Chronicles of cosmic dawn: high redshift quasars as probes of supermassive black holes and the intergalactic medium

Onorato, S.

1 Introduction

What were the first light sources in our Universe? How did they emerge from a sea of neutral hydrogen and begin to shape the cosmic structures we observe today? And how did the earliest supermassive black holes form and grow so quickly in such a young cosmos? These are among the most fundamental questions in modern astrophysics, cutting to the heart of our attempts to understand how the Universe evolved from a nearly uniform state to the richly structured web of galaxies, stars, and black holes that we see in the present. While the cosmic microwave background offers a snapshot of the infant Universe, and local galaxies reveal its mature state, the first billion years of cosmic history remain one of the least understood epochs. During this formative era, the intergalactic medium transitioned from a cold, neutral state to a hot, ionized one in a process known as *cosmic reionization*. In this thesis, we focus on one of the most powerful beacons from the early Universe: *high-redshift quasars*, actively accreting supermassive black holes that illuminate their surroundings and serve as unique probes of the intergalactic medium during the Epoch of Reionization. By studying their spectra, we can extract clues about the ionization state of the Universe, the lifetimes of quasars, and the processes governing black hole growth at cosmic dawn.

1.1 The Universe in the first billion years

According to our current understanding (see Fig. 1.1), the Universe began roughly 13.8 billion years ago with a space-time singularity called the Big Bang. Before even a second had passed, it is believed to have experienced a brief but extraordinary period of exponential expansion, the cosmic inflation (Ellis & Wands 2023). This expansion stretched quantum fluctuations to macroscopic scales, seeding the density variations that would eventually grow into cosmic structures. In the first few minutes, conditions were extreme enough for nuclear reactions to take place: protons and neutrons fused to form the first light elements in a process known as primordial nucleosynthesis (Fields & Sarkar 2006; Cyburt et al. 2016), producing mostly hydrogen and helium, with small traces of deuterium, lithium, and beryllium. However, the Universe remained too hot and ionized for these nuclei to bind with electrons and form neutral atoms. Photons were constantly scattered off free electrons, keeping the cosmos opaque.

Only after about 380 000 years, the Universe cooled sufficiently for electrons to combine with nuclei and form stable atoms, primarily hydrogen. This period, known as recombination, marked a pivotal transition: with

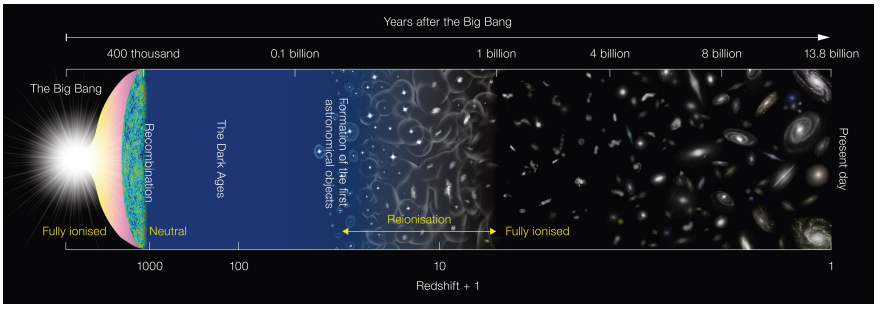

Figure 1.1: Summary of the key stages in the evolution of the Universe since the Big Bang, approximately 13.8 billion years ago (not to scale). Around 400 000 years after the Big Bang, the Universe became neutral and remained so throughout the cosmic dark ages. It was only with the emergence of the first stars that ultraviolet radiation began to ionize the surrounding hydrogen. Within a few hundred million years, this process progressed until the intergalactic medium was almost entirely ionized. *Credit:* NAOJ.

electrons now bound into atoms, the Universe became transparent to light. The photons released during this epoch have traveled largely undisturbed ever since, forming the cosmic microwave background (CMB; Penzias & Wilson 1965), the earliest observable light, still detectable today as a nearly perfect blackbody radiation.

Following recombination, the Universe entered a long, dark phase. With no stars or galaxies yet formed, the cosmos was filled with neutral gas, primarily hydrogen, and devoid of any luminous sources. These cosmic dark ages lasted for several hundred million years, during which the Universe continued to expand and cool. Throughout, small fluctuations in the dark matter density (White et al. 1994; Scott et al. 1995), amplified by gravity, gradually drew gas into the first star-forming regions.

Eventually, the densest pockets of gas collapsed under their own gravity, igniting nuclear fusion in their cores and giving rise to the first generation of stars (Bromm & Larson 2004; Bromm 2013), marking the onset of the cosmic dawn. These early stars were likely massive and short-lived, emitting large amounts of ultraviolet (UV) radiation. Over time, they assembled into the first galaxies (Bromm & Yoshida 2011), which began to shape their surroundings through radiation and feedback processes. Around this epoch, supermassive black holes (SMBHs), often residing at the centers of galaxies (Lynden-Bell 1969; Kormendy & Richstone 1995), also started to form. Importantly, the energetic photons from these sources began to ionize the surrounding neutral hydrogen, initiating the process known as cosmic reionization (see Section 1.1.1).

Reionization represents the next major phase transition of the Universe.

As UV light from early stars and galaxies propagated outward, it carved out growing ionized bubbles in the neutral intergalactic medium (IGM). These bubbles expanded and merged, gradually transforming the IGM from a cold, opaque sea of neutral atoms into a hot, transparent plasma of ionized gas. By the time the Universe was about a billion years old, most of the hydrogen in the IGM had been reionized, completing this global phase change (Fan et al. 2006c). A second reionization, that of helium, would follow later, driven by harder radiation from even more energetic sources such as quasars (see Section 1.3), which became prominent at later times (McQuinn 2016).

This narrative of cosmic evolution is encapsulated within the *standard cosmological model*, ΛCDM (Peebles 1993; Planck Collaboration et al. 2020). This model assumes that the Universe is flat, homogeneous, and isotropic on scales larger than about 100 megaparsecs (Mpc[1]), and explains its expansion history through a specific energy budget: roughly 70% dark energy, 25% cold dark matter (CDM), and 5% baryonic matter, which includes atoms, molecules, and radiation (Tanabashi et al. 2018). In this framework, dark matter supplies the gravitational potential wells that enable structure formation, while dark energy drives the accelerated expansion of the Universe (Riess et al. 1998; Perlmutter et al. 1999; Spergel et al. 2003). Despite its success in explaining a wide range of observations, the physical nature of both dark matter and dark energy remains one of the most profound open questions in modern cosmology. At the same time, many details about the first billion years, such as the sources and timing of reionization and the rapid emergence of supermassive black holes, are still uncertain.

1.1.1 The Epoch of Reionization

The Epoch of Reionization (EoR) denotes the cosmic interval during which the intergalactic hydrogen transitioned from a mostly neutral state to an almost fully ionized one, as a consequence of ionizing radiation produced by the first luminous sources. A convenient global measure of the progress of reionization is the volume-averaged ionized hydrogen fraction $Q_{\rm HII}(t)$, which evolves from $Q_{\rm HII} \simeq 0$ (neutral IGM) to $Q_{\rm HII} \simeq 1$ (ionized IGM). The time evolution of this quantity can be captured in a simple, widely used balance equation that compares the rate at which ionizing photons are supplied to the IGM with the rate at which free electrons and protons recombine (Madau et al. 1999):

$$\frac{dQ_{\rm HII}}{dt} = \frac{\langle \dot{n}_{\rm ion} \rangle}{\langle n_{\rm H} \rangle} - \frac{Q_{\rm HII}}{\bar{t}_{\rm rec}}. \tag{1.1}$$

The first term on the right-hand side is the ionizing photon production rate per hydrogen atom, set by the nature and abundance of sources and by

[1] Namely, $1\,{\rm pc} \simeq 3.09 \times 10^{18}\,{\rm cm}$.

the fraction of ionizing photons that escape galaxies $f_{\rm esc}$. The term $\bar{t}_{\rm rec}$ is an effective recombination timescale that depends on the IGM temperature and on small-scale density inhomogeneity (Kaurov & Gnedin 2015; Gnedin & Madau 2022).

Understanding which sources powered reionization remains an area of active study. Current evidence supports a dominant role for star-forming galaxies (Robertson et al. 2015), especially numerous faint systems (Bouwens et al. 2012), supplemented by contributions from X-ray sources and, possibly, a minority population of active galactic nuclei (AGN; see Section 1.2) in some models (Haardt & Madau 2012; Xu et al. 2014). The total ionizing photon production rate per unit volume contributed by galaxies (Mesinger 2016) can be parametrized as:

$$\dot{n}_{\rm ion} \propto \rho_{\rm UV} \xi_{\rm ion} f_{\rm esc}, \tag{1.2}$$

where $\rho_{\rm UV}$ is the galaxy UV luminosity density measured from the galaxy's UV luminosity function within a specific magnitude range, $\xi_{\rm ion}$ is the number of ionizing photons yield per unit UV luminosity, and $f_{\rm esc}$ is the already defined escape fraction. Recent estimates place $\xi_{\rm ion}$ around $\simeq 10^{25.3}$ erg^{-1} Hz (e.g., Tang et al. 2019; Naidu et al. 2022), while observational and theoretical constraints on $f_{\rm esc}$ are still broad ($f_{\rm esc} \lesssim 50\%$; e.g., Mascia et al. 2023; Endsley et al. 2023), and highly model dependent.

An independent way to estimate the total ionizing emissivity[2] ϵ is by exploiting its connection to the photo-ionization rate $\Gamma_{\rm bg}$ and the mean free path of ionizing photons $\lambda_{\rm mfp}$ via the relation $\Gamma_{\rm bg} \propto \epsilon \lambda_{\rm mfp}$. Both $\Gamma_{\rm bg}$ and $\lambda_{\rm mfp}$ can be constrained from quasar spectra, and have been measured in several studies (e.g., Worseck et al. 2014; D'Aloisio et al. 2018; Becker et al. 2021; Gaikwad et al. 2023).

Recombinations act as an important sink of ionizing photons. Dense, self-shielded structures, such as Lyman-Limit Systems (LLSs, with column densities $N_{\rm HI} \simeq 10^{17-20}$ cm^{-2}), can shorten the effective mean free path of ionizing photons and slow reionization locally; their abundance and structure must therefore be modelled or inferred from observations and simulations (e.g., McQuinn et al. 2011; Davies & Furlanetto 2016; Kulkarni et al. 2019).

Reionization is expected to be *patchy*: ionized bubbles first form around ionizing sources, and then grow and merge following an "inside-out" topology (Furlanetto et al. 2004), so that overdense regions typically reionize earlier than the deepest voids. The degree of patchiness encodes information about the spectral sources and their hardness (X-ray photons are thought to have produced more extended and smoother ionization; see e.g., Chen & Miralda-Escudé 2004), and about the spatial distribution of sinks.

[2]The total number of ionizing photons emitted per unit volume per unit time.

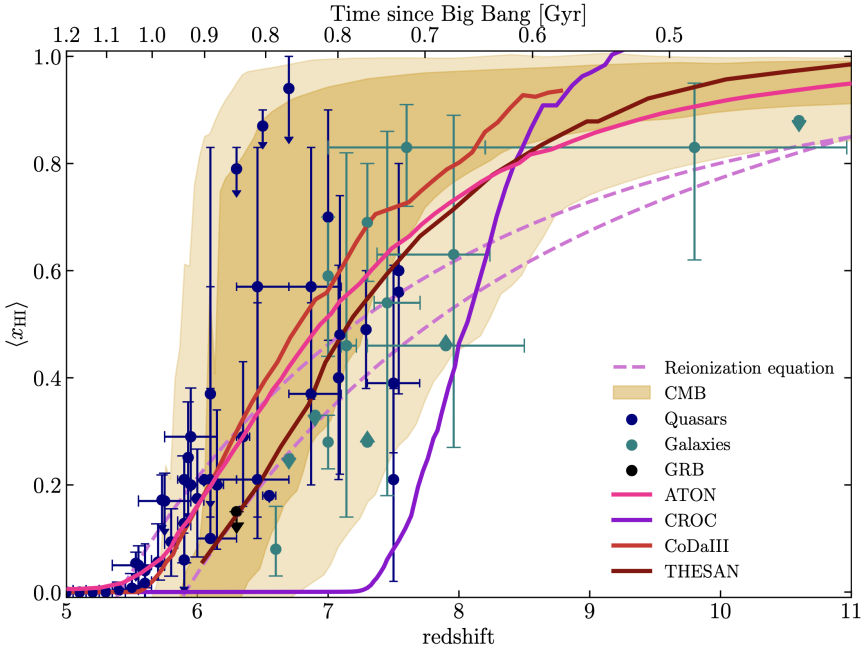

Figure 1.2: Evolution of the volume-averaged neutral hydrogen fraction $\langle x_{\rm HI} \rangle$, from $z = 11$ to $z = 5$. The shaded bands indicate the 1σ and 2σ constraints from CMB measurements (Planck Collaboration et al. 2020). Observational points are grouped by probe. Quasar-derived constraints are plotted in blue; they come from damping wings (Greig et al. 2017, 2019, 2022, 2024a; Bañados et al. 2018; Davies et al. 2018a; Wang et al. 2020; Yang et al. 2020a; Ďurovčíková et al. 2024), dark-pixels (McGreer et al. 2015; Jin et al. 2023), dark-gap statistics (Zhu et al. 2021, 2022), and Lyα opacity cumulative distribution function (Gaikwad et al. 2023). Galaxy-based constraints are shown in green, and are taken from studies of Lyα emitters (Mason et al. 2019; Goto et al. 2021; Morales et al. 2021; Wold et al. 2022; Bruton et al. 2023), Lyman-break galaxies (Bolan et al. 2022), and damping wings (Umeda et al. 2024). Finally, a recent gamma-ray burst (GRB) damping wing measurement is plotted in black (Fausey et al. 2025). Solid model curves display the neutral fraction evolution predicted by several radiative transfer simulations: i.e., ATON (pink; Kulkarni et al. 2019), CROC (purple; Gnedin & Kaurov 2014), CoDaIII (red; Lewis et al. 2022), and THESAN (brown; Garaldi et al. 2022). The dashed violet curves show predictions based on equation 1.1, for fixed and redshift-dependent ionizing photon production rate and effective recombination timescale (see Section 1.1.1). *Credit:* Sindhu Satyavolu.

Observational constraints on the timing and topology of reionization come from multiple, complementary probes (see Fig. 1.2). Measurements of Thomson scattering optical depth $\tau_{\rm re}$ from CMB constrain the integrated column of free electrons produced by reionization. Planck Collaboration et al. (2020) report $\tau_{\rm re} \simeq 0.054 \pm 0.007$, implying a midpoint reionization $z \sim 7.7 \pm 0.7$, which favors a relatively late and fairly rapid process (McQuinn 2016). High-redshift quasar and galaxy spectra provide direct IGM probes through distinct properties, such as damping wings (e.g., Davies et al. 2018b; Greig et al. 2022; Ďurovčíková et al. 2024; Umeda et al. 2024), dark pixels (McGreer et al. 2015; Jin et al. 2023), or dark gaps (Zhu et al. 2021, 2022). These studies indicate that IGM becomes largely ionized by $z \simeq 5-6$, with large fluctuations at $z \gtrsim 6$. Recent measurements of the ionizing mean free path and the Lyman-α (Lyα) opacity support a picture in which the end of hydrogen reionization occurs around $z \simeq 5.3$ (e.g., Becker et al. 2015b; McGreer et al. 2015; Bosman et al. 2022).

Theoretical studies play a central role in advancing our understanding of reionization. Given its inherently multiscale nature, the process is explored through a hierarchy of methods. Fully coupled cosmological radiation-hydrodynamics simulations provide the most complete physical description, but are computationally expensive and limited in volume and parameter exploration. Post-processing radiative transfer on large N-body/hydrodynamic simulations, hybrid schemes, and semi-numerical models are therefore widely used to explore parameter space and generate large mock datasets for comparison with observations (e.g., Iliev et al. 2006; Mesinger et al. 2011; Kulkarni et al. 2019; Gnedin & Madau 2022).

Finally, another unique window into the EoR is the 21 cm hyperfine transition of neutral hydrogen. Measurements of the global 21 cm signal, the 21 cm power spectrum, and 21 cm absorption toward bright radio sources (the 21 cm forest) can probe different aspects of the thermal and ionization history, and the topology of reionization (e.g., Furlanetto et al. 2006; Morales & Wyithe 2010; Pritchard & Loeb 2012; Šoltinský et al. 2021).

1.2 Active Galactic Nuclei

As already introduced in Section 1.1, at the centers of most massive galaxies reside SMBHs with masses ranging from millions to billions of solar masses (e.g., Kormendy & Richstone 1995; Kormendy & Gebhardt 2001; Häring & Rix 2004). When gas is accreted onto them, it forms a rotating accretion disk and releases a substantial fraction of its rest-mass energy as radiation. This process produces an active galactic nucleus (AGN), one of the most luminous phenomena in the Universe (e.g., Fabian 1999), which can outshine its host galaxy across the electromagnetic spectrum, from radio to X-rays. In this section, we describe the main structural components com-

monly invoked in AGN models, the physical processes that occur in each of them, and the observational features through which these components are identified. A schematic showing these components is reported in Fig. 1.3.

1.2.1 Black hole length scales and Eddington limit

We define the gravitational radius and the Schwarzschild radius (Schwarzschild 1916) of a black hole of mass M, respectively, as:

$$r_g \equiv \frac{GM}{c^2}, \tag{1.3}$$

$$r_s \equiv 2r_g \simeq 9.57 \times 10^{-6} \left(\frac{M}{10^8\,\mathrm{M}_\odot} \right) \mathrm{pc}, \tag{1.4}$$

where $G = 6.67 \times 10^{-8}\ \mathrm{cm^3\,g^{-1}\,s^{-2}}$ is Newton's gravitational constant, and $c \simeq 3 \times 10^{10}\ \mathrm{cm\,s^{-1}}$ is the speed of light in vacuum. For reference, $\mathrm{M}_\odot = 1.99 \times 10^{33}$ g is the solar mass. A frequently used physical limit is the *Eddington luminosity* (Eddington 1926), the luminosity at which outward radiation pressure (for electron scattering opacity) balances gravity:

$$L_{\mathrm{Edd}} = \frac{4\pi G M m_p c}{\sigma_T} \simeq 1.26 \times 10^{38} \left(\frac{M}{\mathrm{M}_\odot} \right) \mathrm{erg\ s^{-1}}, \tag{1.5}$$

where $m_p = 1.67 \times 10^{-24}$ g is the proton mass, and $\sigma_T = 6.65 \times 10^{-25}\ \mathrm{cm^2}$ is the Thomson scattering cross section. The *Eddington ratio* is defined as $\lambda_{\mathrm{Edd}} \equiv L_{\mathrm{bol}}/L_{\mathrm{Edd}}$ and provides a simple measure of the accretion power relative to this limit. Here, L_{bol} is the bolometric luminosity, i.e., the total radiative power emitted across all wavelengths by the accretion process. The accretion mode is tightly connected to λ_{Edd} and to the structure of the inner accretion flow. These concepts are further expanded in Section 1.4.2.

1.2.2 Main pieces of the AGN puzzle

Black hole (BH): The central SMBH is described by its mass M and by a dimensionless spin parameter a (the angular momentum per unit mass, with $0 \leq a \leq 1$). The value of the spin sets the location of the innermost stable circular orbit (ISCO) and therefore affects the radiative efficiency of accretion and the ability of the system to produce relativistic jets (e.g., Bardeen et al. 1972; Blandford & Znajek 1977; Blandford & Payne 1982). Observational SMBH mass estimates rely on virial methods that use broad-line widths together with size estimates from reverberation mapping (see e.g., Kaspi et al. 2000, 2005; Peterson et al. 2004), or on single-epoch calibrations based on the radius–luminosity relation (e.g., Vestergaard & Peterson 2006; Bentz et al. 2013). Alternatively, masses can be inferred from host galaxy scaling relations (e.g., Magorrian et al. 1998; Ferrarese & Merritt 2000; Häring & Rix 2004; Kormendy & Ho 2013).

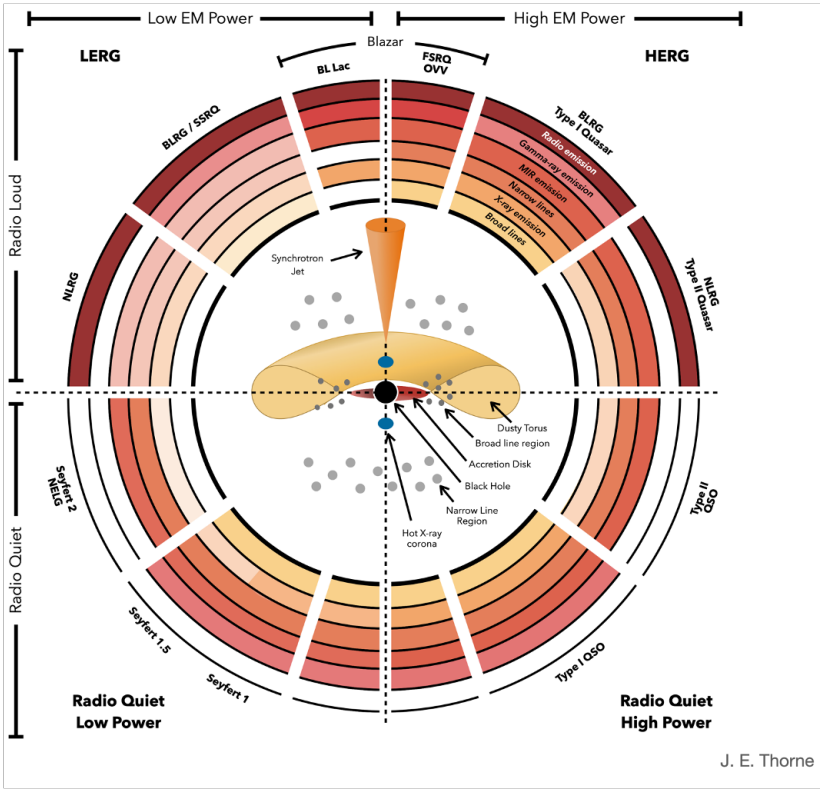

Figure 1.3: Schematic illustration of the AGN orientation-based unified model. The observed class of AGN depends on the line of sight, on whether the central engine drives powerful jets (radio-loud vs radio-quiet), and on the accretion power of the SMBH (low vs high electromagnetic power). The central inset shows the typical components of an AGN, although the geometry of many of these structures remains uncertain. Bipolar jets are typically associated with radio-loud systems. "Radio-quiet" does not imply an absence of radio emission: many hosts display low-level radio flux consistent with star formation, so it is not always clear if weak jets are present. The upper left and upper right panels correspond to low and high excitation radio galaxies (LERG and HERG, respectively). Common class names are included, such as Seyferts, broad line radio galaxy (BLRG), narrow line radio galaxy (NLRG), narrow emission line galaxy (NELG), flat spectrum radio quasar (FSRQ), steep spectrum radio quasar (SSRQ), optically violent variables (OVV), and quasi-stellar objects (QSO). The labels "quasar" and "QSO" are frequently used interchangeably, and do not always indicate radio properties uniquely. Concentric rings surrounding the schematic indicate which combinations of power, jet activity, and geometry favour the production of broad or narrow emission lines, or emission in the mid-infrared, radio, X-ray, or gamma-ray bands; the opacity of each ring reflects the relative strength or likelihood of the corresponding emission. *Credit:* Thorne, J., et al. AGN Unification Diagram. Zenodo, 24 marzo 2022.

Accretion disk: Gas with angular momentum settles into an accretion flow around the SMBH. For luminous, radiatively efficient AGN, the flow is commonly described by the standard thin disk model (Shakura & Sunyaev 1973), which behaves as a multi-temperature blackbody from a few r_g out to $\sim 10^2 - 10^3\, r_g$ and produces the UV/optical "big blue bump" in the spectral energy distribution (SED; Elvis et al. 1994; Richards et al. 2006). The accretion disk supplies the ionizing continuum that photoionizes the broad and narrow line gas.

X-ray corona: Immediately above the inner accretion disk, we can find the *corona*, a region of hot ($T \sim 10^8$ K) electrons which Compton up-scatters the UV/optical disk photons, producing the observed hard X-ray power-law continuum. X-ray reflection signatures, including the Fe Kα line, and rapid X-ray variability are powerful diagnostics of the corona's size, location, and coupling to the inner disk. Models of disk–corona Comptonization provide a framework to interpret these observables (e.g., Haardt & Maraschi 1991).

Broad Line Region (BLR): The BLR is a dense, relatively compact ensemble of clouds located at radii of order pc from the SMBH in luminous AGN, where their scales increase with luminosity following an approximately $R \propto L^{1/2}$ relation established by reverberation mapping. Gas in the BLR is photoionized by the accretion disk continuum, and moves with velocities of order $v \sim 10^3 - 10^4$ km s^{-1}, producing Doppler-broadened permitted emission lines (e.g., Mg II, C IV, Hβ) used as virial BH mass estimators. As already said, reverberation mapping provides the size-luminosity relation for the BLR and is the basis for single-epoch virial black hole mass estimation (e.g., Vestergaard & Peterson 2006; Bentz et al. 2013).

Narrow Line Region (NLR): At larger radii (pc to kpc scales), lower-density gas produces narrow forbidden and permitted emission lines with typical widths $v \sim 10^2 - 10^3$ km s^{-1}. The NLR traces the host ISM photoionized by the AGN and often displays spatially resolved structures such as ionization cones. Because of its extent, the NLR encodes information about AGN influence on galactic scales and can be used to study feedback and the connection between the nucleus and the host galaxy (e.g., Ferland & Osterbrock 1986).

The dusty torus: Many AGN host a geometrically and optically thick circumnuclear distribution of dust and gas, commonly referred to as the "torus", which surrounds the accretion disk and the BLR. The torus efficiently absorbs UV/optical radiation from the nucleus and re-emits it thermally in the mid- and far-infrared (IR), producing a characteristic "IR

bump" in the AGN SED. In unified models, the torus provides orientation-dependent obscuration: sightlines intersecting the torus yield obscured (Type 2) spectra, while direct sightlines to the nucleus produce unobscured (Type 1) spectra with broad emission lines (e.g., Antonucci 1993; Urry & Padovani 1995).

Warm absorbers and winds: A large fraction of AGN exhibits blue-shifted absorption in the UV and soft X-ray bands, signatures of ionized outflows often called "warm absorbers". Typical warm absorbers show velocities of a few $10^2 - 10^3$ km s^{-1}, ionization stratification, and column densities in the approximate range $N_{\rm H} \sim 10^{20} - 10^{23}$ cm^{-2}. These winds, likely launched from the accretion disk or innermost torus by radiation or magnetic forces, can carry substantial mass, momentum, and energy and are a key ingredient of AGN feedback models (e.g., Murray et al. 1995; Crenshaw et al. 2003).

Jets, lobes and large-scale radio structures: In a subset of AGN, a portion of the accretion energy and black hole rotational energy is converted into highly collimated, relativistic jets that emerge perpendicular to the accretion disk. Jets emit synchrotron radiation (from radio up to X-rays) and can produce high-energy emission via inverse Compton scattering; on tens of kpc to Mpc scales, jets inflate radio lobes, drive shocks into the circumgalactic or intracluster medium, and form hotspots where they terminate. Relativistic beaming[3] causes Doppler boosting of jet emission when the jet is closely aligned with the observer's line of sight (i.e., small viewing angles), producing the *blazar* class (see Section 1.2.3 and e.g., Urry & Padovani 1995). Jet–environment interactions are a major channel of mechanical AGN feedback, capable of heating halo gas and regulating star formation on galactic and cluster scales (e.g., Begelman et al. 1984; Blandford & Znajek 1977; McNamara & Nulsen 2007).

1.2.3 A few beasts in the "AGN zoo"

AGN exhibit an extraordinary diversity of observational properties, often described as the "AGN zoo" (see Fig. 1.3 and e.g., Padovani et al. 2017). This diversity can be broadly understood as the result of three main physical axes: the instantaneous accretion power (or Eddington ratio, see Section 1.2.1), the presence and strength of relativistic jets, and the viewing angle of the observer with respect to the system's symmetry axis. Within this framework, a wide variety of AGN subclasses have been identified, reflecting the complexity of the central engine and its interplay with the host galaxy.

[3]The set of special relativistic effects that concentrate and amplify radiation emitted by a source moving toward the observer at speeds close to the speed of light.

The full picture is therefore much richer than what can be captured in a short overview, with many additional categories and nuances that remain the subject of ongoing research. In the following, we summarize some of the most representative classes with their properties.

Seyfert galaxies: Seyfert galaxies (Seyfert 1943) are relatively low- luminosity ($L_{\rm bol} \simeq 10^{43} - 10^{45}\,{\rm erg\ s^{-1}}$) AGN typically found in nearby spiral hosts. In the orientation-based unification picture (Antonucci 1993; Urry & Padovani 1995), Seyfert I objects display both broad and narrow emission lines because the line of sight intercepts the nucleus directly, whereas Seyfert II objects appear dominated by narrow lines when the central regions are obscured by the torus.

Radio galaxies: When powerful relativistic jets inflate extended radio structures, the large-scale radio morphology is commonly described by the Fanaroff–Riley (FR; Fanaroff & Riley 1974) classes. FR I sources ($L_{\rm bol} \simeq 10^{41} - 10^{44}\,{\rm erg\ s^{-1}}$) are edge-darkened with brightest jets near the core and typically correspond to lower radio power, whereas FR II sources ($L_{\rm bol} \simeq 10^{44} - 10^{47}\,{\rm erg\ s^{-1}}$) are edge-brightened with terminal hotspots and generally higher jet power. The FR dichotomy correlates with environment and jet kinetic power and is an important way to connect central engine properties to feedback on galaxy/cluster scales.

Blazars, beamed jet sources: When a relativistic jet is closely aligned with the line of sight, Doppler boosting strongly enhances the jet emission, and the AGN appears as a blazar (Urry & Padovani 1995; Padovani et al. 2017). These objects are observationally split into:

- Flat-Spectrum Radio-loud Quasars (FSRQs): they show strong broad emission lines and high luminosity ($L_{\rm bol} \simeq 10^{46} - 10^{48}\,{\rm erg\ s^{-1}}$), with jet-dominated SEDs.

- BL Lacertae (BL Lac): they have weak or absent emission lines, with highly variable continuum dominated by the jet, and are generally less luminous ($L_{\rm bol} \simeq 10^{44} - 10^{46}\,{\rm erg\ s^{-1}}$).

Relativistic beaming explains rapid variability, high apparent brightness, and often high-energy gamma-ray emission.

Quasars: Quasars (or Quasi-Stellar Objects, QSOs) are the most luminous manifestation of accreting SMBHs, characterized by bright UV continua and broad permitted emission lines (e.g., Schmidt 1963; Matthews & Sandage 1963; Greenstein 1963). A representative luminosity range for

classical quasars is $L_{\rm bol} \simeq 10^{45} - 10^{47}$ erg s^{-1}, with the most extreme objects exceeding 10^{47} erg s^{-1}. They commonly accrete at moderate to high Eddington ratios and can be detected to very large cosmological distances (i.e., high-redshift, such as $z \gtrsim 7$), which makes them powerful probes of the early Universe. To use them for observational studies of the IGM, accurate systemic redshifts are essential. They are commonly obtained from low-ionization broad lines (e.g., Mg II), narrow rest-frame optical/infrared forbidden lines when available, or from sub-mm/far-IR lines of the host galaxy (e.g., [C II], CO), each with their different typical accuracies.

Historically, quasars have been separated into two populations according to their radio-to-optical flux ratio ($F(5\,{\rm GHz})/{\rm F_R}$; where $F(5\,{\rm GHz})$ is the flux at 5 GHz, and F_R is the flux in the R band, centered at ≈ 658 nm):

- Radio-loud ($F(5\,{\rm GHz})/{\rm F_R} \geq 10$): they launch powerful jets and often show extended radio structures.
- Radio-quiet ($F(5\,{\rm GHz})/{\rm F_R} < 10$): they either lack such jets or host weak/compact jets that do not dominate the radio output.

Roughly, $\sim 10\% - 20\%$ of optically selected quasars are classified as radio-loud in bright samples (e.g., Kellermann et al. 1989). The radio-loud fraction may depend on luminosity, host properties, and redshift, with radio-bright quasars becoming increasingly rare at the highest redshifts.

Based on optical emission line properties and accretion signatures, radio AGN can also be divided into low- and high-excitation classes, respectively:

- LERGs: they are associated with low $\lambda_{\rm Edd}$, radiatively inefficient accretion, and energetic output dominated by jets (the so-called "radio/jet mode").
- HERGs: they are radiatively efficient, higher-$\lambda_{\rm Edd}$ systems, with strong ionizing continua ("quasar/radiative mode").

This dichotomy provides a useful physical classification that links accretion physics to observed emission line spectra and radio morphology (e.g., Alexander & Hickox 2012; Best & Heckman 2012; Heckman & Best 2014).

Luminous quasars have been discovered at redshifts $z \gtrsim 6 - 7$, implying the existence of SMBHs with masses of order $10^8 - 10^9\,{\rm M_\odot}$ when the Universe was less than 1 Gyr old. These objects, therefore, provide tight constraints on black hole seeding and early growth. They act as bright background sources whose spectra[4] encode the ionization and thermal state of the IGM via several features, such as the Lyα forest (Fan et al. 2006c; Becker et al. 2015b), Gunn-Peterson troughs (Gunn & Peterson 1965), damping wings (Miralda-Escudé 1998), and proximity zones (Fan et al. 2006c).

[4]We focus on the rest-frame UV/optical spectra that include the Lyα emission line and the spectral regions both redward and blueward of Lyα, where IGM absorption is the strongest.

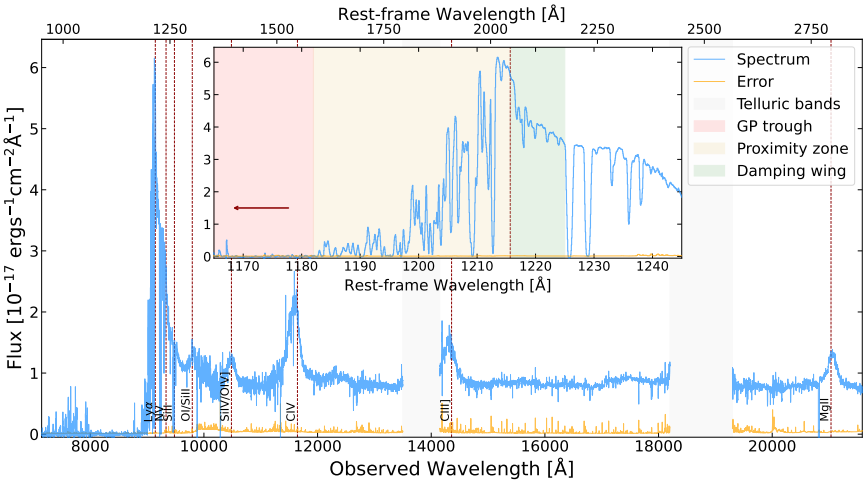

Figure 1.4: Spectrum 1D of the quasar J022426.540−471129.400 (Reed et al. 2017) at $z = 6.5222 \pm 0.0001$ (Wang et al. 2021b), obtained with VLT/X-Shooter (Vernet et al. 2011), and included in the sample of Onorato et al. (2025b). The main panel shows flux as a function of both the observed near-IR (bottom axis) and rest-frame UV wavelengths (top axis) in Å. The spectrum (light blue) and its uncertainty (orange) are smoothed for display purposes. Shaded grey bands mark strong telluric absorption at $[13\,500, 14\,150]$ Å and $[18\,200, 19\,300]$ Å. Red dashed vertical lines indicate the main rest-frame UV emission lines at the systemic redshift. The inset panel shows a zoom on the rest-frame $[1165, 1245]$ Å interval, where we highlight with red, yellow, and green bands the wavelength regions in which we can typically observe the Gunn-Peterson trough, the quasar proximity zone, and the damping wing, respectively. The leftward red arrow indicates that the Gunn-Peterson trough extends blueward (to shorter rest-frame wavelengths).

1.3 High-redshift quasars

High-redshift ($z \gtrsim 6$) quasars remain the most powerful class of background sources for studying the IGM during and immediately after the epoch of hydrogen reionization. Their extreme UV radiation illuminates the intervening gas and creates ionized regions around the sources themselves. The rest-frame UV spectra (see example in Fig. 1.4) act as dual diagnostics: they constrain the quasar's continuum and emission line regions and, simultaneously, probe the integrated ionization, temperature, and metal content of every absorber along the line of sight. Below, we summarize the observational diagnostics most relevant to studies of the EoR, emphasizing their physical interpretation and the main modeling caveats.

1.3.1 Lyα absorption and Gunn-Peterson troughs

Resonant Lyα absorption[5] by neutral hydrogen is extraordinarily sensitive to small neutral fractions due to the large Lyα cross section (at line center, $\sigma_{\rm Ly\alpha} \simeq 6 \cdot 10^{-14}\,{\rm cm}^2 \gg \sigma_{912} = 6.3 \cdot 10^{-18}\,{\rm cm}^2$, where σ_{912} is the hydrogen photo-ionization cross section at 13.6 eV; see Draine 2011). For a uniformly distributed neutral fraction $\langle x_{\rm HI}\rangle$, the Ly$\alpha$ optical depth[6] $\tau_{\rm Ly\alpha}$ grows very rapidly with redshift (see Becker et al. 2015a):

$$\tau_{\rm Ly\alpha} \propto \langle x_{\rm HI}\rangle (1+z)^{3/2}, \tag{1.6}$$

so that even $\langle x_{\rm HI}\rangle \sim 10^{-4}$ produces near-zero transmitted flux F at $z \gtrsim 6$ (i.e., $\tau_{\rm Ly\alpha} \sim 2 \cdot 10^1$). At lower redshifts ($z \lesssim 4$) this resonant absorption appears as the well-known Lyα forest, a dense series of narrow absorption lines produced by clouds of neutral hydrogen at different redshifts along the line of sight. However, as redshift increases, the Universe becomes denser and the mean transmission declines, so that by $z \gtrsim 6$ the Lyα forest gets opaque, saturating for $\langle x_{\rm HI}\rangle \gtrsim 10^{-4}$. This extreme sensitivity to $\langle x_{\rm HI}\rangle$ is the basis of the Gunn-Peterson (GP) effect: a saturated Gunn–Peterson trough (i.e., vanishing Lyα transmission; see Fig. 1.4) signals the presence of non-negligible neutral gas along the sightline (Gunn & Peterson 1965; Fan et al. 2006c).

Because Lyα saturates quickly, useful information in highly neutral regimes must be extracted statistically or with complementary diagnostics. These approaches include:

- Measuring the effective optical depth $\tau_{\rm eff} = -\ln\langle F\rangle$ averaged over fixed redshift intervals to quantify the evolving opacity and its sightline-to-sightline scatter (e.g., Fan et al. 2006c; Becker et al. 2015b; Bosman et al. 2022);

- Using higher-order Lyman-series lines (i.e., Lyβ, Lyγ) whose smaller oscillator strengths[7] can remain unsaturated when Lyα is black, thereby providing tighter constraints on column densities and neutral fractions along particularly opaque sightlines (Fan et al. 2006c; Zhu et al. 2022);

- Analyzing the statistics of transmission 'spikes' and dark gaps (length distributions and frequency of zero-transmission regions) to probe the

[5] Here, "resonant" means that the photon energy matches precisely the energy difference between the ground state and the first excited energy state of the atom. A photon with that frequency can be absorbed with very high probability.

[6] It quantifies how much Lyα light is absorbed by neutral hydrogen in the IGM. Higher values imply that more photons are absorbed or scattered, and less are transmitted.

[7] The oscillator strength f is a dimensionless atomic parameter that measures the probability (or "strength") of an allowed radiative transition. Larger f imply larger σ, and hence stronger, more saturated lines.

topology of residual neutral islands and large-scale UV background (UVB) fluctuations (e.g., Becker et al. 2018; Zhu et al. 2021, 2022).

Interpretation of Lyα statistics requires careful modeling of cosmic density fluctuations, a spatially varying UVB, and temperature fluctuations introduced by patchy reionization and inhomogeneous heating (see e.g., Fan et al. 2006a; McQuinn et al. 2011; D'Aloisio et al. 2015; Davies & Furlanetto 2016; Kulkarni et al. 2019). As a result, Lyα constraints on the global neutral fraction are typically model-dependent and must be assessed jointly with radiative transfer simulations or forward-modelling frameworks that capture the relevant scales of ionized bubbles and self-shielded absorbers.

1.3.2 Damping wings

When the neutral hydrogen fraction becomes large (i.e., $\langle x_{\rm HI} \rangle \gtrsim 0.1$), the Lorentzian wings of the Lyα line profile imprint broad, extended absorption that can extend redward of the quasar systemic Lyα emission (see Fig. 1.4 and 1.5). These so-called *damping wings* are therefore a direct probe of large H I column densities and tracer of neutral gas in ambient IGM for $z \gtrsim 6.5$ (Miralda-Escudé 1998; Davies et al. 2018b; Greig et al. 2022; Ďurovčíková et al. 2024; Hennawi et al. 2025; Kist et al. 2025b).

Extracting robust constraints on $\langle x_{\rm HI} \rangle$ from damping wings require three main ingredients:

1. Accurate modeling the intrinsic unabsorbed quasar emission around Lyα (i.e., both continuum shape and emission lines), since uncertainties here map directly into the inferred damping wing amplitude;

2. Radiative transfer models that include the 3D distribution of neutral islands, peculiar velocities, and large-scale gas densities, as damping wings arise from absorption over extended velocity intervals;

3. Treatments of possible contaminating absorbers, such as proximate Damped Lyα Systems (pDLAs, with $N_{\rm HI} \gtrsim 10^{20}$ cm^{-2}) whose absorption can mimic an IGM feature.

Damping wing measurements can deliver strong, object-specific constraints on the local neutral hydrogen distribution, even though all the elements listed above can lead to significant uncertainties. Robust inferences therefore require a hybrid strategy that pairs carefully modeled, physically motivated analyses of individual spectra with statistical/ensemble approaches (see e.g., Hennawi et al. 2025; Kist et al. 2025b).

Recent works have successfully detected damping wing signatures in a handful of $z \gtrsim 7$ quasars (Mortlock et al. 2011; Bañados et al. 2018; Wang et al. 2020; Yang et al. 2020b; Greig et al. 2022) and have begun to translate these detections into quantitative limits on $\langle x_{\rm HI} \rangle$ at the earliest epochs

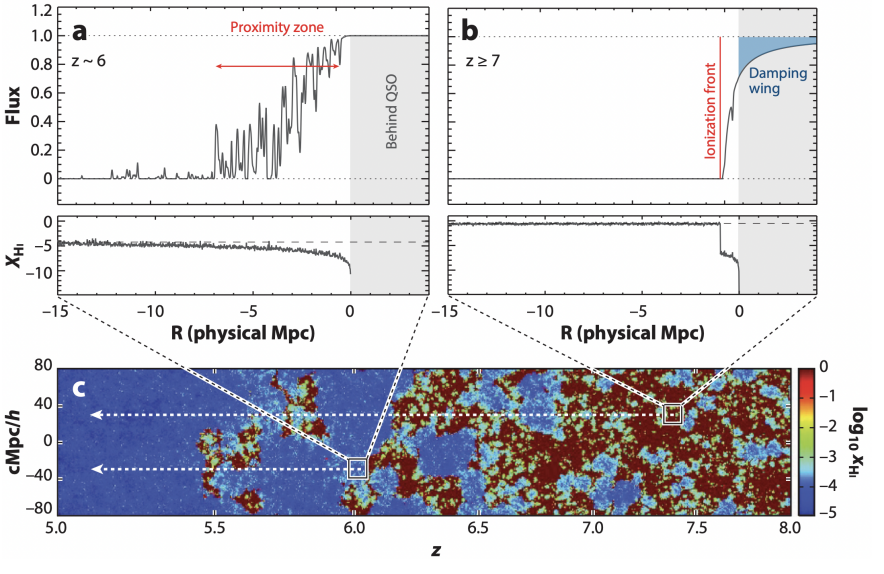

Figure 1.5: Schematic illustration of the observable signatures and physical interpretation of quasar proximity zones (a) and damping wings (b), with a visualization from large-scale simulations (c), where each small box marks the position of a background quasar whose sightline extends to the left. Proximity zones appear as regions of enhanced transmission produced by the quasar ionizing the surrounding, already largely ionized, IGM; no significant absorption is present redward of the systemic Lyα wavelength (gray shaded band), and a low optical depth region of order $\sim 5-10$ pMpc is visible. Damping wings arise when a substantial volume-averaged neutral fraction ($\langle x_{\rm HI}\rangle \sim 0.1$, indicated by the red color scale in the simulation) contributes extended, Lorentzian profile absorption. In this case, on-resonance Lyα absorption is saturated up to the quasar systemic redshift, while the off-resonance damping wing (blue shaded area) extends redward in the quasar rest-frame and encodes information on the IGM neutral fraction. Observational constraints, therefore, come both from the redward damping wing region and from the short run of unsaturated pixels that can remain between the systemic redshift and the end of the proximity zone. *Credit:* Fan et al. (2023).

(Ďurovčíková et al. 2024; Greig et al. 2024a), though the results remain sensitive to the continuum and lifetime priors adopted in the modeling.

1.3.3 Proximity zones

The quasar *proximity zone* (or near-zone) is the region close to the source where its ionizing flux dominates over the metagalactic UVB, and produces enhanced transmission relative to the ambient IGM for $5.5 \lesssim z \lesssim 7.5$ (see Fig. 1.4 and 1.5; Fan et al. 2006c). Assuming a predominantly neutral, homogeneous IGM and neglecting both recombinations and inhomogeneities, the physical size of the overionized "bubble" powered by a quasar that has emitted ionizing photons for a time $t_{\rm Q}$ is set by:

$$R_{\rm ion} \propto \left(\frac{\dot{N}_\gamma t_{\rm Q}}{n_{\rm H} x_{\rm HI}} \right)^{1/3} , \tag{1.7}$$

where $\dot{N}_\gamma$ is the ionizing photon emission rate, and $n_{\rm H}$ is the local hydrogen density (see Cen & Haiman 2000; Haiman & Cen 2001; Bolton & Haehnelt 2007a,b). In a mostly ionized IGM, the proximity zone instead marks the region over which the gas has had time to approach photo-ionization equilibrium with the quasar's radiation. Consequently, the proximity zone size is set by the IGM equilibration timescale $t_{\rm eq} \simeq 3{\cdot}10^4$ yr, and only becomes a sensitive probe of the quasar age when $t_{\rm Q} \lesssim t_{\rm eq}$; for longer $t_{\rm Q}$ the ionization structure reaches steady state and the measured size no longer increases with quasar lifetime (Davies et al. 2020a). Additionally, the spectroscopic proximity zone radius $R_{\rm p}$ (Fan et al. 2006c) does not always equal the physical ionized extent $R_{\rm ion}$: residual H I inside the quasar's bubble (and associated high-density clumps) can absorb Lyα photons and extinguish transmission well interior to the true ionization front, so that $R_{\rm p}$ tends to underestimate $R_{\rm ion}$ (Bolton & Haehnelt 2007a).

Proximity zone measurements have been used to constrain both the ambient neutral fraction and the quasar lifetime, which has important implications on SMBH growth (see Section 1.4 and e.g., Eilers et al. 2017, 2020, 2021b; Davies et al. 2019, 2020a). The key complications are:

1. Degeneracy between quasar lifetime and hydrogen neutral fraction: a small observed proximity zone can arise from a young quasar (i.e., which had insufficient time to ionize a large volume), a neutral-rich surrounding IGM, or an associated optically thick absorber (LLS or pDLA) that truncates the transmission;

2. Line of sight density fluctuations and quasar host overdensity: quasars preferentially live in biased, overdense regions that may be reionized earlier or host more absorbers, altering the proximity zone evolution relative to uniform background expectations;

3. Systemic redshift and continuum uncertainties: precise systemic redshifts (from either low-ionization lines, such as Mg II, or host CO or [C II] measurements) and quasar continuum reconstructions are required to define the proximity zone boundary accurately.

Because of these degeneracies, proximity zone analyses are most robust when combined with detailed radiative transfer forward models that place priors on quasar luminosity and environmental density (e.g., Eilers et al. 2017; Davies et al. 2018b, 2020a; Satyavolu et al. 2023b).

1.4 Supermassive black hole growth

As already introduced, evidences of luminous quasars at $z \gtrsim 6$ demonstrate that black holes with masses around $10^8 - 10^9\,\mathrm{M}_\odot$ already existed less than 1 Gyr after the Big Bang (e.g., Fan et al. 2000; Mortlock et al. 2011; Wu et al. 2015; Bañados et al. 2018; Matsuoka et al. 2019a; Wang et al. 2021a). Together with understanding the EoR, explaining how such extreme masses are assembled in a short cosmic time is another central challenge of early Universe astrophysics (see Volonteri 2010, 2012; Inayoshi et al. 2020; Di Matteo et al. 2023; Fan et al. 2023). In the following, we summarize the principal formation channels and the basic physical limits that govern subsequent growth, together with the principal observational constraints.

1.4.1 Seeding channels

The simplest division of formation scenarios (see Fig. 1.6) distinguishes between *light* and *heavy* seeds. Their differences are detailed below:

1. Light seeds: formed as remnants of the first (Population III) stars, with characteristic high masses $\simeq 10^2\,\mathrm{M}_\odot$ (Madau & Rees 2001), although the timing of the formation of the first stars is still unknown. They would require prolonged, often near–Eddington (or episodic super-Eddington) growth to reach $\simeq 10^9\,\mathrm{M}_\odot$ by $z \simeq 6$. Another possible channel discussed is runaway stellar collisions in dense star clusters, leading to the formation of intermediate mass back holes (IMBHs) with masses $10^3 - 10^4\,\mathrm{M}_\odot$ (Omukai et al. 2008);

2. Heavy seeds: formed by more exotic, direct-collapse black holes[8] (DCBHs) that can produce initial masses $\simeq 10^4 - 10^6\,\mathrm{M}_\odot$. They require rarer, special environments (e.g., strong Lyman–Werner backgrounds, dynamical heating, or synchronized halo pairs), but would ease the time constraint by starting from a much larger mass (e.g., Yoshida et al.

[8] Formed by the direct collapse of a large cloud of molecular gas without intermediate star formation.

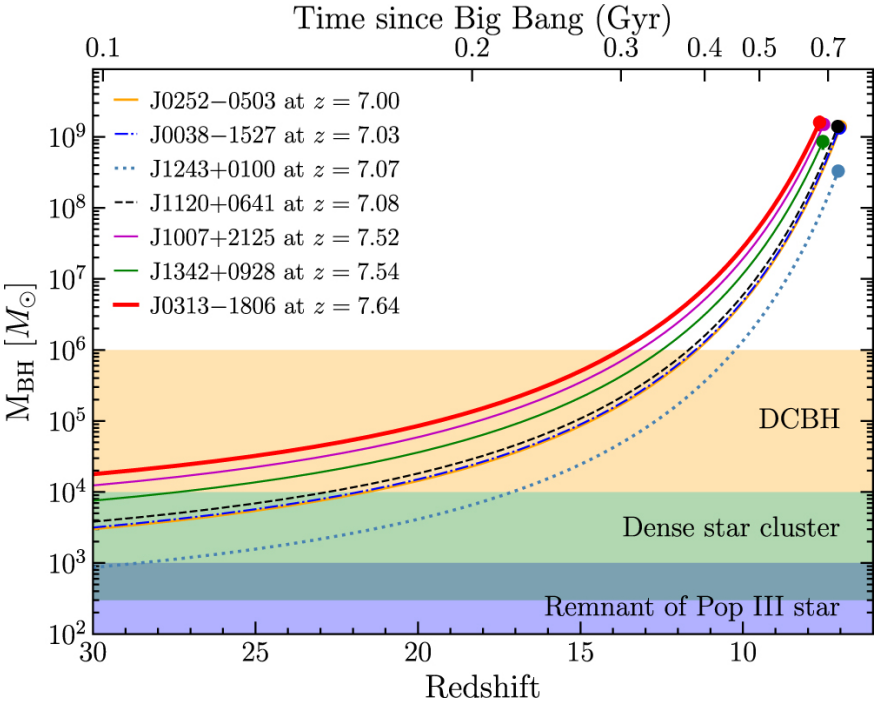

Figure 1.6: Required seed masses to grow the observed SMBHs of $z \gtrsim 7$ quasars under the assumption of continuous, Eddington–limited accretion with a constant radiative efficiency $\varepsilon = 0.1$. Each curve shows the initial seed mass that must be in place to reach the measured SMBH mass by the quasar's observed redshift. The quasars considered are J0252−0503, J0038−1527, J1243+0100, J1120+0641, J1007+2125, J1342+0928, and J0313−0318 (Mortlock et al. 2011; Bañados et al. 2018; Wang et al. 2018; Matsuoka et al. 2019a; Yang et al. 2019, 2020a; Wang et al. 2020, 2021a,b). Colored bands indicate approximate seed mass regimes expected for Pop III stars (blue), dense stellar clusters (green), and direct-collapse black holes (yellow). *Credit:* Wang et al. (2021a).

2003; Visbal et al. 2014; Fernandez et al. 2014; Wise et al. 2019; Regan et al. 2020). Another channel invoking primordial black holes (Choquette et al. 2019) has also been discussed, but its cosmological prevalence is uncertain.

New *James Webb Space Telescope* (JWST) observations of SMBHs at very high-z ($z \sim 10$) and evidences of over-massive black holes, compared to the host galaxies, seem to give more credit to the heavy seed scenario (Larson et al. 2023; Maiolino et al. 2024; Jeon et al. 2025).

1.4.2 Accretion physics and exponential growth

When gas accretes onto a black hole, a fraction ε of the rest–mass energy is released as radiation (the *radiative efficiency*), and the remainder increases the black hole mass. The radiated bolometric luminosity is $L_{\rm bol} = \varepsilon \dot{M}_{\rm in} c^2$ for an inflow rate $\dot{M}_{\rm in}$, so that the black hole mass grows at rate $\dot{M}_{\rm BH} = (1-\varepsilon)\dot{M}_{\rm in}$. The maximum (steady) accretion rate is then the *Eddington accretion rate* (also see Section 1.2.1):

$$\dot{M}_{\rm Edd} \equiv \frac{L_{\rm Edd}}{\varepsilon c^2} \simeq 2.1 \left(\frac{0.1}{\varepsilon}\right) \left(\frac{M}{10^8\,{\rm M}_\odot}\right) {\rm M}_\odot\,{\rm yr}^{-1}. \tag{1.8}$$

In this regime, the black hole mass grows exponentially as $M(t) = M_0 \cdot e^{t/t_{\rm S}}$, with M_0 being the initial (seed) black hole mass, and $t_{\rm S}$ the e-folding (*Salpeter*; Salpeter 1964) timescale[9] given by:

$$t_S = \frac{\varepsilon c \sigma_T}{(1-\varepsilon)\lambda_{\rm Edd} 4\pi G m_p} \simeq 4.5 \times 10^7 \left(\frac{\varepsilon}{0.1}\right) {\rm yr}, \tag{1.9}$$

for $\lambda_{\rm Edd} = 1$. Assuming radiatively efficient accretion dominates SMBH growth, the Soltan argument (Soltan 1982) implies a typical radiative efficiency $\varepsilon \approx 0.1$, consistent with constraints from quasar demographics (e.g., Yu & Tremaine 2002; Marconi et al. 2004; Shankar et al. 2009; Davis & Laor 2011). From equation 1.9, we see that continuous Eddington-limited accretion can, in principle, increase a $10^2\,{\rm M}_\odot$ seed to $10^9\,{\rm M}_\odot$ in a few hundred Myr. In practice, maintaining $\lambda_{\rm Edd} \simeq 1$ for such long times is challenging due to many factors, such as feedback or angular momentum barriers (e.g., Anglés-Alcázar et al. 2017; Weinberger et al. 2018).

1.4.3 Mergers and other mass accretion channels

Black hole mergers make an important contribution to mass build-up and are a natural consequence of hierarchical galaxy assembly. Some important aspects to be taken into account include:

[9] It is the characteristic timescale over which a black hole mass grows by a factor $e \simeq 2.718$ if it accretes continuously at Eddington limit with fixed radiative efficiency ε.

- Mass addition vs gas accretion: mergers add mass directly, but their overall contribution to the highest end of the mass function depends on the merger rate and on whether the merged remnant retains access to dense gas to continue accreting.

- Gravitational wave (GW) recoil: GW emission during coalescence can impart recoil velocities that eject or displace the remnant from shallow potential wells, temporarily suppressing accretion, and delaying growth (e.g., Haiman 2004).

- Dark matter accretion: dark matter is collisionless and inefficiently accreted; simulations indicate dark matter capture contributes only a small fraction ($\lesssim 10\%$) of SMBH growth compared to baryonic accretion (e.g., Peirani & de Freitas Pacheco 2008).

- Modeling limitations: cosmological simulations cannot resolve accretion physics down to event-horizon scales and therefore adopt subgrid prescriptions, such as Bondi accretion (Bondi 1952) or torque-driven inflow (Anglés-Alcázar et al. 2017), and parametrized feedback; different prescriptions can yield substantially different duty cycles and growth histories (e.g., Weinberger et al. 2018; Wellons et al. 2023).

1.4.4 Episodic growth and duty cycle

The accretion process is likely intermittent, with SMBHs alternating between active luminous phases and quieter intervals set by gas supply, feedback, and environmental conditions. Two timescales are commonly distinguished. The first one is the *episodic timescale* $t_{\rm ep}$, defining the duration of an individual bright phase during which the quasar emits the ionizing flux that shapes its local IGM. Line of sight proximity zone size measurements typically probe this time (Davies et al. 2020a; Satyavolu et al. 2023a).

The second one is the integrated *quasar lifetime* $t_{\rm Q}$, the total time a black hole spends accreting in luminous states over cosmic history. These times are connected to the *duty cycle* $f_{\rm duty}$, a widely used quantity that indicates the fraction of time a quasar is actively accreting gas and emitting radiation across cosmic time.

Cosmological models and high-resolution simulations commonly find highly variable accretion with repeated Eddington or super–Eddington episodes interspersed with longer quiescent periods (e.g., Novak et al. 2011; Weinberger et al. 2018; Bustamante & Springel 2019). Recent observations suggest short duty cycle $f_{\rm duty} < 1$ (typically $f_{\rm duty} \simeq 10^{-2} - 10^{-1}$; e.g., Arita et al. 2023; Eilers et al. 2024), increasing the time required to reach a given black hole mass compared to continuous growth.

Observational constraints: Multiple, complementary observational approaches can constrain different aspects of SMBH activity (see Fig. 1.7). The most relevant ones include:

- Clustering: measuring quasar clustering constrains typical host halo masses, and combining halo abundance with the quasar space density yields average duty cycles (and hence integrated-averaged lifetimes). These methods typically infer $t_{\rm Q} \simeq 10^7 - 10^9$ yr, though systematic uncertainties in halo occupation and obscured fractions remain significant (e.g., Haiman & Hui 2001; Martini & Weinberg 2001; Martini 2004; Shen et al. 2007, 2009; White et al. 2012; Pizzati et al. 2024a).

- Proximity zones: the ionization structure around bright quasars is sensitive to the recent ionizing history. H I proximity zone sizes of $z \gtrsim 6$ quasars often imply episodic times $t_{\rm ep} \simeq 10^5 - 10^6$ yr for the emission episode that created the observed ionization pattern (e.g., Eilers et al. 2017, 2020, 2021a; Davies et al. 2019, 2020a), whereas He II proximity measurements (probing longer equilibration times) indicate longer times $\simeq 10^6 - 10^7$ yr (e.g., Khrykin et al. 2016, 2019, 2021; Worseck et al. 2021).

- Transverse proximity, fluorescence, and nebular emission: searches for enhanced transmission in background sightlines (transverse proximity; e.g., Worseck et al. 2007; Keel et al. 2012; Schmidt et al. 2017, 2018; Bosman et al. 2021), or for fluorescent Lyα emission around quasars (e.g., Trainor & Steidel 2013; Hada et al. 2024) provide additional constraints that depend on geometry and IGM conditions; these typically favor episodic lifetimes $\lesssim 10^7 - 10^8$ yr in many systems.

Taken together, these approaches reveal tension, as clustering inferences point to relatively long integrated duty cycles, while proximity zone and fluorescence diagnostics commonly imply shorter luminous episodes. Reconciliation would require complex growth histories, in which SMBHs accumulate mass through many short, luminous episodes (so that $\sum t_{\rm ep} \sim t_{\rm Q}$), or models in which luminous emission is anisotropic, or heavily obscured for much of the growth period (e.g., DiPompeo et al. 2014; Satyavolu et al. 2023a).

1.5 This thesis

This thesis presents a coordinated observational and modeling effort to exploit high-z quasars as probes of the IGM and early SMBH growth. The next three chapters each address a different, but complementary, aspect of this problem: (i) assembling and characterizing a large, uniform sample

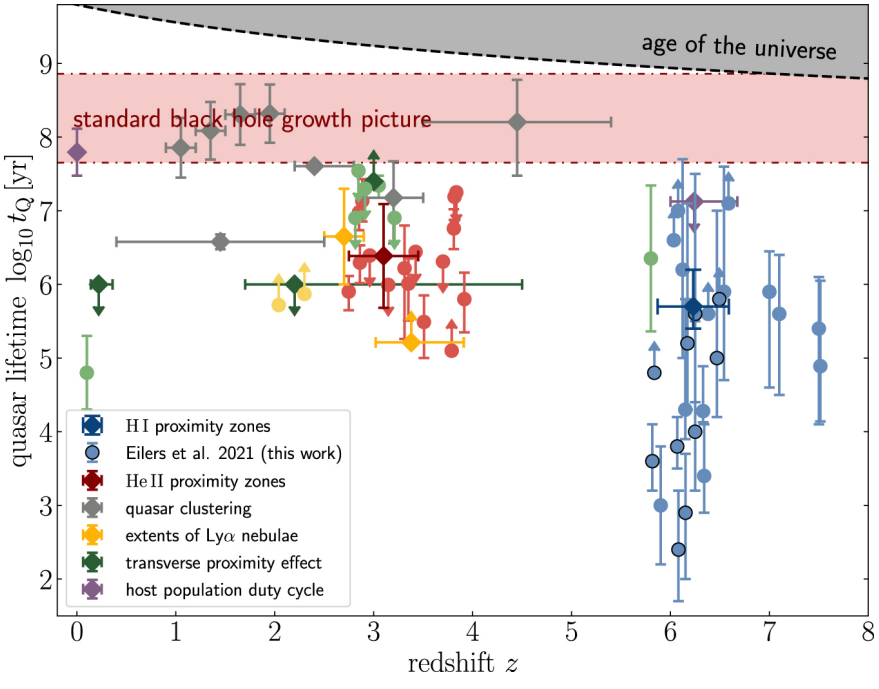

Figure 1.7: Compilation of quasar lifetime and duty cycle estimates from the literature as a function of redshift. Light shaded markers denote constraints from individual objects, while the darker diamonds indicate averaged values. Blue symbols show measurements based on H I line of sight proximity zones (individual points include results from Eilers et al. 2018b, 2021b; Davies et al. 2019, 2020b; Andika et al. 2020; the population average is from Morey et al. 2021). Red symbols correspond to He II proximity zone analyses (Khrykin et al. 2019; Worseck et al. 2021; population average from Khrykin et al. 2021). Gray points give duty cycle constraints derived from quasar clustering (Shen et al. 2007; Shankar et al. 2010; White et al. 2012; Laurent et al. 2017), and yellow symbols summarize constraints based on the spatial extents of Lyα nebulae (individual measurements: Cantalupo et al. 2014; Hennawi et al. 2015; averages: Trainor & Steidel 2013; Borisova et al. 2016). Green markers show lifetimes inferred from transverse proximity and fluorescence searches (Keel et al. 2012; Schmidt et al. 2018; Bosman et al. 2021; population estimates from Kirkman & Tytler 2008; Schmidt et al. 2017; Oppenheimer et al. 2018). Finally, the purple diamond indicates host population duty cycle inferences (Yu & Tremaine 2002; Chen & Gnedin 2018). The shaded red region marks the range of lifetimes expected under the simple exponential growth picture (roughly $1-16\,t_S$), assuming continuous, Eddington-limited accretion with a fiducial radiative efficiency $\varepsilon = 0.1$ (see equation 1.9). The dashed black curve represents the age of the Universe. *Credit:* Eilers et al. (2021b).

of medium/moderate-resolution rest-frame UV spectra of $z \gtrsim 6.5$ quasars, (*ii*) using those spectra to measure and interpret quasar proximity zones as a diagnostic of quasar lifetimes and the fraction of neutral hydrogen, and (*iii*) forward-modeling the Lyα transmission probability distribution function (PDF) to constrain the ionization and thermal state of the IGM during the end stages of reionization. Below, we summarize the contents and key outcomes of each chapter.

Chapter 2: Spectroscopy of $6.50 < z \leq 7.64$ quasars and composite spectrum

Chapter 2 presents optical and near-infrared spectroscopic observations for a sample of 45 quasars in the redshift interval $6.50 < z \leq 7.64$ and in the absolute magnitude at 1450 Å interval $-28.82 \leq M_{1450} \leq -24.13$ (with median values $z_{\rm median} = 6.71$, $M_{1450,\rm median} \simeq -26.1$). The data were obtained with a combination of echelle and long-slit spectrographs, and reduced in a uniform way using the open-source Python-based pipeline `PypeIt`. The chapter documents the sample selection, data reduction, quality assessment (median signal-to-noise ratio SNR in the J, H, K bands), comparison with other high-z quasar samples, and constructs a composite rest-UV spectrum from the dataset.

The spectroscopic catalog and composite serve two main purposes. First, they provide a good quality reference for quasar continua and emission line properties at $z \gtrsim 6.5$, useful to achieve several goals, such as the study of the proximity zones. Second, the composite spectrum is compared to other high- and low-z templates to isolate evolutionary trends in emission line strengths and continuum shape, and we find that all composites are consistent across cosmic times.

Chapter 3: Measurements of proximity zone sizes and their interpretation

Chapter 3 compiles homogeneous measurements of proximity zone sizes $R_{\rm p}$, for a sample of 59 quasars spanning $5.77 \leq z \leq 7.54$ (median $z = 6.59$), including first-time measurements for 15 objects. Using this enlarged dataset, we investigate the dependence of $R_{\rm p}$ on both quasar luminosity and redshift. A bivariate power-law fit (applied to a combined sample of ~ 100 objects) shows that the luminosity scaling is consistent with simple photo-ionization expectations ($R_{\rm p} \propto 10^{-0.4 M_{1450}/2.87}$), while the redshift evolution is steeper than previously reported in the literature, $R_{\rm p} \propto (1+z)^{-2.44}$ from our fit.

The chapter emphasizes the large scatter found in $R_{\rm p}$ at fixed z and M_{1450}, which we interpret as the combined effect of (*i*) variations in the quasar lifetime $t_{\rm Q}$, (*ii*) line of sight density fluctuations and host environment effects, and (*iii*) the ambient neutral fraction and the presence of

optically thick absorbers. We identify 13 objects with anomalously small $R_{\rm p}$; for all but one we rule out proximate optically thick absorbers and suggest that very short lifetimes ($t_{\rm Q} \lesssim 10^4$ yr) can be a plausible explanation. The chapter, therefore, provides both a catalog of uniformly measured $R_{\rm p}$ values and a possible physical interpretation that links the observed proximity zone distribution to quasar activity timescales and the surrounding IGM.

Chapter 4: Forward-modeling the Lyα transmission PDF

Chapter 4 develops a forward-modeling framework to compare the observed Lyα transmission PDFs with those obtained from simulated mock spectra that include uncertainties in the intrinsic continuum and instrumental noise. We apply this method to 21 spectra of the E-XQR-30 sample in the redshift range $5.76 < z < 6.59$, computing the PDFs in two rest-frame intervals: the Lyα forest ($1170 - 1190$ Å) and the proximity zone ($1190 - 1218$ Å). For each object, we generate thousands of mock spectra drawn from IGM models and evaluate goodness-of-fit.

The main findings are: (*i*) the transmission PDFs inside proximity zones is generally well reproduced by the models, indicating that quasar-dominated regions are reasonably captured by current radiative–transfer prescriptions and continuum treatments; (*ii*) in contrast, the Lyα forest PDFs in several redshift bins show statistically significant mismatches with the fiducial, spatially homogeneous IGM models. These discrepancies point to missing physics in the simple models, for example, patchy reionization, elevated IGM temperatures from recent heating events, or large-scale fluctuations in the UVB, and motivate more realistic simulations and hierarchical forward-modeling to interpret Lyα opacity at $z \sim 6$.