prominence spikes that show a redshift coincidence across the list of transitions are reinforced, confirming the detection and providing its identification as a candidate absorption system. In the specific case of the UM673, the automatic procedure was facilitated by pre-selecting the absorption-affected regions and by visually inspecting the cases of multiple identification.
2. **Definition of the model components.** Each system is modelled with a variable number of components. The initial number is the number of local maxima in the prominence profile. After a first fit, further components are iteratively added to adjust the model until no significant residuals are left (see below).

3. **Fitting of the model**: a Voigt profile is defined for each component. Each profile is parametrised by an atomic species, with corresponding atomic parameters, and three free parameters. Atomic parameters (such as laboratory wavelength and oscillator strength) are provided within the Astrocook package, and are a compilation of several literature sources. The free parameters are the redshift of the absorber, z, its column density, generally expressed as $\log N$[6], and a Doppler parameter describing the line broadening b ($\mathrm{km\,s^{-1}}$) accounting for thermal broadening. Astrocook also requires information about the instrumental profile, which is convolved with the intrinsic line profile in the observed spectrum. Here it is assumed to be a Gaussian, with a full width at half maximum of 4.28 $\mathrm{km\,s^{-1}}$, corresponding to the nominal resolution of the adopted ESPRESSO configuration. Astrocook optimises the free parameters using nonlinear least squares minimisation. It then reports the best-fit value for the free parameters, the corresponding errors, the χ^2 and the reduced χ^2_ν (i.e., the χ^2 divided by the number of degrees of freedom). While building the model, changes were made only based on χ^2_ν (that is, a component was added or removed only if it was deemed necessary to lower the χ^2_ν of the fit). A composite Voigt profile is then created by combining the overlapping profiles of all components and fitted to the spectrum.

In practice, we started from the secure identifications of metal doublets in the region redward of the Lyman-α emission peak, and looked for associated absorption both outside and within the Lyman-α forest. We then masked all the identified metal absorbers and looked for differences in the Lyman-α forest absorption pattern along the two sightlines. Under the assumption that Lyman-α and Lyman-β could not be too dissimilar in the two spectra, we looked for significant spikes in the difference spectrum (i.e. contiguous regions of several wavelength bins where the absolute value of the difference was larger than 5 times the local error) and added them to the mask of metal absorbers, even when a secure identification was not possible. Lyman-α and Lyman-β lines were first modelled on a combination of the two masked spectra and then re-fitted to the individual sightlines. The procedure was iterated several times, adding new identifications at each iteration, until a solid assessment of the large majority of features was achieved.

A total of 332 Lyman-α /β components were modelled on both spectra, while we identified 371 metal components on A and 232 metal components on B.

6.5 The Lens Model

Determining the physical separation between the lines of sight associated with images A and B requires a model of the gravitational lens. In this regard, we have employed the methodology presented by Koptelova et al. (2014), which incorporates constraints

[6] To simplify the notation, throughout we use $\log N$ in place of $\log_{10}(N/\mathrm{cm^{-2}})$

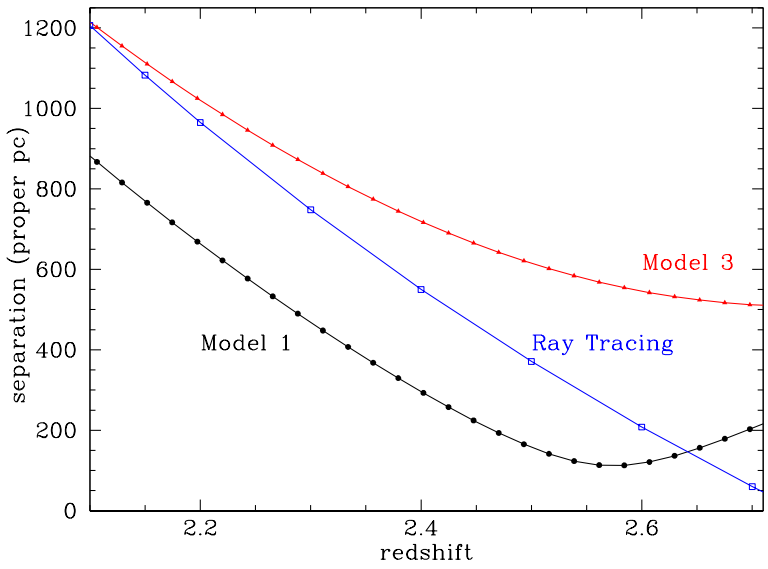

Figure 6.6: *Separation of the lines of sight to the images A and B as a function of the redshift, computed according to the models 1 and 3 of Koptelova et al. (2014). The blue line is derived from a simple ray-trace equation (Smette et al., 1992) assuming for the redshift of the QSO* $z_S = 2.7434$ *(Cooke et al., 2010), for the redshift of the lens* $z_L = 0.493$ *and a separation* $\theta = 2.22$ *arcsec.*

derived from the positional information of the two quasar components, their flux ratio, and a time delay of 89 ± 11 days, modelling the lensing galaxy as a singular isothermal ellipsoid. Nevertheless, further possibilities exist regarding the presence of additional factors, such as an external shear or a shear at the location of one of the galaxies observed with the Hubble Space Telescope (HST) in close proximity to the line of sight (see Koptelova et al. (2014) for more detailed information). While these specific details are not crucial for the analysis carried out in this work, we have adopted the average value between Model 1 and Model 3[7] as the physical separation between the lines of sight associated with images A and B. The uncertainty is estimated as the half-difference between the two aforementioned models. Figure 6.6 shows the range of separations encompassed within the redshift range relevant to the Lyman-α forest ($2.2 < z < 2.7$) where the average separation spans approximately $\langle S \rangle \sim 500\ h_{69.6}^{-1}$ physical parsecs. The blue solid line reports the sightline separation as modelled by simple geometrical optics, following Eq. 6.3.

6.6 The small-scale structure of the IGM traced by the Lyman-α forest

Figure 6.3 shows the transmission of the Lyman-α forest in a section of the spectra of UM 673 A and B. Notably, the spectra exhibit minimal discernible differences, facilitating a straightforward correspondence between the identified absorption features in both images. In the subsequent analysis, we will examine the similarity of the two lines of sight by conducting a global correlation analysis and directly comparing individual absorption features.

[7] Model 3 presented by Koptelova et al. (2014), uses seven parameters and is therefore under-constrained.

6.6.1 Global Correlation

A first statistical estimator of the differences/similarities in the two lines of sight is the Cross Correlation Function (CCF) between the observed transmission and its model, $\xi_{\rm CC}$, computed over the global region between the Lyman-α and the Lyman-β in emission:

$$\xi_{\rm CC}(\Delta v) \equiv \frac{\langle (T(v) - \overline{T}) \cdot (T_M(v + \Delta v) - \overline{T}_M) \rangle}{\sqrt{\langle (T(v) - \overline{T})^2 \rangle \cdot \langle (T_M(v + \Delta v) - \overline{T}_M)^2 \rangle}}, \tag{6.4}$$

where T is the observed transmission along either A or B, T_M is the corresponding modelled transmission, Δv is the velocity lag, and the overline denotes average along the whole wavelength range. The function is defined so as to satisfy $\xi_{\rm CC}(0) = 1$.

Given the similarity between the two forests, we produced a single *reference model* of the Lyman-α and Lyman-β lines (obtained by fitting the spectrum of sightline A) and used AstroCook to compute the CCF along both sightlines in the observed wavelength range $380\,\mathrm{nm}$ to $455\,\mathrm{nm}$ (including only the Lyman-α forest). We masked the regions affected by metal contamination, which were determined with the procedure described in Section 6.4.1 (examples of such regions are shown as shaded areas in Fig. 6.3). The resulting profile of $\xi_{\rm CC}(\Delta v)$ are shown with solid black lines in the top panels of Fig. 6.7. Note that, since the CCF compares pixelated spectral data with an analytical Voigt model of the forest, rebinning or subpixelisation are not required, as we analytically shift the model and re-evaluate it on the data pixel grid, thereby avoiding interpolation and subpixelisation issues.

The CCF peaks at a Δv very close to 0, which would indicate a null lag between the sightline and the model. To estimate the uncertainty of this measurement, we shifted the model around the best fit position (ranging between -200 and $+200\,\mathrm{m\,s^{-1}}$, with a step of $20\,\mathrm{m\,s^{-1}}$) and computed the χ^2 at all positions, after resampling the model on the wavelength grid of the data. We then modelled the resulting profile with a 2nd-degree polynomial around the χ^2 minimum. The 1-σ region around this minimum is defined by the portion of the profile that lies below the minimum χ^2 plus one. With this definition, we obtain a model-data shift of $\Delta v_{\rm min} = -6 \pm 22\,\mathrm{m\,s^{-1}}$ for A and $\Delta v_{\rm min} = -18 \pm 43\,\mathrm{m\,s^{-1}}$ for B. Both values are compatible with zero and with each other. Combining the two values, a velocity difference of $\Delta v = 12 \pm 48\,\mathrm{m\,s^{-1}}$ is measured, indicating an absence of a detectable lag between A and B.

The height of the CCF peak is lower than 1, due to the presence of noise in the spectra. Accounting quantitatively for the effect of noise in Eq.6.4 is difficult, because of the presence of random and systematic noise. Therefore, to assess the effect of noise on the measured CCF we resorted to simulated spectra. The "mocks" were constructed from the AstroCook model of the Lyman-α absorption lines in the non-masked regions (see Sect. 6.4.1) and artificially degraded to match the S/N observed along the A and B sightlines, respectively, in order to account for random errors. We took into account possible systematic errors in the continuum estimation by allowing for a relative tilt in the continuum level of the two mocks. The profile of $\xi_{\rm CC}(\Delta v)$ obtained from the

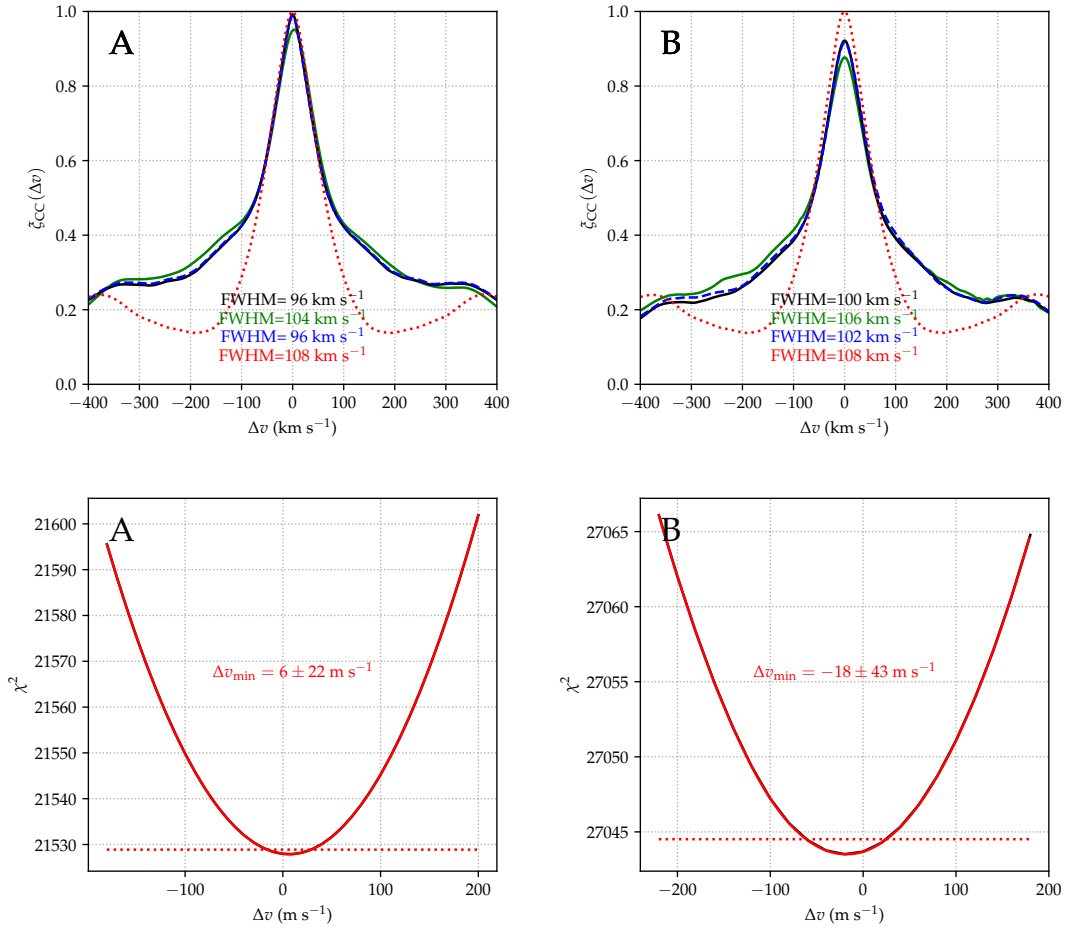

Figure 6.7: *Cross-correlation function (CCF) $\xi_{CC}(\Delta v)$, between data and model in the Lyman-α forest (380 nm to 454 nm observed) for the two sightlines to UM 673. Top panels: CCF profiles for the ESPRESSO observations (solid black line), the HIRES observations (solid green line), mock spectra with the same S/N of the ESPRESSO observations (dashed blue line), and mock spectra with random systems correlated with themselves at infinite S/N (dotted red line). The FWHM of the three distributions are printed with matching colours. Bottom panels: χ^2 profiles around the CCF maxima. The red line (2-degree polynomial fit) almost perfectly overlaps the black line (sampled values). The shift between data model with its uncertainty, computed from the 1-σ region around the χ^2 minimum, is printed in red. The letter in the top left corner of each plot identifies one of the two sightlines.*

mocks is shown with a dashed blue line in Fig. 6.7, adopting a relative continuum adjustment ranging from 0 at the red end to ~20 per cent at the noisier blue end. Other choices of the parameters used to construct the mocks give similar results, showing that the combined effect of observational noise and modelling uncertainties is enough to explain the decrement in the CCF peak value.

The observed FWHM of $\xi_{\rm CC}(\Delta v)$ is about $110\,{\rm km\,s^{-1}}$ for both the observed spectra and the mocks. The broadening of the CCF is a consequence of the typical width of the absorption features in the Lyman-α forest in the velocity space (with the stronger, saturated lines that tend to widen the shape of the CCF). To assess its value, we created a synthetic sightline at (virtually) infinite S/N and populated it with 100 absorption systems. Redshifts, column densities, and Doppler broadening of the systems were chosen randomly in the ranges $2.126 < z < 2.734$, $13 < \log(N_{\rm HI}/{\rm cm^2}) < 14$, $10 < b/{\rm km\,s^{-1}} < 100$. The $\xi_{\rm CC}(\Delta v)$ profile obtained by correlating the synthetic sightline with itself is shown with a dotted red line in Fig. 6.7. The distribution peaks at 1, as expected in the absence of noise, and its FWHM is $112\,{\rm km\,s^{-1}}$, very close to the one of the observed distribution. It is noteworthy that the synthetic sightline does not account for clustering of Lyman-α lines, which is known to have typical scales of $100 - 200$ ${\rm km\,s^{-1}}$ (Cristiani et al., 1997; Saitta et al., 2008; Maitra et al., 2022), thus its CCF does not show the wider wings that are evident in the observed profiles.

The same procedure has been applied to the HIRES spectra and is shown in Fig. 6.7 as a solid green line. The FWHM of this profile is comparable with the one obtained with ESPRESSO, confirming that the width of the CCF is only defined by the typical width of the Lyman-α lines. However, the position of the peak of the HIRES CCF shows a shift of $2.9\,{\rm km\,s^{-1}}$ between A and B (mostly due to a positive offset of A's peak with respect to 0), suggesting the presence of possible systematic effects in the wavelength calibration of the HIRES spectra.

6.6.2 Comparison between individual Lyman-α features

Given the similarity between the Lyman-α forests of image A and B (see Fig. 6.3), it is straightforward to carry out an unambiguous, feature-by-feature comparison of the absorbers along the two lines of sight. To this end, we used the reference model described in Sect. 6.6.1. The *absorption features* to be compared between images A and B have then been defined as the regions enclosed between two adjacent maxima in this fitted model. This specific definition was chosen to avoid potential ambiguities linked to the decomposition of features into Voigt components, and captures the characteristics that are commonly recognised as a "feature" through visual examination. These features represent the distinct signatures of individual absorbers, which can exhibit a simple structure, such as a single Voigt profile, or frequently consist of multiple overlapping components. By employing this definition, a total of 270 absorption features associated with Lyman-α are identified. Among these, in order to avoid the effects of contamination from adjacent absorbers and low S/N, only the 40 *most reliable features*

– i.e. those with well defined boundaries (maxima > 0.85 the continuum level) and large enough equivalent width ($> 0.03\,\mathrm{nm}$) – were used in the following analysis.

Differences in velocity

For each feature, we compared the model with the normalised flux in spectra A and B, sliding the model in velocity space to find the velocity shifts $\Delta v^A_{\mathrm{best}}$ and $\Delta v^B_{\mathrm{best}}$ that gave the best correlation between the model and the data. The resulting distribution of $\Delta v^A_{\mathrm{best}}$ and $\Delta v^B_{\mathrm{best}}$, together with the distribution of $\Delta v^{A-B}_{\mathrm{best}} = \Delta v^A_{\mathrm{best}} - \Delta v^B_{\mathrm{best}}$, are shown in Fig. 6.8. The distribution of the differences in velocity between the absorption features

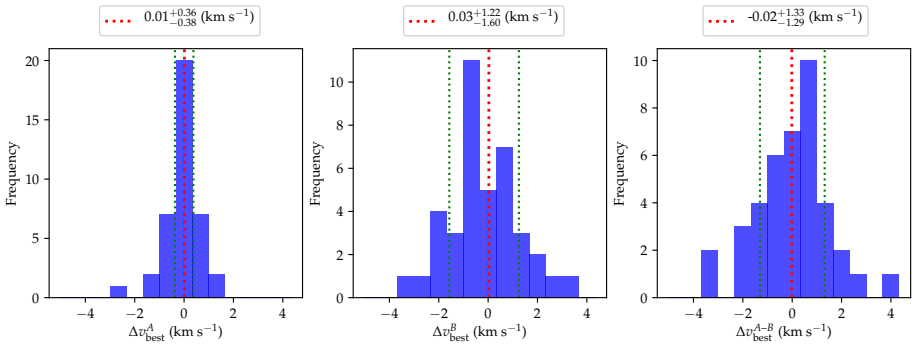

Figure 6.8: *Distribution of the velocity shifts that maximise the cross correlation between model and data, for selected Lyman-α features (see text). Left: sightline A; centre: sightline B; right: difference between the two sightlines. The vertical dotted bars show the average and the 16-84 percentile region, corresponding to the values and confidence interval in the boxes.*

is compatible with a null average and a mean absolute deviation of $0.93\,\mathrm{km\,s^{-1}}$.

If we model the Lyman-α absorbers as slabs following the universal expansion, inclined at an angle θ with respect to the line of sight, the difference of radial velocity, Δv, between two lines of sight with a separation l is expected to be:

$$\Delta v = H(z) \cdot l / \tan(\theta) \tag{6.5}$$

with an expected median $\Delta v \sim 120\,\mathrm{m\,s^{-1}}$, over a random distribution of θ, and a typical separation of $0.5\,\mathrm{kpc}$ at a $\langle z \rangle \sim 2.45$. Clearly, such a signal, if present, is still below our measurement error.

Differences in column density

Under the assumption that the Jeans length is a good estimate of the typical scale of the region where the density is of the order of the maximum density, along any sightline through an absorbing cloud, it is possible to express the observed $N_{\mathrm{H\,I}}$ column density

at a given redshift as a function of the local density contrast, δ (Schaye, 2001):

$$N_{\rm H\,I} \sim 2.7 \times 10^{13}\,{\rm cm}^{-2}\,(1+\delta)^{1.5-0.26\alpha}\,T_0^{-0.26}\,\Gamma_{12}^{-1}\times \\ \times\left(\frac{1+z}{4}\right)^{9/2}\left(\frac{\Omega_b h^2}{0.02}\right)^{3/2}\left(\frac{f_g}{0.16}\right)^{1/2} \tag{6.6}$$

where $\delta \equiv (n_H - \bar{n}_H)/\bar{n}_H$, T_0 is the temperature of the IGM at the mean density, α is the index of the equation of state, $T = T_0(1+\delta)^{\alpha}$ (Hui et al., 1997b), $\Gamma \equiv \Gamma_{12} \times 10^{-12}$ s^{-1} is the hydrogen photoionization rate (Bolton et al., 2007), and f_g is the fraction of the mass in gas. In this way we obtain $N_{\rm H\,I} \propto \rho_b^{\beta}$, with $\beta \equiv 1.5 - 0.26\,\alpha$ and we expect to have (Rauch et al., 2001b):

$$\left\langle(\Delta \log \rho_b)^2\right\rangle \simeq \beta^{-2}\left\langle(\log N_{\rm H\,I}{}^{A} - \log N_{\rm H\,I}{}^{B})^2\right\rangle \tag{6.7}$$

In the literature various values for α have been reported: 1.2 (Gaikwad et al., 2021), 1.5 (Telikova et al., 2019), 1.7 (Walther et al., 2019). Therefore, we make the assumption that the parameter β falls within the range of values between 1.06 and 1.19. In order to measure the differences in the total column densities of the 40 most reliable Lyman-α features we re-fitted their total column density, $\log N_{\rm H\,I,tot}$, independently in the two lines of sight[8]. The differences in logarithmic column density after re-fitting

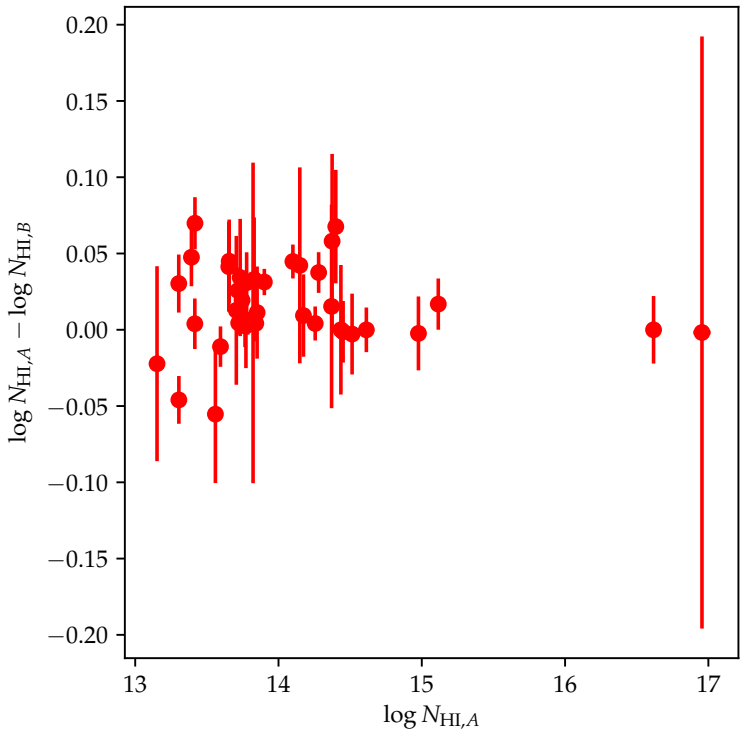

Figure 6.9: *Differences in the logarithmic column density for pairs of Lyman-α absorption features, plotted as a function of the model logarithmic column density.*

[8] To this end, for one of the H I components of each feature in each sightline we defined the column density as the difference between the total column density of the feature and the sum of the column densities of the other components. As a result, $\log N_{\rm H\,I,tot}$ was fitted as one parameter of the composite Voigt model of the feature. The resulting value and its uncertainty are largely insensitive to the choice of the

are displayed in Fig. 6.9 as a function of $\log N_{\mathrm{H\,I}}{}^{A}$. The distribution does not show a significant deviation from zero nor a significant dependence on column density. We obtained $\left\langle (\log N_{\mathrm{H\,I}}{}^{A} - \log N_{\mathrm{H\,I}}{}^{B})^2 \right\rangle = 9.4 \times 10^{-4}$. The mean variance of $\log N_{\mathrm{H\,I}}{}^{A}$ is 3.0×10^{-4} and that of $\log N_{\mathrm{H\,I}}{}^{B}$ is 1.0×10^{-3}. Taking into account the variances of the $N_{\mathrm{H\,I}}$ measurements, the $\left\langle (\Delta \log \rho_b)^2 \right\rangle$ is compatible with a zero value. In order to put an upper limit to the fluctuations in baryon density, $\Delta\rho_b/\rho_b$, we consider that the standard deviation of a variance, for a normally distributed quantity, is expected to be $\sigma_{\mathrm{var}} = \sigma^2\sqrt{2/(n-1)}$, where n is the number of measurements. Then the upper limit to the typical logarithmic change in density can be computed as $\sqrt{\left\langle (\Delta \log \rho_b)^2 \right\rangle} \leq \sqrt{\sigma_{\mathrm{var}}}\,\beta^{-1} = 1.4$ or 1.2×10^{-2} (depending on the assumed value for α, 1.7 or 1.2, respectively) on a scale of $0.2 - 1$ proper kpc; i.e., the RMS fluctuation in the baryon density, $\Delta\rho_b/\rho_b$, is less than $3.1 - 2.8\%$.

Differences in normalised flux

The Lyman-α forests observed in the images UM673A and UM673B present an opportunity for investigating the extent of feedback phenomena, such as galactic winds or supernova explosions. As discussed in Section 6.6.2, an analysis of the fitted Lyman-α features in the spectra of both images reveals no significant discrepancies. However, in order to identify subtle variations in Lyman-α absorption that may have eluded detection during the model fitting process, we conducted a pixel-to-pixel comparison of the normalised flux along the two lines of sight. We considered only the Lyman-α forest, excluding the region of the spectra blueward of $\simeq 381.5\,\mathrm{nm}$, and we masked both metal lines and saturated lines, to isolate the regions most sensitive to potential small-scale disturbances arising from stellar feedback (Rauch et al., 2001b). We masked only metal lines that were surely identified from a redshift coincidence with absorption systems outside the Lyman-α forest. The regions to be masked were defined as the regions where the model of these lines falls below the continuum by more than 1 per cent. Saturated lines were masked by rejecting the regions where the model of Lyman-α absorbers was lower than 0.05 in continuum units. The resulting usable portion of Lyman-α forest amounts to a total wavelength range of $61.48\,\mathrm{nm}$.

Within this region, we computed the normalised flux differences $\Delta f_{A-B} = f_A - f_B$ and $\Delta f_{B-A} = f_B - f_A$, and we look for absorption-like features in both these difference spectra. Absorption features were detected with the same prominence approach described in Sect. 6.4.1 and fitted with Voigt profiles using Astrocook . We selected only features with $\log N_{\mathrm{H\,I}} > 12$ and $b > 5\,\mathrm{km\,s^{-1}}$ that were sufficiently well fitted ($\chi^2_r < 2$), to avoid including a handful of false detections (single-pixel residuals of data reduction mistakenly detected as lines). We visually inspected the detected features to exclude that they were contaminated by nearby metal lines whose tails might have escaped the masking procedure.

Only three features passed these criteria, all three along sightline A (i.e. detected

components used to model the system and provide a robust estimate.

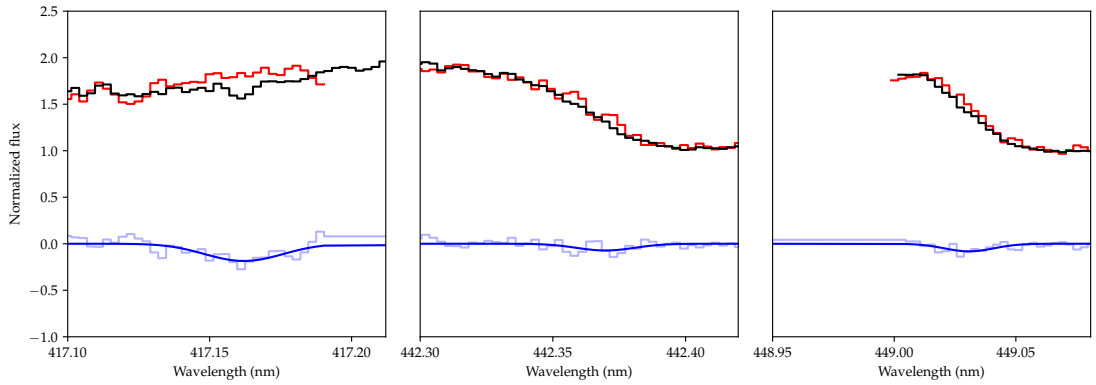

Figure 6.10: *Features detected in the difference spectra. The normalised spectra of sightlines A and B are shown at the top (black and red lines, respectively), vertically offset by 1 unit for displaying purposes. The difference spectra Δf_{A-B} are shown in light blue.*
The corresponding best-fitting Voigt profiles are also shown with their equivalent width (dark blue). The shaded areas correspond to the 2-σ regions around zero.

on Δf_{A-B}). They are shown in Fig. 6.10. The collective rest equivalent width of these features is $1{\times}10^{-2}$ nm, or 0.015 per cent of the usable Lyman-α forest, putting a stringent limit on the fraction of regions affected by local disturbances along the two sightlines. An assessment of these features is uncertain, as they barely extend beyond the 2σ contour obtained by propagating the local error on flux of Δf_{A-B}.

The identification of faint galaxies in close proximity to these features holds the potential to yield insights into feedback mechanisms. Future integral field observations, e.g. with the Multi Unit Spectroscopic Explorer (MUSE) spectrograph (Bacon et al., 2010), may serve as a critical means to establish and elucidate this connection.

6.7 Metal Lines

While the Lyman-α forests of the two images of UM673 appear remarkably similar, the metal lines show significant differences and a global analysis like the one carried out in Section 6.6.1 would be neither effective nor unambiguous. In order to study the differences and similarities of the metal absorptions along the two sightlines, we need, first of all, an operational way to define an "individual metallic feature". Keeping in mind what is typically identified as a "feature" or a "complex of related components" by visual inspection, we have defined metallic absorption features as the spectral regions clearly bounded by two maxima in the flux transmission that can be unambiguously related to the same metallic transition. This definition is different from the one used for the Lyman-α features in Section 6.6.2, due to the inherent complexity of metal absorptions, which poses challenges in providing a comprehensive characterisation. Figure 6.11 shows some examples of these features. The global traits of every feature can then be characterised by the following parameters:

1. Ionic transition;
2. Total column density, computed as the sum of the components' column densities;

3. Redshift-space position, as a weighted mean of the components' redshifts, with weights proportional to the measured column density;
4. Equivalent width.

We grouped together the features that have roughly the same redshift position and are likely to belong to the same physical structure, defining the so-called *Metallic Systems*. Each system is characterised by the presence of one or more ionic transitions, whose velocity profiles resemble each other in shape and number of components. We report in Table 6.1 the metallic systems visible in both images, their redshift as the mean of the values on the two sightlines and the transitions belonging to each system with their respective total column density. The H I column density is also reported for systems whose Lyman-α falls within ESPRESSO's spectral range (i.e. $z > 2.1$). From the initial pool of all the features that were found, we selected the sample used for the rest of the analysis, based on the following criteria:

1. A feature must be present in both sightlines and must have an equivalent width of at least 0.03 Å (observed).
2. A feature must not be contaminated by other transitions.
3. A system must be at least $3000\,\mathrm{km\,s^{-1}}$ away from the quasar emission redshift, in order to minimise the influence of the proximity effect and ensure the reliability of lens model. Effectively, we chose systems with $z < 2.68$.

For ions that produce multiple transitions (e.g. C IV, Fe II), since all of them are modelled with the same parameters, only one feature has been considered. As a rule, we have picked up the strongest transition from the group or the feature that is not contaminated by other lines. With these practical definitions, we defined a sample of 18 features belonging to 9 distinct metallic systems, reported in Table 6.1 with bold letters. Due to the second selection criterion, only three features fall within the Lyman-α forest. This fact reduces the size of our sample but also avoids systematic effects that would be difficult to address during the analysis. From the feature sample, we looked for significant relations amongst their parameters and their differences between the sightlines (e.g. $\Delta \log N = \log N_A - \log N_B$ vs. $\log N_{A,B}$). These relations are dominated by noise, and no significant trend has been found.

6.7.1 Cross-Correlation Analysis

Similarly to the Lyman-α Forest, we computed the CCF for every metallic feature in the selected sample, highlighted with bold characters in Table 6.1 as a function of velocity lag between the sightlines. As the metal features differ significantly in the two spectra and the distribution of the residuals is expected to be non-Gaussian, we have adopted a bootstrap method (Efron, 1979), which is useful when the theoretical distribution of a statistic of interest is complex or unknown, in order to estimate the position error of the peak of the CCF, as shown in Fig. 6.12.

The method is similar to the one described by Peterson et al. (1998). Using Astrocook , we created an ensemble of 100 realisations for the two sightlines by randomly

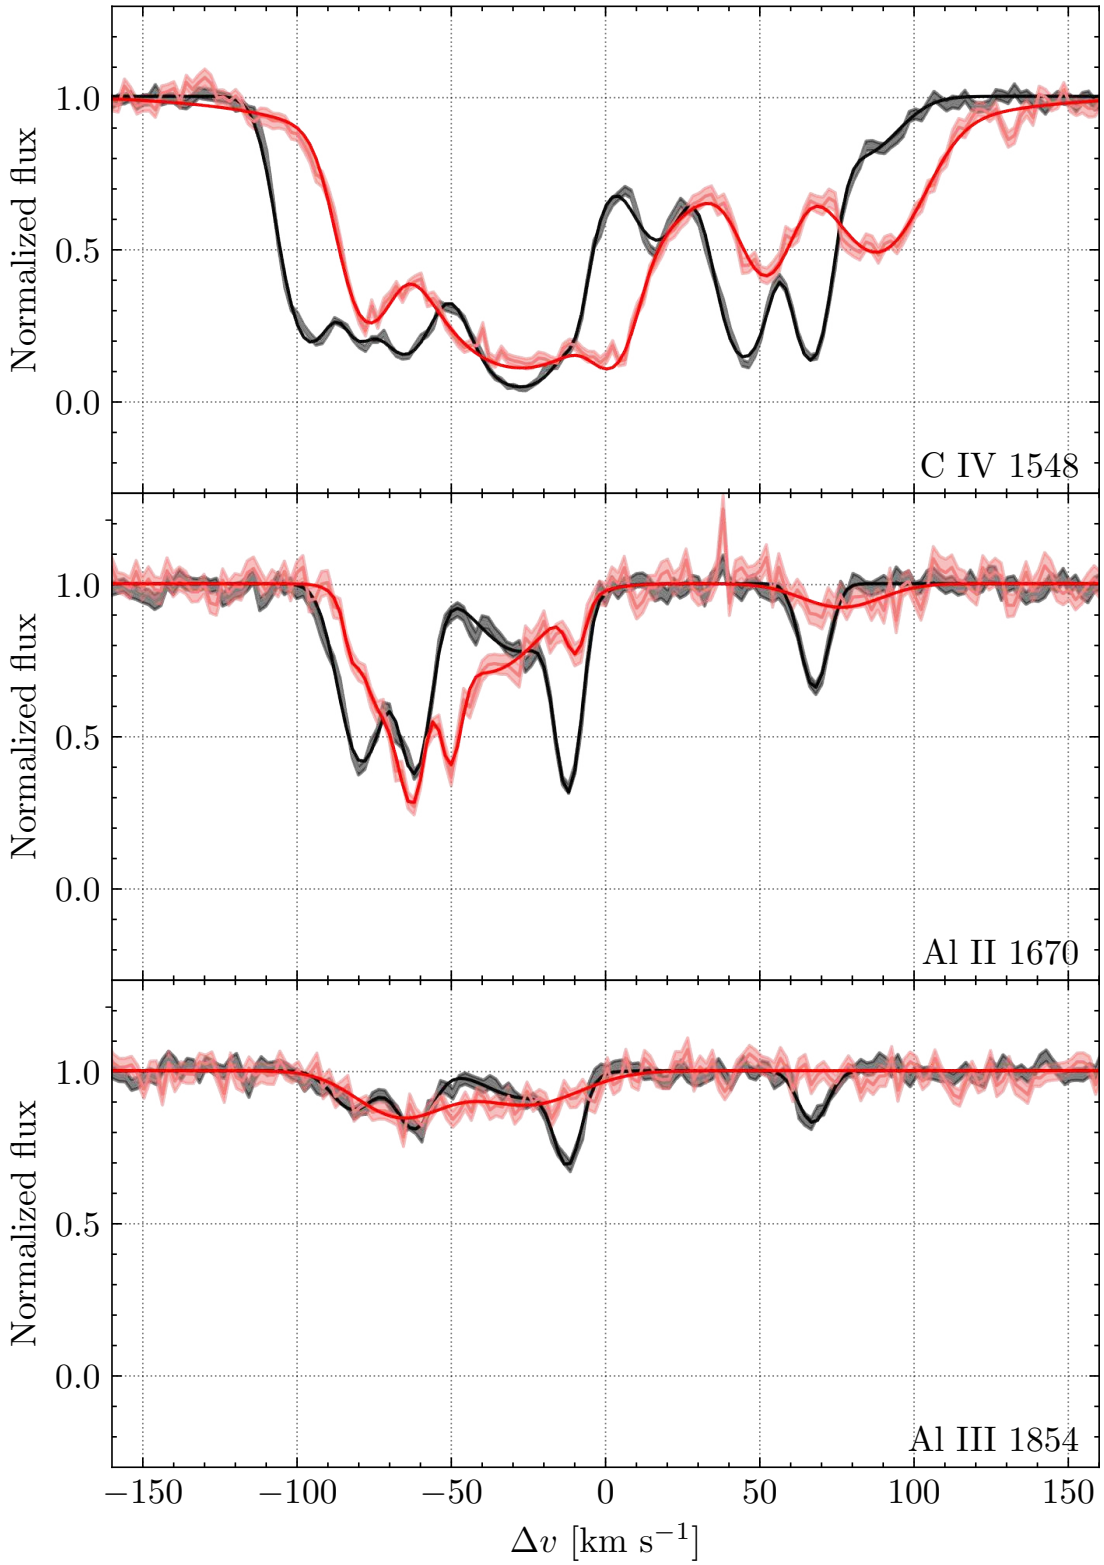

Figure 6.11: *Detail of the C IV, Al II and Al III features of the metallic system at z=1.941 in velocity space. The solid lines define the absorption models of the two spectra. Fluxes are shown as a light-coloured solid line, with a shaded area describing the noise. Image A is black; image B is red.*

z	$\log N_{\mathrm{H\,I}}$	Ion	$\log N_A$	$\log N_B$	Δv [km s^{-1}]
0.564	—	**Mg II**†	15.980	13.940	30.73 ± 0.66
		Ca II	11.808	12.134	32.18 ± 4.72
		Fe II†	13.908	15.471	
1.357	—	**Mg II**	12.601	12.346	17.94 ± 0.84
1.626	—	Fe II	14.433	12.176	7.35 ± 1.45
		Mg II	15.169	12.048	-7.23 ± 1.16
		C IV†	13.250	13.405	
1.772	—	**C IV**†	13.842	13.873	18.69 ± 1.55
1.941	—	**C II**†	14.756	14.990	-2.44 ± 2.11
		C IV	14.539	14.470	-18.56 ± 0.67
		Al II	12.794	12.740	-6.83 ± 1.74
		Al III	12.691	12.579	-0.95 ± 3.68
		Si IV†	14.150	13.996	
1.944	—	**C IV**	14.239	13.834	-1.59 ± 0.35
2.060	—	**C IV**	13.236	12.915	1.77 ±0.42
2.066	—	**C IV**	13.421	13.421	-8.45 ± 2.61
2.356	16.617	**Si IV**	13.925	13.967	1.29 ± 1.62
		Si II	13.649	13.663	0.26 ± 0.28
		S III†	14.201	15.194	
		C II†	14.495	14.428	
		C IV	14.586	14.552	-1.54 ± 0.42
		Al II	12.475	12.425	0.43 ± 0.23
		Al III	12.573	12.504	-0.54 ± 0.34
2.665	15.592	**Si IV**	11.892	12.131	-1.02 ±1.12
		C IV	13.092	13.113	-1.05 ± 0.69
2.707	15.014	C IV	12.967	12.891	1.18 ± 1.35
2.736	15.120	C IV	13.230	13.149	-0.44 ± 0.34
		O VI	13.791	16.639	
2.739	13.285	C IV	12.631	12.556	-0.66 ± 0.79
		O VI	14.148	14.015	

Table 6.1: *List of all metallic systems that are detected on both sightlines. For each system, we indicate the average redshift position between the sightlines and the total H I column density (only for $z > 2.1$ for which the Lyman-α falls in ESPRESSO's range) and the total column density on each sightline. For some of the features was possible to compute the velocity shift that maximises ξ_{CC}. The † indicates transitions that fall within the Lyman-α forest, while features belonging to the final model sample (Sect.6.7) are highlighted in bold letters.*

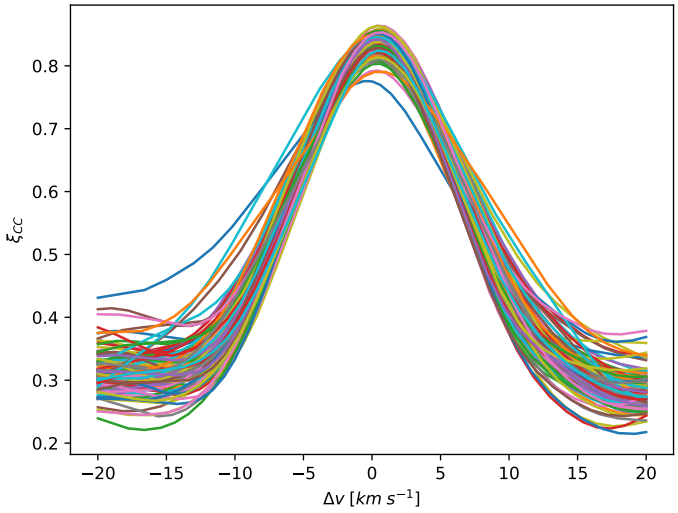

Figure 6.12: *Cross-correlation function ξ_{CC} profiles computed using the bootstrap method for the Al II feature at $z = 2.356$. Every realisation of ξ_{CC} is shown with a different colour and has been fitted by a Gaussian profile. The distribution of the peaks position defines a velocity shift $\Delta v = 0.43 \pm 0.23\,\mathrm{km\,s^{-1}}$.*

sampling the pairs of flux values from the observed spectra. The sampling was done with replacement, but counting just once the pairs that were selected multiple times; due to this approach, the typical number of spectral bins in the realisations was a factor approximately $1 - 1/e$ smaller than the number of spectral bins in the observed spectra. For each pair of realisations in our ensemble, we computed the CCF and fitted a Gaussian profile to its core to determine the position of the maximum. The resulting Δv that maximises the CCF is reported in Table 6.1.

For each system, we can estimate an average redshift position, which directly relates to a separation between the sightlines given by the lens models (see Section 6.5). In order to take into account the uncertainties related to the lens model, we assume for every redshift position a sightline separation given by $l(z) = (l_3(z) + l_1(z))/2$, where l_1 and l_3 are the separations given by Model 1 and Model 3 by Koptelova et al. (2014), respectively. Then, the uncertainty in the separation has been computed as $\sigma_l(z) = (l_3(z) - l_1(z))/2$. The redshift uncertainty of the systems, being negligible, has not been considered.

Figure 6.13 shows the absolute value of the velocity lag that maximises the CCF for each feature as a function of sightline separation. In the figure, the velocity difference is below $2\,\mathrm{km\,s^{-1}}$ for transversal separations $\lesssim 800$ pc, and increases up to $\sim 35\,\mathrm{km\,s^{-1}}$ at higher separations. The observational pattern indicates that the morphology of high-redshift metal absorbers lacks discernible features at scales below a few hundred parsecs, whereas notable variations in the absorption characteristics become evident when examining separations on the scale of kiloparsecs. Besides, it is plausible that this relationship extends above the 10 kpc scale probed by our observation and thus should be investigated with different lenses or pairs of QSOs in future studies.

6.7.2 Characterization of the metal absorbers

To investigate the implications of the metal analysis results, we devised a toy model designed to simulate the absorption of two sightlines by a gaseous structure. It should be noted that the primary objective of our simplified approach was to gain insights into the typical environments responsible for the observed trend and provide an order-of-magnitude estimation of their mass and dimensions. Future analyses, in particular when more data will be available on other lensed systems at different redshifts and separations, will require cosmological simulations (Shin et al., 2021) to better capture the complexity of the astrophysical processes governing galactic feedback in the Circumgalactic Medium (CGM) and IGM, along with fluctuations in the ionising UVB due to He II reionisation.

From the absorption line list obtained in Section 6.4.1, we computed the total H I column densities associated with the selected metallic systems and found values within $13.2 < \log\left(N_{\mathrm{H\,I}}\,/\mathrm{cm}^{-2}\right) < 17$. The system at $z = 1.626$, as described by Cooke et al. (2010), belongs to a Damped Lyman-α System and has not been considered in the subsequent analysis, due to its low redshift that places most of its metal lines in the Lyman-α forest.

The typical H I column densities of the systems listed in Table 6.1 suggest that they do not belong to the general IGM, but to galactic halo structures or to the circumgalactic medium. This is corroborated by the metallicity of the system at $z = 2.356$, the only one with a number of transitions that allows a comprehensive analysis with the CLOUDY package (Ferland et al., 2017). Its metallicity is found to range between solar and one-tenth solar, primarily depending on the assumptions about the UV ionising background (Haardt et al., 1996; Haardt et al., 2012).

A Toy Halo Model

We developed a simple toy model of a DM halo hosting the absorbing structure, defined by a virial mass M that we assumed constant in redshift as a zeroth-order approximation. In a ΛCDM cosmology, a virialised halo has an average density given by (Mo et al., 2010)

$$\overline{\rho_h}(z) = \Delta_v(z)\,\Omega_m(z)\frac{3H^2(z)}{8\pi G}, \tag{6.8}$$

where the virial overdensity Δ_v can be approximated by (Bryan et al., 1998)

$$\Delta_v(z) = \frac{18\pi^2 + 82(\Omega_m(z)-1) - 39(\Omega_m(z)-1)^2}{\Omega_m(z)}. \tag{6.9}$$

The halo's virial radius has been computed as the radius of a sphere of mass M and average density $\overline{\rho_h}$. We assumed a Navarro–Frenk–White (NFW) circular velocity profile

(Navarro et al., 1996)

$$V_h(r,z) = V_{vir}(z) \cdot \sqrt{\frac{f(c \cdot r/R_{vir}(z))}{r/R_{vir}(z) \cdot f(c)}}, \tag{6.10}$$

where

$$f(x) = \ln(x+1) - x/(x+1), \tag{6.11}$$

while V_{vir} is defined as the virial velocity and $c \sim 5$ is the typical concentration parameter for a dark-matter halo at $z \sim 2$ (Zhao et al., 2009). Since we are interested in the velocity of the baryonic component of the halo, we add a factor 1.5 in Eq. 6.10 to account for the rotation of the baryons that, as numerical simulations show (Mo et al., 2010), spin about 1.5 times faster than the hosting DM. We used relation 9 of Vale et al. (2006) to link halo mass to the luminosity of the hosted absorbing structure and computed the radius of the absorber through the Holmberg scaling relation (Holmberg, 1975)

$$R_g = R_\star \left(\frac{L}{L_\star}\right)^\beta, \tag{6.12}$$

where $\beta = 0.4$ (Chen et al., 2001; Nielsen et al., 2013) and $R_\star$ is the effective absorbing radius of a typical $L_\star$ galaxy, that for absorption features of the typical Equivalent Width (EW) limit of our observation is assumed to be $R_\star = 250$ kpc, linearly extrapolating the results of Hasan et al. (2020). As a first-order approximation, we neglected the redshift evolution of $R_\star$. Finally, we pierced the halo model with 1000 sightline pairs separated by l_i, with $0.1 < \log l_i < 1.1$ (l_i in units of kpc). The position of each sightline pair is randomly chosen from a uniform probability circle of radius R_g around the projected sky position of the halo centre. The separation l_i defines the halo's redshift through the lens model. Every pair pierces a different halo with a randomly oriented spin. For the sake of simplicity, we assumed that the velocity of an absorption feature is related to the point along the sightline closest to the halo centre. Therefore, we computed the projected velocity of the baryonic matter along the sightlines at the closest point to the halo centre and retrieved the distribution of the velocity differences between the pairs as a function of the sightline separation.

Best-fit Model

We compared the halo model results to the data in order to estimate the typical halo mass that gives rise to the observed absorption trend. For consistency, we only used the seven C IV absorbers with $|\Delta v| > 3000\,\mathrm{km\,s^{-1}}$ with respect to the emission redshift (blue crosses in Fig. 6.13) to constrain the halo model. Given a mass value, we estimated the probability of such a model to recreate the observed data points by computing the difference between the percentiles of the model-generated velocity distribution that fall between $\Delta v_i \pm \sigma_i$ for every observed data point. The product of the percentile differences on all the data points gives an estimate of the probability of such a mass

recreating the observed trend.

The probability estimate has been computed on 50 realisations of the velocity distribution for 50 log-scaled mass values within $5 \times 10^9 \leq M/M_\odot \leq 10^{12}$. The average of all 50 runs has been smoothed over a 2-bins wide Gaussian window and peaks at $M = 1.9 \times 10^{10} M_\odot$. Figure 6.13 shows the 16-86 and 2-98 percentile intervals of the velocity difference distribution computed by the model assuming the best-fitting mass value. The C IV systems used to constrain the model, as well as the other features from the metal sample, are shown.

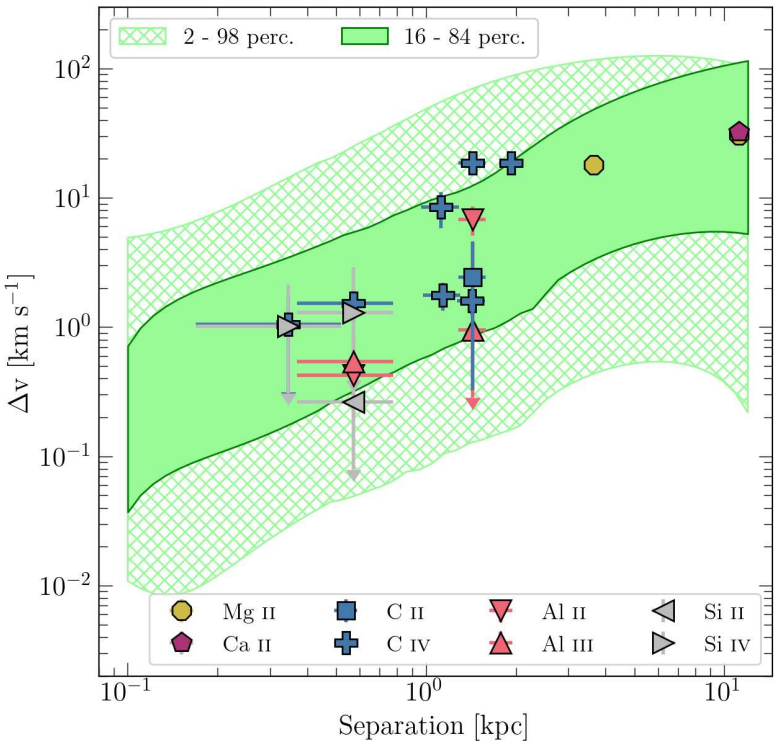

Figure 6.13: *Velocity shifts of metallic spectral features versus the separation between sightlines. The absolute values of the velocity shifts, determined through CCF analysis, are displayed for all metal features within the sample, with distinct ions represented in different colours. Areas shaded in green denote the smoothed 16-84 and 2-98 percentile intervals of the model's velocity distribution that better fits the C IV absorbers. This model is derived using the best-fit mass of $M = 1.9 \times 10^{10} M_\odot$. To enhance visual clarity, instances of low-separation systems exhibiting velocity shifts Δv that are consistent with zero are represented as upper limits.*

Cosmic Incidence

The consistency of the previous mass estimate can be roughly checked by computing the halo's cosmic incidence (i.e. the typical number of halos pierced by a sightline per unit of comoving length; see Sect. 2.3.2) and comparing it to the number of C IV features found on the spectra that fall redwards of the Lyman-α emission. We found $N = 6$ systems over a redshift range $\Delta z = 0.724$ $(1.941 < z < 2.665)$. We can estimate the comoving path density of C IV absorbers as

$$\frac{dN}{dX} = \frac{dN}{dz}\frac{dz}{dX} \approx \frac{N}{\Delta z}\frac{dz}{dX}, \tag{6.13}$$

where dX/dz is the comoving path length along the sightline per unit redshift and is given by Eq. 2.31.

$$\frac{dX}{dz} = \frac{(1+z)^2}{\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}}. \tag{6.14}$$

At $z \sim 2$ this estimate yields $dN/dX \sim 2.67$.

On the other hand, the cosmic incidence of our typical halo structure can be computed from Eq. 2.3.2 as

$$\frac{dN}{dX} = \frac{c}{H_0} n\sigma, \tag{6.15}$$

where $\sigma = \pi R_g^2$ is the spherical cross-sectional area of the absorbing gas (assuming a unity covering fraction $f_c = 1$). The cosmic number density of gas structures n can be obtained by integrating the galaxy luminosity function derived from Parsa et al. (2016)

$$\Phi(L)dL = \Phi_\star \left(\frac{L}{L_\star}\right)^\alpha \exp\left(\frac{L}{L_\star}\right) \frac{dL}{L_\star}, \tag{6.16}$$

with parameters, evaluated at $z = 2$, $\Phi_\star = 0.002$, $\alpha = -1.4$ and $L_\star = 2.2 \times 10^{10} L_\odot$, as fitted from observed UV luminosity functions over the range $0 \leq z \leq 8$. We integrated the product $n\sigma$ over luminosity from a minimum luminosity L_{min} and retrieved the halo cosmic incidence at $z = 2$.

$$\frac{dN}{dX} = \frac{\pi c \Phi_\star}{H_0} \Gamma[x, l] R_\star^2, \tag{6.17}$$

where $\Gamma[x, l]$ is the upper incomplete gamma function

$$\Gamma[x, l] = \int_l^\infty t^{x-1} e^{-t} dt, \tag{6.18}$$

while $l = L_{min}/L_\star$ and $x = 2\beta + \alpha + 1$. Assuming that the lower luminosity bound of the integration $L_{min}/L_\star$ is related to the estimated halo mass $M = 1.9 \times 10^{10} M_\odot$ through relation 9 of Vale et al. (2006), we compute the halo's cosmic incidence $dN/dX = 3.73$, compatible with the C IV estimate.

6.8 Future prospects – Another Lens

The results presented so far rely on only one quasar, and therefore on a limited number of absorber couples and are bound by the geometry of the double lens. The statistical significance of the result is also limited by the redshift range probed by the specific lens-source configuration. The immediate and natural extension of the study, to give more precise and stringent constraints on the dynamical nature of different absorbers, is to look for another target.

Keeping in mind the selection criteria reported in Sect. 6.2, only one other candidate was found in the southern sky – so that it could be observed with ESPRESSO – the quadruply lensed QSO 2M J113440.50-210321.71 ($z_{em} = 2.77$). This object has

been studied in multiple works, especially in the context of time-delay cosmography (Schmidt et al., 2023), and has been observed with HST (see Fig. 6.14,Schmidt et al., 2023), HIRES at KECK (Rusu et al., 2019), ESO Faint Object Spectrograph and Camera 2 (EFOSC2) mounted at the New Technology Telescope (NTT) (Buzzoni et al., 1984), and MUSE at VLT.

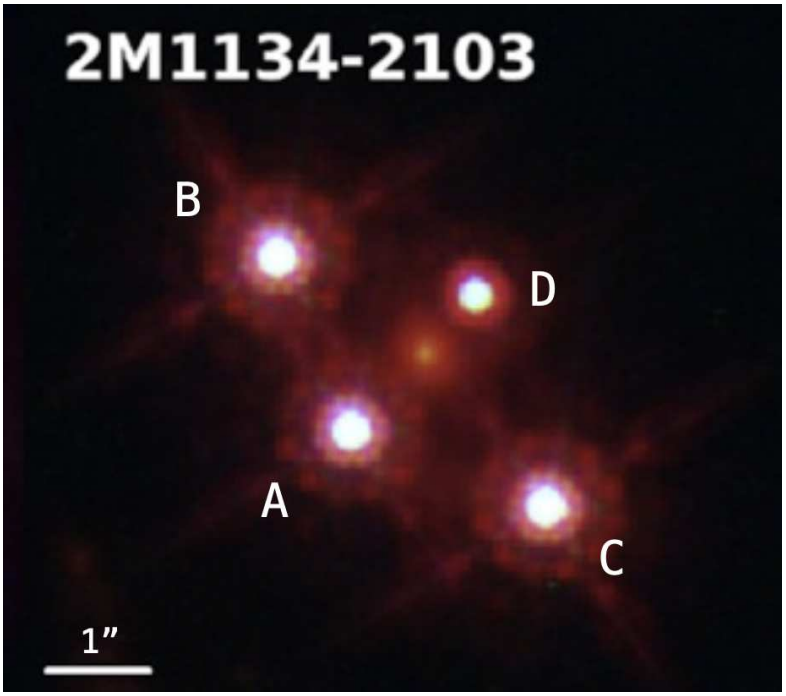

Figure 6.14: *Composite red–green–blue (RGB) image of 2M J113440.50-210321.71, generated from HST observation in bands F160W (red channel), F475X (blue channel), and F814W (green channel). The four images of the quasar are labelled. Adapted from Schmidt et al. (2023).*

The three brightest images[9] of the quasar have $G \sim 17.2$ and an angular separation of 1.89″ (A-B), 2.15″ (A-C), and 3.68″ (B-C). The lens geometry probes the IGM at transversal separations ranging between tens of kpc to sub-kpc (proper) at the redshift of interest, following a simple ray-tracing model, as shown in Fig. 6.15.

Having three distinct usable images, and therefore three distinct sightlines, this target can largely enhance the sample of probed absorbers of the previous study, both for Lyman-α and metal absorbers, and enables a tomographic – or three-dimensional – study of the absorbers. For this reason, we carried out a preliminary analysis of this object on archival MUSE observations, extracting low-resolution ($R \sim 3000$) spectra from the data cube at the position of the three images (on a 3 pixel aperture skewer) and confirmed the presence of 16 metallic absorbing systems along the QSO sightlines, between $z \sim 0.85$ and 2.7, as shown in Fig. 6.16, probing such systems as scales below $\sim 30\,\mathrm{kpc}$, depending on redshift.

Most systems show great variability in EW, gas density and line profile shape across the sightlines, with greater differences at larger separations. Few peculiar systems at high separation show absorption only in one sightline, whilst being undetectable within the noise in the two other images, hinting at the patchy nature of the absorbers. We also looked for galaxies in the MUSE field using SExtractor (Bertin et al., 1996), ex-

[9] Image D is too faint for our purposes.

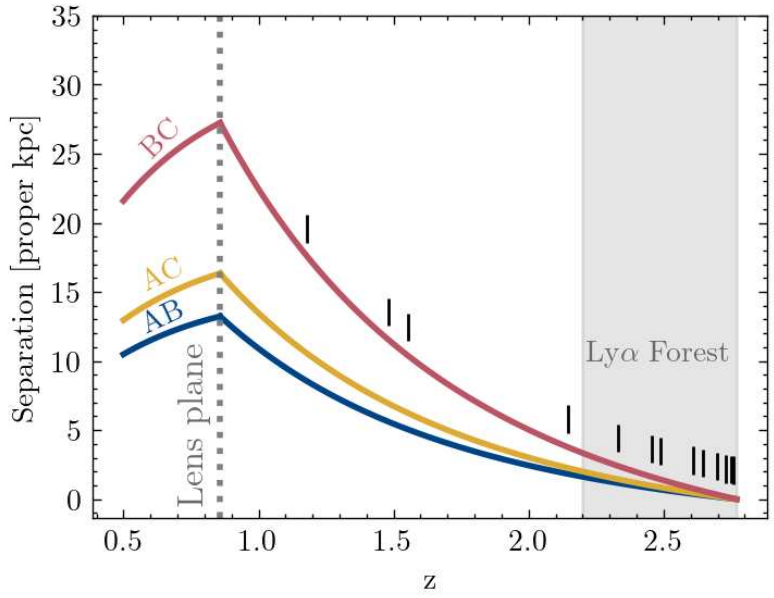

Figure 6.15: *Transversal separation of the three sightlines of images A, B and C as a function of redshift, computed following simple ray tracing (Eq. 6.3). The shaded region highlights the redshift of interest of the Lyman-α forest. Vertical ticks define the redshifts of strong metal absorbers found in the MUSE spectra.*

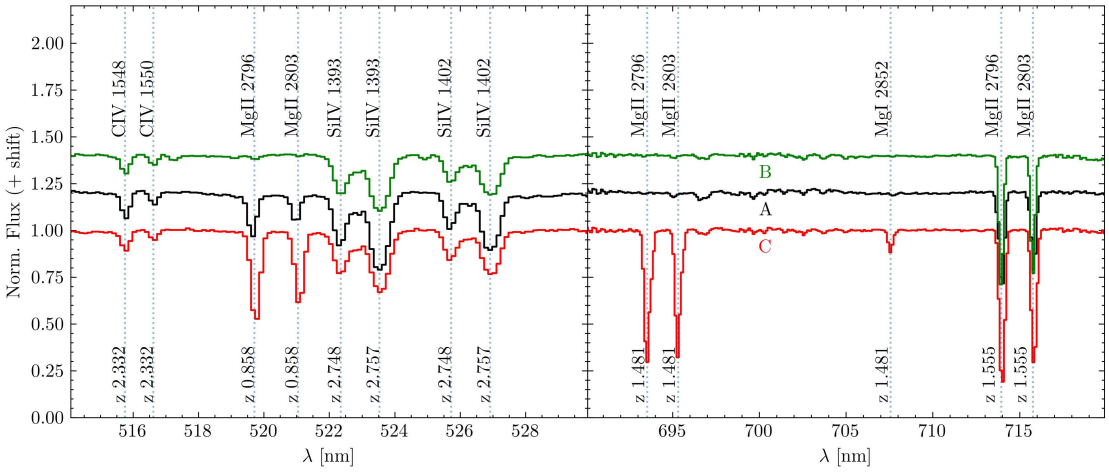

Figure 6.16: *Two sections of the QSO spectra extracted from the MUSE datacube. Metal lines are highlighted with vertical lines. Note that low-z lines show great variability across sightlines.*

tracting 1D spectra at the source positions to identify their redshift by template matching using the Marz software (Hinton et al., 2016) and found 3 galaxies at the same redshift (0.858 and 1.181) as strong absorption systems in the QSO sightlines. This will also constrain the absorber's velocity field and spatial extent if the absorber-galaxy match is confirmed (Augustin et al., 2021; Dutta et al., 2024).

The MUSE data offer valuable insight into the target's unexplored potential; however, they lack the resolution and stability needed for a thorough dynamical examination of the absorber structures, as well as not encompassing the Lyman-α forest of the target. Consequently, we applied for and were granted 19.5 hours of ESPRESSO time, scheduled for the ESO period P116 (April - September 2025). At the time of writing, only part of the allocated time had been executed, with ESPRESSO in 1UT mode, 4x2 binning, with a resolving power of $R \sim 140000$, resulting in a total integration of 4.9 hours, 3.9 hours, and 58 minutes for images A, B, and C, respectively. It

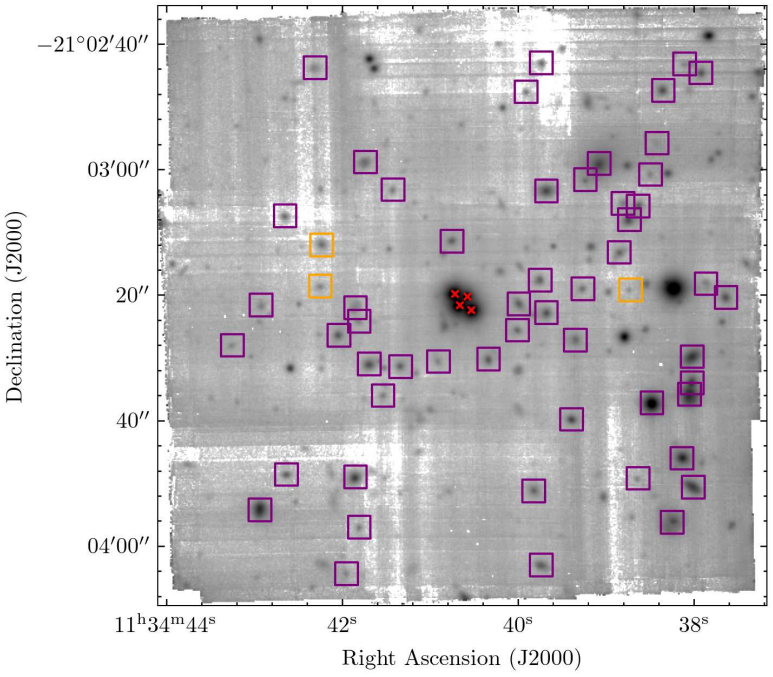

Figure 6.17: *White light MUSE field centred on 2M J1134-2103. The 4 images of the lensed quasar are shown with red crosses. Purple squares show all galaxies found in the field, where the three galaxies having the same redshift as strong absorbers found in the spectra are highlighted in orange.*

is evident that a full analysis is infeasible with these incomplete datasets, particularly because image C's limited integration time would lead to noise dominating the error budget. Nonetheless, we performed a preliminary analysis of the spectra that we have obtained thus far. The data were reduced using ESPRESSO DRS v3.3.10, with wavelengths calibrated to the Solar System barycentre using ThAr lines. The order-by-order spectra were flat-fielded and sky-subtracted. Relative flux calibration was performed to remove the instrument response and ease the continuum fitting process. Order-by-order equalisation was carried out before combining all spectra of the same image to a common wavelength grid of step $1\ \mathrm{km\,s^{-1}}$.

For a quick preliminary analysis, we identified and modelled Voigt profiles to the metal absorption systems redward of the Lyman-α emission, finding the same strong absorbers seen in the MUSE spectra, plus many weaker systems that were undetectable in the low-resolution data. On a first visual inspection, these absorbers show great differences among the spectra at low redshift – i.e. at larger separations – being more similar closer to the quasar emission. However, the forest Lyman-α does not show notable differences between the sightlines at any redshift at this level of noise, as shown in Fig. 6.18.

In the future, when all allocated observations are carried out, these spectra will enable us to significantly improve the limits reported in the previous section. Not only will these data increase the statistical significance of our findings, increasing, for example, the amount of data points in Fig. 6.13 and therefore better constraining our models, but the triple geometry of the lens opens up the possibility of a three-dimensional tomography of the single absorbers. Jointly with ad hoc simulations of galaxies and exploiting the MUSE field, these data can help to tightly constrain galactic environments and turbulent phenomena at the smallest scales.

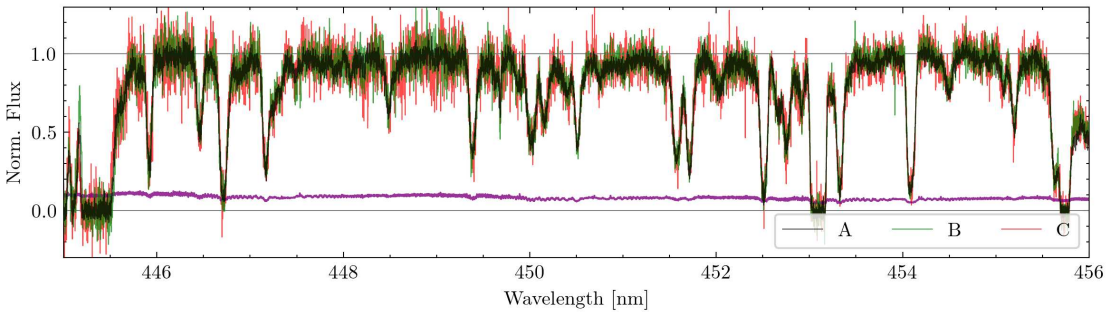

Figure 6.18: *A section of the Lyman-α forest of 2M J1134-2103 as seen in the spectra of the three images. The purple solid line reports the flux error of image C, dominant with respect to images A and B.*

Fig. 6.19 shows the Mg II, Fe II, Al II and Al III lines of the strong absorbing system at $z = 1.555$, where the absorber is probed on scales of $\sim 5 - 15\,\mathrm{kpc}$ by the three sightlines and shows great variability among the spectra.

6.9 Discussion on line stability

In this chapter, we have shown the potential provided by the spectroscopic analysis of lensed quasars in the context of the study of the dynamical properties of the IGM at different scales. We provided proof of the capabilities of the ESPRESSO spectrograph in this field, which can be exploited by the flexible analysis pipeline of the Astrocook environment. Here we outline the results of our analysis of UM673 on the small-scale structure of the IGM.

6.9.1 Lyman-α forest

The analysis of the Lyman-α features highlighted the absence of any global systematic drift effects between the two sightlines and showed that the width of the CCF profile is due to the width of the lines and the typical clustering scale of the Lyman-α forest in velocity space. With a feature-by-feature approach, we also showed that the Lyman-α clouds motions are coherent on sub-kpc scales with a confidence level of $\sim 1\,\mathrm{km\,s^{-1}}$.

We investigated whether, with the present observational setup, the effects of the universal expansion on the Lyman-α absorbers, modelled as slabs following the expansion, could be detected. The expected median $\Delta v \sim 120\,\mathrm{m\,s^{-1}}$, for a typical separation of $0.5\,\mathrm{kpc}$ at a $\langle z \rangle \sim 2.45$, turns out to be below the sensitivity of the present observations, but a detection is not beyond reach, if a higher S/N or/and a wider separation of sightlines is achieved.

In the context of the Sandage test, the observed global shift of the Lyman-α forest of $\Delta v = 12 \pm 48\,\mathrm{m\,s^{-1}}$ can be compared with the velocity accuracy scaling relation presented by Liske et al. (2008) and described in Eq. 5.8. In our case, images A and B must be considered as separated sources and, with an average S/N of ~ 26 and ~ 16, respectively, per ~ 0.0028 nm pixel in the Lyman-α forest ($\sim 90\%$ of which has been

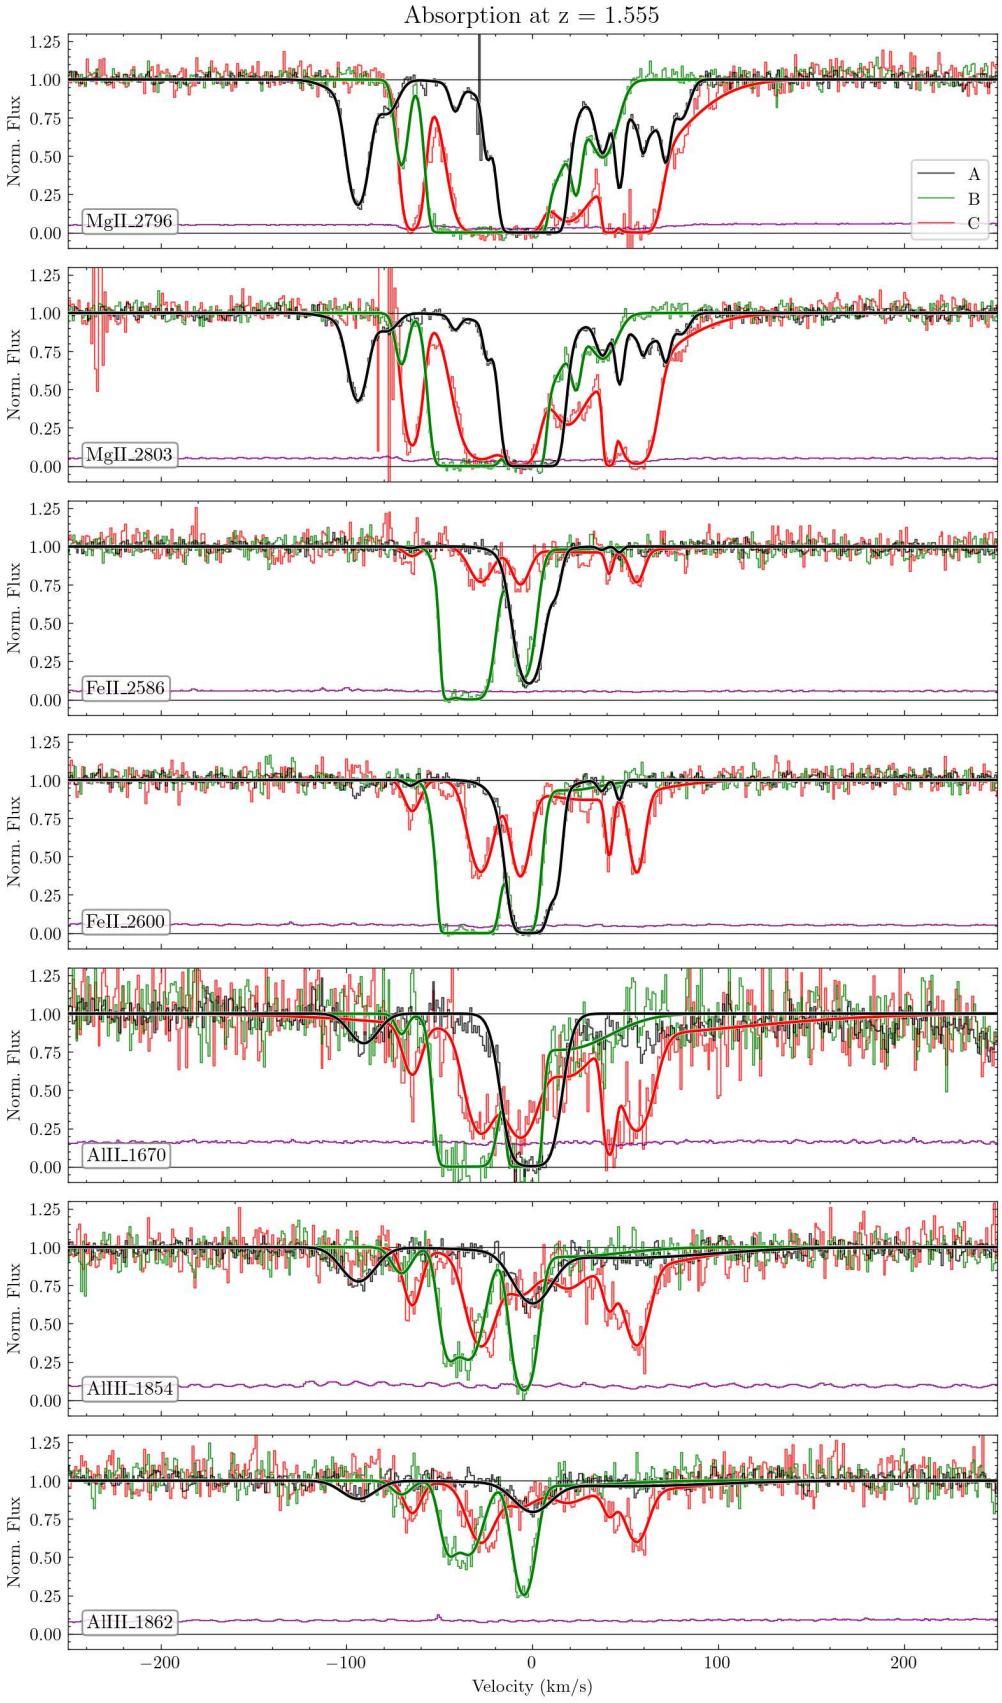

Figure 6.19: *Detail of the Mg II, Fe II, Al II and Al III lines of the metallic system at z=1.555 identified on the spectra of 2M J1134-2103 in velocity space. The normalised fluxes of the three sightlines are shown: A (black), B (green), and C (red). Voigt profile models fitted to the three spectra are overplotted, with the same colour coding. The solid purple line describes the normalised flux error of spectrum C, the noisiest of the set.*

used for the CCF), we obtain an expected total accuracy in the velocity position of the features

$$\sigma_v = \sqrt{\sigma_{v,A}^2 + \sigma_{v,B}^2} \simeq 50\,\mathrm{m\,s^{-1}} \tag{6.19}$$

which is the same order of magnitude uncertainty recovered by the CCF process developed in Section 6.6.1, corroborating the proposed scaling relation on real data, even at a relatively low S/N.

Lastly, we can check whether the infinitesimal changes in the Lyman-α absorbers over a given separation will produce systematic effects in the measure of the redshift drift. From D'Odorico et al. (2006) we assume that the Lyman-α clouds have peculiar velocities likely smaller than $\lesssim 100\,\mathrm{km\,s^{-1}}$. Over a 20-year baseline for a hypothetical redshift drift experiment, these clouds would cover a distance of about 0.002 pc in proper units. In Sect. 6.6.2 we observed a null velocity shift of the Lyman-α features on 1 kpc scales, with an upper bound of $1\,\mathrm{km\,s^{-1}}$. Linearly extrapolating this upper bound to milliparsec scales, we retrieve a velocity lag of $\sim 0.2\,\mathrm{cm\,s^{-1}}$, which is about two orders of magnitude smaller than the expected redshift drift effect of $12\,\mathrm{cm\,s^{-1}}$ over the same 20-year baseline. Therefore, the peculiar velocities of the Lyman-α forest appear to be too weak to induce significant effects on experiments aimed at measuring $\dot{z}$.

6.9.2 Metals

In contrast to the Lyman-α clouds, significant differences are observed in the metallic absorption pattern between the two images. We conducted an analysis of these dissimilarities utilising the cross-correlation function of the fluxes as a function of velocity lag and evidenced a relationship between the observed velocity shift along the sightlines and the spatial separations of the metallic absorbers.

We have observed that the velocity lag of these absorbers remains close to zero for scales below 800 pc. Although one might be tempted to infer coherence among the metal absorbers below this scale, our sample lacks strong metallic absorption instances at $z > 2.5$ (the typical redshifts sampling the shorter separation scales). Consequently, the higher-redshift features with shorter separations may appear indistinguishable within the noise, primarily due to their relatively small column densities. Through the use of a toy model, we have shown how the relation between velocity difference (Δv) and separation can be induced by rotating virialised dark matter halos with a virial mass of approximately $M_{\mathrm{vir}} \sim 1.9 \times 10^{10} M_{\odot}$. It is important to note that our analysis provides only an order-of-magnitude estimate. Nonetheless, even with this rudimentary model, we successfully replicated the observed trend and found consistency with the cosmic incidence of the absorbers.

In relation to the Sandage test of the cosmological redshift drift, considering the halo model discussed above, it can be inferred that the metallic absorbers exhibit peculiar velocities of the order of $\sim 200\,\mathrm{km\,s^{-1}}$. Consequently, over a duration of 20 years, these absorbers are anticipated to traverse a proper distance of approximately 0.004 pc. By linearly scaling the observed velocity shear of $\Delta v \sim 10\,\mathrm{km\,s^{-1}}$ at 1 kpc to the 0.004

pc scale, we can estimate the expected velocity change of a metallic absorber over a 20-year interval, which amounts to $\sim 4\,\mathrm{cm\,s^{-1}}$. The velocity difference resulting from the motion of metallic absorbers over a 20-year baseline is thus of the same order of magnitude as the velocity shift induced by the accelerated expansion of the Universe at the typical redshifts considered for the test. These velocity differences introduce a non-negligible noise effect that needs to be taken into account when measuring $\dot{z}$ using metallic lines. Moreover, this noise effect may pose a substantial limitation for specific approaches such as the *Lyman-α cell* technique proposed by Cooke (2020).

These results are limited to the two UM673 sightlines. With the completion of the observation of 2M J1134-2103, we will be able to probe a larger range of separations, with a statistically greater number of absorbers, providing tighter constraints on these conclusions and expanding our understanding of the 3D dynamical processes that shape the intergalactic gas.

In the meantime, having shown that the Lyman-α forest lines are stable over many decades, we can exploit them as a tracer of the cosmological expansion. In the following Chapter, we set ourselves to start the Sandage-Loeb test using current facilities such as ESPRESSO, in order to study the systematic effects related to the experiment and properly forecast when a detection of the acceleration will be reached.

7

The first Lyman-α forest Redshift Drift Experiment

The previous chapter demonstrated the physical stability of Lyman-α lines in quasar spectra across large spatial and temporal baselines, confirming their suitability for measuring cosmic acceleration. In contrast, metal lines originate from clumpy, dynamic systems, which could introduce noise on the same order of magnitude as the cosmological signal. Building on our findings, this chapter addresses the critical challenges that precede a real-world experiment. Our primary goal is to develop, test, and validate a comprehensive, end-to-end analysis pipeline for measuring the redshift drift.

We will begin by using the existing ESPRESSO spectrograph to collect 31 hours of integration time on two bright quasars from the QUBRICS Golden Sample. These observations will be spread across three epochs. This initial campaign will allow us to estimate the currently achievable precision and improve upon existing upper limits from the literature.

A key objective is to identify and characterise any systematic effects, whether astrophysical or instrumental, that emerge at the current signal-to-noise ratio. We will then propose methods to mitigate these effects. Finally, we will extrapolate our findings to a future ELT/ANDES observational campaign to estimate the total integration time required to achieve a significant detection of the cosmological expansion.

7.1 Target selection

For the present case study and the first epoch observations, we considered the brightest object from the Golden Sample, J052915.80-435152.0, hereafter called Super Bright 2 (SB2). Parallelly, in Sect. 7.8 we discuss the observation of the second-most luminous object of the sample J212540.97-171951.4 (hereafter called SB1)[1]. We set the stage for

[1] The terminology "Super Bright quasars" used in this paper is purely chronological and does not reflect hierarchy based on luminosity, originating from the observation sequence: J212540.97 was identified at an earlier stage in the QUBRICS survey and is designated as SB1, whereas the brighter J052915.80 was

the long-term experiment, starting with two objects in order to maximise our abilities to investigate systematic effects due to cosmic variance, while still requiring a limited amount of telescope time. SB2 has the identifier QID 1128023 in the QUBRICS database (Calderone et al., 2019) with coordinates (J2000) RA 05:29:15.81 and Dec -43:51:52.1 and magnitudes $r = 16.3005$, $i = 16.1184$ (Onken et al., 2024) and $G = 16.3452$ (Gaia Collaboration et al., 2021)

Besides being particularly suited for the Sandage test, SB2 is an outstanding object, the most luminous quasar known in the Universe. Medium resolution spectra taken with X-Shooter (Vernet et al., 2011) reveal that SB2 has a bolometric luminosity of $\log(L_{\rm bol}/{\rm erg\ s^{-1}}) = 48.27$, supported by an accretion of $370 M_\odot$ per year onto a SMBH of mass $\log M/M_\odot \sim 10.24$ or, in other terms, about one solar mass per day accreting onto a SMBH of ~ 17 billion solar masses. From the same data, Wolf et al. (2024) report an emission redshift for SB2 of $z_{em} = 3.962$.

7.2 Data acquisition and treatment

ESPRESSO observations have been carried out using the single UT mode (Pepe et al., 2021b) and achieve a resolving power of $R = \lambda/\Delta\lambda \sim 135000$, where $\Delta\lambda$ is the full width at half maximum of resolution element ($\sim 2.2\ {\rm km\,s^{-1}}$ in velocity space), with a 4x2 detector binning, chosen to optimise the S/N in the region covering the Lyman-α forest. A summary of the observations is given in Table 7.1.

The observations were reduced using the ESPRESSO Data Reduction Software[2] (DRS, Di Marcantonio et al. 2018), version 3.3.0. LFC frames were used for wavelength calibration of the spectra. We also produced spectra calibrated on the Fabry-Pérot etalon (FP) etalon combined with a ThAr lamp frames to investigate the presence of possible instrumental systematics due to calibration stability and accuracy (see Sect. 7.7.5).

The final products of the pipeline were flat-fielded, blaze-corrected, sky-subtracted, and wavelength-calibrated. The spectra were also corrected to the Solar System barycentre using standard pipeline procedures, employing INPOP13c ephemeris (Fienga et al., 2014). Earth's precession is taken into account in the correction, consistently with the requirements for achieving sub-meter-per-second precision (Wright et al., 2014).

Order-by-order spectra were optimally extracted along detector rows (in the cross-dispersion direction), retaining for each pixel the original calibrated wavelength of the rows themselves. Relative flux calibration is carried out by removing the instrument's response function. This is done to ease continuum estimation and reduce systematics related to its position. To compensate for spurious distortions in the spectra due to differences in instrumental response, atmospheric conditions and fibre positioning on the sky, all orders of each spectrum were equalised to match the median flux values of

discovered later.

[2] https://www.eso.org/sci/software/pipelines/espresso/

Table 7.1: *Summary of single Observing Blocks (OBs) of SB2 taken during observing periods P110 and P112. The first column specifies the ESO period in which the observations were collected. The second column gives the date and hour of the start of observations (UT), which is also reported in MJD format in the third column. The fourth column gives the exposure time in seconds. The fifth column reports the OB grade, given by ESO on a set of constraints on observing parameters (image quality, airmass, sky transparency, fractional moon illumination, moon distance, precipitable water vapour, twilight allowance): 'A' means that all constraints are fulfilled; 'B' and 'C' mean that at least one constraint is exceeded by less than 10% or more than 10%, respectively. The sixth and seventh columns report the seeing and airmass, respectively, as quoted from the FITS files' headers. Upon inspection, grade B and C observations did not show a decrease in the achieved S/N. Exceptionally, we consider them for our purposes since we do not expect any strong systematic influence on the final results at our level of S/N. In the future, only grade A observations will be considered.*

Period	Exposure Start (UTC)	MJD	t_{exp} [s]	Grade	Seeing [arcsec]	Airmass
P110	2022-10-29 07:21:01	59881.31068001	3438.0	C	0.93	1.066
	2023-01-22 01:04:40	59966.04843679	3438.0	A	0.62	1.071
	2023-01-22 02:13:07	59966.09453567	3438.0	B	0.72	1.075
	2023-01-23 01:14:56	59997.04166665	3438.0	C	1.28	1.066
	2023-02-22 00:54.53	59967.05436154	3438.0	A	0.60	1.116
P112	2023-08-25 08:41:12	60285.19820316	3438.0	A	0.42	1.278
	2023-11-11 05:22:38	60284.18013713	3453.0	C	0.99	1.087
	2023-11-19 05:02:35	60345.18147687	3453.0	A	0.64	1.078
	2023-12-06 04:15:02	60259.22743581	3453.0	A	0.49	1.069
	2023-12-07 04:41:30	60267.21411237	3453.0	A	0.56	1.065
	2023-12-11 04:47:35	60289.20252708	3453.0	A	0.57	1.070
	2023-12-13 05:22:56	60291.22794520	3453.0	A	0.54	1.103
	2024-02-05 04:15:20	60181.36712803	3453.0	C	0.50	1.512

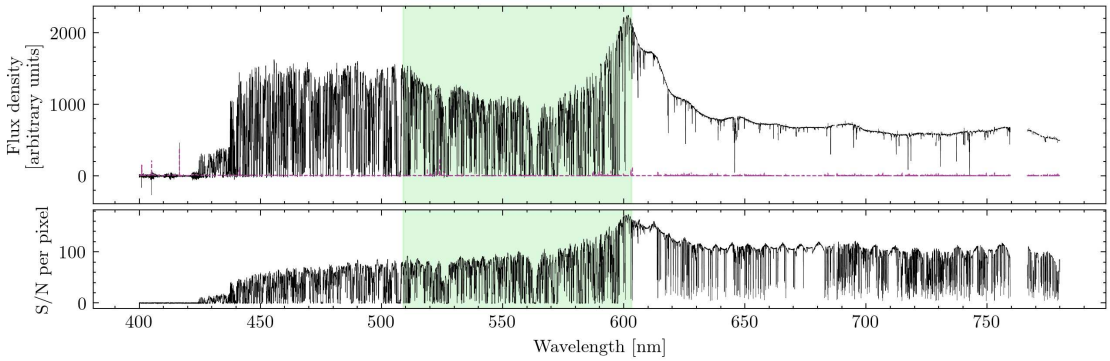

Figure 7.1: *Combined spectrum of SB2. Top panel: Flux density in arbitrary units is shown in black, with the dashed purple line denoting the flux density error. Bottom panel: S/N per* $1\ \mathrm{km\,s^{-1}}$ *pixel. The shaded green area highlights the Lyman-α forest considered in the redshift drift measurement, bound by the Lyman-α and Lyman-β emissions of the quasar, namely between* $509 - 603\ \mathrm{nm}$.

the orders of the first observation[3]. Subsequently, a sky mask was applied to the single exposures while retaining an order-by-order format, shifted by the specific Barycentric Earth Radial Velocity (BERV).

The coaddition of the extracted spectra was performed using the Astrocook software[4] (Cupani et al., 2020), following the 'drizzling'-like approach described in Cupani et al. (2016). A wavelength grid was defined with a fixed step of $1\ \mathrm{km\,s^{-1}}$, and the pixels contributing to each bin in the grid were selected from the order-by-order spectra, to combine them all together in a single rebinning procedure. Contributing pixels were weighted by their associated error and by how much they overlapped with the wavelength range of the bins themselves.

Coaddition of individual exposures resulted in a single spectrum, with a median S/N ~ 86 per $1\ \mathrm{km\,s^{-1}}$ pixel at the continuum level in the forest, and ~ 105 redward of the Lyman-α emission, as shown in Fig. 7.1. The Lyman-α forest highlighted in the figure is defined as the spectral interval between the Lyman-α and Lyman-β emission of the quasar, namely $509 - 603\ \mathrm{nm}$, shown in more detail in Fig. 7.2. This section of the spectrum contains Lyman-α absorption lines with $3.18 \leq z \leq 3.96$, and no Lyman-β absorptions.

Observations taken in the two ESO observing periods P110 and P112 (see Table 7.1), respectively, were also combined separately to provide two independent spectra used as the two epochs of the SL test. The first (P110) and second (P112) epoch spectra have a median S/N level at the continuum in the Lyman-α forest of ~ 47 and ~ 72, respectively. We defined the temporal baseline occurring between the two epochs' spectra as the difference between the mean observational date of each period $\Delta t = 0.875\ \mathrm{yr}$.

7.2.1 Removing metal transitions from the Lyman-α forest

As shown in the previous Chapter of this thesis, Lyman-α and metal lines have significantly different dynamical behaviours on the temporal and physical scales probed

[3] The choice of the reference spectrum in the equalisation does not bias our results.
[4] https://github.com/DAS-OATs/astrocook

by the experiment. We showed that absorption lines due to neutral hydrogen on adjacent sightlines (with sub-kiloparsec separations) appear identical within the noise, whereas metal lines show significant differences in their velocity structure. These intrinsic dynamical effects can be a source of strong systematic effects, inducing a velocity shift in the metal lines up to a few centimetres per second per year, of the same order of magnitude as the expected cosmic signal. We therefore removed the metal-polluted regions from our analysis to reduce possible systematics.

This process was done using the AstroCook software (Cupani et al., 2020) to provide a secure identification of metal lines in the spectrum. First, the quasar continuum emission was estimated by applying an iterative sigma-clipping procedure to remove absorption features while masking sky emission lines, telluric absorption lines in the barycentric frame of reference, or spurious, narrow spikes in flux (e.g. 'cosmic rays' that were not removed by the coadding process), and manually correcting the continuum level where necessary. Sky and telluric lines were masked using a reference SkyCorr model (Noll et al., 2014) and empirically defining proper cuts to conservatively remove even the line wings. We performed a search for metal absorption lines throughout the spectrum, starting from line doublets (e.g. C IV, Si II and Mg II) redward of the Lyman-α emission and following up with associated lines, i.e. other ionic transitions at the same redshift as the found doublets. Table 7.2 reports the absorption systems, i.e. the groups of ionic transitions with the same redshift, found along the spectrum of SB2.

Once identified, we fitted composite Voigt profiles to the metal absorption lines that fall redwards of the Lyman-α emission, estimating the absorber's redshift, column density, and Doppler broadening. Metal lines bluewards of the Lyman-α emission were not fitted, as this would require modelling multiple blended H I components, making the analysis significantly more complex. Since this is beyond the scope of the present study, we left it unaddressed and masked out such lines, removing the spectral regions in velocity space associated with higher wavelength transitions.

7.2.2 A sub-DLA at $z_{\rm abs} = 3.63$

As is visible in Fig. 7.1, the spectrum of SB2 presents a strong H I absorption at $\lambda \sim 563$ nm or $z \sim 3.63$. We study this absorbing system in more detail to understand, based on its column density, whether it might belong to a dense environment prone to proper motions related to galactic dynamical processes, or it stems from the sparse cosmic web following the Hubble flow and thus can be considered in the redshift drift measurement. We found strong and complex associated metal absorption redward of the Lyman-α emission, stemming from ions with both low (O I, Al II, C II, Fe II, Fe II) and high (C IV, Si IV) ionisation states. We fitted the Lyman-α absorber within AstroCook using 3 components, tying the redshifts of H I and low ionisation species. The total H I column density of the system is $\log N_{\rm H\,I} \approx 19.8$, corresponding to a sub-Damped Lyman-α System (sub-DLA, Péroux et al. (2003)).

z	Ions
0.96	Fe II, Mg I, Mg II
1.13	Fe II, Mg I, Mg II
1.44	Fe II, Mg II
1.63	Fe II, Mg II
1.84	Al II, Al III, FeII
2.12	Al II, Fe II, SiII
2.30	Al II, Al III, C II, C IV, Si II, Si IV
2.33	C IV
2.46	C IV, Si IV
2.91	C IV
3.01	C IV
3.13	C IV, Si III
3.17	Al II, C II, C IV, Si II, Si III, Si IV
3.29	C IV
3.33	Al II, C II, C IV, O I, Si II, Si IV
3.29	C IV
3.49	C III, C IV, Si III, Si IV
3.59	C IV
3.60	C III, C IV, Si III, Si IV
3.63	Al II, C II, C IV, Fe II, O I, S IV, Si II, Si IV
3.67	C IV, Fe III
3.71	C IV, Si IV
3.78	C IV, Si III, Si IV
3.80	C IV, Si III, Si IV
3.84	C III, C IV, Si IV
3.86	C IV
3.90	C III, C IV

Table 7.2: *Metallic absorption systems identified in the spectrum of SB2. The left column reports the system redshift, and the right column the identified atomic species.*

Given this value, we decided to exclude the spectral range tied to this absorber from the analysis, as it is probably related to an intervening galactic system and could potentially induce systematics in the measurement due to its peculiar motions. We thus mask the region of the spectrum spanning ± 2000 km s^{-1} from the column density weighted average redshift position of the system. The final masked spectra used in our analysis have 39220 pixels (1 km s^{-1} wide) in the Lyman-α forest, compared to 50805 pixels without mask. Meaning we consider only 77% of the total Lyman-α forest. Fig. 7.2 shows the normalised flux of the Lyman-α forest in the spectrum of SB2, highlighting the masked-out regions.

7.3 Expected velocity precision

In a ΛCDM Universe (Planck Collaboration et al., 2020), the expected cosmological redshift drift at redshift $z = 3.573$ (i.e. the middle of SB2's Lyman-α forest) is $\dot{z} = -6.5 \times 10^{-11}$ yr^{-1}, equivalent to $\dot{v} = -0.43$ cm s^{-1} yr^{-1} in velocity space. On the baseline of our experiment, the expected velocity shift between the two epoch spectra is therefore $\Delta v = -0.38$ cm s^{-1}.

Liske et al. (2008) provided a practical formula to assess the precision expected from an experiment performed on the ELT based on an ensemble of quasars at different redshifts (See Sect. 5.3.2, Eq. 5.8). We adapted this formula to estimate the precision expected from our data. Given the SB2 ESPRESSO spectra at the two epochs, we estimated the predicted precision in the velocity shift measurement as

$$\sigma_v = 2 \cdot g \cdot \left(\frac{S/N}{4058}\right)^{-1} \left(\frac{1 + z_{\rm QSO}}{5}\right)^{-1.7} \left(\frac{N}{30}\right)^{-1/2} f_{Ly\alpha}^{-1/2} \ \text{cm s}^{-1}, \tag{7.1}$$

where S/N = 86 is the median S/N at the continuum level per 1 km s^{-1} pixel per object, accumulated over all observations (i.e. the S/N level of the combined spectrum). $N_{\rm QSO}$ is the number of quasars in the sample ($N_{\rm QSO} = 1$ in our case), and $z_{\rm QSO} = 3.962$ is the redshift of the quasar. The γ exponent is 1.7 for $z_{\rm QSO} \leq 4$ and 0.9 above. The factor $f_{\rm Ly_\alpha} = 0.77$ is the fraction of the Lyman-α forest considered after masking. The "form factor", g, depends on the observational strategy and the distribution of the single observations throughout the program. In general terms, assuming N_e observational epochs, the j-th epoch being separated by the beginning of the program by Δt_j, we can define $h_j = \Delta t_j / \Delta t$, where Δt is the total temporal baseline of the experiment. Similarly, by specifying the fraction of the total integration time related to each epoch f_j, we define the "form factor" as:

$$g(N_e, \mathbf{h}, \mathbf{f}) = \frac{1}{2} \left[\sum_j^{N_e} h_j^2 f_j - \left(\sum_j^{N_e} h_j f_j \right)^2 \right]^{-1/2}. \tag{7.2}$$

In our case, given the fact that 38% of the total integration time is collected at the first

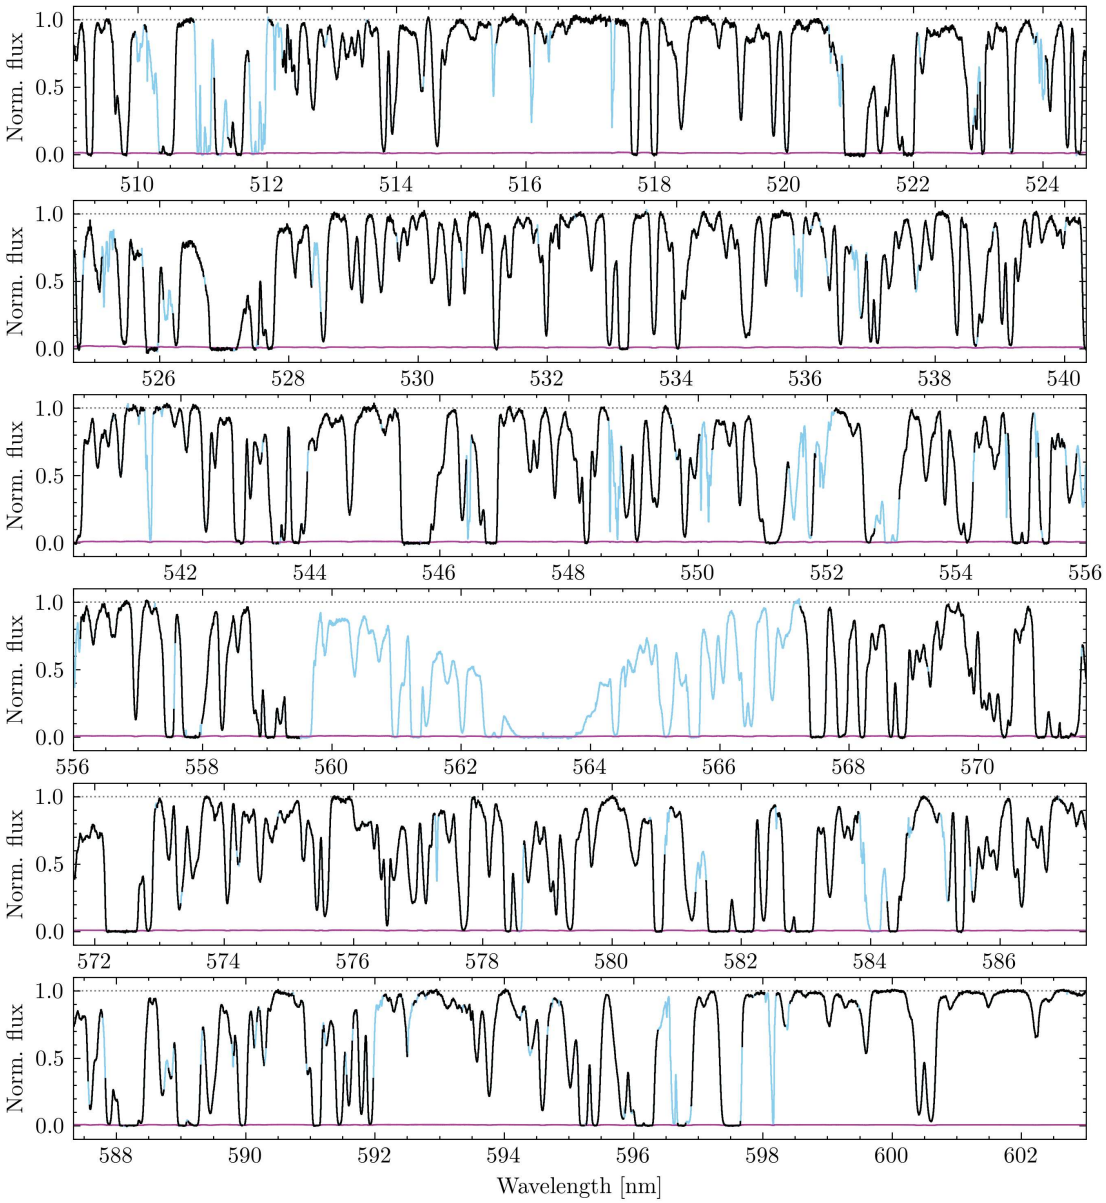

Figure 7.2: *Normalised transmitted flux of the Lyman-α forest in the combined spectrum of SB2 as a function of wavelength. The solid light blue line highlights the spectral regions masked due to metal absorption, the sub-DLA and bad pixels. The solid purple line reports the flux error.*

epoch, and 62% at the second, the form factor has a value of $g = 1.03$. From Eq. 7.1, we can estimate the expected maximum precision allowed by photon noise on our current data to be $\sigma_v = 5.99\,\mathrm{m\,s^{-1}}$ or, in terms of acceleration $\sigma_{\dot{v}} = \sigma_v/\Delta t = 6.89\,\mathrm{m\,s^{-1}\,yr^{-1}}$.

Note that in this work we have analysed only one sightline, whereas the relation from which Eq. 7.1 is derived has been calibrated on a measurement based on an ensemble of spectra. Sightline-to-sightline deviations from the predicted measurement precision due to cosmic variance are expected due to the varying number of Lyman-α lines and strong absorbers found in the spectra of different objects.

7.4 A model of the Lyman-α forest

To begin our measurement, we must first establish a clear methodological approach. In essence, our goal is to detect a systematic shift in spectral lines across a set of multiple spectra. This task is conceptually similar to the detection of exoplanets through Radial Velocity (RV) measurements of their host stars (e.g. Zechmeister et al., 2020; Silva et al., 2022; Figueira et al., 2025). For this reason, we drew inspiration from planetary searches when designing the procedures for carrying out the Sandage–Loeb Test.

From a technical standpoint, one can either compare spectra directly or adopt a spectral template as a fixed reference in velocity space. In RV studies, stellar templates are widely used because the absorption features of stars of a given spectral type can be modelled a priori. Quasar absorption spectra, however, are unique to each sightline, making the construction of a universal template unfeasible. Despite this, a model-based approach offers significant advantages: it reduces errors in line-position estimation, yields precise evaluations of flux derivatives, and is generally less sensitive to spectral defects than direct spectrum-to-spectrum comparison.

There are several ways to model the Lyman-α forest. The most common and physically motivated approach is Voigt profile fitting. Yet, in our case, we are not concerned with the detailed physical interpretation of Lyman-α transmission or the properties of the intervening gas. What we require is simply a function that reproduces the absorption profiles accurately, without overfitting the noise. A full physical description would not only be time-consuming but could also introduce systematic biases due to degeneracies and ambiguities in component placement.

At the same time, as emphasised earlier, it is crucial to disentangle Lyman-α lines from foreground metal absorptions. The latter, when identified, are modelled with Voigt profiles and masked. Following the same reasoning, we aim to develop a modelling procedure that captures the characteristic shape of Lyman-α lines while minimising the risk of inadvertently fitting unrecognised metal lines or noise.

On these grounds, we now introduce our model-construction strategy, which defines the reference framework for the redshift-drift measurement. The main steps of this procedure are:

- Production of realistic mock sightlines that reproduce the properties of the ESPRESSO

spectra of SB2 considered in the measurement.
- Calibration of the modelling procedure on the mock data.
- Construction of an ensemble of models of the Lyman-α forest of SB2.
- Evaluation of the velocity shift between the model ensemble and the first and second epoch data.

We stress that the mock data were only used to calibrate the modelling procedure and were not involved in the final measurement.

7.4.1 Mock spectra

We used realistic mock spectra that reproduce the properties of the three SB2 spectra (first epoch, second epoch, and combined) to calibrate the modelling pipeline and validate the measurement results. The procedure to produce realistic mocks follows the one used in the previous part of this thesis, in the context of Dark Photon searches in the IGM (see Sect. 4.2.2), and is here adapted to the case of ESPRESSO and the redshift drift experiment.

Summarising, starting from the hydrodynamical cosmological simulations of the IGM of the Sherwood Simulation Suite (Bolton et al., 2017b) (with a box size of 10 Mpc/h and 2×1024^3 dark matter and gas particles), we extracted the gas particle data and computed the H I optical depth along 5000 skewers piercing the box at every redshift step $\Delta z = 0.1$, within $3.1 \leq z \leq 4$. The optical depth of the spectra was rescaled to match the mean optical depth observed at $z \sim 3.57$ (centre of the forest) from Becker et al. (2013). The sightlines (10 Mpc/h long) are randomly selected from the appropriate redshift interval, shifted so that they begin and end with regions of no absorption, and stitched together to form 100 spectra that match the length of the Lyman-α of SB2 (~ 421 Mpc/h long, or $50321\ \mathrm{km\,s^{-1}}$).

The long spectra were then convolved with a Gaussian of FWHM $\sim 2.2\ \mathrm{km\,s^{-1}}$ to mimic ESPRESSO's resolution and re-binned to a grid of pixels $1\ \mathrm{km\,s^{-1}}$ wide (the average native pixel size extracted from the simulation was $0.593\ \mathrm{km\,s^{-1}}$), matching the binning of our data. The noise properties of the original spectrum were added to the mocks following the prescription of Rorai et al. (2017). Summarising, the original first epoch spectrum was divided into ten chunks, ~ 10 nm wide. For each of these chunks, we:

- divided the pixels into 50 bins according to their normalised flux value, producing 50 corresponding noise distributions;
- assigned a noise value to each simulated pixel, where the value was randomly sampled from the distribution associated with the pixel's flux value, as defined in the point above.

This procedure was repeated for each chunk of the first epoch spectrum and its corresponding mock spectra, to preserve the relation between S/N at the continuum level and wavelength.

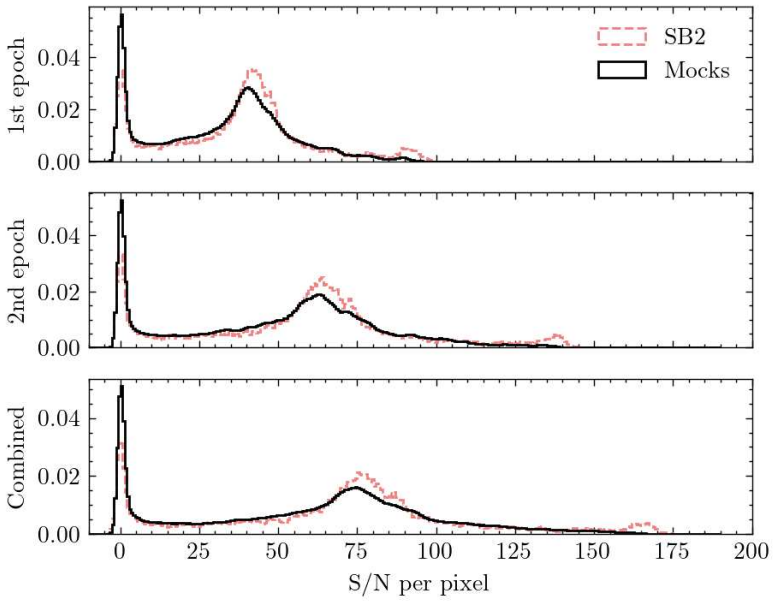

Figure 7.3: *Distribution of the S/N per pixel in the Lyman-α forest of the three spectra of SB2 (dashed red) and of the corresponding mock spectra (solid black). Upper panel: First epoch. Middle panel: Second epoch. Bottom panel: Combined spectrum.*

The second epoch spectra were generated in the same way from the same sightlines. However, before sightline extraction, we artificially added a velocity drift of $-0.38\,\mathrm{cm\,s^{-1}}$, i.e. the expected drift at $z \sim 3.57$ for a ΛCDM Universe during our baseline $\Delta t = 0.875\,\mathrm{yr}$, to the simulated particles' velocities. These sightlines were then processed in the same way, and the noise of the second epoch was added to the flux, as explained above.

We combined the first and second epochs' mock datasets via a weighted average, with weights proportional to the inverse of the pixels' flux variance, to create the mock combined spectrum. Fig. 7.3 shows the distribution of the S/N per pixel in the three mock datasets and the original spectra. To further validate the analysis (see Sect. 7.7), we also built second epoch mock datasets assuming either no drift between the two epochs or 10^3 and 10^4 years of temporal separation, adding a velocity drift to the simulation's particles of $-3.8\,\mathrm{m\,s^{-1}}$ and $-37.6\,\mathrm{m\,s^{-1}}$, respectively. On all spectra, we applied the same pixel masking of the original data, derived from metal, sub-DLA and bad pixels' removal, to ensure the absence of systematics in the final measurement due to pixel masking and to conserve the total amount of pixels involved in the measurement.

7.4.2 Model calibration

We relied on mock data to calibrate the modelling procedure before building the best possible model of SB2. Multiple sightlines were considered in the calibration to properly account for cosmic variance effects in the modelling procedure. Given the short timescales between the two epochs and that the total S/N is insufficient to detect the cosmic drift signal, we simplified the procedure by creating a fiducial model of the combined spectrum instead of modelling the two epoch spectra individually. A further simplification regards our decision to model the spectrum using cubic splines[5]

[5] Using higher order splines does not affect the results of the analysis.

instead of the more traditional approach using Voigt profile decomposition. This was done because we were not interested in measuring physical properties of the gas but rather differences between the SB2's spectra in the two epochs, for which the spline method was sufficient.

In quasar spectra, Lyman-α lines have a typical width of $\sim 30\ \mathrm{km\,s^{-1}}$ (Kim et al., 2002), while outlier features (e.g. cosmic rays) and metal transitions are typically narrower (down to a few kilometres per second). We calibrated the spline knot spacing on mock spectra containing only Lyman-α absorption, so that our procedure was 'informed' on the scales that we want to probe. Then, by generating an ensemble of models for SB2, we recovered a higher variance among models in the regions that might be affected by sharp non-Lyman-α absorption and could de-prioritise such features in the final measurement of Δv. The described procedure also prevents over-fitting of the noise.

In practice, let $S^k(\lambda|\mathbf{n}(A,\phi))$ be the spline model fitted on the k-th spectrum of the combined mock dataset described in the previous section, using the square inverse of the flux noise as weights in the fit, where $\mathbf{n}$ is a vector defining the positions of the spline knots along the spectrum. The knots are equispaced in velocity space, separated by a constant velocity, A, starting from an initial position (or phase) $0 < \phi < A$ with respect to the beginning of the Lyman-α forest.

Goodness of fit for the model was determined from its residuals to the data, $R^k(\lambda|A,\phi) = \left(F - S^k(\lambda|\mathbf{n}(A,\phi))\right)/\sigma_F$. We considered a model to be good if the residuals are centred at zero with a variance equal to unity, $R^k \sim \mathcal{N}(0,1)$.

One hundred mock spectra were fitted in this way, varying both the inter-nodal distance A and the initial phase ϕ, and their residuals were calculated as above. For each value of A and ϕ, we calculated the mean and the variance of the residuals for all 100 mock spectra. The optimal value of A was determined by examining which combination of parameters results in residuals most closely matching values drawn from $\mathcal{N}(0,1)$. Mathematically, we searched for a value of A that minimises:

$$\varepsilon(A) = \frac{\left(\sigma_{\mathrm{R}}(A) - 1\right)^2}{\sigma^2_{\sigma_{\mathrm{R}}}(A)}, \tag{7.3}$$

where $\sigma_R(A)$ is the standard deviation of the fit residuals of splines with internodal distance A, averaged over all spectra and all considered initial node phases. $\sigma^2_{\sigma_{\mathrm{R}}}(A)$ is the variance of the quantity $\sigma_R(A)$ due to sightline variance and variations in the initial node phase. Fig. 7.4 shows $\sigma_R - 1$ and its variance as a function of internodal distance A. Minimising Eq. 7.3 yields a best inter-nodal distance of $A = 8.37^{+0.26}_{-0.24}\ \mathrm{km\,s^{-1}}$. Note that this value of A is valid only for this precise S/N level, resolution, pixel size, and redshift. For any change in the spectral properties, A should be re-evaluated.

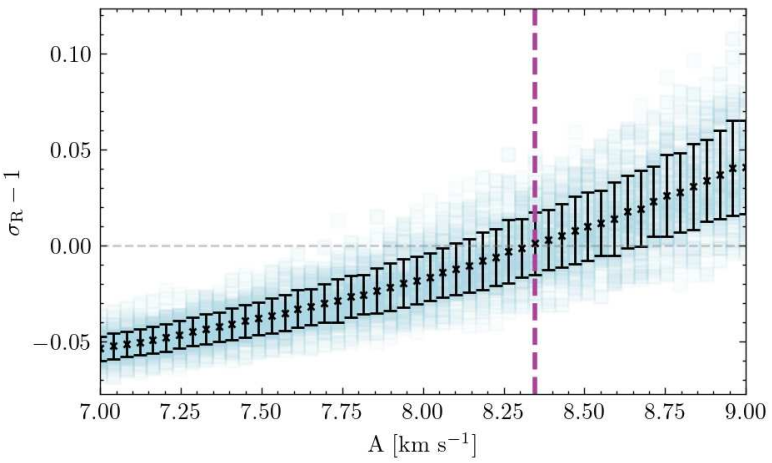

Figure 7.4: *Standard deviation of the fit residuals between the combined mock spectra's flux and the spline model with internodal distance A. The blue squares describe the σ_R of models on different mock spectra and randomly drawn initial node phase ϕ. The black error bars define the average and standard deviation of $\sigma_R(A) - 1$. The vertical dashed purple line denotes the optimal value of A.*

7.4.3 Building the forest model

Given the optimal knot spacing and its uncertainty, A and σ_A, we were able to construct a large number of models of the Lyman-α forest of SB2 that describe the data equally well in terms of a goodness-of-fit statistic (e.g. χ^2). This allowed us to more thoroughly assess the impact of model non-uniqueness on the final result. This is conceptually similar to the approach taken by, for example, Lee et al. (2021), Milaković et al. (2024b), and Webb et al. (2025) in the context of fundamental constant and isotopic ratio measurements. Additionally, in our quest to study systematic effects, we exploit the model non-uniqueness to understand whether the model-making procedure brings significant systematic uncertainties and biases in the measurement.

We started by fitting the combined SB2 spectrum with $N = 500$ splines, $S_j(\lambda \,|\, \mathbf{n}_j)$. The node positions of the j-th model, $\mathbf{n}_j$, were generated starting from an equispaced sequence of nodes in velocity space, $\mathbf{n}_{j,0}$, with inter-nodal separation, A_j, drawn from a normal distribution, $\mathcal{N}(A, \sigma_A)$, and initial phase $0 \leq \phi_j < A_j$, drawn with uniform probability. The node positions were then perturbed by adding random noise in velocity space as $\mathbf{n}_j = \mathbf{n}_{j,0} + \mathbf{m}$, where each element of the $\mathbf{m}$ array was randomly drawn from a normal distribution, $\mathcal{N}(0, \sigma_A)$.

From the ensemble of spline models, we evaluated the mean model, $\bar{S}$, and the variance between the models, σ_S^2.

Fig. 7.5 shows a region of the first, second, and combined spectra, alongside the mean spline model. In the figure, flux noise and model variance are shown in the bottom panel. The periodic pattern found in the flux noise results from the coaddition of several exposures and the subsequent rebinning to a fixed grid, which creates an aliasing effect in the error arrays, whose frequency depends on the position of the pixels within the spectral order.

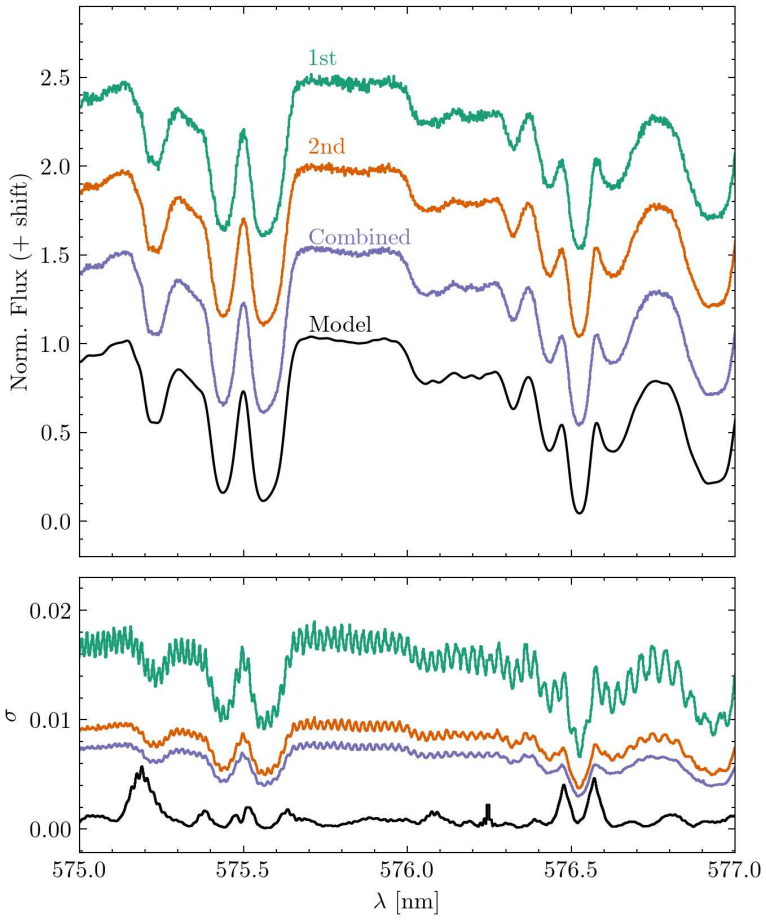

Figure 7.5: *Section of the SB2 spectrum. Top panel: Normalised flux of the first epoch (green), second epoch (orange), and combined (purple) spectra. The mean model, $\bar{S}$, is shown as a solid black line. The spectra are shifted vertically for clarity. Bottom panel: Normalised flux errors, σ_F, of the three spectra shown in the panel above, with the same colour coding. The solid black line is the standard deviation of our model, $\sigma_{\bar{S}}$, i.e. it shows where our ensemble of models exhibits large variation, and hence uncertainty.*

7.5 Redshift drift measurement

7.5.1 Pixel-by-pixel comparison

To measure the redshift drift that occurred between the two observational epochs, we built on the techniques developed for measuring radial velocity shifts between repeated observations of stars for exoplanet search (Bouchy et al., 2001), which we refer to as the pixel-by-pixel method. This method was adapted to the redshift drift measurement and used in Liske et al. (2008) to derive the scaling relation reported in Eq. 7.1. However, differently from Liske et al. (2008), instead of directly comparing the two epoch spectra, we exploit the model of the Lyman-α forest transmission to evaluate the flux derivatives and compare the spectra to the same common model, similarly to what is done with stellar RV with stellar spectral templates.

Summarising, the normalised flux at the i-th pixel of the k-th epoch spectrum, $F_{k,i}$, can be defined as a small perturbation of the model of the Lyman-α forest transmission $\bar{S}$, due to a small velocity shift, $\delta v_{k,i}$, as:

$$F_{k,i} = \bar{S}(\lambda_i) - \frac{\partial \bar{S}}{\partial \lambda_i} \frac{\delta v_{k,i}}{c} \lambda_i, \tag{7.4}$$

where each pixel of the spectrum defines a contribution to the velocity shift of the k-

epoch

$$\delta v_{k,i} = \frac{\bar{S}(\lambda_i) - F_{k,i}}{\partial \bar{S}_i / \partial \lambda_i} \frac{c}{\lambda_i}. \tag{7.5}$$

Note that the above equations hold only for δv_i much smaller than the typical line width ($\sim 30\,\mathrm{km\,s^{-1}}$). This assumption is fulfilled in the case of the redshift drift, as the expected drift is of the order of a few centimetres per second per decade.

The velocity shift occurring between the model and the spectrum can then be estimated as a weighted average of the contributions of all pixels

$$\delta v_k = \frac{\sum_i \delta v_{k,i} w_i}{\sum_i w_i}, \tag{7.6}$$

where the weights are defined as the inverse of the variance of the single pixel velocity contribution

$$w_i = \sigma_{v_{k,i}}^{-2} = \left[\frac{c}{\lambda_i \left(\partial \bar{S}_i / \partial \lambda_i \right)} \sigma_{F_{k,i}} \right]^{-2}, \tag{7.7}$$

where $\sigma_{F_{k,i}}$ is the error on the normalised flux at pixel i and epoch k. Note that, differently from a spectrum-to-spectrum comparison, we compare data to a model, thus only the flux error appears in Eq. 7.7. The uncertainty on the final δv_k estimate is

$$\sigma_{v_k} = \left[\sum_i \sigma_{v_{k,i}}^{-2} \right]^{-1/2}. \tag{7.8}$$

We apply this approach, comparing the Lyman-α transmission model to normalised fluxes in the Lyman-α forest of the two spectra, where we have masked metal lines, the strong H I absorber at $z \sim 3.63$, and bad pixels. This analysis yields the velocity shift between the model and the first epoch spectrum $\delta v_1 = 0.99 \pm 6.45\,\mathrm{m\,s^{-1}}$, and between the model and the second epoch spectrum $\delta v_2 = -0.55 \pm 4.12\,\mathrm{m\,s^{-1}}$. The shift occurring between the two epochs, Δv, was measured indirectly as

$$\Delta v = \delta v_2 - \delta v_1 = 0.45 \pm 7.66\,\mathrm{m\,s^{-1}}. \tag{7.9}$$

Taking into account the temporal baseline between the two epochs, $\Delta t = 0.867\,\mathrm{yr}$, we estimated a cosmic acceleration of

$$\dot{v} \simeq \frac{\Delta v}{\Delta t} = 0.51 \pm 8.80\,\mathrm{m\,s^{-1}\,yr^{-1}}, \tag{7.10}$$

or, equivalently, in units of redshift

$$\dot{z} = (0.08 \pm 1.34) \times 10^{-7}\,\mathrm{yr^{-1}}. \tag{7.11}$$

7.5.2 Model likelihood correlation

Another way to tackle the measurement is by maximising the likelihood between model and data, as a function of a velocity shift between the two. Strictly speaking, $|\dot{z}|$ is larger at the red end of the Lyman-α forest than at its blue end, causing a redshift-dependent compression of features in our spectrum. However, this effect is tiny and completely negligible at the S/N of our data. We therefore ignored the redshift dependence of $\dot{z}$ and assumed that only a rigid translation Δv, in velocity space, occurred between the two epochs.

For each epoch k, we determined the relative velocity shift δv_k between the model and the observed spectrum by maximising the likelihood function

$$\ln \mathcal{L}(\delta v_k) = -\frac{1}{2} \sum_{i=1}^{N_{\mathrm{pix}}} \frac{\left[F_{k,i} - \bar{S}(\lambda_i, \delta v_k)\right]^2}{\sigma_{F_{k,i}}^2}, \tag{7.12}$$

where $F_{k,i}$ and $\sigma_{F_{k,i}}$ denote the normalised flux and its associated uncertainty at pixel i of the k-th epoch spectrum. The model $\bar{S}(\lambda_i, \delta v_k)$ represents the mean Lyman-α forest transmission as constructed in Sect. 7.4.3, rigidly shifted in velocity space by δv_k and evaluated on the same wavelength grid $\{\lambda_i\}$ as the data. The summation in Eq. 7.12 is performed over all non-masked pixels within the Lyman-α forest region.

We evaluated Eq. 7.12 over a dense grid of velocity shifts in the range $-50\ \mathrm{m\,s^{-1}} \leq \delta v_k \leq 50\ \mathrm{m\,s^{-1}}$ and fitted a parabola to the resulting $\ln \mathcal{L}(\delta v_k)$ profile to obtain a precise estimate of its maximum. The statistical uncertainty on δv_k was derived from the curvature of the log-likelihood at the maximum,

$$\sigma_{\delta v_k} \approx \left[-\left. \frac{d^2 \ln \mathcal{L}}{d(\delta v_k)^2} \right|_{\delta v_k, \max} \right]^{-1/2}, \tag{7.13}$$

following standard maximum likelihood theory.

The analysis yields the relative shift between the model and the first epoch spectrum, $\delta v_1 = -0.78 \pm 6.55\ \mathrm{m\,s^{-1}}$, and between the model and the second epoch spectrum $\delta v_2 = -0.29 \pm 6.19\ \mathrm{m\,s^{-1}}$, leading to an estimate of the velocity drift between the two datasets

$$\Delta v = \delta v_2 - \delta v_1 = 0.49 \pm 7.78\ \mathrm{m\,s^{-1}}, \tag{7.14}$$

and therefore a cosmic acceleration measurement of

$$\dot{v} = 0.57 \pm 8.94\ \mathrm{m\,s^{-1}\,yr^{-1}} \quad \leftrightarrow \quad \dot{z} = (0.09 \pm 1.36) \times 10^{-7}\ \mathrm{yr^{-1}}. \tag{7.15}$$

Both measurement methods give results that are consistent with the expected cosmological signal ($-0.43\ \mathrm{cm\,s^{-1}\,yr^{-1}}$) and are consistent with each other. Also, the measurement uncertainties estimated by the two independent methods are coherent with each other, with the likelihood approach yielding a measurement only $\sim 2\%$ less precise than the pixel-by-pixel method.

One might be puzzled by the fact that, for both methods, both estimates of δv_i have negative values. While this may first seem surprising, there is an easily understandable explanation, and a quick test can be carried out to confirm the results. Both methods compare a spectrum to the spline model of the Lyman-α forest, where the latter has been built on the combination of the two epochs' spectra. The same techniques used to evaluate δv_i can be applied to the combined spectrum to test the goodness of the model and its proper description of the spectrum.

Therefore, with both methods, we computed δv_0 by comparing the model to the combined spectrum, yielding

$$\begin{aligned} \delta v_0 &= -0.59 \pm 3.48 \,\mathrm{m\,s^{-1}} \quad (\mathrm{Pixel-by-pixel}) \\ \delta v_0 &= -0.57 \pm 3.54 \,\mathrm{m\,s^{-1}} \quad (\mathrm{Likelihood}) \end{aligned} \tag{7.16}$$

As expected, since the model is meant to faithfully represent the combined spectrum, both measurements of δv_0 are consistent with a null drift. Moreover, note that both values fall between the corresponding δv_i estimated from the two epochs, where, since the second epoch spectrum has a higher S/N, the value of δv_0 is closer to δv_2 than to δv_1, since the second epoch spectrum has a stronger impact in the combination than the first one. Lastly, the uncertainties of δv_1, δv_2 and δv_0 scale accordingly to the three spectra S/N, In that sense, the obtained values of δv_1 and δv_2 are not surprising.

7.6 A third epoch!

As shown in the previous analysis, the SL test requires time – a lot of it – and photons – even more of them. The natural extension of the former analysis is therefore to acquire new spectroscopic data of the same source, extending the temporal baseline, and subsequently reducing the measurement uncertainties. This approach, even without the aim of a significant detection of the cosmological signal, has multiple advantages with respect to an experiment in which half of the data are taken at the beginning and half a few decades later. In particular, with continuous monitoring, we are capable of properly tracking the status of the instrument throughout the experiment, the possible variability of the quasar and its absorption lines, and, most importantly, fine-tune and thoroughly validate the analysis pipeline at different levels of noise, slowly uncovering the various systematic effects that might spoil our measurement, such as wavelength calibration errors. On the other hand, a continuous campaign is statistically less optimal with respect to a half-and-half campaign, where the former has an uncertainty form factor $g = 1.7$, whilst $g = 1$ for the latter (see Eq. 7.1). We nevertheless do not consider a half-and-half campaign since this would require an unrealistic allocation of a massive amount of telescope time ($10^3 - 10^4$ hours, depending on total baseline and instrument) in only two concentrated epochs, and would not enable a rigorous control of systematics of instrumental nature and the swift development of methods to tackle such effects.

To this end, in December 2024, we carried out 9.5 additional hours of ESPRESSO integration of SB2, with the same instrumental setup adopted before, defining the third observational epoch of the SL test on this target. Table 7.3 reports the details of the new observations.

Table 7.3: *Summary of single OBs of SB2 taken during observing period P114, as the third epoch of the ESPRESSO Redshift Drift Experiment.*

Exposure Start (UTC)	MJD	t_{exp} [s]	Grade	Seeing [arcsec]	Airmass	S/N
2024-12-26 01:05:00	60670.04514905	3453.0	A	0.54	1.215	21.4
2024-12-26 02:11:26	60670.09127605	3453.0	A	0.53	1.101	22.6
2024-12-26 03:18:08	60670.13759482	3453.0	A	0.54	1.065	21.8
2024-12-26 04:22:38	60670.18239184	3453.0	A	0.38	1.091	23.1
2024-12-27 01:10:19	60671.04883566	3453.0	A	0.73	1.194	18.8
2024-12-27 03:08:49	60671.13113179	3453.0	B	0.86	1.065	17.3
2024-12-29 02:59:01	60673.12431786	3453.0	B	0.70	1.065	16.0
2024-12-29 05:03:38	60673.21086665	3453.0	A	0.44	1.163	22.4
2024-12-30 01:29:23	60674.06208193	3453.0	A	0.48	1.137	22.9
2024-12-30 02:31:39	60674.10531465	3453.0	A	0.53	1.073	26.7

The data reduction and treatment of the new observations (see Sect. 7.2) follow the same procedures shown previously for the two-epoch experiment. The data collected in period P114 are combined in the third epoch spectrum, with median $\mathrm{S/N} \sim 75$ per $1\,\mathrm{km\,s^{-1}}$ pixel at continuum in the Lyman-α forest. All observations (about 22 hours in total) are combined into the updated combined spectrum, with $\mathrm{S/N} \sim 113$. With the new observations, the total temporal baseline of the experiment is $\Delta t = 1.958\,\mathrm{yr}$.

Following the analysis procedure presented in the sections above, we re-computed the model on the updated combined spectrum. Since the S/N has notably increased (by about 30% from the 2-epoch case), before computing the ensemble of models, we produced a new set of mocks of the Lyman-α forest resembling the three spectra in terms of noise and temporal separation, following Sect. 7.4.2 and computed the best internodal distance for the modelling procedure at this level of S/N, $A = 7.34^{+0.33}_{-0.13}\,\mathrm{km\,s^{-1}}$ (see Sect. 7.4.2).

Precision prediction

For reference, and to check the consistency of our new results, we evaluated the expected precision of a measurement based on the three epochs collected. For the latest sequence of observations, considering all three epochs (reported in Tables 7.1 and7.3), the current form factor is $g = 1.27$ (Eq. 7.2). Evaluating Eq. 7.1 with the new details, we estimated to reach a measurement precision of $\sigma_v = 5.7\,\mathrm{m\,s^{-1}}$ in velocity space. Accounting for the temporal baseline of the current experiment, the measurement precision for the cosmological acceleration is $\sigma_{\dot{v}} = \sigma_v/\Delta t = 2.9\,\mathrm{m\,s^{-1}\,yr^{-1}}$, equivalent to $\sigma_{\dot{z}} = 4.4 \times 10^{-8}\,\mathrm{yr^{-1}}$.

The new measurement

Following the procedure of Sect. 7.5.1, we measured the velocity shift between the new model, built on the updated combined spectrum, and the three single-epochs' spectra, adopting the Pixel-by-pixel approach, yielding

$$\begin{aligned} \delta v_1 &= -1.19 \pm 6.48\ \mathrm{m\,s^{-1}}, \\ \delta v_2 &= \ \ 2.58 \pm 4.14\ \mathrm{m\,s^{-1}}, \\ \delta v_3 &= -3.29 \pm 3.93\ \mathrm{m\,s^{-1}}. \end{aligned} \tag{7.17}$$

Similarly, we measure the following single epoch shifts with the Likelihood correlation of the model (see Sect. 7.5.2)

$$\begin{aligned} \delta v_1 &= -0.43 \pm 6.59\ \mathrm{m\,s^{-1}}, \\ \delta v_2 &= \ \ 3.06 \pm 4.21\ \mathrm{m\,s^{-1}}, \\ \delta v_3 &= -3.82 \pm 4.00\ \mathrm{m\,s^{-1}}. \end{aligned} \tag{7.18}$$

Note that the new single δv_k values estimated here differ from those reported in Sects. 7.5.1 - 7.5.2, since the model has been updated to the latest version of the combined spectrum. For reference, the autocorrelation of the new model onto the final 3-epochs combined spectrum yields $\delta v_0 = -0.74 \pm 2.61\ \mathrm{m\,s^{-1}}$ (Pixel-by-pixel) and $\delta v_0 = -0.74 \pm 2.65\ \mathrm{m\,s^{-1}}$ (Likelihood), with proportionally smaller uncertainties with respect to the two epoch case (See Eq. 7.16).

Differently from the 2-epoch case, with the current three data points, we estimated the cosmic acceleration by means of a simple linear fit of the velocity shifts between the model and the three epochs' spectra, δv_k, as a function of time, or rather, as a function of the time difference between the first set of observations, as shown in Figure 7.6. We fit the linear function

$$\delta v = \dot{v}\left(\Delta t - \overline{\Delta t}\right) - q, \tag{7.19}$$

where $\overline{\Delta t} = 1.207\ \mathrm{yr}$ is the weighted average Δt of the three epochs. This pivoted linear relation was chosen to minimise correlation between the fitted parameters $\{\dot{v}, q\}$, as shown in Fig. 7.7.

Fitting the δv estimates computed through the two methods yields

$$\begin{aligned} &\dot{v} = -2.28 \pm 3.58\ \mathrm{m\,s^{-1}\,yr^{-1}}; \quad \dot{z} = (-3.46 \pm 5.46) \times 10^{-8}\ \mathrm{yr^{-1}} \quad (\mathrm{Pixel-by-pixel}), \\ &\dot{v} = -2.84 \pm 3.64\ \mathrm{m\,s^{-1}\,yr^{-1}}; \quad \dot{z} = (-4.34 \pm 5.55) \times 10^{-8}\ \mathrm{yr^{-1}} \quad (\mathrm{Likelihood}). \end{aligned} \tag{7.20}$$

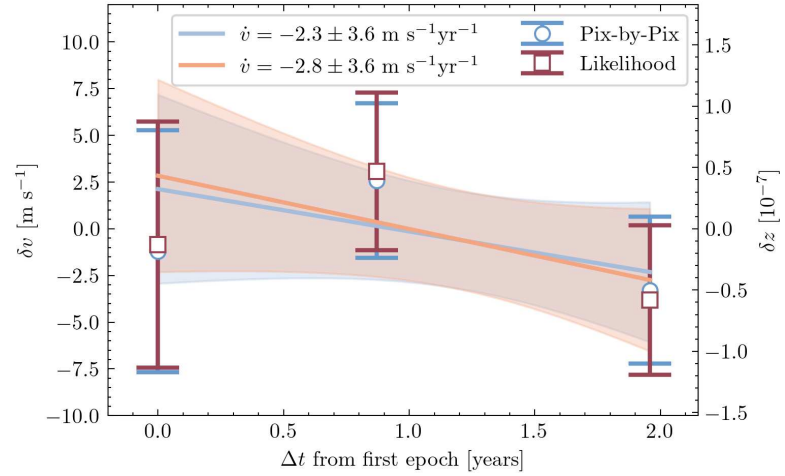

Figure 7.6: *Velocity shifts of the three epochs spectra with respect to the combined mean model, as a function of elapsed time from the start of observations, obtained via the Pixel-by-Pixel (blue circles) and the Likelihood Correlation approaches (red squares). For both approaches, the cosmic acceleration is computed as a linear fit of the data points and shown as a solid line. Shaded areas report the uncertainty in the fit.*

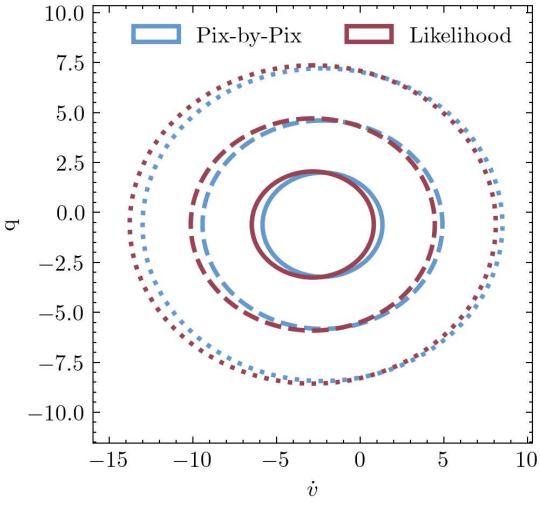

Figure 7.7: *Confidence ellipses of the fit parameters $\{\dot{v}, q\}$ estimated fitting the velocity shifts computed through the Pixel-by-pixel (blue) and Likelihood (red) methods. Solid, dashed and dotted lines define 1σ, 2σ and 3σ contours, respectively.*

7.7 Validation and systematics

7.7.1 Measurement validation

We validated the measurement procedures by applying the pixel-by-pixel (see Sect. 7.5.1) and the model-based (see Sect. 7.5.2) methods to a sample of synthetic spectra resembling the three SB2 spectra in noise level, resolution, pixel size, spectral range, and masked regions. Upon construction, the synthetic spectra of the second and third epoch have been shifted by -0.38 cm s^{-1} and -0.84 cm s^{-1}, respectively, to mimic the expected velocity shift occurring in the respective baselines in a ΛCDM universe. For each sightline, a model was computed using the combination of the three epoch spectra and used as a reference to compute $\dot{v}$ using both methods described above. The results are shown in Fig. 7.8, where the top panels relate to the Pixel-by-pixel analysis, whilst the bottom panels to the Likelihood correlation analysis. The right-hand panels report the distribution of measured $\dot{v}$ over all 100 sightlines, where the black solid line represents a Normal distribution centred at 0 and with the average uncertainty on $\dot{v}$, $\sigma_{\dot{v}}$, as standard deviation. The bright green dashed line reports a Normal distribution

with the RMS among the data points as the standard deviation. The two profiles are remarkably similar, showing how both methods properly estimate the uncertainty on $\dot{v}$, as this faithfully reflects the statistical scatter of the data.

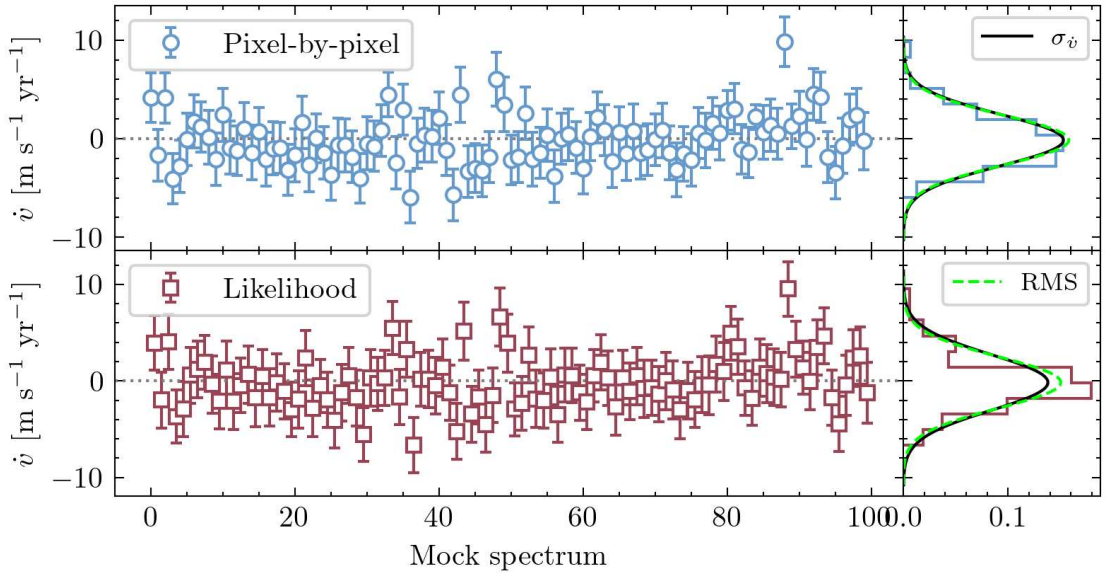

Figure 7.8: *Cosmic acceleration, $\dot{v}$, measured on one hundred random mock sightlines, each with 3 epochs, with both the pixel-by-pixel (top left panel, blue circle scatter points, see Sect. 7.5.1) and Likelihood correlation methods (bottom left panel, red squares scatter points, see Sect. 7.5.2). The right panels report the distribution of the 100 $\dot{v}$ measurements (histogram), with the Normal distribution defined by their average error $\sigma_{\dot{v}}$ (solid black) and the Normal distribution defined by the RMS of the data points (dashed bright green).*

7.7.2 Uncertainty scaling

Having executed the experiment both with two and three epochs separately, we can investigate whether the resulting measurement precision scales as expected. Following Eq. 7.1, the uncertainty in the measurement of the cosmic accelerations should scale as

$$\sigma_{\dot{v}} \propto g\, \mathrm{S/N}^{-1}\, \Delta t^{-1}, \tag{7.21}$$

where S/N refers to the total signal-to-noise. Scaling the results reported for the two-epoch experiment, and substituting the S/N, temporal baseline and form factor of the three-year experiment, we get

$$\begin{aligned} &\sigma_{\dot{v},\mathrm{expected}} = 3.74\ \mathrm{m\,s^{-1}\,yr^{-1}}; \quad \sigma_{\dot{v},\mathrm{measured}} = 3.58\ \mathrm{m\,s^{-1}\,yr^{-1}}, \quad (\mathrm{Pixel-by-pixel}),\\ &\sigma_{\dot{v},\mathrm{expected}} = 3.81\ \mathrm{m\,s^{-1}\,yr^{-1}}; \quad \sigma_{\dot{v},\mathrm{measured}} = 3.64\ \mathrm{m\,s^{-1}\,yr^{-1}}. \quad (\mathrm{Likelihood}). \end{aligned} \tag{7.22}$$

Both methods show results that properly scale as expected with S/N, temporal baseline and distribution of the observations, although improving the precision level by $\sim 5\%$ with respect to expectations.

On the other hand, we compared the results from our analysis with the scaling relation of Liske et al. (2008) (Eq. 7.1) and found an excess in uncertainty on the real

data of $\sim 20-25\%$ with respect to the theoretical expectations. This discrepancy is not surprising nor worrisome when we recall the fact that the relation reported in Eq. 7.1 is calibrated on an ensemble measurement, performed on several synthetic spectra. Our measurement, instead, not only relies on real data, prone to errors and pixel defects that might enhance uncertainties, but is also based on a single sightline. As shown in Liske et al. (2008), sightline-to-sightline differences in measurement precision are expected up to $\sim 50\%$ mainly due to Lyman-α lines clustering and cosmic variance, thus the small deviation that we detect is not unexpected nor unexplained.

7.7.3 Uncertainty across the spectrum

It is possible that a narrow segment of the Lyman-α forest contains unidentified contaminants that spoil the measurement on $\dot{v}$. We investigated this possibility by dividing the Lyman-α forest into ten equally sized chunks (of width $9.4\ \mathrm{nm}$) and measuring $\dot{v}$ over each one, using both the pixel-by-pixel and the likelihood correlation methods, following the usual procedures described in Sects. 7.5.1 and 7.5.2.

Fig. 7.9 shows the measured $\dot{v}$ for each chunk, as a function of the chunk's central wavelength λ_c, in the upper panel. The middle panel shows the acceleration uncertainty, $\sigma_{\dot{v}}$, computed on each chunk with the pixel-by-pixel (blue) and likelihood (red) methods. In the absence of systematics, the measurement uncertainty is expected to depend solely on the median S/N at continuum and the number of pixels (or the fraction of the forest f_{Ly_α}) in each chunk, scaling as $\sigma_{\dot{v}} \propto \mathrm{S/N}^{-1} f_{\mathrm{Ly}_\alpha}^{-0.5}$. The expected measurement uncertainty is shown in the middle panel as a dash-dotted (dashed) grey line, scaled by the uncertainty obtained on the full spectrum. The estimated uncertainties closely follow the scaling expectations for both methods, highlighting the absence of strong wavelength-dependent systematic effects.

7.7.4 Influence of the quasar

Given the extraordinary nature of SB2 (Wolf et al., 2024), with its high black hole accretion rate ($\dot{M} \sim 370\ M_\odot \mathrm{yr}^{-1}$) and luminosity variability ($\sim 15\%$ over the last 6 years), one would expect to find a strong signature of IGM photoionisation in the Lyman-α forest close to the quasar emission redshift. Similarly, since QSOs are known to live in halos of high matter overdensity, the peculiar motions of the Lyman-α lines close to QSO emission might be stronger than in the general IGM. However, our data do not show an important decrease in strength and number of Lyman-α lines close to z_{em}, nor a significant variability of such lines, which could be related to the quasar's strong activity. Therefore, in our previous analysis, we based the redshift drift measurement on the whole forest, without excluding the proximity region, about $5000\ \mathrm{km\,s^{-1}}$ from the quasar emission redshift.

We checked for systematic effects due to the proximity to the quasar by performing the same measurement as before, but excluding such a region and measuring a velocity drift between the two epochs of $\dot{v} = -3.40 \pm 3.91\ \mathrm{m\,s^{-1}\,yr^{-1}}$. Having excluded the

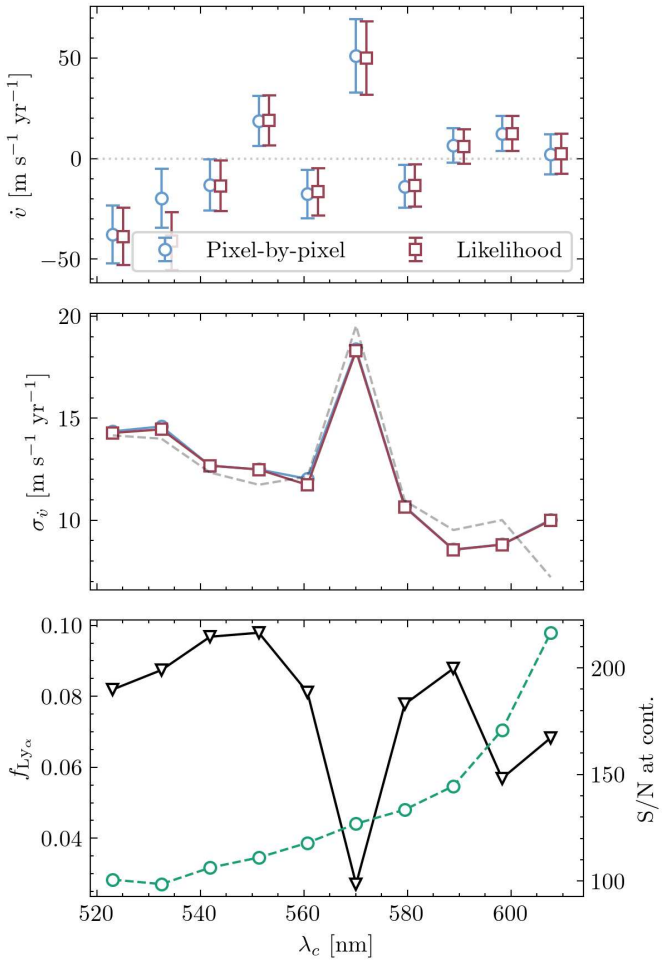

Figure 7.9: *Estimated cosmic acceleration, $\dot{v}$, computed on ten equispaced sections of the Lyman-α forest. Top panel: $\dot{v}$ estimated on each spectral chunk with the pixel-by-pixel (blue circles) and likelihood correlation (red squares) methods, as a function of the section's central wavelength λ_c. Middle panel: Measurement uncertainty, $\sigma_{\dot{v}}$, for the two methods, shown with the same colour coding as the panel above. The dashed grey line reports the expected uncertainty, based on the median S/N and pixel number of each chunk. Bottom panel: Fraction of the Lyman-α forest f_{Ly_α} considered in the measurement within each chunk (solid black line, left-hand axis), taking into account the masking of metals, strong absorbers and bad pixels. Note that the spectral chunk centred at 560.7 nm contains the sub-DLA absorber and has a significantly smaller fraction of usable pixels than the other regions. The right-hand axis of the same panel reports the median S/N at continuum level in each chunk in the combined spectrum (green dashed line).*

proximity region reduces the number of pixels used in the measurement to ∼ 35600, or ∼ 70% of the total amount found in the forest. When ignoring the proximity region, we excluded the range close to the Lyman-α emission where the spectra have the highest S/N levels at the continuum, with a median of ∼ 217 per pixel in the combined spectrum, and performed the measurement on spectra with effectively lower S/N. Therefore, the uncertainty grows not only due to a reduced pixel sample, but also due to a smaller median S/N level (∼ 109 per pixel at the continuum in the combined spectrum) outside of the proximity region. These two effects add up to a scaling factor $(109/113)^{-1}\,(0.70/0.77)^{-1/2} \sim 1.08$, explaining the new uncertainty. Still, the measurement is compatible with the expected value, and no systematic effect due to excluding or taking into account the proximity region is clearly visible.

7.7.5 Instrumental systematics

ESPRESSO's wavelength calibration is a crucial aspect for robustly measuring the redshift drift signal. Previous studies identified differences as large as 20 $\mathrm{m\,s^{-1}}$ between the two calibration methods available on ESPRESSO, the LFC and the Fabry-Pérot (FP) etalon combined with a ThAr lamp (Schmidt et al., 2021). Their impact is expected to decrease by considering a large number of lines over a broad spectral range,

but the true quantification on a real measurement has so far lacked. We therefore repeated our analysis on SB2 observations calibrated using ThAr+FP frames instead of LFC, created specifically for this purpose, following the same procedures used previously in Sect. 7.2. Adopting the Likelihood correlation method, we measured $\dot{v} = 0.79 \pm 3.65\ \mathrm{m\,s^{-1}\,yr^{-1}}$ from these data, a change of $+3.63\ \mathrm{m\,s^{-1}\,yr^{-1}}$ with respect to LFC-calibrated data. The two measurements agree within uncertainties.

Another source of uncertainty relates to assumptions on the shape of the line-spread function (LSF), used for determining line centres during wavelength calibration. The standard ESPRESSO DRS wavelength calibration procedures assume a Gaussian LSF shape and, because this cannot be modified without significant changes to the entire procedure, we were unable to directly quantify the impact of this assumption on our results. However, considering that using a Gaussian LSF should introduce wavelength calibration errors similar to the differences observed between using LFC and ThAr+FP (i.e. $20\ \mathrm{m\,s^{-1}}$, Schmidt et al. (2021) and Schmidt (2024)), a reasonable expected additional systematic uncertainty is $\sim 0.5\ \mathrm{m\,s^{-1}}$ through analogy with the previous paragraph.

7.7.6 Model non-uniqueness

One of the main interests at this stage of the experiment is the effective and accurate description of systematic effects that might hinder the final measurement of the redshift drift. Since the latter relies on the production of a transmission model of the Lyman-α forest, we ought to determine whether this procedure induces systematic uncertainties in the procedure.

In Sect. 7.4.3 we built a model of the Lyman-α transmission as the average of 500 independent realisations of spline models, calibrated on physical simulations of the Lyman-α forest at $z = 3 - 4$, that statistically fit the data equally well. From this ensemble, we can estimate the spread between independent models at each pixel, highlighting how absorption line modelling suffers from intrinsic model non-uniqueness. The model variance at each pixel can be taken into account in both measurement methods proposed beforehand, as an additional noise term, to investigate whether model non-uniqueness adds an impactful systematic uncertainty to the measurement.

In the Pixel-by-pixel method (see Sect. 7.5.1), the model variance term was accounted for by modifying the velocity contribution weights of each pixel (Eq. 7.7) to

$$w_i = \sigma_{v_{k,i}}^{-2} = \left[\frac{c}{\lambda_i \left(\partial \bar{S}_i / \partial \lambda_i\right)}\right]^{-2} \left[\sigma_{F_{k,i}}^2 + \sigma_{S_i}^2 + \frac{\left(\bar{S}(\lambda_i) - F_{k,i}\right)^2}{\left(\partial \bar{S}_i / \partial \lambda_i\right)^2} \sigma_{S_i'}^2\right]^{-1}, \tag{7.23}$$

where we propagated the model variance σ_S^2 and the model derivative variance $\sigma_{S'}^2$, to the velocity contribution variance.

Similarly, the same variance term can be taken into account when correlating the transmission model to the data via the likelihood maximisation, adding a noise term

at the denominator of Eq 7.12, as

$$\ln \mathcal{L}(\delta v_k) = -\frac{1}{2} \sum_{i=1}^{N_{\rm pix}} \frac{\left[F_{k,i} - \bar{S}(\lambda_i, \delta v_k)\right]^2}{\sigma^2_{F_{k,i}} + \sigma^2_{S_i}}. \tag{7.24}$$

We performed the measurement again on the three epochs, following the same procedures as reported above for both methods, adopting the model non-uniqueness noise terms and recovering $\dot{v} = -2.02 \pm 3.60\,\mathrm{m\,s^{-1}\,yr^{-1}}$ (pixel-by-pixel) and $\dot{v} = -2.84 \pm 3.72\,\mathrm{m\,s^{-1}\,yr^{-1}}$ (likelihood correlation) with a minimal increase of the uncertainty ($\lesssim 2\%$). The systematic uncertainty due to model non-uniqueness is strongly subdominant with respect to the statistical uncertainty and does not provide an important source of systematic uncertainty in the measurement. Moreover, since the transmission model is built on the combination of all exposures, in future developments of the measurement with higher S/N, the model variance is expected to decrease and remain subdominant to the photon noise, never providing a significant source of systematic uncertainty.

7.7.7 Single exposures analysis

In order to investigate systematic effects and spurious velocity contribution in the single ESPRESSO exposures, we carried out the measurement described in the previous sections, comparing the Lyman-α transmission model to the single rebinned exposures, instead of the three combined epoch spectra. As before, in both methods the model is compared to each spectrum to estimate δv_k, where in this case $k = \{1, \ldots, 23\}$ refers to the single exposure. Again, a linear fit is carried out in the $\delta v - \Delta t$ plane to estimate the cosmological acceleration. This analysis yielded $\dot{v} = -3.02 \pm 4.19\,\mathrm{m\,s^{-1}\,yr^{-1}}$ (pixel-by-pixel) and $\dot{v} = -3.55 \pm 4.19\,\mathrm{m\,s^{-1}\,yr^{-1}}$ (likelihood).

The results of the two approaches are compatible with each other and with the measurement carried out on the combined epoch spectra. The uncertainty on the final measurement is slightly larger ($\sim 20\%$) with respect to the three epochs measurements presented in Sects. 7.5.1-7.5.2. This increase is tied to the non-linear processes that take place during spectral combination (e.g. sigma-clipping) and increase the overall S/N level of the final spectrum while removing outlier pixels.

Fig. 7.10 shows the velocity shift between the Lyman-α transmission model and the normalised flux of the single exposures. The shaded areas represent the weighted average velocity shift across the single epochs, with their uncertainty. The median S/N per pixel at continuum level for each exposure is reported in the figure, as well as the measurement uncertainty obtained on each spectrum.

7.7.8 Influence of local motions

As is shown by Inoue et al. (2020), the mass of the Milky Way (MW) and the other galaxies in the Local Group (LG), especially the Large Magellanic Cloud and M31,

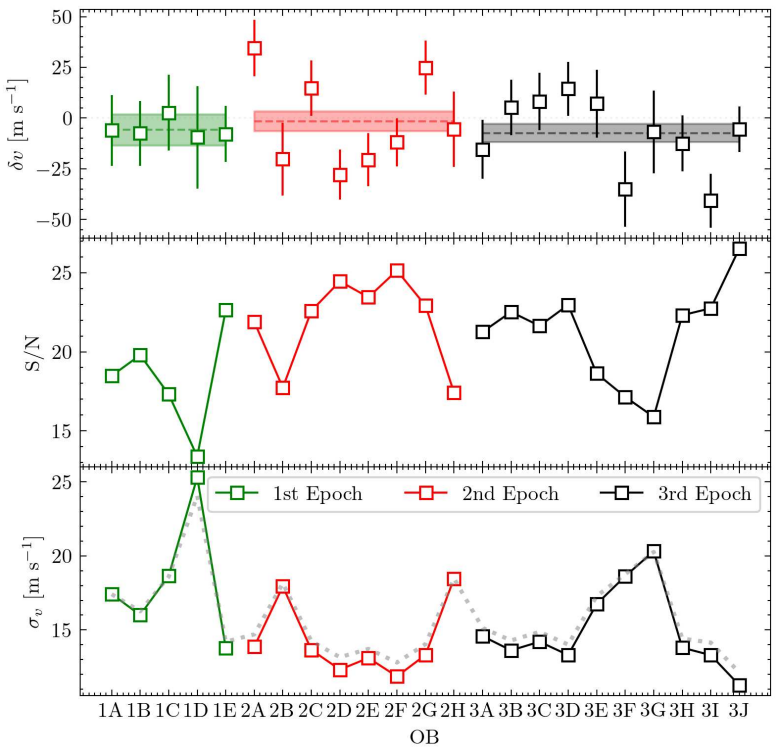

Figure 7.10: *Top panel: Velocity shift between transmission model and single exposure flux, subdivided in the three observational epochs: first (green), second (red) and third (black). The horizontal dashed lines report the average velocity shift for each epoch, weighted by single exposure velocity shift uncertainty, and their uncertainty as shaded areas. Middle panel: median S/N per 1 km s^{-1} pixel at continuum in the forest for each exposure. Bottom panel: velocity shift uncertainty for each exposure. The dotted grey line reports the expected precision computed as a simple S/N^{-1} scaling.*

have a non-negligible impact on the proper motion of the Sun, providing a local source of acceleration along a certain sightline. Left uncorrected, this effect will impact the cosmological redshift drift signal measurement. The amplitude of this effect is sightline dependent and, in the case of SB2, induces an additional systematic drift in the Lyman-α lines of $\dot{v}_\odot = -0.16 \pm 0.04$ cm s^{-1} yr^{-1}. This is of the same order of magnitude as the cosmological signal we want to measure. Thus, the expected velocity shift between two epochs of SB2 grows effectively to $\dot{v}_{\rm SB2} \sim -0.59$ cm s^{-1} yr^{-1}, implying it can be measured on a shorter temporal baseline, at fixed total integration time.

However, such a shift is the composite of the local and cosmological effects, where the former should be removed from the measured shift to infer the actual redshift drift. The required duration of the experiment is not affected by the local component, if not for the fact that its amplitude is inferred from measurements of the MW and LG's galaxies' masses, and has an uncertainty that will be propagated to the redshift drift measurement's uncertainty. Such an uncertainty is one order of magnitude smaller than the expected redshift drift, of the order of 0.04 cm s^{-1} yr^{-1}, where the cosmological signal is ~ -0.429 cm s^{-1} yr^{-1}, and its propagation in the actual measurement is sub-dominant with respect to the measurement uncertainty related to the spectral analysis.

7.8 Another Quasar – the case of J212540.97-171951.4

In parallel to the analysis of the brightest SB2, the group has begun observations and analysis of the quasar J212540.97-171951.4 ($z_{\rm em} = 3.897$), called SB1. Targeting two different quasars has the potential to validate the measurement procedures and unveil

hidden systematic effects due to, for example, cosmic variance or telescope pointing. Moreover, we took the opportunity to test and develop a different modelling approach on SB1 to strengthen the measurement with an independent validation.

The first observational epoch of SB1 has been carried out between July and August 2023, with the same ESPRESSO set-up used in the case of SB2. Unfortunately, we were not able to collect the data for the second epoch as ESPRESSO's LFC experienced technical problems with the intensity of the comb lines, resulting in a non-stable wavelength calibration. Since one of the main focuses of this preliminary ESPRESSO experiment is to study the systematic effects related to wavelength calibration, comparing LFC and ThAr calibrated spectra, the observations of SB1 have been put on hold and will resume once the LFC's performances enable an accurate calibration. Table 7.4 reports the details of the observations of the first observing epoch of SB1.

Table 7.4: *Summary of single Observing Blocks (OB) composing the first epoch of SB1, carried out during the P111 ESO period.*

Exposure Start (UTC)	MJD	t_{exp} [s]	Grade	Seeing [arcsec]	Airmass
2023-07-21 05:57:20	60146.24815289	3438.0	A	0.395	1.017
2023-07-21 07:01:15	60146.29253746	3438.0	A	0.420	1.069
2023-07-24 03:57:33	60149.16497110	3438.0	A	0.550	1.101
2023-07-25 04:48:03	60150.20004087	3438.0	A	0.450	1.031
2023-07-25 05:53:47	60150.24568803	3438.0	A	0.535	1.021
2023-07-25 06:58:29	60150.29062090	3438.0	A	0.450	1.089
2023-08-11 05:06:01	60167.21251990	3438.0	A	0.605	1.032
2023-08-15 04:32:51	60171.18947932	3438.0	A	0.485	1.021
2023-08-15 05:36:39	60171.23379563	3438.0	A	0.525	1.091

Nevertheless, although an estimate of the cosmic acceleration along SB1's line-of-sight is not possible with a single epoch, we can still investigate the properties of its spectrum, identify metal absorption lines and develop the required analytical procedures. Also, with just one epoch, we want to answer the question: Does SB1's Lyman-α forest provide for a better or worse tracer of the expansion, compared to SB2?

Data products of the observations were reduced and processed following the same procedures adopted for SB2 and described in Sect. 7.2. All exposures were combined in a first-epoch spectrum, with median S/N$\sim$ 43 per $1\ \mathrm{km\,s^{-1}}$ pixel at continuum level in the forest. Figure 7.11 shows the combined spectrum of all exposures and its S/N. Metal absorption lines outside and inside the Lyman-α forest region were recognised and masked, together with bad pixels and telluric lines. Finally, a fraction of the Lyman-α forest $f_{Ly\alpha} = 0.771$ was used for the subsequent analysis, coincidentally similar to the usable fraction of SB2's forest.

The analysis is similar to what was already presented in the case of SB2. We generated a model for the normalised SB1 spectrum and compared the model to the spectrum itself through the Likelihood Correlation method. Even though such a correlation does not aim to measure any velocity differences, it can be used to estimate the

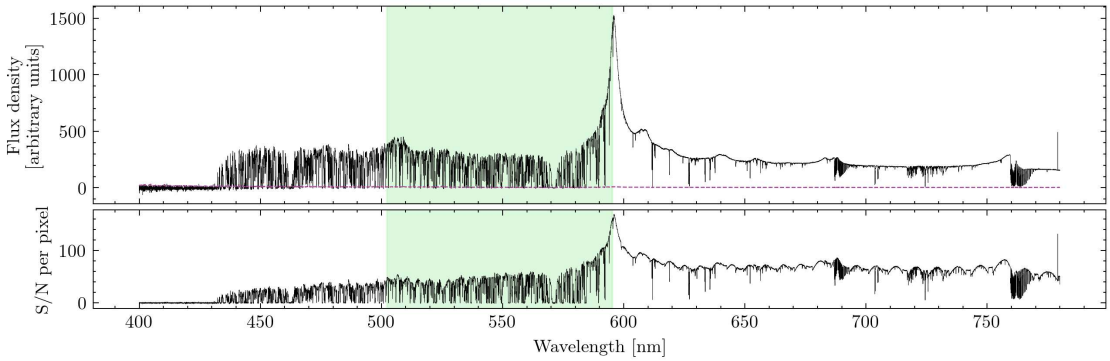

Figure 7.11: *Top panel: combined spectrum of the first epoch of observations of QSO J212540.96-171951.32. Flux density and flux density error are shown in black and purple, respectively. Bottom panel: S/N of the combined spectrum per* $1\ \mathrm{km\,s^{-1}}$ *pixel. The green shaded area encompasses the Lyman-α forest used for the analysis.*

precision of a future measurement that targets the first epoch. We call this precision *intrinsic velocity dispersion*, θ_v, where the name is meant to stress that we are measuring an intrinsic property of the spectrum that depends on its S/N and on the amount and disposition of usable lines. In the case of the first epoch spectrum of SB1, we estimate $\theta_v = 8.63\ \mathrm{m\,s^{-1}}$.

We can compare this estimate with the uncertainties in δv measured for SB2 to understand which of the two sightlines is "richer" in information and can better constrain the acceleration term. Fig. 7.12 reports the intrinsic velocity dispersion of SB1 compared to the measurement precision achieved on the three epochs of SB2, as a function of the spectra's S/N. Both sets of points are fitted with a simple power law[6] $\theta_v \propto \mathrm{S/N}^{-1}$, showing that, at given S/N, SB1's sightline provides a measurement 22% less precise than SB2's sightline. This difference is not unexpected, as the two sightlines pierce through different media and can have significant differences in the number and shape of the Lyman-α lines. Nevertheless, a 22% difference in precision is only a small factor compared to the fact that analysing two sightlines enables us to get two independent measurements on opposite sides of the sky, thus providing a joint constraint on the final estimation of $\dot{v}$, and significantly enhancing the detection of systematic effects related to each single source.

7.9 Discussion and Future Prospects

In the previous sections, we have started the effort towards the cosmic redshift drift measurement by taking limited amounts of data with ESPRESSO to properly study our techniques and the limiting factors that might hinder the experiment. We have shown that no significant systematic effect rises above the photon-noise level we currently reach, while the measurement precision follows expectations. Having done so, one

[6] A more detailed comparison should also take into account the scaling of the uncertainty on the fraction of analysed forest and the emission redshift of the QSO, $\theta_v \propto f_{Ly\alpha}^{-1/2}(1+z)^{-1.7}$, however the difference in the usable fractions in the spectra of SB1 and SB2 and their emission redshifts are negligible.

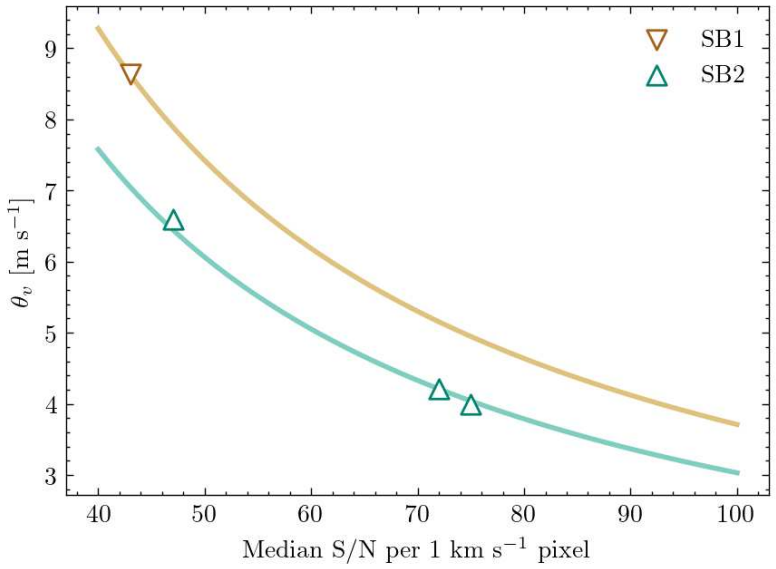

Figure 7.12: *Intrinsic velocity dispersion of SB1's first epoch spectrum (brown downward triangle) compared to the measurement precision achieved on SB2's three epochs (green upward triangle). Simple scaling relations are fitted to the two datasets ($\theta_v \propto \mathrm{S/N}^{-1}$) and shown with the same colour coding.*

crucial question remains unanswered: How long will it take to measure the signal? Here, we forecast the timeline of the experiment depending on the instrument used and the observational strategy adopted, on the basis of simple extrapolations of the results found so far.

Assume an observing strategy of N_e epochs evenly spread throughout the temporal baseline of the experiment, Δt, so that $\Delta t = (N_e - 1)\delta t$, where $\delta t = 1$ yr is the time interval between epochs, and consider carrying out T hours of observations of SB2 each epoch. In this framework, the uncertainty on a measurement of the cosmic acceleration, $\dot{v}$, scales as

$$\sigma_{\dot{v}}(\Delta t, T) \propto g(\Delta t)\, \mathrm{S/N_{tot}}(\Delta t, T)^{-1}\, \Delta t^{-1}, \tag{7.25}$$

where $\mathrm{S/N_{tot}}(\Delta t, T)$ is the total S/N achieved over all $N_e = \Delta t + 1$ epochs and scales as $\mathrm{S/N_{tot}}(\Delta t, T) \propto \sqrt{(\Delta t + 1)T}$. The form factor g follows from Eq. 7.2 and, in the case of the uniform observing strategy that we assume, takes the analytical form $g = \sqrt{3(N_e - 1)/(N_e + 1)} \approx 1.7$ for $N_e \gg 1$. Fig. 7.13 shows the achievable precision in a cosmic acceleration measurement carried out with ESPRESSO integrating SB2 for 10, 100, 1000 hours each year, as a function of total experimental baseline.

Similarly, we extrapolated the expected temporal baseline required for a significant detection in an ELT/ANDES-like experiment by simply considering that the ELT will have a collecting area a factor x20 larger than one VLT unit. Thus, given that the proposed efficiency of ANDES is similar to ESPRESSO's ($\sim 10\%$), one hour of ANDES integration is equivalent to 20 hours of ESPRESSO integration time, in terms of the number of photons collected by the telescope. Thus, the temporal baselines to achieve significant detection are shortened by a factor[7] $20^{1/3} \sim 2.7$, as shown in Fig. 7.14. Table 7.5 reports the required baseline to achieve a 1σ, 3σ and 5σ detection following a

[7] The cube-root scaling comes as a direct consequence of the fact that $\sigma_{\dot{v}} \propto \mathrm{S/N}^{-1}\Delta t^{-1} \propto (T\Delta t)^{-1/2}\, \Delta t^{-1} \propto T^{-1/2}\Delta t^{-3/2}$. Fixing an uncertainty level yields $\Delta t \propto T^{-1/3}$, where, due to the collecting area, 1 hour of ANDES integration is equivalent to 20 hours of ESPRESSO integration. Thus, $T_{\mathrm{ANDES}} = 20 T_{ESPRESSO}$.

10, 100, or 1000 hours-per-year campaign with either ESPRESSO or ANDES.

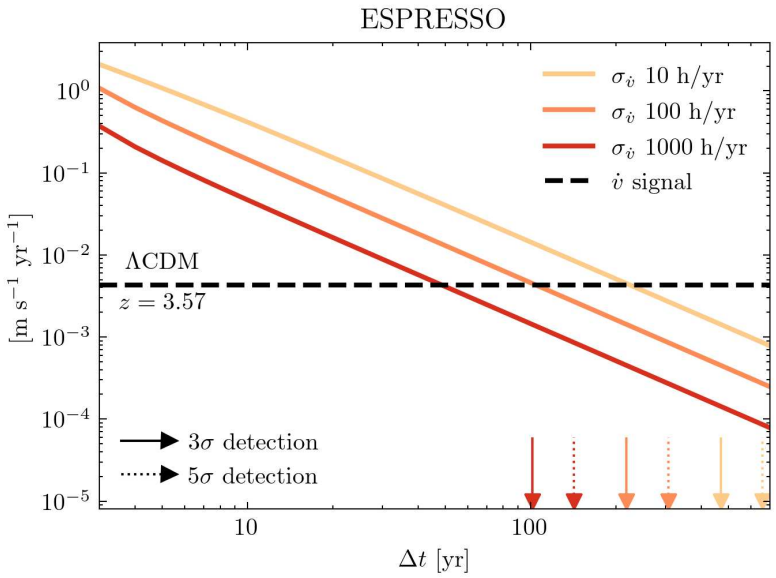

Figure 7.13: *Measurement uncertainty of an ESPRESSO redshift drift experiment targeting SB2 with 10, 100, and 1000 hours of integration per year, as a function of experiment baseline. Horizontal black dashed line defines the magnitude of the cosmological signal in a ΛCDM scenario at $z = 3.57$. Vertical arrows denote the timestamp at which a 3σ (solid) and 5σ (dotted) detection is achieved for each time allocation, with the same colour coding.*

	ESPRESSO			ANDES		
T [h/yr]	1σ [yr]	3σ [yr]	5σ [yr]	1σ [yr]	3σ [yr]	5σ [yr]
10	225	469	659	83	174	243
100	104	218	306	38	80	112
1000	48	101	142	18	27	52

Table 7.5: *Temporal baselines required to achieve a 1σ, 3σ, or 5σ detection following up the SB2 spectra with a total integration time of 10, 100, or 1000 hours per year, if performed with ESPRESSO or ANDES (assuming a 10% efficiency).*

7.9.1 A realistic ESPRESSO + ANDES campaign

Moving to a realistic scenario for a future observing campaign, we shall consider that ANDES will likely not start observations before 2040. In the meantime, carrying out the programme with ESPRESSO is crucial to extend the experiment's baseline and to characterise systematic effects at the centimetre per second per year level. Moreover, a realistic scenario cannot be based solely on one target, but requires a sample of bright quasars that cover the sky uniformly, to maximise observable time and to reduce cosmic variance effects related to a single sightline approach. The Golden Sample of 7 bright quasars ($2.9 \leq z \leq 4.7$) presented in the QUBRICS Survey (Cristiani et al., 2023, see Sect. 5.3.2) was proposed specifically with this aim.

We can extend the forecasts presented in the previous section to a more complex, multi-target, multi-instrument campaign, assuming to continue the experiment with ESPRESSO with a yearly cadence of T_E hours of integration, and follow up with T_A yearly hours of ANDES integration as the instrument comes online. Moreover, assuming that the pressure on the VLTs will significantly drop after the commissioning of the ELT, we can foresee allocating a large amount of yearly time on ESPRESSO, such as

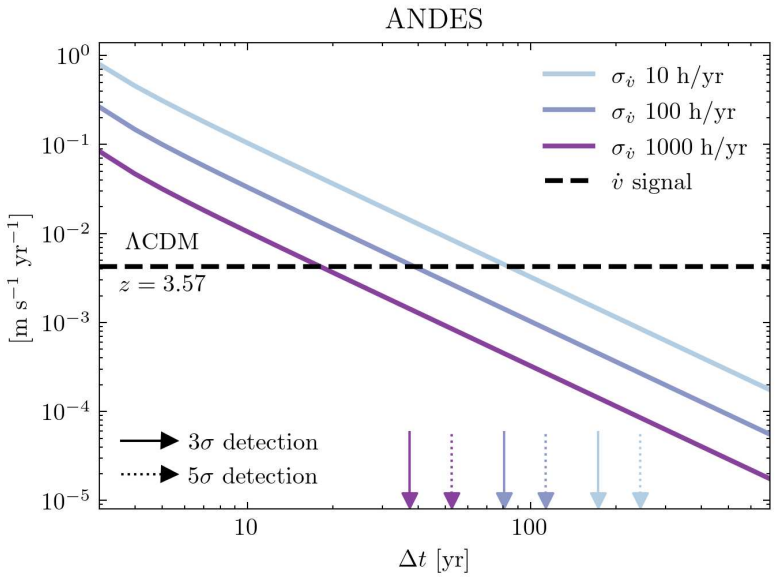

Figure 7.14: *Measurement uncertainty of an ANDES redshift drift experiment targeting SB2 with 10, 100, and 1000 hours of integration per year, as a function of experiment baseline, assuming 10% efficiency.*

1000 h/yr equivalent to all the dark time of one dedicated VLT UT. On the other hand, ELT's high demand will limit the amount of ANDES time devoted to the experiment. Here we assume an allocation of 100 h/yr (i.e. 12.5 nights per year).

Given the results obtained on three epochs of SB2, we extrapolated the measurement precision to the other objects of the GS, as

$$\sigma_{\dot{v}} \propto g \, \mathrm{S/N}^{-1} \Delta t^{-1} \, (1+z)^{-\gamma} \, , \tag{7.26}$$

where g is the form factor computed following Eq. 7.2, Δt is the experiment's baseline, $\gamma = 1.7$ for $z < 4$ and 0.9 otherwise, and S/N is the total signal-to-noise reached by the target. Assuming shot-noise limited observations, S/N is extrapolated from the results of SB2 as

$$\mathrm{S/N} \propto 10^{-0.2\Delta m} \sqrt{(\Delta t + 1) \; T / N_{\mathrm{obj}}}, \tag{7.27}$$

where $\Delta m = m_j - m_{\mathrm{SB2}}$ is the difference in apparent magnitude between the j-th target and SB2, Δt is the experiment baseline, N_{obj} is the number of targets considered. The total amount of integration time per year, T, depends on the time allocations of both ESPRESSO and ANDES as $T = T_E$ before 2040, and $T = T_E + 20 \cdot T_A$ after 2040, assuming that we will not discontinue the ESPRESSO observations to better control systematic effects.

For a given amount of integration time and experiment duration, we estimate the detection significance (in units of σ) of the drift signal as

$$S[\sigma] = \sqrt{\sum_j \left(\frac{\dot{v}(z_j)}{\sigma_{\dot{v},j}} \right)^2} \, , \tag{7.28}$$

where $\dot{v}(z_j)$ is the expected acceleration signal due to cosmic expansion in a ΛCDM universe at the average redshift of the Lyman-α forest of the j-th target, whereas $\sigma_{\dot{v},j}$ is the

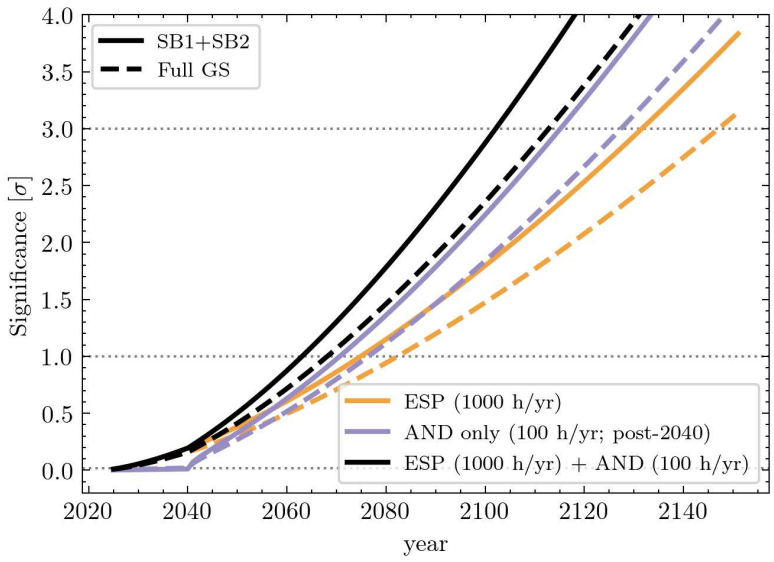

Figure 7.15: *Detection significance as a function of time for an experiment based only on ESPRESSO (1000 h/yr, orange), only on ANDES (100 h/yr after 2040, purple), or both instruments (black) for different choices of targets: SB1+SB2 (solid) or the full Golden Sample (dashed).*

measurement precision reached on such target, as computed in Eq 7.26.

Fig. 7.15 shows the measurement significance in a scenario in which the experiment is carried out adopting only ESPRESSO, only ANDES (after 2040) or both instruments in parallel, allocating 1000 h/yr to the VLT and 100 h/yr to the ELT. Moreover, the figure shows the effects on the measurement significance of targeting all 7 quasars in the GS, or observing only the two brightest ones: SB1 and SB2. These two bright targets have opposite positions in the sky. With a programme of such kind, leveraging both ESPRESSO and ANDES, exploiting the time until the commissioning of the ELT, the measurement precision would reach the cosmic signal level by 2068 when targeting all 7 objects of the GS, and a 3σ by 2112. Note, however, that about 11% less time is required to reach the same measurement significance if focusing only on the two brightest targets, i.e. 1σ in 2063 and 3σ in 2102, although a continuous monitoring of SB1 and SB2 alone is not easily carried out. On the other hand, the 7 quasars of the GS are selected precisely to homogeneously cover the sky so that at least one target is observable with limited airmass at all times, ensuring the feasibility of a continuous program.

The forecasts also show that an ESPRESSO-only experiment can reach significant precision, reaching the 1σ level by 2074 (SB1+SB2), highlighting the fact that the experiment will be carried out even if the scheduling and technical aspects of ANDES would change before the instrument sees first light. In the meantime, a strong effort in the upcoming years before such commissioning is crucial to reach a precision level of $\sim 2\ \mathrm{cm\,s^{-1}\,yr^{-1}}$ where most systematic effects will become dominant, detectable and analysable, providing a smooth transition into the ELT-era.

8

Conclusion, Discussion and Future Perspectives

8.1 The Lyman-α forest as a particle detector

In Part II of this Thesis, we have shown how the Lyman-α forest can be used as a particle detector and as a cosmic laboratory to study fundamental particle physics and put constraints on different theories of particle Dark Matter. In particular, we have tackled the Dark Photon Dark Matter model to put limits on the particle's properties based on its thermal effects on the intergalactic hydrogen gas. We have carried out cosmological hydrodynamical simulations modelling the energy injection in the IGM due to dark photon dark matter resonantly converting into standard photons. From these simulations, we have extracted synthetic Lyman-α skewers and compared them to the ultra-high S/N VLT/UVES spectrum of quasar HE0940-1050 through the regulated and transformed flux PDF, setting robust constraints on the A' DM mass and kinetic mixing parameters.

For the first time, we utilised the full Lyman-α forest in quasar spectra to constrain IGM heating from a DM process; almost all previous works do not account for inhomogeneities in the IGM and do not directly make use of Lyman-α spectra, simply comparing a single, predicted IGM temperature with the gas temperature at mean density inferred from the Lyman-α forest (see e.g. Cirelli et al. (2009), Diamanti et al. (2014), Liu et al. (2016), Muñoz et al. (2017), and Liu et al. (2021)).

We have obtained robust 95% confidence limits on the kinetic mixing parameter of A' DM with mass around $m_{A'} \sim 10^{-13} - 10^{-12}\,\mathrm{eV/c^2}$. This represents a significant improvement compared to previous work (Caputo et al., 2020a; Witte et al., 2020), which provided approximate limits. These limits were determined by requiring that the energy injected by dark photons during the He II reionisation epoch ($2 \leq z \leq 6$) did not cause an increase in the IGM temperature of more than $\Delta T \lesssim 10^4$ K.

Our new results tested the presence of A' DM with significantly more informa-

tion than Bolton et al. (2022a), since we compared simulation results directly with Lyman-α forest data at $z \sim 3$. We found that the best-fit point of Bolton et al. (2022a) is still consistent with the data; therefore, A' DM heating remains a viable solution to the discrepancy between HST/COS data and simulations at $z \sim 0$, although we saw no preference for the model either.

Low and high mass A' DM models remain unconstrained by our data within the interval of kinetic mixing we considered, as these induce energy injection in a density regime not probed by the Lyman-α forest. Nonetheless, these models can be investigated with the current data by expanding the parameter space to very strong mixing scenarios where large amounts of energy are injected into the IGM and can have a non-negligible effect on the Lyman-α forest statistics, even if the oscillations occur in environments that are not well captured in the Lyman-α forest. Similarly, A' DM models with heavy mass and strong mixing can alter the thermal state of the IGM at an earlier epoch and, if enough energy is injected, such a trace will remain detectable at later times. We did not extend the parameter space to high enough mixing values to tackle this occurrence.

Future deep spectroscopic observations will be required to explore the A' DM mass parameter space further, especially targeting the low redshift Lyman-α forest down to $z \sim 2$, where the best fit A' DM model of Bolton et al. (2022a) is expected to affect gas close to the mean density and leave stronger traces in the Lyman-α forest, eventually leading to either the confirmation or ruling out of the model. Such observations can already be carried out with current facilities and instrumentation (such as VLT/UVES Dekker et al., 2000, VLT/ESPRESSO Pepe et al., 2021a and KECK/HIRES Vogt et al., 1994), requiring, however, a substantial amount of integration time (of the order of a hundred hours), due to the low efficiency in the blue part of the spectrum of such instruments. In the future, spectrographs such as ANDES (Marconi et al., 2022) will be more powerful and reach high S/N levels faster, being able to exploit the larger collecting area of the ELT, whereas, projects directly targeting the ultraviolet wavelengths such as VLT/CUBES (Zanutta et al., 2023) could push the search for A' DM down to $z \sim 1.5$, probing even lighter models. In particular, the latter instrument, meant to see first light by 2029, could be used to create a collection of high-S/N spectra of $1.9 < z_{\rm em} < 2.4$ quasars, a low-redshift counterpart of other legacy QSO surveys such as XQR-30 (D'Odorico et al., 2023), XQ-100 (López et al., 2016) and EQUALS (ESO ESPRESSO Proposal 112.25NR, PI T. Berg). A small survey of this kind, even with a limited number of targets (~ 30), would have incredible legacy value for low redshift IGM studies and pave the way to stronger constraints on light A' DM models.

We also note that dark photons in the mass range of interest for us can be produced around spinning black holes through the superradiance instability, then affecting the spin-mass distribution of black holes or producing gravitational waves, for $m_{A'} \sim 8 \times 10^{-14}\,\mathrm{eV}\,c^{-2}$ (Baryakhtar et al., 2017; Cardoso et al., 2018; Ghosh et al., 2021). However, such constraints are currently subject to significant uncertainties (Belczynski et al., 2024; Ghosh et al., 2021) and model dependences (Fukuda et al., 2020; Baryakhtar

et al., 2017; Caputo et al., 2021b; Cannizzaro et al., 2022). This highlights the importance of having complementary probes with distinct systematics in this mass range, such as the one presented in this work.

Finally, although in this Thesis we have focused on a specific DM candidate, our pipeline can be applied to many other cases in which a DM process imprints a thermal effect on the intergalactic medium, such as decaying and annihilating dark matter, evaporating primordial black holes or axions converting in extra-galactic magnetic fields.

8.2 The Lyman-α forest as a tracer of expansion

In Part III, we demonstrated how the Lyman-α forest can serve as a direct tracer of the cosmic expansion. By monitoring the temporal evolution of the redshift of individual Lyman-α absorption lines, it is possible to reconstruct the expansion history of the Universe without assuming a specific cosmological model. The expected signal, however, is extremely small, of order $5\ \mathrm{cm\,s^{-1}}$ per decade, and thus requires a remarkably stable tracer to minimise both statistical noise and systematic effects.

In Chapter 6, we investigated the stability of Lyman-α and metal absorption lines in the spectra of gravitationally lensed quasars. By comparing closely separated sightlines, we found that the Lyman-α absorbers show no measurable variations within the noise level, implying coherence over scales $\gtrsim 1\ \mathrm{kpc}$ and negligible peculiar motions ($\sim 0.01\ \mathrm{cm\,s^{-1}\,yr^{-1}}$). This noise contribution is an order of magnitude smaller than the expected redshift drift signal, confirming that Lyman-α absorbers are sufficiently stable for use as cosmological tracers. Conversely, metal lines exhibit significant variability across separations $\lesssim 30\ \mathrm{kpc}$, consistent with gas dynamics within virialised dark-matter haloes of $M \sim 2 \times 10^{10}\ M_\odot$. The inferred peculiar velocity noise of $\sim 0.2\ \mathrm{cm\,s^{-1}\,yr^{-1}}$ is comparable to the cosmological signal itself, rendering metal lines unsuitable for the Sandage test.

Building on these results, in Chapter 7 we initiated the first redshift-drift experiment based solely on the Lyman-α forest. Using ESPRESSO at the VLT, we obtained high-resolution and high-S/N spectra of J052915.80-435152.0 (SB2, $z_{\mathrm{em}} = 3.962$, total S/N ~ 85), the most luminous known quasar. This source is ideally suited for a decades-long experiment aimed at detecting the cosmic redshift drift. The first two observational epochs provided a baseline of one year, allowing us to quantify intrinsic systematics and validate the theoretical predictions for achievable precision.

The measurement is based on the comparison of a spectral model, computed via spline reconstruction of the Lyman-α forest transmission, to the single epoch datasets. We adopted two different independent methods to estimate the velocity shift occurring between model and data, and therefore between the two epochs. Firstly, we applied a direct Pixel-by-pixel comparison based on the concept of maximum radial velocity information, as employed in RV measurements of stars (Bouchy et al., 2001). With this approach, each pixel of the spectrum contributes to the total estimated velocity

shift. Over the two epochs, we measure $\Delta v = 0.45 \pm 7.66\ \mathrm{m\,s^{-1}}$, which translates to a cosmic acceleration of $\dot{v} = 0.51 \pm 8.80\ \mathrm{m\,s^{-1}\,yr^{-1}}$ ($\dot{z} = (0.08 \pm 1.34) \times 10^{-7}\ \mathrm{yr^{-1}}$). Parallelly, we employed a correlation method, where the model is rigidly shifted over the data and the best fitting δv is estimated through the maximisation of a Gaussian Likelihood. With this alternative approach, we measured over the two epochs a velocity drift of $\Delta v = 0.49 \pm 7.78\ \mathrm{m\,s^{-1}}$, or in acceleration terms $\dot{v} = 0.57 \pm 8.94\ \mathrm{m\,s^{-1}\,yr^{-1}}$ ($\dot{z} = (0.09 \pm 1.39) \times 10^{-7}\ \mathrm{yr^{-1}}$).

A third observational epoch extended the temporal baseline and raised the total S/N per $1\ \mathrm{km\,s^{-1}}$ pixel to ~ 113. Reapplying our modelling and measurement framework, we obtained $\dot{v} = -2.27 \pm 3.58\ \mathrm{m\,s^{-1}\,yr^{-1}}$ (pixel-by-pixel) and $\dot{v} = -2.84 \pm 3.64\ \mathrm{m\,s^{-1}\,yr^{-1}}$ (Likelihood). Both measurements are mutually consistent with each other and with the predicted ΛCDM drift at $z = 3.5$, representing the most precise Lyman-α forest-based constraint to date. Systematic uncertainties related to solar-system motion, quasar variability, and wavelength calibration were found to be subdominant, at the level of $\lesssim 1\ \mathrm{m\,s^{-1}}$.

A complementary analysis was performed on the second target of the QUBRICS Golden Sample, J212540.97-171951.4 (SB1). The first epoch of ESPRESSO high-S/N observations, obtained in 2023, allowed us to construct the spectral model of its Lyman-α forest and assess its suitability for future drift measurements. Owing to the temporary unavailability of a second epoch, no velocity shift could yet be derived; however, the present data have been used to compare the two quasar sightlines, highlighting that, at given S/N, SB2 provides a 22% more precise measurement of the velocity drift, with respect to SB1, due to cosmic variance effects such as the variation of the number of strong and faint absorbers and their distribution along the two individual sightlines. Regardless of this difference, being the two quasars at similar redshifts, once a second epoch of SB1 is collected, a new stringent joint constraint will be put on the redshift drift at $z \sim 3.5$.

From the results presented above, we forecasted the observational timeline required for a definitive detection of the cosmic redshift drift. ESPRESSO alone could, in principle, achieve a 3σ detection within ~ 100 years assuming a continuous allocation of $1000\ \mathrm{h\ yr^{-1}}$. A more realistic and efficient approach exploits the synergy with ELT/ANDES, enabling a 1σ (3σ) detection by 2068 (2112) when observing the seven quasars of the QUBRICS Golden Sample. Limiting the programme to the two brightest quasars, SB1 and SB2, achieves a 1σ detection by 2063 and a 3σ detection by 2102. These results highlight the feasibility of the Sandage test through Lyman-α -forest spectroscopy and confirm that the method, while technically demanding, rests on a robust physical foundation established in this work.

In the short term, new results from the redshift–drift experiment are expected within the next few years. Continued observations, amounting to a few tens of hours per year on both SB1 and SB2 (already allocated for ESO Period 116, 2025–2026), will soon achieve a precision of $\sim 1\ \mathrm{m\,s^{-1}\,yr^{-1}}$, at which level the first systematic effects are expected to emerge. Among these, wavelength calibration is likely to represent the dom-

inant source of uncertainty. In parallel, optical Lyman-α -based measurements show a strong complementarity with low-redshift H I 21cm absorption-based experiments such as that of Darling (2012). Repeating those observations in the near future would effectively double the radio temporal baseline (to $> 20\,\mathrm{yr}$), tightening the low-redshift constraint to $\sim 1\,\mathrm{m\,s^{-1}\,yr^{-1}}$. The combination of high- and low-redshift measurements will elevate the redshift–drift experiment to a new level, enabling a much stronger joint constraint on the dynamics of cosmic expansion.

Bibliography

Aad, G. et al. (2012). "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Physics Letters B* 716.1, pp. 1–29. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2012.08.020. URL: https://www.sciencedirect.com/science/article/pii/S037026931200857X.

Aaij, R. et al. (May 2025). "Measurement of the Z-boson mass". In: arXiv: 2505.15582 [hep-ex].

Abbott, T. M. C. et al. (Aug. 2021). "The Dark Energy Survey Data Release 2". In: *ApJS* 255.2, 20, p. 20. DOI: 10.3847/1538-4365/ac00b3. arXiv: 2101.05765 [astro-ph.IM].

Adelberger, E. G. et al. (Jan. 2011). "Solar fusion cross sections. II. The pp chain and CNO cycles". In: *Reviews of Modern Physics* 83.1, pp. 195–246. DOI: 10.1103/RevModPhys.83.195. arXiv: 1004.2318 [nucl-ex].

Agrawal, P., N. Kitajima, M. Reece, T. Sekiguchi, and F. Takahashi (2020). "Relic Abundance of Dark Photon Dark Matter". In: *Phys. Lett.* B801, p. 135136. DOI: 10.1016/j.physletb.2019.135136. arXiv: 1810.07188 [hep-ph].

Almgren, A., J. Bell, M. Lijewski, Z. Lukic, and E. Van Andel (2013). "Nyx: A Massively Parallel AMR Code for Computational Cosmology". In: *Astrophys. J.* 765, p. 39. DOI: 10.1088/0004-637X/765/1/39. arXiv: 1301.4498 [astro-ph.IM].

Alonso-Álvarez, G., D. Curtin, A. Rasovic, and Z. Yuan (2024). "Baryogenesis through asymmetric reheating in the mirror twin Higgs". In: *JHEP* 05, p. 069. DOI: 10.1007/JHEP05(2024)069. arXiv: 2311.06341 [hep-ph].

Alonso-Álvarez, G., T. Hugle, and J. Jaeckel (2020). "Misalignment & Co.: (Pseudo-)Scalar and Vector Dark Matter with Curvature Couplings". In: *JCAP* 02, p. 014. DOI: 10.1088/1475-7516/2020/02/014. arXiv: 1905.09836 [hep-ph].

An, H., X. Chen, S. Ge, J. Liu, and Y. Luo (2024). "Searching for ultralight dark matter conversion in solar corona using Low Frequency Array data". In: *Nature Commun.* 15.1, p. 915. DOI: 10.1038/s41467-024-45033-4. arXiv: 2301.03622 [hep-ph].

An, H., F. P. Huang, J. Liu, and W. Xue (2021). "Radio-frequency Dark Photon Dark Matter across the Sun". In: *Phys. Rev. Lett.* 126.18, p. 181102. DOI: 10.1103/PhysRevLett.126.181102. arXiv: 2010.15836 [hep-ph].

Antonucci, R. (1993). "Unified Models for Active Galactic Nuclei and Quasars". In: *Annual Review of Astronomy and Astrophysics* 31.1, pp. 473–521. DOI: 10.1146/annurev.aa.31.090193.002353. eprint: https://doi.org/10.1146/annurev.aa.31.090193.002353. URL: https://doi.org/10.1146/annurev.aa.31.090193.002353.

Arcadi, G., A. Djouadi, and M. Kado (July 2021). "The Higgs-portal for dark matter: effective field theories versus concrete realizations". In: *European Physical Journal C* 81.7, 653, p. 653. DOI: 10.1140/epjc/s10052-021-09411-2. arXiv: 2101.02507 [hep-ph].

Arias, P. et al. (2012). "WISPy Cold Dark Matter". In: *JCAP* 1206, p. 013. DOI: 10.1088/1475-7516/2012/06/013. arXiv: 1201.5902 [hep-ph].

Armengaud, E., N. Palanque-Delabrouille, C. Yèche, D. J. E. Marsh, and J. Baur (2017). "Constraining the mass of light bosonic dark matter using SDSS Lyman-α forest". In: *Mon. Not. Roy. Astron. Soc.* 471.4, pp. 4606–4614. DOI: 10.1093/mnras/stx1870. arXiv: 1703.09126 [astro-ph.CO].

Augustin, R. et al. (Aug. 2021). "Clumpiness of observed and simulated cold circumgalactic gas". In: *MNRAS* 505.4, pp. 6195–6205. DOI: 10.1093/mnras/stab1673. arXiv: 2105.11480 [astro-ph.GA].

Bacon, R. et al. (July 2010). "The MUSE second-generation VLT instrument". In: *Ground-based and Airborne Instrumentation for Astronomy III*. Ed. by I. S. McLean, S. K. Ramsay, and H. Takami. Vol. 7735. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 773508, p. 773508. DOI: 10.1117/12.856027. arXiv: 2211.16795 [astro-ph.IM].

Bahcall, J. N. and P. J. E. Peebles (Apr. 1969). "Statistical Tests for the Origin of Absorption Lines Observed in Quasi-Stellar Sources". In: *ApJ* 156, p. L7. DOI: 10.1086/180337.

Baryakhtar, M., R. Lasenby, and M. Teo (2017). "Black Hole Superradiance Signatures of Ultralight Vectors". In: *Phys. Rev. D* 96.3, p. 035019. DOI: 10.1103/PhysRevD.96.035019. arXiv: 1704.05081 [hep-ph].

Basant, R. et al. (Mar. 2025). "Four Sub-Earth Planets Orbiting Barnard's Star from MAROON-X and ESPRESSO". In: *ApJ* 982.1, L1, p. L1. DOI: 10.3847/2041-8213/adb8d5. arXiv: 2503.08095 [astro-ph.EP].

Bastero-Gil, M., J. Santiago, L. Ubaldi, and R. Vega-Morales (2019). "Vector dark matter production at the end of inflation". In: *JCAP* 1904.04, p. 015. DOI: 10.1088/1475-7516/2019/04/015. arXiv: 1810.07208 [hep-ph].

Beadle, C., A. Caputo, and S. A. R. Ellis (May 2024). "Resonant Conversion of Wave Dark Matter in the Ionosphere". In: *arXiv e-prints*, arXiv:2405.13882, arXiv:2405.13882. DOI: 10.48550/arXiv.2405.13882. arXiv: 2405.13882 [hep-ph].

Becker, G. D. and J. S. Bolton (Dec. 2013). "New measurements of the ionizing ultraviolet background over $2 < z < 5$ and implications for hydrogen reionization". In: *MNRAS* 436, pp. 1023–1039. DOI: 10.1093/mnras/stt1610. arXiv: 1307.2259 [astro-ph.CO].

Becker, G. D., M. Rauch, and W. L. W. Sargent (June 2007). "The Evolution of Optical Depth in the Lyα Forest: Evidence Against Reionization at z~6". In: *ApJ* 662, pp. 72–93. DOI: 10.1086/517866. eprint: arXiv:astro-ph/0607633.

Beckmann, V. and C. R. Shrader (2012). *Active Galactic Nuclei*.

Belczynski, K., C. Done, S. Hagen, J.-P. Lasota, and K. Sen (2024). "Common origin for black holes in both high mass X-ray binaries and gravitational-wave sources". In: *Astron. Astrophys.* 690, A21. DOI: 10.1051/0004-6361/202450229. arXiv: 2111.09401 [astro-ph.HE].

Bergeron, J. et al. (Dec. 2004). "The large programme "Cosmic Evolution of the IGM"". In: *The Messenger* 118, pp. 40–44.

Berlin, A., J. A. Dror, X. Gan, and J. T. Ruderman (2023). "Millicharged relics reveal massless dark photons". In: *JHEP* 05, p. 046. DOI: 10.1007/JHEP05(2023)046. arXiv: 2211.05139 [hep-ph].

Bertin, E. and S. Arnouts (June 1996). "SExtractor: Software for source extraction." In: *A&AS* 117, pp. 393–404. DOI: 10.1051/aas:1996164.

Blandford, R., D. Meier, and A. Readhead (Aug. 2019). "Relativistic Jets from Active Galactic Nuclei". In: *ARA&A* 57, pp. 467–509. DOI: 10.1146/annurev-astro-081817-051948. arXiv: 1812.06025 [astro-ph.HE].

Bolton, J. S. and M. G. Haehnelt (Nov. 2007). "The observed ionization rate of the intergalactic medium and the ionizing emissivity at $z > 5$: evidence for a photon-starved and extended epoch of reionization". In: *MNRAS* 382, pp. 325–341. DOI: 10.1111/j.1365-2966.2007.12372.x. eprint: arXiv:astro-ph/0703306.

Bolton, J. S., M. Viel, T. .-. Kim, M. G. Haehnelt, and R. F. Carswell (May 2008). "Possible evidence for an inverted temperature-density relation in the intergalactic medium from the flux distribution of the Lyα forest". In: *Mon. Not. Roy. Astron. Soc.* 386.2, pp. 1131–1144. DOI: 10.1111/j.1365-2966.2008.13114.x. arXiv: 0711.2064 [astro-ph].

Bolton, J. S. et al. (Jan. 2017a). "The Sherwood simulation suite: overview and data comparisons with the Lyman α forest at redshifts $2 \leq z \leq 5$". In: *Mon. Not. Roy. Astron. Soc.* 464, pp. 897–914. DOI: 10.1093/mnras/stw2397. arXiv: 1605.03462.

Bolton, J. S., A. Caputo, H. Liu, and M. Viel (2022a). "Comparison of Low-Redshift Lyman-α Forest Observations to Hydrodynamical Simulations with Dark Photon Dark Matter". In: *Phys. Rev. Lett.* 129.21, p. 211102. DOI: 10.1103/PhysRevLett.129.211102. arXiv: 2206.13520 [hep-ph].

Bolton, J. S. et al. (Jan. 2017b). "The Sherwood simulation suite: overview and data comparisons with the Lyman α forest at redshifts $2 \leq z \leq 5$". In: *MNRAS* 464.1, pp. 897–914. DOI: 10.1093/mnras/stw2397. arXiv: 1605.03462 [astro-ph.CO].

Bolton, J. S. et al. (June 2022b). "Limits on non-canonical heating and turbulence in the intergalactic medium from the low redshift Lyman α forest". In: *Mon. Not. Roy. Astron. Soc.* 513.1, pp. 864–885. DOI: 10.1093/mnras/stac862. arXiv: 2111.09600 [astro-ph.CO].

Bond, J. R., L. Kofman, and D. Pogosyan (Apr. 1996). "How filaments of galaxies are woven into the cosmic web". In: *Nature* 380.6575, pp. 603–606. DOI: 10.1038/380603a0. arXiv: astro-ph/9512141 [astro-ph].

Bosman, S. E. I. et al. (2022). "Hydrogen reionization ends by $z = 5.3$: Lyman-α optical depth measured by the XQR-30 sample". In: *Mon. Not. Roy. Astron. Soc.* 514.1, pp. 55–76. DOI: 10.1093/mnras/stac1046. arXiv: 2108.03699 [astro-ph.CO].

Bouchy, F., F. Pepe, and D. Queloz (Aug. 2001). "Fundamental photon noise limit to radial velocity measurements". In: *A&A* 374, pp. 733–739. DOI: 10.1051/0004-6361:20010730.

Broderick, A. E., P. Chang, and C. Pfrommer (June 2012). "The Cosmological Impact of Luminous TeV Blazars. I. Implications of Plasma Instabilities for the Intergalactic Magnetic Field and Extragalactic Gamma-Ray Background". In: *ApJ* 752.1, 22, p. 22. DOI: 10.1088/0004-637X/752/1/22. arXiv: 1106.5494 [astro-ph.CO].

Bryan, G. L. and M. L. Norman (Mar. 1998). "Statistical Properties of X-Ray Clusters: Analytic and Numerical Comparisons". In: *ApJ* 495.1, pp. 80–99. DOI: 10.1086/305262. arXiv: astro-ph/9710107 [astro-ph].

Bryan, G. L. et al. (Apr. 2014). "ENZO: An Adaptive Mesh Refinement Code for Astrophysics". In: *ApJS* 211.2, 19, p. 19. DOI: 10.1088/0067-0049/211/2/19. arXiv: 1307.2265 [astro-ph.IM].

Buchner, J. et al. (Apr. 2015). "Obscuration-dependent Evolution of Active Galactic Nuclei". In: *ApJ* 802.2, 89, p. 89. DOI: 10.1088/0004-637X/802/2/89. arXiv: 1501.02805 [astro-ph.HE].

Burbidge, E. M., G. R. Burbidge, W. A. Fowler, and F. Hoyle (Jan. 1957). "Synthesis of the Elements in Stars". In: *Reviews of Modern Physics* 29.4, pp. 547–650. DOI: 10.1103/RevModPhys.29.547.

Burkhart, B. et al. (July 2022). "The Low-redshift Lyα Forest as a Constraint for Models of AGN Feedback". In: *ApJ* 933.2, L46, p. L46. DOI: 10.3847/2041-8213/ac7e49. arXiv: 2204.09712 [astro-ph.GA].

Buzzoni, B. et al. (Dec. 1984). "The ESO Faint Object Spectrograph and Camera / EFOSC". In: *The Messenger* 38, p. 9.

Calderone, G. et al. (Dec. 2019). "Finding the Brightest Cosmic Beacons in the Southern Hemisphere". In: *ApJ* 887.2, 268, p. 268. DOI: 10.3847/1538-4357/ab510a. arXiv: 1909.06391 [astro-ph.IM].

Caldwell, R. R., R. Dave, and P. J. Steinhardt (Feb. 1998). "Cosmological Imprint of an Energy Component with General Equation of State". In: *Phys. Rev. Lett.* 80.8, pp. 1582–1585. DOI: 10.1103/PhysRevLett.80.1582. arXiv: astro-ph/9708069 [astro-ph].

Caldwell, R. R., M. Kamionkowski, and N. N. Weinberg (2003). "Phantom energy and cosmic doomsday". In: *Phys. Rev. Lett.* 91, p. 071301. DOI: 10.1103/PhysRevLett.91.071301. arXiv: astro-ph/0302506.

Calura, F. et al. (June 2012). "The Lyman α forest flux probability distribution at z>3". In: *MNRAS* 422.4, pp. 3019–3036. DOI: 10.1111/j.1365-2966.2012.20811.x. arXiv: 1201.5121 [astro-ph.CO].

Cannizzaro, E., L. Sberna, A. Caputo, and P. Pani (2022). "Dark photon superradiance quenched by dark matter". In: *Phys. Rev. D* 106.8, p. 083019. DOI: 10.1103/PhysRevD.106.083019. arXiv: 2206.12367 [hep-ph].

Caputo, A., H. Liu, S. Mishra-Sharma, and J. T. Ruderman (2020a). "Dark Photon Oscillations in Our Inhomogeneous Universe". In: *Phys. Rev. Lett.* 125.22, p. 221303. DOI: 10.1103/PhysRevLett.125.221303. arXiv: 2002.05165 [astro-ph.CO].

Caputo, A., H. Liu, S. Mishra-Sharma, and J. T. Ruderman (2020b). "Modeling Dark Photon Oscillations in Our Inhomogeneous Universe". In: *Phys. Rev. D* 102.10, p. 103533. DOI: 10.1103/PhysRevD.102.103533. arXiv: 2004.06733 [astro-ph.CO].

Caputo, A., A. J. Millar, C. A. J. O'Hare, and E. Vitagliano (2021a). "Dark photon limits: A handbook". In: *Phys. Rev. D* 104.9, p. 095029. DOI: 10.1103/PhysRevD.104.095029. arXiv: 2105.04565 [hep-ph].

Caputo, A., S. J. Witte, D. Blas, and P. Pani (2021b). "Electromagnetic signatures of dark photon superradiance". In: *Phys. Rev. D* 104.4, p. 043006. DOI: 10.1103/PhysRevD.104.043006. arXiv: 2102.11280 [hep-ph].

Cardoso, V. et al. (2018). "Constraining the mass of dark photons and axion-like particles through black-hole superradiance". In: *JCAP* 03, p. 043. DOI: 10.1088/1475-7516/2018/03/043. arXiv: 1801.01420 [gr-qc].

Carroll, S. M., W. H. Press, and E. L. Turner (Jan. 1992). "The cosmological constant." In: *ARA&A* 30, pp. 499–542. DOI: 10.1146/annurev.aa.30.090192.002435.

Carroll, S. M. et al. (2005). "The Cosmology of generalized modified gravity models". In: *Phys. Rev. D* 71, p. 063513. DOI: 10.1103/PhysRevD.71.063513. arXiv: astro-ph/0410031.

Chabanier, S. et al. (July 2019). "The one-dimensional power spectrum from the SDSS DR14 Lyα forests". In: *JCAP* 2019.7, 017, p. 017. DOI: 10.1088/1475-7516/2019/07/017. arXiv: 1812.03554 [astro-ph.CO].

Chatrchyan, S. et al. (2012). "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". In: *Physics Letters B* 716.1, pp. 30–61. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2012.08.021. URL: https://www.sciencedirect.com/science/article/pii/S0370269312008581.

Chekhovsky, V. et al. (Dec. 2024). "High-precision measurement of the W boson mass with the CMS experiment at the LHC". In: arXiv: 2412.13872 [hep-ex].

Chen, H.-W., K. M. Lanzetta, and J. K. Webb (July 2001). "The Origin of C IV Absorption Systems at Redshifts z < 1: Discovery of Extended C IV Envelopes around Galaxies". In: *ApJ* 556.1, pp. 158–163. DOI: 10.1086/321537. arXiv: astro-ph/0104403 [astro-ph].

Choi, E. et al. (Nov. 2020). "The Impact of Outflows Driven by Active Galactic Nuclei on Metals in and around Galaxies". In: *ApJ* 904.1, 8, p. 8. DOI: 10.3847/1538-4357/abba7d. arXiv: 2009.10184 [astro-ph.GA].

Cieplak, A. M. and A. Slosar (Oct. 2017). "Characterizing the Lyα forest flux probability distribution function using Legendre polynomials". In: *JCAP* 2017.10, 013, p. 013. DOI: 10.1088/1475-7516/2017/10/013. arXiv: 1608.08808 [astro-ph.CO].

Cirasuolo, M. et al. (Mar. 2024). "The Rise of the Giant: ESO's Extremely Large Telescope". In: *The Messenger* 192, pp. 3–3. DOI: 10.18727/0722-6691/5346.

Cirelli, M., F. Iocco, and P. Panci (Oct. 2009). "Constraints on Dark Matter annihilations from reionization and heating of the intergalactic gas". In: *JCAP* 2009.10, 009, p. 009. DOI: 10.1088/1475-7516/2009/10/009. arXiv: 0907.0719 [astro-ph.CO].

Co, R. T., A. Pierce, Z. Zhang, and Y. Zhao (2019). "Dark Photon Dark Matter Produced by Axion Oscillations". In: *Phys. Rev.* D99.7, p. 075002. DOI: 10.1103/PhysRevD.99.075002. arXiv: 1810.07196 [hep-ph].

Cole, S., C. G. Lacey, C. M. Baugh, and C. S. Frenk (Nov. 2000). "Hierarchical galaxy formation". In: *MNRAS* 319, pp. 168–204. eprint: arXiv:astro-ph/0007281.

Cooke, R. (Feb. 2020). "The ACCELERATION programme: I. Cosmology with the redshift drift". In: *MNRAS* 492.2, pp. 2044–2057. DOI: 10.1093/mnras/stz3465. arXiv: 1912.04983 [astro-ph.CO].

Cooke, R. et al. (Dec. 2010). "A newly discovered DLA and associated Lyα emission in the spectra of the gravitationally lensed quasar UM673A,B". In: *MNRAS* 409.2, pp. 679–693. DOI: 10.1111/j.1365-2966.2010.17331.x. arXiv: 1007.1409 [astro-ph.CO].

Cowan, G., K. Cranmer, E. Gross, and O. Vitells (2011). "Asymptotic Formulae for Likelihood-Based Tests of New Physics". In: *Eur. Phys. J. C* 71. [Erratum: *Eur. Phys.*

J. C **73**, 2501 (2013)], p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0,10.1140/epjc/s10052-013-2501-z. arXiv: 1007.1727 [physics.data-an].

Cristiani, S. et al. (Feb. 1997). "The clustering properties of the Lyman <e6>alpha</e6> clouds". In: *MNRAS* 285.1, pp. 209–217. DOI: 10.1093/mnras/285.1.209. arXiv: astro-ph/9610006 [astro-ph].

Cristiani, S., L. M. Serrano, F. Fontanot, E. Vanzella, and P. Monaco (Nov. 2016). "The spectral slope and escape fraction of bright quasars at z ∼ 3.8: the contribution to the cosmic UV background". In: *MNRAS* 462.3, pp. 2478–2485. DOI: 10.1093/mnras/stw1810. arXiv: 1603.09351 [astro-ph.CO].

Cristiani, S. et al. (June 2023). "Spectroscopy of QUBRICS quasar candidates: 1672 new redshifts and a golden sample for the Sandage test of the redshift drift". In: *MNRAS* 522.2, pp. 2019–2028. DOI: 10.1093/mnras/stad1007. arXiv: 2304.00362 [astro-ph.CO].

Croft, R. A. C., D. H. Weinberg, N. Katz, and L. Hernquist (Mar. 1998). "Recovery of the Power Spectrum of Mass Fluctuations from Observations of the Lyα Forest". In: *ApJ* 495.1, pp. 44–62. DOI: 10.1086/305289. arXiv: astro-ph/9708018 [astro-ph].

Crotts, A. P. S. and Y. Fang (July 1998). "Reobservation of Close QSO Groups: The Size Evolution and Shape of Lyα Forest Absorbers". In: *ApJ* 502.1, pp. 16–47. DOI: 10.1086/305872. arXiv: astro-ph/9702185 [astro-ph].

Cupani, G. et al. (July 2016). "Integrated data analysis in the age of precision spectroscopy: the ESPRESSO case". In: *Software and Cyberinfrastructure for Astronomy IV*. Ed. by G. Chiozzi and J. C. Guzman. Vol. 9913. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 99131T, 99131T. DOI: 10.1117/12.2231379.

Cupani, G. et al. (Oct. 2019). "Field Tests for the ESPRESSO Data Analysis Software". In: *Astronomical Data Analysis Software and Systems XXVI*. Ed. by M. Molinaro, K. Shortridge, and F. Pasian. Vol. 521. Astronomical Society of the Pacific Conference Series, p. 362. arXiv: 1808.04249 [astro-ph.IM].

Cupani, G. et al. (2020). "Astrocook: your starred chef for spectral analysis". In: *Software and Cyberinfrastructure for Astronomy VI*. Ed. by J. C. Guzman and J. Ibsen. Vol. 11452. International Society for Optics and Photonics. SPIE, pp. 372–388. DOI: 10.1117/12.2561343. URL: https://doi.org/10.1117/12.2561343.

Cyncynates, D. and Z. J. Weiner (Oct. 2023). "Detectable, defect-free dark photon dark matter". In: arXiv: 2310.18397 [hep-ph].

Cyncynates, D. and Z. J. Weiner (Oct. 2024). "Experimental targets for dark photon dark matter". In: arXiv: 2410.14774 [hep-ph].

D'Odorico, V., P. Petitjean, and S. Cristiani (July 2002). "High matter density peaks from UVES observations of QSO pairs: Correlation properties and chemical abun-

dances". In: *A&A* 390, pp. 13–25. DOI: 10.1051/0004-6361:20020737. arXiv: astro-ph/0205299 [astro-ph].

D'Odorico, V. et al. (Nov. 1998). "The size and geometry of the Lyα clouds". In: *A&A* 339, pp. 678–686. arXiv: astro-ph/9806204 [astro-ph].

D'Odorico, V. et al. (Dec. 2016). "Metals in the z ~ 3 intergalactic medium: results from an ultra-high signal-to-noise ratio UVES quasar spectrum". In: *MNRAS* 463.3, pp. 2690–2707. DOI: 10.1093/mnras/stw2161. arXiv: 1608.06116 [astro-ph.GA].

D'Odorico, V. et al. (July 2023). "XQR-30: The ultimate XSHOOTER quasar sample at the reionization epoch". In: *MNRAS* 523.1, pp. 1399–1420. DOI: 10.1093/mnras/stad1468. arXiv: 2305.05053 [astro-ph.GA].

D'Odorico, V. et al. (2006). "Tomography of the intergalactic medium with Lyα forests in close QSO pairs". In: *Monthly Notices of the Royal Astronomical Society* 372, pp. 1333–1344.

da Fonseca, V. et al. (Oct. 2022). "Fundamental physics with ESPRESSO: Constraining a simple parametrisation for varying α". In: *A&A* 666, A57, A57. DOI: 10.1051/0004-6361/202243795. arXiv: 2204.02930 [astro-ph.CO].

Danforth, C. W. et al. (Feb. 2016). "An HST/COS Survey of the Low-redshift Intergalactic Medium. I. Survey, Methodology, and Overall Results". In: *ApJ* 817, 111, p. 111. DOI: 10.3847/0004-637X/817/2/111. arXiv: 1402.2655.

Darling, J. (2012). "Toward a Direct Measurement of the Cosmic Acceleration". In: *Astrophys. J.* 761, p. L26. DOI: 10.1088/2041-8205/761/2/L26. arXiv: 1211.4585 [astro-ph.CO].

Davé, R., L. Hernquist, N. Katz, and D. H. Weinberg (Feb. 1999). "The Low-Redshift Lyα Forest in Cold Dark Matter Cosmologies". In: *ApJ* 511.2, pp. 521–545. DOI: 10.1086/306722. arXiv: astro-ph/9807177 [astro-ph].

Davis, M., G. Efstathiou, C. S. Frenk, and S. D. M. White (May 1985). "The evolution of large-scale structure in a universe dominated by cold dark matter". In: *ApJ* 292, pp. 371–394. DOI: 10.1086/163168.

Deffayet, C., G. R. Dvali, and G. Gabadadze (2002). "Accelerated universe from gravity leaking to extra dimensions". In: *Phys. Rev. D* 65, p. 044023. DOI: 10.1103/PhysRevD.65.044023. arXiv: astro-ph/0105068.

Dekker, H., S. D'Odorico, A. Kaufer, B. Delabre, and H. Kotzlowski (Aug. 2000). "Design, construction, and performance of UVES, the echelle spectrograph for the UT2 Kueyen Telescope at the ESO Paranal Observatory". In: *Optical and IR Telescope Instrumentation and Detectors*. Ed. by M. Iye and A. F. Moorwood. Vol. 4008. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 534–545. DOI: 10.1117/12.395512.

DESI Collaboration et al. (Mar. 2025a). "DESI DR2 Results I: Baryon Acoustic Oscillations from the Lyman Alpha Forest". In: *arXiv e-prints*, arXiv:2503.14739, arXiv:2503.14739. DOI: 10.48550/arXiv.2503.14739. arXiv: 2503.14739 [astro-ph.CO].

DESI Collaboration et al. (Mar. 2025b). "DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints". In: *arXiv e-prints*, arXiv:2503.14738, arXiv:2503.14738. DOI: 10.48550/arXiv.2503.14738. arXiv: 2503.14738 [astro-ph.CO].

Di Marcantonio, P. et al. (July 2018). "ESPRESSO data flow in operations: results of commissioning activities". In: *Observatory Operations: Strategies, Processes, and Systems VII*. Vol. 10704. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 107040F, 107040F. DOI: 10.1117/12.2311285.

Diamanti, R., L. Lopez-Honorez, O. Mena, S. Palomares-Ruiz, and A. C. Vincent (Feb. 2014). "Constraining dark matter late-time energy injection: decays and p-wave annihilations". In: *JCAP* 2014.2, 017, p. 017. DOI: 10.1088/1475-7516/2014/02/017. arXiv: 1308.2578 [astro-ph.CO].

Doyle, L. et al. (Oct. 2022). "The Hot Neptune WASP-166 b with ESPRESSO - I. Refining the planetary architecture and stellar variability". In: *MNRAS* 516.1, pp. 298–315. DOI: 10.1093/mnras/stac2178. arXiv: 2207.10127 [astro-ph.EP].

Dutta, R. et al. (Feb. 2024). "Probing coherence in metal absorption towards multiple images of strong gravitationally lensed quasars". In: *MNRAS* 528.2, pp. 1895–1905. DOI: 10.1093/mnras/stae048. arXiv: 2401.03024 [astro-ph.GA].

East, W. E. and J. Huang (2022). "Dark photon vortex formation and dynamics". In: *JHEP* 12, p. 089. DOI: 10.1007/JHEP12(2022)089. arXiv: 2206.12432 [hep-ph].

Efron, B. (1979). "Bootstrap Methods: Another Look at the Jackknife". In: *The Annals of Statistics* 7.1, pp. 1–26. DOI: 10.1214/aos/1176344552. URL: https://doi.org/10.1214/aos/1176344552.

Efstathiou, G. and R. I. Jedrzejewski (Jan. 1984). "Observational constraints on dark matter in the universe". In: *Advances in Space Research* 3.10-12, pp. 379–386. DOI: 10.1016/0273-1177(84)90119-4.

Efstathiou, G., W. J. Sutherland, and S. J. Maddox (Dec. 1990). "The cosmological constant and cold dark matter". In: *Nature* 348.6303, pp. 705–707. DOI: 10.1038/348705a0.

Eigenbrod, A., F. Courbin, and G. Meylan (Apr. 2007). "COSMOGRAIL: the COSmological MOnitoring of GRAvItational Lenses. VI. Redshift of the lensing galaxy in seven gravitationally lensed quasars". In: *A&A* 465.1, pp. 51–56. DOI: 10.1051/0004-6361:20066939. arXiv: astro-ph/0612419 [astro-ph].

ESO (2022). *ESPRESSO Pipeline User Manual*. VLT-MAN-ESO-XXXXX-XXXX, European Southern Observatory. Available at https://www.eso.org/sci/software/pipelines/espresso/.

Fabbrichesi, M., E. Gabrielli, and G. Lanfranchi (May 2020). "The Dark Photon". In: *arXiv e-prints*, arXiv:2005.01515, arXiv:2005.01515. DOI: 10.48550/arXiv.2005.01515. arXiv: 2005.01515 [hep-ph].

Fabjan, D. et al. (Jan. 2010). "Simulating the effect of active galactic nuclei feedback on the metal enrichment of galaxy clusters". In: *MNRAS* 401.3, pp. 1670–1690. DOI: 10.1111/j.1365-2966.2009.15794.x. arXiv: 0909.0664 [astro-ph.CO].

Faucher-Giguère, C.-A. (Apr. 2020). "A cosmic UV/X-ray background model update". In: *Mon. Not. Roy. Astron. Soc.* 493.2, pp. 1614–1632. DOI: 10.1093/mnras/staa302. arXiv: 1903.08657 [astro-ph.CO].

Ferland, G. J. et al. (Oct. 2017). "The 2017 Release Cloudy". In: *Rev. Mexicana Astron. Astrofis.* 53, pp. 385–438. DOI: 10.48550/arXiv.1705.10877. arXiv: 1705.10877 [astro-ph.GA].

Fernandez, M. A., S. Bird, and M.-F. Ho (July 2024). "Cosmological constraints from the eBOSS Lyman-α forest using the PRIYA simulations". In: *JCAP* 2024.7, 029, p. 029. DOI: 10.1088/1475-7516/2024/07/029. arXiv: 2309.03943 [astro-ph.CO].

Fernández-Martínez, E., M. González-López, J. Hernández-García, M. Hostert, and J. López-Pavón (2023). "Effective portals to heavy neutral leptons". In: *JHEP* 09, p. 001. DOI: 10.1007/JHEP09(2023)001. arXiv: 2304.06772 [hep-ph].

Fienga, A., H. Manche, J. Laskar, M. Gastineau, and A. Verma (May 2014). "INPOP new release: INPOP13b". In: *arXiv e-prints*, arXiv:1405.0484, arXiv:1405.0484. DOI: 10.48550/arXiv.1405.0484. arXiv: 1405.0484 [astro-ph.EP].

Figueira, P. et al. (Aug. 2025). "A comprehensive study on radial velocity signals using ESPRESSO: Pushing precision to the 10 cm/s level". In: *A&A* 700, A174, A174. DOI: 10.1051/0004-6361/202553869. arXiv: 2507.07514 [astro-ph.EP].

Foot, R. and S. Vagnozzi (2015). "Dissipative hidden sector dark matter". In: *Phys. Rev. D* 91, p. 023512. DOI: 10.1103/PhysRevD.91.023512. arXiv: 1409.7174 [hep-ph].

Foreman-Mackey, D., D. W. Hogg, D. Lang, and J. Goodman (Mar. 2013). "emcee: The MCMC Hammer". In: *PASP* 125.925, p. 306. DOI: 10.1086/670067. arXiv: 1202.3665 [astro-ph.IM].

Fukuda, H. and K. Nakayama (2020). "Aspects of Nonlinear Effect on Black Hole Superradiance". In: *JHEP* 01, p. 128. DOI: 10.1007/JHEP01(2020)128. arXiv: 1910.06308 [hep-ph].

Fukugita, M. and P. J. E. Peebles (Dec. 2004). "The Cosmic Energy Inventory". In: *ApJ* 616.2, pp. 643–668. DOI: 10.1086/425155. arXiv: astro-ph/0406095 [astro-ph].

Gaia Collaboration et al. (May 2021). "Gaia Early Data Release 3. Summary of the contents and survey properties". In: *A&A* 649, A1, A1. DOI: 10.1051/0004-6361/202039657. arXiv: 2012.01533 [astro-ph.GA].

Gaia Collaboration et al. (June 2023). "Gaia Data Release 3. The extragalactic content". In: *A&A* 674, A41, A41. DOI: 10.1051/0004-6361/202243232. arXiv: 2206.05681 [astro-ph.GA].

Gaikwad, P., R. Srianand, T. R. Choudhury, and V. Khaire (May 2017). "VoIgt profile Parameter Estimation Routine (VIPER): H I photoionization rate at z < 0.5". In: *Mon. Not. Roy. Astron. Soc.* 467.3, pp. 3172–3187. DOI: 10.1093/mnras/stx248. arXiv: 1610.06572 [astro-ph.CO].

Gaikwad, P., R. Srianand, M. G. Haehnelt, and T. R. Choudhury (Sept. 2021). "A consistent and robust measurement of the thermal state of the IGM at $2 \leq z \leq 4$ from a large sample of Ly α forest spectra: evidence for late and rapid He II reionization". In: *MNRAS* 506.3, pp. 4389–4412. DOI: 10.1093/mnras/stab2017. arXiv: 2009.00016 [astro-ph.CO].

Garcia-Gallego, O., V. Iršič, M. G. Haehnelt, M. Viel, and J. S. Bolton (Aug. 2025). "Constraining mixed dark matter models with high-redshift Lyman-alpha forest data". In: *Phys. Rev. D* 112.4, 043502, p. 043502. DOI: 10.1103/4k29-h99l. arXiv: 2504.06367 [astro-ph.CO].

Gelman, A. and D. B. Rubin (Jan. 1992). "Inference from Iterative Simulation Using Multiple Sequences". In: *Statistical Science* 7, pp. 457–472. DOI: 10.1214/ss/1177011136.

Ghosh, D. and D. Sachdeva (2021). "Constraining Light Dark Photons from GW190517 and GW190426_152155". In: *Phys. Rev. D* 103.9, p. 095028. DOI: 10.1103/PhysRevD.103.095028. arXiv: 2102.08857 [astro-ph.HE].

Gnedin, N. Y. and L. Hui (May 1998). "Probing the Universe with the Lyalpha forest - I. Hydrodynamics of the low-density intergalactic medium". In: *MNRAS* 296.1, pp. 44–55. DOI: 10.1046/j.1365-8711.1998.01249.x. arXiv: astro-ph/9706219 [astro-ph].

Graham, P. W., J. Mardon, and S. Rajendran (2016). "Vector Dark Matter from Inflationary Fluctuations". In: *Phys. Rev.* D93.10, p. 103520. DOI: 10.1103/PhysRevD.93.103520. arXiv: 1504.02102 [hep-ph].

Grazian, A. et al. (Oct. 2024). "What Are the Pillars of Reionization? Revising the AGN Luminosity Function at $z \sim 5$". In: *ApJ* 974.1, 84, p. 84. DOI: 10.3847/1538-4357/ad6980. arXiv: 2407.20861 [astro-ph.CO].

Green, J. C., E. Wilkinson, and J. A. Morse (Feb. 2003). "The Cosmic Origins Spectrograph". In: *Future EUV/UV and Visible Space Astrophysics Missions and Instrumentation*. Ed. by J. C. Blades and O. H. W. Siegmund. Vol. 4854. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 72–80. DOI: 10.1117/12.459769.

Green, J. C. et al. (Jan. 2012). "The Cosmic Origins Spectrograph". In: *ApJ* 744.1, 60, p. 60. DOI: 10.1088/0004-637X/744/1/60. arXiv: 1110.0462 [astro-ph.IM].

Guarneri, F. et al. (Apr. 2024). "Fundamental physics with ESPRESSO: a new determination of the D/H ratio towards PKS1937-101". In: *MNRAS* 529.2, pp. 839–854. DOI: 10.1093/mnras/stae452. arXiv: 2402.05586 [astro-ph.CO].

Gunn, J. E. and B. A. Peterson (Nov. 1965). "On the Density of Neutral Hydrogen in Intergalactic Space." In: *ApJ* 142, pp. 1633–1636. DOI: 10.1086/148444.

Haardt, F. and P. Madau (Feb. 2012). "Radiative Transfer in a Clumpy Universe. IV. New Synthesis Models of the Cosmic UV/X-Ray Background". In: *ApJ* 746, 125, p. 125. DOI: 10.1088/0004-637X/746/2/125. arXiv: 1105.2039 [astro-ph.CO].

Haardt, F. and P. Madau (Apr. 1996). "Radiative Transfer in a Clumpy Universe. II. The Ultraviolet Extragalactic Background". In: *ApJ* 461, p. 20. DOI: 10.1086/177035. arXiv: astro-ph/9509093 [astro-ph].

Hasan, F. et al. (Nov. 2020). "Evolution of C IV Absorbers. I. The Cosmic Incidence". In: *ApJ* 904.1, 44, p. 44. DOI: 10.3847/1538-4357/abbe0b. arXiv: 2007.11751 [astro-ph.GA].

Higgs, P. (1964). "Broken symmetries, massless particles and gauge fields". In: *Physics Letters* 12.2, pp. 132–133. ISSN: 0031-9163. DOI: https://doi.org/10.1016/0031-9163(64)91136-9. URL: https://www.sciencedirect.com/science/article/pii/0031916364911369.

Hinton, S. R., T. M. Davis, C. Lidman, K. Glazebrook, and G. F. Lewis (Apr. 2016). "MARZ: Manual and automatic redshifting software". In: *Astronomy and Computing* 15, pp. 61–71. DOI: 10.1016/j.ascom.2016.03.001. arXiv: 1603.09438 [astro-ph.IM].

Holdom, B. (1986). "Two U(1)'s and Epsilon Charge Shifts". In: *Phys. Lett. B* 166, pp. 196–198. DOI: 10.1016/0370-2693(86)91377-8.

Holmberg, E. (1975). "Magnitudes, Colors, Surface Brightness, Intensity Distributions Absolute Luminosities, and Diameters of Galaxies". In: *Galaxies and the Universe*. Ed. by A. Sandage, M. Sandage, and J. Kristian, p. 123.

Hook, A. (2023). "New Solutions to the Gauge Hierarchy Problem". In: *Ann. Rev. Nucl. Part. Sci.* 73.1, pp. 23–39. DOI: 10.1146/annurev-nucl-102422-080830.

Hoyle, F. (Jan. 1946). "The synthesis of the elements from hydrogen". In: *MNRAS* 106, p. 343. DOI: 10.1093/mnras/106.5.343.

Hubble, E. P. (Dec. 1926). "Extragalactic nebulae." In: *ApJ* 64, pp. 321–369. DOI: 10.1086/143018.

Hui, L. and N. Y. Gnedin (Nov. 1997a). "Equation of state of the photoionized intergalactic medium". In: *Mon. Not. Roy. Astron. Soc.* 292, pp. 27–+.

Hui, L. and N. Y. Gnedin (Nov. 1997b). "Equation of state of the photoionized intergalactic medium". In: *MNRAS* 292.1, pp. 27–42. DOI: 10.1093/mnras/292.1.27. arXiv: astro-ph/9612232 [astro-ph].

Inoue, T. et al. (Feb. 2020). "The effect of our local motion on the Sandage-Loeb test of the cosmic expansion". In: *PASJ* 72.1, L1, p. L1. DOI: 10.1093/pasj/psz131. arXiv: 1911.01467 [astro-ph.CO].

Irsic, V. et al. (2017). "New Constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman-α forest data". In: *Physical Review D* 96, p. 023522. DOI: 10.1103/PhysRevD.96.023522.

Iršič, V., M. Viel, M. G. Haehnelt, J. S. Bolton, and G. D. Becker (2017). "First constraints on fuzzy dark matter from Lyman-α forest data and hydrodynamical simulations". In: *Phys. Rev. Lett.* 119.3, p. 031302. DOI: 10.1103/PhysRevLett.119.031302. arXiv: 1703.04683 [astro-ph.CO].

Ivanov, M. M., M. W. Toomey, and N. G. Kara çayl□□ (Mar. 2025). "Fundamental Physics with the Lyman-Alpha Forest: Constraints on the Growth of Structure and Neutrino Masses from SDSS with Effective Field Theory". In: *Phys. Rev. Lett.* 134 (9), p. 091001. DOI: 10.1103/PhysRevLett.134.091001. URL: https://link.aps.org/doi/10.1103/PhysRevLett.134.091001.

Janknecht, E., D. Reimers, S. Lopez, and D. Tytler (Nov. 2006). "The evolution of Lyman α absorbers in the redshift range $0.5 < z < 1.9$". In: *A&A* 458.2, pp. 427–439. DOI: 10.1051/0004-6361:20065372. arXiv: astro-ph/0608342 [astro-ph].

Jones, A. M., S. Noll, W. Kausch, C. Szyszka, and S. Kimeswenger (2013). "An advanced scattered moonlight model for Cerro Paranal". In: *A&A* 560, A91. DOI: 10.1051/0004-6361/201322433.

Kajino, T. et al. (July 2019). "Current status of r-process nucleosynthesis". In: *Progress in Particle and Nuclear Physics* 107, pp. 109–166. DOI: 10.1016/j.ppnp.2019.02.008. arXiv: 1906.05002 [astro-ph.HE].

Kamenshchik, A., U. Moschella, and V. Pasquier (July 2001). "An alternative to quintessence". In: *Physics Letters B* 511.2-4, pp. 265–268. DOI: 10.1016/S0370-2693(01)00571-8. arXiv: gr-qc/0103004 [gr-qc].

Kannan, R. et al. (2022). "Introducing the thesan project: radiation-magnetohydrodynamic simulations of the epoch of reionization". In: *Mon. Not. Roy. Astron. Soc.* 511.3, pp. 4005–4030. DOI: 10.1093/mnras/stab3710. arXiv: 2110.00584 [astro-ph.GA].

Käppeler, F., R. Gallino, S. Bisterzo, and W. Aoki (Jan. 2011). "The s process: Nuclear physics, stellar models, and observations". In: *Reviews of Modern Physics* 83.1, pp. 157–194. DOI: 10.1103/RevModPhys.83.157. arXiv: 1012.5218 [astro-ph.SR].

KATRIN Collaboration et al. (2025). "Direct neutrino-mass measurement based on 259 days of KATRIN data". In: *Science* 388.6743, pp. 180–185. DOI: 10.1126/science.adq9592. eprint: https://www.science.org/doi/pdf/10.1126/science.adq9592. URL: https://www.science.org/doi/abs/10.1126/science.adq9592.

Kauffmann, G. and S. D. M. White (Apr. 1993). "The merging history of dark matter haloes in a hierarchical universe." In: *MNRAS* 261, pp. 921–928. DOI: 10.1093/mnras/261.4.921.

Kim, T. .-., R. F. Carswell, S. Cristiani, S. D'Odorico, and E. Giallongo (Sept. 2002). "The physical properties of the Lyα forest at z > 1.5". In: *MNRAS* 335.3, pp. 555–573. DOI: 10.1046/j.1365-8711.2002.05599.x. arXiv: astro-ph/0205237 [astro-ph].

Kim, T. .-. et al. (Mar. 2021). "The evolution of the low-density H I intergalactic medium from z = 3.6 to 0: data, transmitted flux, and H I column density". In: *Mon. Not. Roy. Astron. Soc.* 501.4, pp. 5811–5833. DOI: 10.1093/mnras/staa3844. arXiv: 2012.05861 [astro-ph.CO].

Kobayashi, T., R. Murgia, A. De Simone, V. Iršič, and M. Viel (2017). "Lyman-α constraints on ultralight scalar dark matter: Implications for the early and late universe". In: *Phys. Rev. D* 96.12, p. 123514. DOI: 10.1103/PhysRevD.96.123514. arXiv: 1708.00015 [astro-ph.CO].

Koptelova, E., T. Chiueh, W. P. Chen, and H. H. Chan (June 2014). "New near-infrared observations and lens-model constraints for UM673". In: *A&A* 566, A36, A36. DOI: 10.1051/0004-6361/201322230. arXiv: 1307.2390 [astro-ph.CO].

Koptelova, E., V. L. Oknyanskij, B. P. Artamonov, and O. Burkhonov (Feb. 2010). "Intrinsic quasar variability and time delay determination in the lensed quasar UM673". In: *MNRAS* 401.4, pp. 2805–2815. DOI: 10.1111/j.1365-2966.2009.15872.x.

Kulkarni, G., G. Worseck, and J. F. Hennawi (Sept. 2019). "Evolution of the AGN UV luminosity function from redshift 7.5". In: *Mon. Not. Roy. Astron. Soc.* 488.1, pp. 1035–1065. DOI: 10.1093/mnras/stz1493. arXiv: 1807.09774 [astro-ph.GA].

Lee, C.-C., J. K. Webb, D. Milaković, and R. F. Carswell (Oct. 2021). "Non-uniqueness in quasar absorption models and implications for measurements of the fine structure constant". In: *MNRAS* 507.1, pp. 27–42. DOI: 10.1093/mnras/stab2005. arXiv: 2102.11648 [astro-ph.CO].

Lewis, A., A. Challinor, and A. Lasenby (Aug. 2000). "Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models". In: *ApJ* 538.2, pp. 473–476. DOI: 10.1086/309179. arXiv: astro-ph/9911177 [astro-ph].

Liske, J. et al. (May 2008). "Cosmic dynamics in the era of Extremely Large Telescopes". In: *MNRAS* 386.3, pp. 1192–1218. DOI: 10.1111/j.1365-2966.2008.13090.x. arXiv: 0802.1532 [astro-ph].

Liu, H., W. Qin, G. W. Ridgway, and T. R. Slatyer (2021). "Lyman-α constraints on cosmic heating from dark matter annihilation and decay". In: *Phys. Rev. D* 104.4, p. 043514. DOI: 10.1103/PhysRevD.104.043514. arXiv: 2008.01084 [astro-ph.CO].

Liu, H., T. R. Slatyer, and J. Zavala (2016). "Contributions to cosmic reionization from dark matter annihilation and decay". In: *Phys. Rev. D* 94.6, p. 063507. DOI: 10.1103/PhysRevD.94.063507. arXiv: 1604.02457 [astro-ph.CO].

Loeb, A. (June 1998). "Direct Measurement of Cosmological Parameters from the Cosmic Deceleration of Extragalactic Objects". In: *ApJ* 499.2, pp. L111–L114. DOI: 10.1086/311375. arXiv: astro-ph/9802122 [astro-ph].

Long, A. J. and L.-T. Wang (2019). "Dark Photon Dark Matter from a Network of Cosmic Strings". In: *Phys. Rev. D* 99.6, p. 063529. DOI: 10.1103/PhysRevD.99.063529. arXiv: 1901.03312 [hep-ph].

López, S. et al. (Oct. 2016). "XQ-100: A legacy survey of one hundred $3.5 < z < 4.5$ quasars observed with VLT/X-shooter". In: *A&A* 594, A91, A91. DOI: 10.1051/0004-6361/201628161. arXiv: 1607.08776.

Louis, T. et al. (2025). "The Atacama Cosmology Telescope: DR6 power spectra, likelihoods and ΛCDM parameters". In: *JCAP* 11, p. 062. DOI: 10.1088/1475-7516/2025/11/062. arXiv: 2503.14452 [astro-ph.CO].

Lyman, T. (Apr. 1906). "The Spectrum of Hydrogen in the Region of Extremely Short Wave-Lengths". In: *ApJ* 23, p. 181. DOI: 10.1086/141330.

MacAlpine, G. M. and F. R. Feldman (Oct. 1982). "Discovery and spectrophotometry of high-redshift quasars." In: *ApJ* 261, pp. 412–421. DOI: 10.1086/160353.

Madau, P. and F. Haardt (Nov. 2015). "Cosmic Reionization after Planck: Could Quasars Do It All?" In: *ApJ* 813, L8, p. L8. DOI: 10.1088/2041-8205/813/1/L8. arXiv: 1507.07678.

Madau, P. and M. Dickinson (2014). "Cosmic star formation history". In: *Ann. Rev. Astron. Astrophys.* 52, pp. 415–486. DOI: 10.1146/annurev-astro-081811-125615. arXiv: 1403.0007 [astro-ph.CO].

Madau, P., A. Ferrara, and M. J. Rees (July 2001). "Early Metal Enrichment of the Intergalactic Medium by Pregalactic Outflows". In: *ApJ* 555.1, pp. 92–105. DOI: 10.1086/321474. arXiv: astro-ph/0010158 [astro-ph].

Madau, P., F. Haardt, and M. J. Rees (Apr. 1999). "Radiative Transfer in a Clumpy Universe. III. The Nature of Cosmological Ionizing Sources". In: *ApJ* 514.2, pp. 648–659. DOI: 10.1086/306975. arXiv: astro-ph/9809058 [astro-ph].

Maitra, S., R. Srianand, and P. Gaikwad (Jan. 2022). "Measurement of redshift-space two- and three-point correlation of Lyα absorbers at $1.7 < z < 3.5$: implications on

evolution of the physical properties of IGM". In: *MNRAS* 509.1, pp. 1536–1556. DOI: 10.1093/mnras/stab3053. arXiv: 2107.05651 [astro-ph.CO].

Marconi, A. et al. (Aug. 2022). "ANDES, the high resolution spectrograph for the ELT: science case, baseline design and path to construction". In: *Ground-based and Airborne Instrumentation for Astronomy IX*. Ed. by C. J. Evans, J. J. Bryant, and K. Motohara. Vol. 12184. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1218424, p. 1218424. DOI: 10.1117/12.2628689.

Martins, C. J. A. P. et al. (2024). "Cosmology and fundamental physics with the ELT-ANDES spectrograph". In: *Exper. Astron.* 57.1, p. 5. DOI: 10.1007/s10686-024-09928-w. arXiv: 2311.16274 [astro-ph.CO].

Martins, C. J. A. P., M. A. F. Melo e Sousa, S. Q. Fernandes, and C. M. J. Marques (Feb. 2025). "Impact of spatial curvature on forecast constraints from standard and differential redshift drift measurements". In: *Physics of the Dark Universe* 47, 101775, p. 101775. DOI: 10.1016/j.dark.2024.101775. arXiv: 2412.07532 [astro-ph.CO].

Martins, C. J. A. P. et al. (June 2022). "Fundamental physics with ESPRESSO: Constraints on Bekenstein and dark energy models from astrophysical and local probes". In: *Phys. Rev. D* 105.12, 123507, p. 123507. DOI: 10.1103/PhysRevD.105.123507. arXiv: 2205.13848 [astro-ph.CO].

McDermott, S. D. and S. J. Witte (2020). "Cosmological Evolution of Light Dark Photon Dark Matter". In: *Phys. Rev. D* 101.6, p. 063030. DOI: 10.1103/PhysRevD.101.063030. arXiv: 1911.05086 [hep-ph].

McDonald, P. et al. (Nov. 2000). "The Observed Probability Distribution Function, Power Spectrum, and Correlation Function of the Transmitted Flux in the Lyα Forest". In: *ApJ* 543.1, pp. 1–23. DOI: 10.1086/317079. arXiv: astro-ph/9911196 [astro-ph].

McQuinn, M. (Sept. 2016). "The Evolution of the Intergalactic Medium". In: *Ann. Rev. Astron. Astrophys.* 54, pp. 313–362. DOI: 10.1146/annurev-astro-082214-122355. arXiv: 1512.00086 [astro-ph.CO].

Milaković, D. and P. Jethwa (Apr. 2024a). "A new method for instrumental profile reconstruction of high-resolution spectrographs". In: *A&A* 684, A38, A38. DOI: 10.1051/0004-6361/202348532. arXiv: 2311.05240 [astro-ph.IM].

Milaković, D. et al. (Oct. 2024b). "Isotopic abundance of carbon in the DLA towards QSO B1331+170". In: *MNRAS* 534.1, pp. 12–29. DOI: 10.1093/mnras/stae2056. arXiv: 2407.17953 [astro-ph.GA].

Mirizzi, A., J. Redondo, and G. Sigl (Mar. 2009). "Microwave background constraints on mixing of photons with hidden photons". In: *JCAP* 2009.3, 026, p. 026. DOI: 10.1088/1475-7516/2009/03/026. arXiv: 0901.0014 [hep-ph].

Mo, H., F. van den Bosch, and S. White (2010). *Galaxy Formation and Evolution*. Galaxy Formation and Evolution. Cambridge University Press. ISBN: 9780521857932. URL: https://books.google.it/books?id=Zj7fDU3Z4wsC.

Muñoz, J. A. et al. (June 1998). "The Castles Project". In: *Ap&SS* 263, pp. 51–54. DOI: 10.1023/A:1002120921330. arXiv: astro-ph/9902131 [astro-ph].

Muñoz, J. B. and A. Loeb (Nov. 2017). "Constraints on dark matter-baryon scattering from the temperature evolution of the intergalactic medium". In: *JCAP* 2017.11, 043, p. 043. DOI: 10.1088/1475-7516/2017/11/043. arXiv: 1708.08923 [astro-ph.CO].

Murphy, M. T. et al. (Feb. 2022). "Fundamental physics with ESPRESSO: Precise limit on variations in the fine-structure constant towards the bright quasar HE 0515−4414". In: *A&A* 658, A123, A123. DOI: 10.1051/0004-6361/202142257. arXiv: 2112.05819 [astro-ph.CO].

Navarro, J. F., C. S. Frenk, and S. D. M. White (May 1996). "The Structure of Cold Dark Matter Halos". In: *ApJ* 462, p. 563. DOI: 10.1086/177173. arXiv: astro-ph/9508025 [astro-ph].

Nelson, D. et al. (May 2019). "The IllustrisTNG simulations: public data release". In: *Computational Astrophysics and Cosmology* 6.1, 2, p. 2. DOI: 10.1186/s40668-019-0028-x. arXiv: 1812.05609 [astro-ph.GA].

Nielsen, N. M., C. W. Churchill, G. G. Kacprzak, and M. T. Murphy (Oct. 2013). "MAGIICAT I. The Mg II Absorber-Galaxy Catalog". In: *ApJ* 776.2, 114, p. 114. DOI: 10.1088/0004-637X/776/2/114. arXiv: 1304.6716 [astro-ph.CO].

Noll, S. et al. (2012). "An atmospheric radiation model for Cerro Paranal. I. The optical spectral range". In: *A&A* 543, A92. DOI: 10.1051/0004-6361/201219040.

Noll, S. et al. (July 2014). "Skycorr: A general tool for spectroscopic sky subtraction". In: *A&A* 567, A25, A25. DOI: 10.1051/0004-6361/201423908. arXiv: 1405.3679 [astro-ph.IM].

O'Meara, J. M. et al. (Oct. 2015). "The First Data Release of the KODIAQ Survey". In: *AJ* 150.4, 111, p. 111. DOI: 10.1088/0004-6256/150/4/111. arXiv: 1505.03529 [astro-ph.CO].

O'Meara, J. M. et al. (Sept. 2017). "The Second Data Release of the KODIAQ Survey". In: *AJ* 154.3, 114, p. 114. DOI: 10.3847/1538-3881/aa82b8. arXiv: 1707.07905 [astro-ph.GA].

Ocvirk, P. et al. (2016). "Cosmic Dawn (CoDa): the First Radiation-Hydrodynamics Simulation of Reionization and Galaxy Formation in the Local Universe". In: *Mon. Not. Roy. Astron. Soc.* 463.2, pp. 1462–1485. DOI: 10.1093/mnras/stw2036. arXiv: 1511.00011 [astro-ph.GA].

Onken, C. A. et al. (Oct. 2024). "SkyMapper Southern Survey: Data release 4". In: *PASA* 41, e061, e061. DOI: 10.1017/pasa.2024.53. arXiv: 2402.02015 [astro-ph.CO].

Oppenheimer, B. D. and R. Davé (June 2008). "Mass, metal, and energy feedback in cosmological simulations". In: *MNRAS* 387.2, pp. 577–600. DOI: 10.1111/j.1365-2966.2008.13280.x. arXiv: 0712.1827 [astro-ph].

Padovani, P. et al. (Aug. 2017). "Active galactic nuclei: what's in a name?" In: *A&A Rev.* 25.1, 2, p. 2. DOI: 10.1007/s00159-017-0102-9. arXiv: 1707.07134 [astro-ph.GA].

Padovani, P. and M. Cirasuolo (Jan. 2023). "The Extremely Large Telescope". In: *Contemporary Physics* 64.1, pp. 47–64. DOI: 10.1080/00107514.2023.2266921. arXiv: 2312.04299 [astro-ph.IM].

Pakmor, R. et al. (Sept. 2023). "The MillenniumTNG Project: the hydrodynamical full physics simulation and a first look at its galaxy clusters". In: *MNRAS* 524.2, pp. 2539–2555. DOI: 10.1093/mnras/stac3620. arXiv: 2210.10060 [astro-ph.CO].

Palanque-Delabrouille, N. et al. (Nov. 2015). "Neutrino masses and cosmology with Lyman-alpha forest power spectrum". In: *JCAP* 2015.11, pp. 011–011. DOI: 10.1088/1475-7516/2015/11/011. arXiv: 1506.05976 [astro-ph.CO].

Palanque-Delabrouille, N. et al. (Apr. 2020). "Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman-α and Planck full-survey data". In: *JCAP* 2020.4, 038, p. 038. DOI: 10.1088/1475-7516/2020/04/038. arXiv: 1911.09073 [astro-ph.CO].

Parsa, S., J. S. Dunlop, R. J. McLure, and A. Mortlock (Mar. 2016). "The galaxy UV luminosity function at z $\simeq$ 2-4; new results on faint-end slope and the evolution of luminosity density". In: *MNRAS* 456.3, pp. 3194–3211. DOI: 10.1093/mnras/stv2857. arXiv: 1507.05629 [astro-ph.GA].

Passegger, V. M. et al. (Apr. 2024). "The compact multi-planet system GJ 9827 revisited with ESPRESSO□". In: *A&A* 684, A22, A22. DOI: 10.1051/0004-6361/202348592. arXiv: 2401.06276 [astro-ph.EP].

Peeples, M. S., D. H. Weinberg, R. Davé, M. A. Fardal, and N. Katz (May 2010). "Pressure support versus thermal broadening in the Lyman α forest - I. Effects of the equation of state on longitudinal structure". In: *Mon. Not. Roy. Astron. Soc.* 404.3, pp. 1281–1294. DOI: 10.1111/j.1365-2966.2010.16383.x. arXiv: 0910.0256 [astro-ph.CO].

Pepe, F. et al. (Jan. 2014). "ESPRESSO: The next European exoplanet hunter". In: *Astronomische Nachrichten* 335.1, p. 8. DOI: 12.1002/asna.201312004.

Pepe, F. et al. (Jan. 2021a). "ESPRESSO at VLT. On-sky performance and first results". In: *Astronomy & Astrophysics* 645, A96, A96. DOI: 10.1051/0004-6361/202038306. arXiv: 2010.00316 [astro-ph.IM].

Pepe, F. et al. (Jan. 2021b). "ESPRESSO at VLT. On-sky performance and first results". In: *A&A* 645, A96, A96. DOI: 10.1051/0004-6361/202038306. arXiv: 2010.00316 [astro-ph.IM].

Perlmutter, S. et al. (1999). "Measurements of Ω and Λ from 42 High Redshift Supernovae". In: *Astrophys. J.* 517, pp. 565–586. DOI: 10.1086/307221. arXiv: astro-ph/9812133.

Péroux, C., R. G. McMahon, L. J. Storrie-Lombardi, and M. J. Irwin (Dec. 2003). "The evolution of Ω_{HI} and the epoch of formation of damped Lyman α absorbers". In: *MNRAS* 346.4, pp. 1103–1115. DOI: 10.1111/j.1365-2966.2003.07129.x. arXiv: astro-ph/0107045 [astro-ph].

Péroux, C. and J. C. Howk (Aug. 2020). "The Cosmic Baryon and Metal Cycles". In: *ARA&A* 58, pp. 363–406. DOI: 10.1146/annurev-astro-021820-120014. arXiv: 2011.01935 [astro-ph.GA].

Peterson, B. M. et al. (June 1998). "On Uncertainties in Cross-Correlation Lags and the Reality of Wavelength-dependent Continuum Lags in Active Galactic Nuclei". In: *PASP* 110.748, pp. 660–670. DOI: 10.1086/316177. arXiv: astro-ph/9802103 [astro-ph].

Planck Collaboration et al. (Sept. 2020). "Planck 2018 results. VI. Cosmological parameters". In: *A&A* 641, A6, A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

Press, W. H. and P. Schechter (Feb. 1974). "Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation". In: *ApJ* 187, pp. 425–438. DOI: 10.1086/152650.

Prochaska, J. X., J. M. O'Meara, and G. Worseck (July 2010). "A Definitive Survey for Lyman Limit Systems at z ∼3.5 with the Sloan Digital Sky Survey". In: *ApJ* 718.1, pp. 392–416. DOI: 10.1088/0004-637X/718/1/392. arXiv: 0912.0292 [astro-ph.CO].

Puchwein, E. et al. (2023). "The Sherwood–Relics simulations: overview and impact of patchy reionization and pressure smoothing on the intergalactic medium". In: *Mon. Not. Roy. Astron. Soc.* 519.4, pp. 6162–6183. DOI: 10.1093/mnras/stac3761. arXiv: 2207.13098 [astro-ph.CO].

Puchwein, E., F. Haardt, M. G. Haehnelt, and P. Madau (May 2019). "Consistent modelling of the meta-galactic UV background and the thermal/ionization history of the intergalactic medium". In: *Mon. Not. Roy. Astron. Soc.* 485.1, pp. 47–68. DOI: 10.1093/mnras/stz222. arXiv: 1801.04931 [astro-ph.GA].

Puchwein, E. et al. (July 2015). "The photoheating of the intergalactic medium in synthesis models of the UV background". In: *Mon. Not. Roy. Astron. Soc.* 450.4, pp. 4081–4097. DOI: 10.1093/mnras/stv773. arXiv: 1410.1531 [astro-ph.CO].

Rahmati, A., A. H. Pawlik, M. Raicevic, and J. Schaye (Apr. 2013). "On the evolution of the H I column density distribution in cosmological simulations". In: *Mon. Not. Roy. Astron. Soc.* 430, pp. 2427–2445. DOI: 10.1093/mnras/stt066. arXiv: 1210.7808 [astro-ph.CO].

Ramos Almeida, C. et al. (Apr. 2011). "Testing the Unification Model for Active Galactic Nuclei in the Infrared: Are the Obscuring Tori of Type 1 and 2 Seyferts Different?" In: *ApJ* 731.2, 92, p. 92. DOI: 10.1088/0004-637X/731/2/92. arXiv: 1101.3335 [astro-ph.CO].

Ratra, B. and P. J. E. Peebles (1988). "Cosmological Consequences of a Rolling Homogeneous Scalar Field". In: *Phys. Rev. D* 37, p. 3406. DOI: 10.1103/PhysRevD.37.3406.

Rauch, M. (Jan. 1998). "The Lyman Alpha Forest in the Spectra of QSOs". In: *ARA&A* 36, pp. 267–316. DOI: 10.1146/annurev.astro.36.1.267. arXiv: astro-ph/9806286 [astro-ph].

Rauch, M., W. L. W. Sargent, and T. A. Barlow (June 2001a). "Small-Scale Structure at High Redshift. II. Physical Properties of the C IV Absorbing Clouds". In: *ApJ* 554.2, pp. 823–840. DOI: 10.1086/321402. arXiv: astro-ph/0104216 [astro-ph].

Rauch, M., W. L. W. Sargent, T. A. Barlow, and R. F. Carswell (Nov. 2001b). "Small-Scale Structure at High Redshift. III. The Clumpiness of the Intergalactic Medium on Subkiloparsec Scales". In: *ApJ* 562.1, pp. 76–87. DOI: 10.1086/323523. arXiv: astro-ph/0107516 [astro-ph].

Rauch, M. et al. (Oct. 2005). "Expansion and Collapse in the Cosmic Web". In: *ApJ* 632.1, pp. 58–80. DOI: 10.1086/432904. arXiv: astro-ph/0509262 [astro-ph].

Riess, A. G. et al. (Sept. 1998). "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant". In: *AJ* 116.3, pp. 1009–1038. DOI: 10.1086/300499. arXiv: astro-ph/9805201 [astro-ph].

Rogers, K. K. and H. V. Peiris (2021). "General framework for cosmological dark matter bounds using N-body simulations". In: *Phys. Rev. D* 103.4, p. 043526. DOI: 10.1103/PhysRevD.103.043526. arXiv: 2007.13751 [astro-ph.CO].

Rorai, A. et al. (Nov. 2016). "Exploring the thermal state of the low-density intergalactic medium at z = 3 with an ultrahigh signal-to-noise QSO spectrum". In: *Mon. Not. Roy. Astron. Soc.* 466.3, pp. 2690–2709. ISSN: 0035-8711. DOI: 10.1093/mnras/stw2917. eprint: https://academic.oup.com/mnras/article-pdf/466/3/2690/10901806/stw2917.pdf. URL: https://doi.org/10.1093/mnras/stw2917.

Rorai, A. et al. (Apr. 2017). "Exploring the thermal state of the low-density intergalactic medium at z = 3 with an ultrahigh signal-to-noise QSO spectrum". In: *MNRAS* 466.3, pp. 2690–2709. DOI: 10.1093/mnras/stw2917. arXiv: 1611.03805 [astro-ph.CO].

Rubin, V. C., W. K. Ford Jr., and N. Thonnard (June 1980). "Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc)." In: *ApJ* 238, pp. 471–487. DOI: 10.1086/158003.

Rusu, C. E. et al. (July 2019). "A search for gravitationally lensed quasars and quasar pairs in Pan-STARRS1: spectroscopy and sources of shear in the diamond 2M1134-2103". In: *MNRAS* 486.4, pp. 4987–5007. DOI: 10.1093/mnras/stz1142. arXiv: 1803.07175 [astro-ph.GA].

Saitta, F. et al. (Mar. 2008). "Tracing the gas at redshift 1.7-3.5 with the Lyα forest: the FLO approach". In: *MNRAS* 385.1, pp. 519–530. DOI: 10.1111/j.1365-2966.2008.12860.x. arXiv: 0712.2452 [astro-ph].

Sandage, A. (Sept. 1962). "The Change of Redshift and Apparent Luminosity of Galaxies due to the Deceleration of Selected Expanding Universes." In: *ApJ* 136, p. 319. DOI: 10.1086/147385.

Sartori, L. F. et al. (Feb. 2023). "L 363-38 b: A planet newly discovered with ESPRESSO orbiting a nearby M dwarf star". In: *A&A* 670, A42, A42. DOI: 10.1051/0004-6361/202244347. arXiv: 2210.12710 [astro-ph.EP].

Schaye, J. (Oct. 2001). "Model-independent Insights into the Nature of the Lyα Forest and the Distribution of Matter in the Universe". In: *ApJ* 559.2, pp. 507–515. DOI: 10.1086/322421. arXiv: astro-ph/0104272 [astro-ph].

Schaye, J., T. Theuns, M. Rauch, G. Efstathiou, and W. L. W. Sargent (Nov. 2000). "The thermal history of the intergalactic medium*". In: *Mon. Not. Roy. Astron. Soc.* 318.3, pp. 817–826. DOI: 10.1046/j.1365-8711.2000.03815.x. arXiv: astro-ph/9912432 [astro-ph].

Schaye, J. et al. (Jan. 2015). "The EAGLE project: simulating the evolution and assembly of galaxies and their environments". In: *MNRAS* 446.1, pp. 521–554. DOI: 10.1093/mnras/stu2058. arXiv: 1407.7040 [astro-ph.GA].

Schmidt, T. et al. (Jan. 2023). "STRIDES: automated uniform models for 30 quadruply imaged quasars". In: *MNRAS* 518.1, pp. 1260–1300. DOI: 10.1093/mnras/stac2235. arXiv: 2206.04696 [astro-ph.CO].

Schmidt, T. M. (Aug. 2024). "Accurate estimate of the ESPRESSO fibre-injection losses inferred from integrated field-stabilization images". In: *MNRAS* 532.2, pp. 1635–1645. DOI: 10.1093/mnras/stae1477. arXiv: 2406.07193 [astro-ph.IM].

Schmidt, T. M. et al. (Feb. 2021). "Fundamental physics with ESPRESSO: Towards an accurate wavelength calibration for a precision test of the fine-structure constant". In: *A&A* 646, A144, A144. DOI: 10.1051/0004-6361/202039345. arXiv: 2011.13963 [astro-ph.IM].

Shin, E.-J., J.-H. Kim, and B. K. Oh (Aug. 2021). "How Metals Are Transported Into and Out of a Galactic Disk: Dependence on Hydrodynamic Schemes in Numerical Simulations". In: *ApJ* 917.1, 12, p. 12. DOI: 10.3847/1538-4357/abffd0. arXiv: 2106.04640 [astro-ph.GA].

Silva, A. M. et al. (July 2022). "A novel framework for semi-Bayesian radial velocities through template matching". In: *A&A* 663, A143, A143. DOI: 10.1051/0004-6361/202142262. arXiv: 2205.00067 [astro-ph.EP].

Simcoe, R. A., W. L. W. Sargent, and M. Rauch (May 2004). "The Distribution of Metallicity in the Intergalactic Medium at z$_\sim$2.5: O VI and C IV Absorption in the Spectra of Seven QSOs". In: *ApJ* 606.1, pp. 92–115. DOI: 10.1086/382777. arXiv: astro-ph/0312467 [astro-ph].

Smette, A. et al. (Apr. 1992). "A Spectroscopic Study of UM 673 A and B: On the Size of Lyman-Alpha Clouds". In: *ApJ* 389, p. 39. DOI: 10.1086/171187.

Smette, A. et al. (Oct. 1995). "The gravitational lens candidate HE 1104-1805 and the size of absorption systems." In: *A&AS* 113, p. 199.

Smoot, G. F. et al. (Sept. 1992). "Structure in the COBE Differential Microwave Radiometer First-Year Maps". In: *ApJ* 396, p. L1. DOI: 10.1086/186504.

Springel, V. (Dec. 2005). "The cosmological simulation code GADGET-2". In: *Mon. Not. Roy. Astron. Soc.* 364.4, pp. 1105–1134. DOI: 10.1111/j.1365-2966.2005.09655.x. arXiv: astro-ph/0505010 [astro-ph].

Springel, V., R. Pakmor, O. Zier, and M. Reinecke (Sept. 2021). "Simulating cosmic structure formation with the GADGET-4 code". In: *MNRAS* 506.2, pp. 2871–2949. DOI: 10.1093/mnras/stab1855. arXiv: 2010.03567 [astro-ph.IM].

Suárez Mascareño, A. et al. (July 2020). "Revisiting Proxima with ESPRESSO". In: *A&A* 639, A77, A77. DOI: 10.1051/0004-6361/202037745. arXiv: 2005.12114 [astro-ph.EP].

Surdej, J. et al. (Oct. 1987). "A new case of gravitational lensing". In: *Nature* 329.6141, pp. 695–696. DOI: 10.1038/329695a0.

Surdej, J. et al. (June 1988). "Observations of the new gravitational lens system UM 673 = Q 0142-100." In: *A&A* 198, pp. 49–60.

Tabernero, H. M. et al. (Feb. 2021). "ESPRESSO high-resolution transmission spectroscopy of WASP-76 b". In: *A&A* 646, A158, A158. DOI: 10.1051/0004-6361/202039511. arXiv: 2011.12197 [astro-ph.EP].

Tejos, N. (Apr. 2025). "The Intergalactic Medium". In: *arXiv e-prints*, arXiv:2504.12539, arXiv:2504.12539. DOI: 10.48550/arXiv.2504.12539. arXiv: 2504.12539 [astro-ph.CO].

Telikova, K. N., P. S. Shternin, and S. A. Balashev (Dec. 2019). "Thermal State of the Intergalactic Medium at z ~ 2-4". In: *ApJ* 887.2, 205, p. 205. DOI: 10.3847/1538-4357/ab52fe. arXiv: 1910.13184 [astro-ph.CO].

Tepper Garcia, T. (2006). "Voigt profile fitting to quasar absorption lines: a simple approximation to the Voigt-Hjerting function". In: *Mon. Not. Roy. Astron. Soc.* 369, pp. 2025–2035. DOI: 10.1111/j.1365-2966.2006.10450.x. arXiv: astro-ph/0602124.

Teyssier, R. (Apr. 2002). "Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES". In: *A&A* 385, pp. 337–364. DOI: 10.1051/0004-6361:20011817. arXiv: astro-ph/0111367 [astro-ph].

Theuns, T., A. Leonard, G. Efstathiou, F. R. Pearce, and P. A. Thomas (Dec. 1998). "P3M-SPH simulations of the Lyalpha forest". In: *Mon. Not. Roy. Astron. Soc.* 301.2, pp. 478–502. DOI: 10.1046/j.1365-8711.1998.02040.x. arXiv: astro-ph/9805119 [astro-ph].

Trost, A. et al. (July 2025). "The ESPRESSO Redshift Drift Experiment: I. High-resolution spectra of the Lyman-α forest of QSO J052915.80-435152.0". In: *A&A* 699, A159, A159. DOI: 10.1051/0004-6361/202554502. arXiv: 2505.21615 [astro-ph.CO].

Tyson, J. A., F. Valdes, and R. A. Wenk (Jan. 1990). "Detection of Systematic Gravitational Lens Galaxy Image Alignments: Mapping Dark Matter in Galaxy Clusters". In: *ApJ* 349, p. L1. DOI: 10.1086/185636.

Tytler, D. (Oct. 1987). "The Redshift Distribution of QSO Lyman-Alpha Absorption Systems". In: *ApJ* 321, p. 69. DOI: 10.1086/165616.

Urry, C. M. and P. Padovani (Sept. 1995). "Unified Schemes for Radio-Loud Active Galactic Nuclei". In: *PASP* 107, p. 803. DOI: 10.1086/133630. arXiv: astro-ph/9506063 [astro-ph].

Vale, A. and J. P. Ostriker (Sept. 2006). "The non-parametric model for linking galaxy luminosity with halo/subhalo mass". In: *MNRAS* 371.3, pp. 1173–1187. DOI: 10.1111/j.1365-2966.2006.10605.x. arXiv: astro-ph/0511816 [astro-ph].

Vernet, J. et al. (Dec. 2011). "X-shooter, the new wide band intermediate resolution spectrograph at the ESO Very Large Telescope". In: *A&A* 536, A105, A105. DOI: 10.1051/0004-6361/201117752. arXiv: 1110.1944 [astro-ph.IM].

Viel, M., M. G. Haehnelt, and V. Springel (Nov. 2004). "Inferring the dark matter power spectrum from the Lyman α forest in high-resolution QSO absorption spectra". In: *Mon. Not. Roy. Astron. Soc.* 354, pp. 684–694. DOI: 10.1111/j.1365-2966.2004.08224.x.

Viel, M., M. G. Haehnelt, and V. Springel (Apr. 2006). "Testing the accuracy of the hydrodynamic particle-mesh approximation in numerical simulations of the Lyman α forest". In: *MNRAS* 367.4, pp. 1655–1665. DOI: 10.1111/j.1365-2966.2006.10063.x. arXiv: astro-ph/0504641 [astro-ph].

Viel, M., J. Lesgourgues, M. G. Haehnelt, S. Matarrese, and A. Riotto (2005). "Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman-alpha forest". In: *Phys. Rev. D* 71, p. 063534. DOI: 10.1103/PhysRevD.71.063534. arXiv: astro-ph/0501562.

Viel, M. et al. (May 2017). "Diagnosing galactic feedback with line broadening in the low-redshift Lyα forest". In: *Mon. Not. Roy. Astron. Soc.* 467.1, pp. L86–L90. DOI: 10.1093/mnrasl/slx004. arXiv: 1610.02046 [astro-ph.GA].

Villasenor, B., B. Robertson, P. Madau, and E. Schneider (July 2023). "New constraints on warm dark matter from the Lyman-α forest power spectrum". In: *Phys. Rev. D* 108.2, 023502, p. 023502. DOI: 10.1103/PhysRevD.108.023502. arXiv: 2209.14220 [astro-ph.CO].

Vogt, S. S. et al. (June 1994). "HIRES: the high-resolution echelle spectrometer on the Keck 10-m Telescope". In: *Instrumentation in Astronomy VIII*. Ed. by D. L. Crawford and E. R. Craine. Vol. 2198. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 362. DOI: 10.1117/12.176725.

Walther, M., J. Oñorbe, J. F. Hennawi, and Z. Lukić (Feb. 2019). "New Constraints on IGM Thermal Evolution from the Lyα Forest Power Spectrum". In: *ApJ* 872.1, 13, p. 13. DOI: 10.3847/1538-4357/aafad1. arXiv: 1808.04367 [astro-ph.CO].

Walther, M. et al. (May 2025). "Emulating the Lyman-Alpha forest 1D power spectrum from cosmological simulations: new models and constraints from the eBOSS measurement". In: *JCAP* 2025.5, 099, p. 099. DOI: 10.1088/1475-7516/2025/05/099. arXiv: 2412.05372 [astro-ph.CO].

Wang, F. et al. (Jan. 2021). "A Luminous Quasar at Redshift 7.642". In: *ApJ* 907.1, L1, p. L1. DOI: 10.3847/2041-8213/abd8c6. arXiv: 2101.03179 [astro-ph.GA].

Webb, J. K. et al. (Jan. 2025). "The Mystery of Alpha and the Isotopes". In: *MNRAS*. DOI: 10.1093/mnrasl/slaf012. arXiv: 2401.00888 [astro-ph.CO].

Weinberg, S. (Jan. 1989). "The cosmological constant problem". In: *Rev. Mod. Phys.* 61 (1), pp. 1–23. DOI: 10.1103/RevModPhys.61.1. URL: https://link.aps.org/doi/10.1103/RevModPhys.61.1.

Weinberger, R., V. Springel, and R. Pakmor (2020). "The Arepo public code release". In: *Astrophys. J. Suppl.* 248.2, p. 32. DOI: 10.3847/1538-4365/ab908c. arXiv: 1909.04667 [astro-ph.IM].

Weymann, R. J. and C. B. Foltz (Sept. 1983). "Common Lyman-alpha absorption lines in the triple QSO PG 1115+08." In: *ApJ* 272, pp. L1–L4. DOI: 10.1086/184105.

White, S. D. M. and C. S. Frenk (Sept. 1991). "Galaxy formation through hierarchical clustering". In: *ApJ* 379, pp. 52–79. DOI: 10.1086/170483.

Wilks, S. (1938). "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses". In: *Annals Math. Statist.* 9.1, pp. 60–62. DOI: 10.1214/aoms/1177732360.

Witte, S. J., S. Rosauro-Alcaraz, S. D. McDermott, and V. Poulin (June 2020). "Dark photon dark matter in the presence of inhomogeneous structure". In: *Journal of High Energy Physics* 2020.6, 132, p. 132. DOI: 10.1007/JHEP06(2020)132. arXiv: 2003.13698 [astro-ph.CO].

Wolf, C. et al. (Apr. 2024). "The accretion of a solar mass per day by a 17-billion solar mass black hole". In: *Nature Astronomy* 8, pp. 520–529. DOI: 10.1038/s41550-024-02195-x. arXiv: 2402.15101 [astro-ph.CO].

Wright, J. T. and J. D. Eastman (2014). "Barycentric Corrections at 1 cm/s for Precise Radial Velocity Work". In: *Publications of the Astronomical Society of the Pacific* 126.938, pp. 838–852. DOI: 10.1086/677192. URL: https://iopscience.iop.org/article/10.1086/677192.

Yèche, C., N. Palanque-Delabrouille, J. Baur, and H. du Mas des Bourboux (June 2017). "Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100". In: *JCAP* 2017.6, 047, p. 047. DOI: 10.1088/1475-7516/2017/06/047. arXiv: 1702.03314 [astro-ph.CO].

Zanutta, A. et al. (Feb. 2023). "CUBES phase a design overview". In: *Experimental Astronomy* 55.1, pp. 241–265. DOI: 10.1007/s10686-022-09837-w. arXiv: 2203.15352 [astro-ph.IM].

Zechmeister, M. et al. (June 2020). *SERVAL: SpEctrum Radial Velocity AnaLyser*. Astrophysics Source Code Library, record ascl:2006.011.

Zhao, D. H., Y. P. Jing, H. J. Mo, and G. Börner (Dec. 2009). "Accurate Universal Models for the Mass Accretion Histories and Concentrations of Dark Matter Halos". In: *ApJ* 707.1, pp. 354–369. DOI: 10.1088/0004-637X/707/1/354. arXiv: 0811.0828 [astro-ph].