

Quasars as standard candles

VI. Spectroscopic validation of the cosmological sample

Bartolomeo Trefoloni^{1,2,*}, Elisabeta Lusso^{1,2}, Emanuele Nardini², Guido Risaliti^{1,2}, Alessandro Marconi^{1,2},
Giada Bargiacchi^{3,4}, Andrea Sacchi⁵, Paola Pietrini¹, and Matilde Signorini^{1,2}

¹ Dipartimento di Fisica e Astronomia, Università di Firenze, via G. Sansone 1, 50019 Sesto Fiorentino, Firenze, Italy

² INAF – Osservatorio Astrofisico di Arcetri, Largo Enrico Fermi 5, 50125 Firenze, Italy

³ Scuola Superiore Meridionale, Largo S. Marcellino 10, 80138 Napoli, Italy

⁴ Istituto Nazionale di Fisica Nucleare (INFN), Sez. di Napoli, Complesso Univ. Monte S. Angelo, Via Cinthia 9, 80126 Napoli, Italy

⁵ Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA

Received 13 December 2023 / Accepted 8 April 2024

ABSTRACT

A sample of quasars has been recently assembled to investigate the non-linear relation between their monochromatic luminosities at 2500 Å and 2 keV and to exploit quasars as a new class of ‘standardized candles’. The use of this technique for cosmological purposes relies on the non-evolution with redshift of the UV-optical spectral properties of quasars, as well as on the absence of possible contaminants such as dust extinction and host galaxy contribution. We address these possible issues by analysing the spectral properties of our cosmological quasar sample. We produced composite spectra in different bins of redshift and accretion parameters (black hole mass, bolometric luminosity), to investigate any possible evolution of the spectral properties of the continuum of the composites with these parameters. We found a remarkable similarity amongst the various stacked spectra. Apart from the well known evolution of the emission lines with luminosity (i.e. the Baldwin effect) and black hole mass (i.e. the virial relation), the overall shape of the continuum, produced by the accretion disc, does not show any statistically significant trend with black-hole mass (M_{BH}), bolometric luminosity (L_{bol}), or redshift (z). The composite spectrum of our quasar sample is consistent with negligible levels of both intrinsic reddening (with a colour excess $E(B - V) \lesssim 0.01$) and host galaxy emission (less than 10%) in the optical. We tested whether unaccounted dust extinction could explain the discrepancy between our cosmographic fit of the Hubble–Lemaître diagram and the concordance Λ CDM model. The average colour excess required to solve the tension should increase with redshift up to unphysically high values ($E(B - V) \approx 0.1$ at $z > 3$) that would imply that the intrinsic emission of quasars is much bluer and more luminous than ever reported in observed spectra. The similarity of quasar spectra across the parameter space excludes a significant evolution of the average continuum properties with any of the explored parameters, confirming the reliability of our sample for cosmological applications. Lastly, dust reddening cannot account for the observed tension between the Hubble–Lemaître diagram of quasars and the Λ CDM model.

Key words. galaxies: active – quasars: general – quasars: supermassive black holes

1. Introduction

Quasars are the most luminous persistent sources in the Universe and, as such, they enable its investigation up to redshifts of $z > 6$ (e.g. Mortlock et al. 2011; Bañados et al. 2018; Yang et al. 2020; Wang et al. 2021; Zappacosta et al. 2023). Thanks to their ubiquity and high luminosity, these sources were quickly recognized as powerful cosmological probes. Since the 1970s, correlations between quasars’ observables have been employed to determine their distances. The discovery of the Baldwin effect (i.e. the emission-line equivalent widths anti-correlate with the quasar luminosity; Baldwin 1977) envisaged that quasars could be utilised as cosmological probes. Techniques involving similar relations have been subsequently developed (see Section 6 in Czerny et al. 2018 for a detailed review). Yet the large dispersion as well as the scarcity of suitable sources, have in most cases limited their applicability for measuring the expansion parameters (e.g. H_0). Other attempts to extract cosmological information from quasars are based on the broad line region (BLR; Elvis & Karovska 2002) or dust (Yoshii et al. 2014) reverber-

ation mapping in combination with spectro-astrometric measurements (SARM; Wang et al. 2020), the Eigenvector 1 main sequence (Marziani & Sulentic 2014), or the X-ray excess variance (La Franca et al. 2014).

In the last decade, the efforts of our group have been mostly devoted to developing and refining a novel technique based on the non-linear relation between the X-ray and UV luminosity ($L_X - L_{\text{UV}}$) in quasars (or equivalently, the $\alpha_{\text{OX}} - L_{\text{UV}}$ relation; e.g. Vignali et al. 2003; Lusso et al. 2010), which has long been known (Tananbaum et al. 1979). The $L_X - L_{\text{UV}}$ relation, according to our current knowledge, reflects the interplay between the UV photons coming from the disc (Shakura & Sunyaev 1973; Novikov & Thorne 1973) and the hot ‘corona’ (Haardt & Maraschi 1991, 1993) that boosts their energy producing the X-ray emission. Despite the theoretical efforts (e.g. Nicastro 2000; Merloni 2003; Lusso & Risaliti 2017; Arcodia et al. 2019), a comprehensive description of the physical mechanism underlying the $L_X - L_{\text{UV}}$ relation is still missing. Even the choice of the monochromatic luminosities at 2 keV and 2500 Å can be regarded as rather arbitrary, being mainly due to historical and empirical reasons. Indeed, we have

* Corresponding author: bartolomeo.trefoloni@uni.fi.it

argued that better proxies, such as the Mg II emission line (see Signorini et al. 2023 for further details), of the physical mechanisms responsible for the observed relation can be exploited.

In addition to important implications for accretion physics, the non-linearity of the $L_X - L_{UV}$ relation, combined with the non-evolution with redshift of its parameters (Risaliti & Lusso 2019), makes it a useful tool to determine the luminosity distance of quasars and thus transforms them into a new class of ‘standardizable candles’. The $L_X - L_{UV}$ relation indeed proved to be a fruitful cosmological tool after it was demonstrated that most of the observed scatter is not intrinsic but rather due to observational issues, as the dispersion can be narrowed down to 0.09 dex in case of samples with high quality data (Sacchi et al. 2022). This new category of standard candles becomes particularly convenient when it comes to populating the intermediate redshifts between the farthest supernovae at $z \sim 2$ (Scolnic et al. 2018) and the cosmic microwave background (CMB). Despite a good agreement with the Λ CDM model in the low-redshift regime of the Hubble–Lemaître diagram, a 3–4 σ tension with the model, driven by the quasar data at earlier cosmic epochs ($z > 1.5$), has been reported (Risaliti & Lusso 2019; Bargiacchi et al. 2021; Sacchi et al. 2022; Bargiacchi et al. 2022; Lenart et al. 2023).

In Lusso et al. (2020, henceforth L20), we discussed how we produced our latest compilation of 2421 quasars for cosmological purposes, describing how the measurements of X-ray and UV luminosities were performed. However, in order to validate the selection criteria adopted to build the sample, we still need to address two main issues that could be present to some extent: i) an evolution with redshift of the slope of the optical and UV continuum and ii) the contamination of host galaxy in the optical and flux attenuation due to dust and gas. While we discuss the validity of the cosmological techniques adopted to analyse the sample in a companion paper (Risaliti et al., in preparation), in this work we mainly focus on a spectral confirmation of the robustness of the selection criteria. In particular, we aim at demonstrating, by means of composite spectra, that the continuum of our sample (in which the UV proxy $L_{2500\text{\AA}}$ is evaluated) does not show any appreciable evolution with redshift. At the same time, the spectral composites are a useful tool to reveal the presence of residual (i.e., still present despite the selection cuts) dust extinction and/or host galaxy contribution.

To the end of investigating any possible spectral diversity among the objects constituting our cosmological sample, we employed stacked templates. The lack of a significant evolution of the continuum among the stacked spectra, representative of different regions of the parameter space, ensures the universality of the accretion mechanism powering the optical–UV emission.

In this paper, we analyse the average spectral properties of the UV and optical data of the cosmological sample described in L20, with three main goals: i) directly verifying how close the average spectral properties of our sample are to those of typical blue quasars, by means of composite spectra in different regions of the $M_{\text{BH}} - L_{\text{bol}} - z$ parameter space; ii) estimating the degree of host galaxy contamination; iii) quantifying the effect of intrinsic dust absorption in our sample of quasars. The last two effects would influence our estimates of $f_{2500\text{\AA}}$ and therefore undermine the cosmological application of the $L_X - L_{UV}$ relation. Employing spectral data to understand – and if necessary, account for – such possible observational effects, that can alter both X-ray and UV measurements would definitely corroborate the evidence that the observed discrepancy with the concordance cosmological model is not due to observational issues.

The paper is structured as follows. A brief summary about the selection of the quasar sample is provided in Section 2, whilst

the analysis performed on the sample spectra is described in Section 3. Results are discussed in Section 4, and conclusions are drawn in Section 5.

2. The data set

The quasar sample analysed here was published by L20 and is made of 2421 sources within the redshift range $0.13 < z < 5.42$. The data set was produced by merging seven different subsamples available from the literature.

The Sloan Digital Sky Survey data release 14 (SDSS-DR14, Pâris et al. 2018) was cross-matched with the *XMM-Newton* 4XMM-DR9 (Webb et al. 2020) catalogue of serendipitous sources, the second release of the *Chandra* Source Catalogue (CSC2.0, Evans et al. 2010) and the AGN sample described in Menzel et al. (2016), belonging to the *XMM-Newton* XXL survey (XMM–XXL, PI: Pierre), producing respectively the SDSS-4XMM, the SDSS-Chandra, and the SDSS-XXL samples.

Local quasars ($0.009 < z < 0.1$) with available UV data from the International Ultraviolet Explorer (IUE) in the Minkowski Archive for Space Telescopes (MAST) and X-ray data were included in the sample to anchor the normalisation of the quasar Hubble diagram with Type Ia supernovae.

To improve the coverage at high redshifts, three samples of quasars with pointed observations were added. We added the sample of 29 luminous quasars at redshift $3.0 < z < 3.3$ with dedicated *XMM-Newton* X-ray observations described in Nardini et al. (2019), the 53 high redshift ($z > 4.01$) quasars from Salvestrini et al. (2019), and the 25 quasars at $z > 6$ from Vito et al. (2019).

In order to produce such a compilation, several selection criteria were adopted. We thus redirect the interested reader to L20 (in particular, their Section 5) for a more comprehensive description, while here we briefly highlight the filters adopted to define the sample. All the selected sources are radio quiet¹ and not flagged as broad absorption line (BAL) quasars according to the SDSS-DR14 quasar catalogue (Pâris et al. 2018), which was the latest release at the time of the L20 analysis. A broadband photometric spectral energy distribution (SED) was built for each source, and only quasars with UV-optical colours close to the ones of typical blue quasars (Richards et al. 2006) were retained. In addition, steep X-ray spectra ($1.7 < \Gamma < 2.8$) and deep enough X-ray observations (to avoid the Eddington bias), were required. Such properties make the L20 sample a truly unique sample of blue quasars, sharing a high degree of homogeneity in terms of broadband properties and observational data. The present work is focused on the spectral properties of the sources with available SDSS data, which constitute the bulk (~93%) of the sample.

3. Methods

3.1. The parameter space

Our aim is to study whether the continuum properties of our sample evolve with redshift or any of the governing parameters of the quasar spectrum, namely the bolometric luminosity (L_{bol}), black hole mass (M_{BH}), and accretion rate (parametrised through the Eddington ratio $\lambda_{\text{Edd}} = L_{\text{bol}}/L_{\text{Edd}}$, where L_{Edd} is the luminosity at the Eddington limit). To this end, we considered the latest quasar catalogue by Wu & Shen (2022, henceforth WS22), who

¹ We applied the usual criteria based on the ratio between the flux density at 6 cm and 2500 Å, i.e. AGN with $R = F(6\text{ cm})/F(2500\text{\AA}) < 10$ are considered radio quiet (Kellermann et al. 1989).




Fig. 1. Parameter space of $\log M_{\text{BH}} - \log L_{\text{bol}}$ for the quasar sample published in L20 with colour-coded redshifts. Dot-dashed lines represent constant λ_{Edd} values. Typical uncertainties on $\log(M_{\text{BH}})$ and $\log(L_{\text{bol}})$ are represented in the bottom-right corner.

compiled continuum and emission-line properties for 750,414 broad-line quasars from the SDSS DR16, including physical quantities such as single-epoch virial M_{BH} , L_{bol} , and λ_{Edd} . We also gathered the latest spectral data available for each source in our sample, hence the spectra of $\sim 4\%$ of the sources belong to the DR17 release, where calibration files have been updated (Abdurro'uf et al. 2022).

Fig. 1 presents the $\log L_{\text{bol}} - \log M_{\text{BH}}$ plane for our sample. The partition of the plane was defined by seeking the balance between a large number of bins to improve the sensitivity on the parameters (L_{bol} , M_{BH}) that control the average spectrum, and a statistically robust number of objects in each region. We further divided the sample into five redshift intervals, $z = [0.13, 0.77, 1.12, 1.52, 2.03, 5.42]$, each one containing a similar number of objects (roughly 480 objects per bin). The resulting grid is shown in Fig. 1. The width of the $\log M_{\text{BH}}$ ($\log L_{\text{bol}}$) bins is 0.7 dex (1 dex), large enough to encompass the large systematic uncertainties associated with the estimate of this quantity, typically 0.3 dex (but up to 0.4 dex in case of C IV-based M_{BH} , e.g. Shen et al. 2011; Kratzer & Richards 2015; Rakshit et al. 2020). The points of the grid are $\log M_{\text{BH}} = [7.0, 7.7, 8.4, 9.1, 9.8, 10.5]$ and $\log L_{\text{bol}} = [44.4, 45.4, 46.4, 47.4, 48.4]$. This procedure defines a 3D grid made of $5 \times 5 \times 4 = 100$ bins. The individual cells are labelled using the following nomenclature: increasing capital letters (A, B, C, D) denote decreasing bolometric luminosity across the bins, while increasing numbers (1, 2, 3, 4, 5) mark increasing average black hole masses, and increasing lower-case letters (a, b, c, d, e) represent increasing redshift intervals, which are also colour-coded according to the values in the legend.

Clearly, not all the bins are densely, nor uniformly populated. Lower redshift quasars preferentially occupy the lower luminosity and black hole mass region, whereas higher redshift sources are observed at higher values of both L_{bol} and M_{BH} . Moreover,

there are cells in the $L_{\text{bol}} - M_{\text{BH}}$ plane where the barycentre of the sources is close to the edge of a given interval (e.g. B2, B5, C1). In these cases, low sample statistics would produce not only low signal-to-noise (SNR) composite spectra, but also a stack not fully representative of its nominal locus in the parameter space.

To select the regions in the $\log M_{\text{BH}} - \log L_{\text{bol}}$ plane populated enough to be included in the stacking analysis, we quantified the probability that a statistical fluctuation in terms of $\log M_{\text{BH}}$ and $\log L_{\text{bol}}$ of objects in adjacent bins pushed them into a given region. We simulated 1000 mock samples starting from the actual distribution of $\log M_{\text{BH}}$ and $\log L_{\text{bol}}$ in each bin within the grid. The values of $\log M_{\text{BH}}$ and $\log L_{\text{bol}}$ are then scattered by a value drawn from a Gaussian distribution whose amplitude was set by the total uncertainty on each given data point. In doing so, we conservatively assumed the largest possible systematic uncertainty on both parameters and summed it in quadrature to the statistical uncertainty reported in WS22. For the bolometric corrections, we assumed a systematic uncertainty of $\sim 50\%$ on $\log L_{\text{bol}}$ for each data point in the plane, suitable in case the bolometric correction is performed on the rest-frame 1350 Å luminosity (see Shen et al. 2011), which translates into $\lesssim 0.22$ dex. The systematic uncertainty on $\log M_{\text{BH}}$ was assumed to be 0.40 dex, a value based on the C IV calibration (i.e. the least reliable one). As a consequence of the random scatter, the points of the mock sample distribution could escape the bin where the original data were instead contained. We then computed the average $\log M_{\text{BH}}$ and $\log L_{\text{bol}}$ of the resulting scattered points and retained only the regions within the parameter space where in at least 90% of the mock samples the average $\log M_{\text{BH}}$ and $\log L_{\text{bol}}$ remained within the bin. We refer to bins retained by this procedure as bins meeting the ‘representativeness criterion’. In general, 44 bins out of 100 in the parameter space do not contain any object, while 25 bins are excluded by the aforementioned criterion. The whole selection procedure ultimately excluded 69 scarcely populated bins (with less than nine objects), and it ensures that the composite spectrum is truly representative of its associated region in the parameter space. At the end of this procedure, 2316 objects were retained in the parameter space.

3.2. Building the stacks

The spectra were collected according to their PLATE, MJD and FIBERID reported in WS22². We produced the stacked spectra and adopted a bootstrap recipe for the evaluation of the uncertainty on the stacks (more quantitative assessments on how different stacking prescriptions affect the results are explored in Appendix A) in each of the selected regions following the steps described below.

Each spectrum was corrected for Galactic absorption assuming the value for the colour excess $E(B - V)$ available in the WS22 catalogue (Schlafly & Finkbeiner 2011) and a Fitzpatrick (1999) extinction curve. Then, the de-reddened spectra were shifted to the rest frame.

All the spectra in the same redshift bin were resampled, by means of linear interpolation, onto a fixed wavelength grid. The grid was defined in the interval $\lambda_1 - \lambda_2$ with a fixed dispersion $\Delta\lambda$ so as to cover the shortest and the longest rest-frame wavelengths in each redshift interval. The values for $\lambda_1, \lambda_2, \Delta\lambda$ are listed in Table 1. The value for $\Delta\lambda$ was chosen in order to preserve the

² The spectrum of SDSS J124406.96+113524.2 could not be found. However, this quasar belongs to a region where there are ~ 60 objects, thus we expect that its inclusion would not play any significant difference in the resulting stack.

Table 1. Stacking parameters.

Redshift interval (1)	N_{obj} (2)	λ_1 (3)	λ_2 (4)	$\Delta\lambda$ (5)	$\bar{\lambda}$ (6)
0.13–0.77	484	2062	9200	2.2	5100
0.77–1.12	479	1722	5876	1.7	4200
1.12–1.52	473	1448	4906	1.4	3100
1.52–2.03	483	1205	4127	1.2	2200
2.03–5.42	474	569	3432	0.7	1450

Notes. Redshift interval (1), number of sources within the redshift interval (2), lower (3) and upper (4) wavelengths of the wavelength grid, fixed dispersion (5), and reference wavelength (6). All wavelengths are reported in units of Å.

SDSS observed frame resolution also in the rest frame spectra, similarly to what described in Section 3.2 of [Vanden Berk \(2001, VB01\)](#), assuming an observed frame spectral resolution of ~ 2000 at 6500 Å.

We also applied a few quality cuts. Objects with more than 10% of the pixels flagged as bad (`AND_MASK > 0`) in their SDSS spectra were excluded. Also objects where the normalisation wavelength $\bar{\lambda}$ (the wavelength chosen to scale all the spectra in the cell) was not covered in the rest-frame spectrum were excluded.

The average flux in each spectral channel and the relative uncertainty were evaluated adopting a bootstrap technique. In brief, a number of spectra equal to the total number of objects in the bin is drawn, allowing for replacement. These spectra are normalised by their flux value at the reference wavelength $\bar{\lambda}$, and an average stack is produced. This procedure is repeated with new draws until the desired number of stacks is reached. Subsequently, in each spectral channel the distribution of the normalised fluxes is computed, a 3σ clipping is applied, and the median value (less affected by possible strong outliers) is chosen. The uncertainty on the median value was evaluated as the semi-inter-percentile range between 1st and 99th percentile.

The information about the relevant quantities adopted in each redshift interval is shown in Table 1. We explored different possible ways of combining the spectra, averaging the flux in each spectral channel and estimating the uncertainty. How these assumptions affect the resulting composites and the subsequent results, is presented in Appendices A and B. All the composite spectra are presented in Appendix G.

3.3. Spectral fit of the composites

To estimate the characteristic spectral quantities of the continuum (chiefly the slope α_λ of the power law $f_\lambda \propto \lambda^{-\alpha_\lambda}$) and the emission lines, in particular the rest-frame equivalent width (EW), the full width at half-maximum (FWHM), and the peak offset velocity (v_{off}), we performed the spectral fits of all of our composite spectra. To this end we adopted a custom-made code, based on the IDL MPFIT package ([Markwardt 2009](#)), which takes advantage of the Levenberg-Marquardt technique ([Moré 1978](#)) to solve the least-squares problem. Our main goal is a faithful reproduction of the continuum, in order to investigate whether the continuum slope and consequently the 2500 Å flux exhibits any significant evolution with the accretion parameters.

The baseline model is given by a continuum power law pivoting at different emission-free wavelengths depending on the redshift interval (3000 Å in the three bins at $z < 1.52$, 2200 Å

Table 2. Fit parameters.

Redshift interval (1)	λ_1 (2)	λ_2 (3)	$\bar{\lambda}$ (4)
0.13–0.77	2150	5500	3000
0.77–1.12	2150	5500	3000
1.12–1.52	1700	4500	3000
1.52–2.03	1300	3800	2200
2.03–5.42	1300	3200	2200

Notes. Redshift interval (1), lower (2) and upper (3) wavelengths, and normalisation wavelength (4) of the stacks. All wavelengths are reported in units of Å.

in the two bins at $z \geq 1.52$), a set of emission lines, the Balmer continuum, and a set of iron templates convolved with a Gaussian line-of-sight velocity distribution whose weights are free parameters in the fit. Each line required a different decomposition: the C IV line was generally fitted with two Lorentzian profiles (one broad, one narrow). A single broad Lorentzian profile combined with the iron templates is able to adequately reproduce the Mg II profile. The H β was modelled by considering a core narrow Gaussian component whose FWHM could not exceed ~ 1000 km s $^{-1}$ and a broad component whose analytical shape (Gaussian or Lorentzian) was chosen each time to best reproduce the data. The [O III] doublet was modelled with Gaussian or Lorentzian profiles with a ratio of the [O III] $\lambda 4959$ Å to $\lambda 5007$ Å components fixed at one to three. In addition, other weaker forbidden ([Ne V] $\lambda 3426$, [O II] $\lambda 3727$, [Ne III] $\lambda 3869$) and permitted (H γ , H δ) emission lines were included to improve the residuals. Above $\log(L_{\text{bol}}/(\text{erg s}^{-1})) \geq 46$ the host galaxy contribution at optical wavelengths is expected to be negligible with respect to the quasar emission ([Shen et al. 2011](#); [Jalan et al. 2023](#)). Therefore, we only added an elliptical galaxy template from [Assef et al. \(2010\)](#) to the model, to account for possible galactic emission, only in the two lowest redshift bins, where the bulk of the sample has L_{bol} close or below above value and the rest-frame wavelengths sampled in these redshift ranges could be imprinted by galactic emission. The fits were performed adopting the same wavelength interval within the same redshift bin to avoid mismatches between the adopted continuum windows. The fitting wavelengths are listed in Table 2 and examples of a low- and a high-redshift fits are shown in the left and right panels of Fig. 2, respectively.

The uncertainties on the relevant parameters were estimated by means of mock samples. For each composite we produced 100 mock samples by adding a random value to the actual flux, extracted from a Gaussian distribution whose amplitude was set by the uncertainty value in that spectral channel. We then fitted the mock samples and evaluated the 3σ clipped distribution of the relevant parameters. The final uncertainty on the reported parameters was estimated as the standard deviation of such a distribution.

4. Results

4.1. The full sample and literature composite spectra

As a reference for the whole sample we produced a stacked spectrum collecting all the spectra of the L20 sample matching our quality criteria. The result is shown in Fig. 3. The procedure adopted to produce our template is the same as that described in Sect. 3.2. The normalisation wavelength $\bar{\lambda}$ was set to 2200 Å,




Fig. 2. Examples of a low-redshift bin (left) and a high-redshift bin (right) fit. The continuum power law is marked in blue, Fe II templates in grey, emission lines in different colours (orange, purple, green, cyan). The total model is depicted in red.




Fig. 3. Spectral composite of the L20 sample. *Top panel:* composite spectrum of our cosmological sample described in this work. The Ly α emission line is cut for visualisation purposes. The inset zooms in on the H β -[O III] region. The main emission lines are labelled accordingly. *Middle panel:* relative residuals with respect to the VB01 template. *Bottom panel:* number of objects contributing to each spectral channel. The composite spectrum of the L20 sample is available in its entirety in a machine-readable form at the CDS.

as this is the only relatively emission-line-free continuum window available for most of the objects in our sample. To place our sample in a broader context, we compared our new template with other literature average spectra, namely Francis et al. (1991), Brotherton et al. (2001), VB01, Selsing et al. (2016), and Harris et al. (2016).

The composite spectrum by Francis et al. (1991) was produced using 718 objects from the Large Bright Quasar Survey (LBQS) with absolute magnitude in the B -band $M_B \leq -20.5$ and redshifts in the range $0.05 < z < 3.36$; this sample is strongly biased towards luminous objects at high redshift. The FIRST Bright Quasar Survey (FBQS, Brotherton et al. 2001)

took advantage of FIRST radio observations of known point-like sources with colour indices³ $O - E < 2.0$ and $E < 17.8$ from the Palomar Observatory Sky Survey I (POSS-I). This match yielded 636 sources, here we show only the composite spectrum made of the radio-quiet sources as a comparison.

The first SDSS quasar template was produced by VB01, who selected 2204 sources between $0.004 \leq z \leq 4.789$ based on the SDSS colour difference (Richards et al. 2001), a non-stellar continuum, and at least one broad ($\text{FWHM} > 500 \text{ km s}^{-1}$)

³ More details about the POSS-I photometric system can be found in Evans (1989).

permitted emission line. The sample adopted by Harris et al. (2016) was instead based on the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2012), as part of SDSS-III (Eisenstein et al. 2011), aimed at characterizing the effects of the baryon acoustic oscillations (BAO) on the distribution of luminous red galaxies and studying the Ly α forest in high-redshift quasars. These purposes resulted in a sample of 102,150 objects between $2.1 \leq z \leq 3.5$ (see Harris et al. 2016 for more details about the sample selection). Given the redshift range covered by such sources, the comparison with this composite spectrum was possible only in the UV up to ~ 3000 Å. Lastly, Selsing et al. (2016) observed with XSHOOTER/VLT a sample of seven bright ($r \leq 17$) quasars between $1.0 < z < 2.1$, for which galactic contribution should be regarded as negligible.

The comparison between our template and the literature ones is shown in Fig. 4. Note that, the y-axis represents λf_λ , rather than the usual flux density. The agreement is generally good, especially in the region between 1100–4000 Å. At shorter wavelengths, the intergalactic medium (IGM) can affect the composites differently, according to the distribution in redshift of the individual spectra upon which the stacks are based (e.g. Madau 1995; Faucher-Giguère et al. 2008; Prochaska et al. 2009). On the red side, instead, some intrinsic differences start to become visible. Our composite is almost indistinguishable from the VB01 composite, but these two seem flatter than the others. There are several explanations for this slight difference in the long-wavelength end of the composites. First, the colour selections performed to select the quasar candidates is not homogeneous among the different samples, thus the SDSS objects that constitute the VB01 sample as well as ours could present slightly redder indices in the optical range below $z = 0.8$. Secondly, it is possible that the host galaxy, whose contribution is negligible in the UV (see Sect. 4.5) plays a minor role in flattening the optical side of the spectra. VB01 estimated the galaxy contribution to be roughly 7–15% at the Ca II $\lambda 3934$ wavelengths, and a similar fraction can be estimated in our global composite. On the other hand, the other surveys and the Selsing et al. (2016) sample are made, on average, of more luminous objects for the same redshift (see inset in Fig. 4). It is therefore likely that the galactic contribution in those templates is even lower than in ours.

Throughout this work, we will refer to the full sample template as a reference to compare to the composite spectra of the individual bins in the $M_{\text{BH}} - L_{\text{bol}} - z$ space, unless differently stated.

4.2. Non-evolution with redshift

A key observational feature bridging local and high-redshift quasars is the apparent similarity of their optical–UV continuum up to primordial times (e.g. Bañados et al. 2018; Yang et al. 2020; Wang et al. 2021). However, if the quasar spectra showed even a subtle evolution with redshift, the accretion physics underlying the X-ray to UV relation could also be evolving, making quasars unsuitable candidates as standard candles. Given the wide range in terms of redshift covered by our sample, we expect to be capable of detecting any such possible effect. This would appear as a systematic difference in the slope of the continuum produced by the accretion process at different cosmic times.

As a very basic step to check for any evolution with redshift of the spectral properties within our sample, in Fig. 5 we compare the average spectra in the redshift bins with the global composite (described in Sect. 4.1), obtained by stacking the spectra of all the sources in the sample. We find a very good agreement with the global composite at all redshifts, without any clear evi-




Fig. 4. Comparison between our quasar template and others from the literature. All the templates are scaled by their emission at 2200 Å. The agreement on the UV side, apart from minor differences in the Fe II pseudo-continuum, specifically the 2250–2650 Å and the 3090–3400 Å bumps, is apparent. The Ly α emission line is cut for visualisation purposes. The inset shows the distribution of the absolute R magnitude against the redshift for the different samples reported.

dence of evolution, apart from the known trends of the emission lines, whose FWHM varies with M_{BH} (i.e. the virial relation; e.g. Peterson et al. 2004) and whose EW anti-correlates with L_{bol} (i.e. the Baldwin effect; Baldwin 1977).

The visual agreement is confirmed by a more quantitative assessment presented in the bottom-right panel in Fig. 5. There we show the slope of the continuum evaluated in continuum intervals sampled at different redshifts by the spectral composites. The slopes are generally consistent within 1σ , and no systematic trends are observed. The optimal agreement between the composite spectrum in every redshift bin and the global composite, as well as among the different redshift composites, demonstrates the non-evolution of the average continuum properties of the sample across cosmic time.

In addition, as discussed in Section 4.1, the global sample composite exhibits a remarkable agreement with the average SDSS spectrum described in VB01, meaning that the stacks in the different redshift bins are also similar to the spectrum of a typical blue quasar. Since even this direct comparison reveals that the average continuum properties of our sample do not suggest any redshift evolution, we now investigate and characterise in more detail how our sample behaves across the $\log M_{\text{BH}} - \log L_{\text{bol}}$ parameter space.

4.3. Analysis of the evolution of the continuum slope of the stacks

The fits of the stacked spectra returned the slope of the continuum in each bin of the parameter space. We thus explored whether any possible evolution of the continuum slope exists throughout our sample. With the term slope (α_λ), we denote the slope of the continuum power law underlying the emission lines roughly between 1300 and 5500 Å. This wavelength range is sampled differently in the different redshift bins: for instance in the lowest redshift bin only the interval 2200–5500 Å is available, while in the highest redshift bin the composite spectra were fitted over the interval 1300–3200 Å. We also underline that,

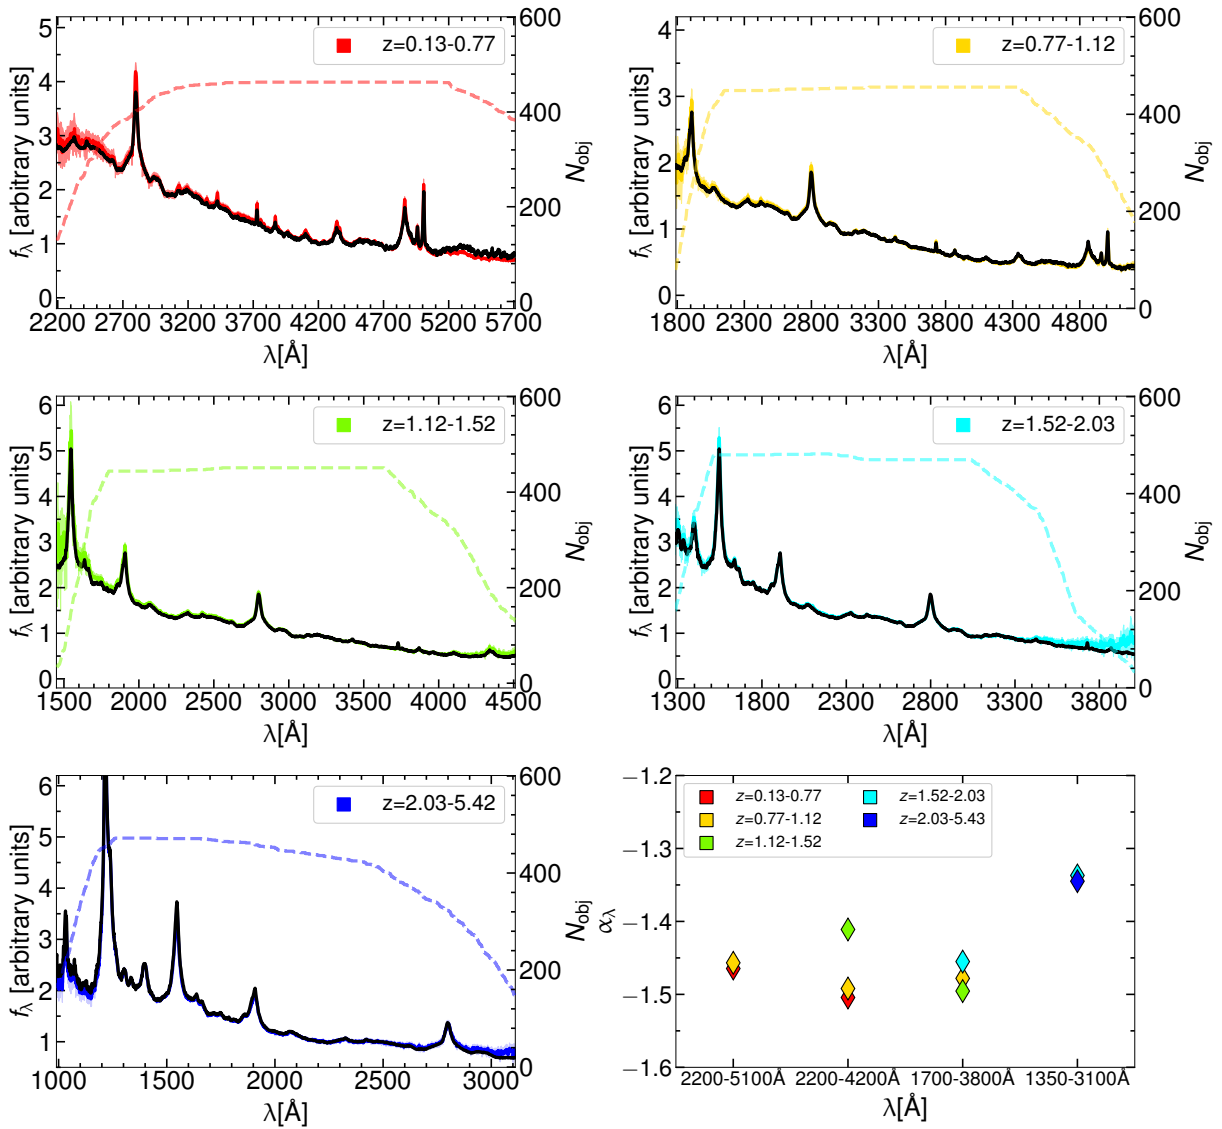


Fig. 5. Composite spectra divided in redshift bins superimposed to the full sample composite. The dashed line denotes the number of objects contributing to the stack in each spectral channel. The composite spectra in each bin are scaled to match the template respectively at 4200 Å for the low redshift bin, 3000 Å for the $z = [0.77-1.12)$, $z = [1.12-1.52)$, $z = [1.52-2.03)$ bins, and 2200 Å for the highest redshift bin. In the bottom-right panel we show the slopes of the continuum evaluated in different pairs of continuum windows, which are generally consistent.

throughout this work, we never assign a physical meaning to the slope of the continuum, adopting it exclusively as a comparative parameter. The physical meaning of such a parameter would require the assumption of a physical continuum model, generally an accretion disc, but see also e.g. Antonucci et al. (1989) for some limitations, whose implications go beyond the scope of this work.

In Fig. 6, we show the values of the slope in the redshift bins as a function of $\log M_{\text{BH}}$, $\log L_{\text{bol}}$, and $\log \lambda_{\text{Edd}}$. To avoid dust reddening or galaxy contamination, a selection in terms of photometric colours has been performed when constructing the cosmological sample (see L20, Section 5.1). In brief, for each object, the slope of the power law in the $\log \nu - \log(\nu f_\nu)$ space was computed in the rest-frame ranges 3000–10000 Å (Γ_1) and 1450–3000 Å (Γ_2). Objects whose photometric colours were compatible with those of typical blue quasars, $\Gamma_1 = 0.82$ and $\Gamma_2 = 0.40$ (Richards et al. 2006), only tolerating mild red-

dening, i.e. $E(B - V) < 0.1$, were selected. Converting the photometric indices Γ in terms of spectral slopes as $\alpha_\lambda = -\Gamma - 1$, the expected slopes in the 3000–10000 Å and the 1450–3000 Å ranges, respectively covered by the low and high redshift bins, are $\alpha_{\lambda,1} = -1.82$ and $\alpha_{\lambda,2} = -1.40$. Although this process operates a sort of pre-selection on the possible continuum slopes, we note that only roughly 20% of the starting sample is excluded according to this colour selection.

In Fig. 6 we show the slopes of the continuum evaluated in the bins of the parameter space. We mark with a hollow circle the values of the bins with a low N_{obj} which would be excluded as a result of the representativeness criterion described in Sect. 3.1. The values cluster around $\alpha_\lambda = -1.50$ with a standard deviation of $\sigma_\alpha = 0.14$, showing little or no evolution from low to high redshift. Considering also the excluded bins, the results are not significantly altered as the mean value negligibly shifts to $\alpha_\lambda = -1.51$ and the standard deviation to $\sigma_\alpha = 0.17$. The slope




Fig. 6. Values of the slope of the continuum produced by our fits. Bins not meeting the representativeness criterion are marked with a circle. The slope of the VB01 template (dot-dashed line) is shown as an external reference. Different pairs of line-free windows were adopted to define the continuum in each bin.




Fig. 7. De-reddened stacks, according to the same colour code as in Fig. 13, scaled by their mean emission between 1770 Å and 1810 Å. For comparison, the VB01 template is also shown in black as an external reference, rather than the usual L20 template. The rising continuum slopes with increasing redshift (and extinction) are clearly observed. Inset: UV continuum slopes of the literature samples described in Section 4.1 and plotted in Fig. 4 (squares), against the slopes of the de-reddened spectra (dashed lines).

average value is close to those expected from the Richards et al. (2006) typical photometric index and from the templates of blue quasars described in Sect. 4.1, whose continuum slopes, evaluated using our fitting routine, are reported in the inset panel of Fig. 7. We recall that the slopes are evaluated using different continuum windows at different redshifts. For a fixed value of M_{BH} (or equivalently L_{bol}), the distribution in the observed slopes is therefore partly due to the spread along the other axis of the parameter space, which is anyway modest, and partly driven by the different continuum windows.

We assessed the degree of correlation between the slope and the axes of the parameter space by means of the Spearman rank correlation index S (correlation increases in modulus from 0 to 1). We do not find any significant trend with any of the accretion parameters, (i.e. $|S| \leq 0.3$ and associated p -values > 0.05). The inclusion of the bins excluded as a result of the representativeness criterion, does not change significantly the correlation index, as we show in Appendix E.

In addition to exploring the evolution of the continuum within our parameter space, we also explored how the number of sources in each bin affects the continuum properties with respect to the global template. In particular, in Fig. 8 we present the relative difference between the continuum slope for each bin (α_{bin})




Fig. 8. Values of $\Delta\alpha$ ($\Delta\alpha = 1 - \alpha_{\text{bin}}/\alpha_{\text{FS}}$) in the bins of our parameter space as a function of the objects in each bin. Bins not meeting the representativeness criterion are marked with a circle. The dashed lines define the $\sim 1/\sqrt{N_{\text{obj}}}$ envelope characterising the standard error on the average slope assuming a standard deviation $\sigma_{\alpha} = 1$.

and the full sample composite (α_{FS}) evaluated in the same wavelength range $\Delta\alpha = 1 - \alpha_{\text{bin}}/\alpha_{\text{FS}}$ against the number of objects in the bin. Regardless of the redshift, the difference decreases in modulus for increasingly populated bins, as if the samples were drawn from a single population. In this picture, the spread in the observed slopes results from sample variance, rather than from an actual evolution with the accretion parameters explored.

4.4. Comparison of the 2500 Å monochromatic fluxes

We produced the normalised monochromatic 2500 Å fluxes ($f_{2500\text{Å}}$) by interpolating the continuum resulting from the best fit slope and assuming a unitary monochromatic flux at 3000 Å. In Fig. 9 we present $f_{2500\text{Å}}$ against $\log M_{\text{BH}}$, $\log L_{\text{bol}}$, and $\log \lambda_{\text{Edd}}$ in order to study any possible evolution. In principle, a systematic change of the $f_{2500\text{Å}}$ with the accretion parameters would require a correction to ‘standardise’ the UV proxy. However, there is evidence against this possibility.

In Signorini et al. (2023) we showed that the monochromatic 2500 Å flux is a good UV proxy for the $L_{\text{X}} - L_{\text{UV}}$ relation, although other wavelengths in the range 1350–5100 Å can perform equally well (see also Jin et al. 2024). Therefore, even




Fig. 9. Values of the normalised 2500 Å fluxes along the axes of the parameter space, assuming a common normalisation of the continuum at 3000 Å in all the bins. Bins not meeting the representativeness criterion are marked with a circle.

assuming different slopes, the UV term of the $L_X - L_{UV}$ relation would hardly be affected.

Here we aimed at quantifying the effect on the monochromatic 2500 Å fluxes of the differences in the spectroscopic slopes within our sample. The slopes are comprised between roughly $[-2.0, -1.1]$, with the extremal values found in scarcely populated bins (respectively 9 and 6 objects), which would be excluded as a result of the procedure described in Sect. 3.1. Since our goal is to corroborate our sample for its cosmological application, we can estimate how much the spread in the observed slopes affects, for instance, the distance modulus (DM) estimates. Here we show that, even assuming a systematic evolution in the continuum slope with one of the explored parameters (an occurrence which is not supported by our analysis), the variation induced in the $f_{2500\text{Å}}$ would propagate into a negligible difference in the DM for our scopes. We can easily check this by expressing DM in terms of $f_{2500\text{Å}}$ and adopting the expression (see Risaliti & Lusso 2015, 2019; Lusso et al. 2020)

$$\text{DM} = 5 \left[\frac{\log f_X - \beta - \gamma (\log f_{2500\text{Å}} + 27.5)}{2(\gamma - 1)} - \frac{1}{2} \log(4\pi) + 28.5 \right] - 5 \log(10 \text{ pc}), \quad (1)$$

where f_X and $f_{2500\text{Å}}$ are, respectively, the X-ray and UV flux densities in units of $\text{erg s}^{-1} \text{cm}^{-2} \text{Hz}^{-1}$. The parameter $\log f_{2500\text{Å}}$ is normalised to the value of 27.5, whereas the logarithm of the luminosity distance is normalised to 28.5. The $\gamma = 0.600 \pm 0.015$ term is the slope of the $L_X - L_{UV}$ relation reported in L20, and β is the intercept. Here we can express the 2500 Å monochromatic flux in a power-law form pivoting at 5100 Å (i.e. the farthest among the explored characteristic wavelengths). By differentiating Eq. (1) with respect to the slope of the continuum α_λ , we get

$$\Delta \text{DM} = \left| -\frac{5\gamma}{2(\gamma - 1)} \log\left(\frac{2500}{5100}\right) \right| \Delta \alpha_\lambda. \quad (2)$$

We can then substitute the standard deviation of the observed slopes ($\sigma_\alpha = 0.17$) as a measure of the spread of the $\Delta \alpha_\lambda$ term. The DM variation ΔDM is ~ 0.2 in modulus, is not sufficient to explain the differences observed with the ΛCDM at high redshifts (see 13). If, as we believe, 2500 Å is a better approximation to the UV pivot point than 5100 Å, the average ΔDM is expected to be even smaller.

4.5. Host galaxy contamination

Although the broadband properties of the objects included in the sample are designed to minimise the additional contribution of the host galaxy in the quasar spectrum, here we perform an a posteriori test based on the spectral data to estimate

the average degree of host galaxy emission in our sample. In Rakshit et al. (2020) a host galaxy–quasar decomposition was performed using a principal component analysis (PCA, Yip et al. 2004) in two windows close to 4200 Å and 5100 Å, obtaining reliable results for $\sim 12\,700$ objects below redshift $z \leq 0.8$. We employed these measurements to assess the level of residual (i.e. not filtered out by our criteria) host galaxy contamination within the quasars of our sample. With this aim, we built a purposefully designed sample of sources below $z = 0.8$ to ensure the coverage of both the host fraction indices, from the latest WS22 catalogue by adopting the following criteria.

We considered only the objects for which the host galaxy fraction indices at 4200 Å (HOST_FR_4200) and 5100 Å (HOST_FR_5100) were both evaluated, with values ranging continuously from 0 (quasar-dominated emission) to 1 (galaxy-dominated emission). If any of the two indices was flagged as not measured (i.e. = -999), the object was discarded.

We defined the host galaxy fraction index (HGF) and created 10 bins of width $\Delta \text{HGF} = 0.1$. A source was a given HGF index when both the 4200 Å and the 5100 Å fractions fell within the same HGF bin in order to define an unambiguous index. The sample so constructed has 5748 objects.

We then stacked the spectra of the objects in each bin according to the same procedure described in Sect. 3.2. In Fig. 10 we show the stacked spectra as a function of their HGF index. To perform a fair comparison, we also produced a stack of all the 529 sources in our sample with $z \leq 0.8$ (shown in black). On average, the quasars in the low redshift tail of our sample retain minimal levels ($\lesssim 10\%$) of host galaxy contamination below 5400 Å, in good agreement with VB01 who report a galaxy flux fraction between 7% and 15% on the basis of the Ca II $\lambda 3934$ absorption feature. At redder wavelengths, the host contribution is expected to become increasingly important. In addition, we also note that the difference is terms of continuum slope and line strength is more evident in the UV region of the spectra, where our template exhibits the best match with the other samples of blue quasars. This test confirms that our UV proxy for the $L_X - L_{UV}$ relation is not systematically affected by host galaxy contamination.

4.6. Intrinsic extinction

Our photometric selection is designed to exclude dust reddened quasars, by retaining only sources near the locus in the $\Gamma_1 - \Gamma_2$ plane where blue unobscured quasars are expected to reside (Richards et al. 2006). We again checked, a posteriori, that the spectra composing our sample only accept minimal levels of extinction, and we bring three arguments in favour of this hypothesis.

First, as a general consideration, we find a very good agreement between our stacks and the VB01 SDSS template, which




Fig. 10. Comparison between the stacks in bins of HGC and our reference spectrum (black) for the $z \leq 0.8$ interval, as colour-coded in the legend. All the spectra are scaled to their 4200 Å flux. Increasingly host-dominated composites exhibit flatter UV continua.

is made of sources whose colour-selection had been designed to select blue quasars (Richards et al. 2001). The agreement is excellent – especially on the bluer UV side – also with the other average spectra shown in Sect. 4.1. This implies that, if the quasars used to build such templates are not redder than the average, the same should hold for our sample. We also recall that, by selection, the VB01 template was built using objects with at least one broad emission line, thus the presence of Type 2 obscured objects should be limited.

As a second test, to get a more quantitative assessment of the effect on the spectral slope of increasing reddening, we reddened the VB01 template under the reasonable hypothesis that this is representative of the intrinsic spectrum of a typical blue quasar. We reddened the template with increasing $E(B-V)$ from 0 to 0.1 with steps of 0.01, using the extinction curve described in Gordon et al. (2016) for an average Small Magellanic Cloud (SMC) extinction curve, which according to several authors is a good approximation of the AGN extinction curve (see e.g. Brotherton et al. 2001; Hopkins et al. 2004), and then checked the agreement with our full sample stack. A more detailed discussion on how a different extinction curve can affect this procedure is presented in Appendix C, where other prescriptions are explored. In brief, we found that the variation of the continuum slope for a given colour excess derived adopting the Gordon et al. (2016), used to account for intrinsic reddening throughout this work, is the closest to the average (see Fig. C.1). The result of the reddening procedure is shown in Fig. 11. The best match is obtained for minimal levels of intrinsic reddening, i.e. $E(B-V) \lesssim 0.01$, an estimate similar to that provided by Reichard et al. (2003) on the basis of spectral data from the SDSS early data release. To quantify this effect, we performed a spectral fit of the reddened templates in the region between 1300 and 4100 Å, and compared the slope of the continuum with that obtained by the fit of our stack in the same region. This test confirmed the finding emerging from a crude comparison of the overall spectral shape, that is a continuum slope consistent with a low degree of intrinsic extinction. The slopes resulting from the spectral fit of the reddened templates are shown in the inset of Fig. 11. We note that, choosing a different template than VB01 would produce analogous results, given the similarity of all the explored templates in the near UV.




Fig. 11. Extinction of the VB01 template according to the colour-coded $E(B-V)$ values. All the spectra are scaled by their 4200-Å flux, where extinction becomes less relevant. The black spectrum is our full sample stack and is in very good agreement with a low degree of extinction. The inset shows the best-fit slope in the range 1300–4100 Å of the VB01 template reddened as a function of the $E(B-V)$ value. The black dot-dashed line represents the slope of the continuum fitted to our stack. Minimal levels of extinction are allowed to match the reddened template with our composite. We note that the visual match between the L20 and the VB01 reddened template depends slightly on the scaling wavelength. The slope derived from the fits, however, does not depend on such a wavelength.

Finally, we estimated the Balmer decrement of the low-redshift sources with available broad $H\alpha$ and $H\beta$ emission lines to assess possible levels of intrinsic extinction of the BLR. Dong et al. (2007) computed the distribution of the broad-line Balmer decrement $H\alpha_{br}/H\beta_{br} \equiv f(H\alpha_{br})/f(H\beta_{br})$ for a sample of 446 low-redshift ($z \lesssim 0.35$) blue AGN, finding that the distribution can be well approximated by a log-normal, with a mean value of $\langle \log(H\alpha_{br}/H\beta_{br}) \rangle = 0.486$ and a standard deviation of 0.03. A value of $\log(3.1) = 0.49$ is generally adopted for the Balmer decrement in the narrow line region in AGN (Halpern & Steiner 1983; Gaskell & Ferland 1984), while it is found to vary widely according to different BLR conditions (e.g. Dong et al. 2007 and references therein).

In this context, large Balmer decrements are interpreted as due to local reddening, as also suggested by large ratios between infrared emission and broad lines (Dong et al. 2005) and significant X-ray absorption (Wang et al. 2009). In Fig. 12, we show the distribution of $\log(H\alpha_{br}/H\beta_{br})$ ratios assuming the values reported in the WS22 catalogue for the 167 objects in our sample with reliable $H\alpha$ and $H\beta$ measurements (see Sect. 4 of WS22 for the detailed criteria). The mean value of the distribution is $\langle \log(H\alpha_{br}/H\beta_{br}) \rangle = 0.48$, and the standard deviation is 0.14. Under the assumption that the sample of blue quasars described in Dong et al. (2007) retains, on average, minimal levels of BLR extinction (see their Sect. 4.1), we find that the intrinsic extinction is negligible, at least for the low redshift sample.

Finally, the combination of the three tests proposed here strongly suggests that intrinsic reddening should be minimal across the entire redshift range covered by our sample.




Fig. 12. Distribution of $\log(H\alpha_{br}/H\beta_{br})$ for our sources with available data. The mean values from other literature samples (Osterbrock 1977, O77, Neugebauer et al. 1979, N79, Goodrich 1990, G90, Greene & Ho 2005, GH05, Zhou et al. 2006, Z06, Dong et al. 2007, D07, Mejía-Restrepo et al. 2022, MR22, Ma et al. 2023, M23) are also shown for reference according to the colour code. The mean of our values is shown in black. The BASS sample described in Mejía-Restrepo et al. (2022) has been split according to the column density N_H inferred from the X-ray analysis.




Fig. 13. Effect of the intrinsic extinction on the Hubble–Lemaître diagram. *Top panel:* Hubble–Lemaître diagram of the sources used to investigate the possible effects of local reddening. The redshift bins used for this purpose are colour-coded in the legend together with the number of objects per bin and the corresponding $E(B - V)$. Coloured stars represent the average values in each redshift bin, while black stars mark the Λ CDM model expectations, with red segments joining these two values. Increasingly high values of $E(B - V)$ are needed in order to produce the observed tension with the concordance cosmological model. *Bottom panel:* difference between the predicted and observed distance modulus (ΔDM). Values are not extrapolated below $z = 1.0$ and above $z = 4.8$.

4.7. Is the Λ CDM tension driven by reddening?

One of the main consequences of reddening would be the underestimate of the 2500 Å monochromatic fluxes. A relatively small amount of extinction, such as $E(B - V) \approx 0.1$, would diminish the 5100-Å flux by $\sim 25\%$, but the extinguished radiation would be roughly double (depending on the assumed extinction curve) at 2500 Å, as a result of the differential cross section of dust.

The net effect of extinction is thus the underestimation of the DM, enhancing the tension between our cosmographic best-fit model and the high-redshift ($z \gtrsim 2-3$) extrapolation of the Λ CDM model (Risaliti & Lusso 2019; Bargiacchi et al. 2023). Here we focus on the possibility that the observed tension with the Λ CDM model is caused by local dust extinction affecting the UV measurements. To this end, we performed the following exercise, neglecting any possible effects on the X-ray fluxes. We assumed that the difference between the DM predicted by Λ CDM and that based on the $L_X - L_{UV}$ relation is entirely driven by local (i.e. at the source redshift) extinction on $f_{2500\text{Å}}$. Therefore, assuming a standard extinction curve, we can evaluate the average $E(B - V)$ by which $f_{2500\text{Å}}$ should be de-extinguished in order to match the predicted DM. With this quantity in hand, we can de-redden the UV spectra and check how much their intrinsic shapes would differ from those observed. In the following, the term “de-reddened” will be used to indicate the extinction-corrected spectral fluxes that produce a DM complying with the Λ CDM model.

Fig. 13 shows the samples used for this purpose. We gathered spectral data of objects at $z \geq 1$, and divided them into 8 logarithmically equi-spaced redshift bins. For each of them, we evaluated the average redshift and DM_{obs} values. The DM can be written in terms of $f_{2500\text{Å}}$ by using Eq. (1) as:

$$DM = -\frac{5}{2} \frac{\gamma}{\gamma - 1} \log f_{2500\text{Å}} + C, \quad (3)$$

where C includes the X-ray flux contribution, the normalisation β , and the other constants. If we assume that the measured DM values are affected by reddening, we can evaluate the amount of extinction needed to place the intrinsic DM exactly onto the Λ CDM expectations and de-redden $f_{2500\text{Å}}$ accordingly. Such a correction can be expressed as (see Appendix F for the explicit derivation):

$$E(B - V) = \Delta DM \frac{1}{R_V} \frac{A_V}{A_{2500\text{Å}}} \frac{1 - \gamma}{\gamma}, \quad (4)$$

where ΔDM is the difference between the average distance modulus in a given redshift bin according to the Λ CDM model and that observed based on the $L_X - L_{UV}$ relation. Assuming a selective to total extinction ratio $R_V = 3.1$ and the SMC extinction curve from Gordon et al. (2016) used above, we retrieve the values listed in the legend of Fig. 13.

Next, we produced the composite spectrum in each bin, de-reddened it by applying the $E(B - V)$ values found in the previous step, and performed the spectral fits to each of them to evaluate the slope of the continuum. We show the so derived spectra and the slopes of the continuum obtained by means of spectral fits in Fig. 7. The slopes characterising the continua of the de-reddened spectra are not compatible with the expectations based on literature samples, as shown in the inset of Fig. 7. Indeed, moving towards higher redshift, under the hypothesis of local reddening, we would select intrinsically bluer objects with increasing amounts of extinction, whose UV slopes are anomalously steeper with respect to those typically found in blue quasars, while still delivering similar spectral shapes at all redshifts.

With the aim to explore the consequences of de-reddening the spectra by such amounts in terms of quasar physics, we built the semi-analytic (i.e. using our bin-averaged data and adopting analytic approximations from the literature in the other bands) intrinsic quasar SED in each redshift bin and computed the hard X-ray (2–10 keV) to bolometric corrections ($k_{\text{bol},X}$) to be compared with other literature samples. Anomalously high $k_{\text{bol},X}$




Fig. 14. Average hard X-ray bolometric corrections in each redshift bin. As a term of comparison, we report the $k_{\text{bol},X}$ for the semi-analytical SED templates from M04 (blue squares) and the empirical ones from D20 (pentagons). The blue dashed line represents the best-fit reported in M04. Black and red continuous lines show, respectively, the best fit of our sample and the D20 objects in the high L_{bol} ($\geq 10^{46}$ erg s $^{-1}$) regime. The red dashed lines represent the 95% confidence interval of the best-fit regression line for the selected high-luminosity objects from D20 (red pentagons). Our values of $k_{\text{bol},X}$ from the high-redshift, de-reddened SEDs are inconsistent with the extrapolation of the lower luminosity trend from M04, and steadily offset with respect to our best fit of the D20 high-luminosity type-1 quasars. Inset: average SEDs in each redshift bin, according to the same colour code as in Fig. 13. The shaded area is given by the individual mock SEDs in each bin.

would further testify the selection of a very peculiar class of objects at high redshift, seemingly at odds with the typical spectral properties upon which the sample is built, and which we have verified so far.

In order to produce the average SEDs we built a semi-analytical parametrisation, similar to that discussed in Marconi et al. 2004 (M04). For wavelengths $\lambda > 1 \mu\text{m}$ we adopted a power-law spectrum $L_\nu \propto \nu^{\alpha_\nu}$, with slope $\alpha_{\nu,IR} = 2$ as in the Rayleigh-Jeans tail of the disc. Between $1 \mu\text{m} > \lambda > 1200 \text{ \AA}$ we adopted the spectral index $\alpha_{\nu,UV}$ found in the fits of the de-reddened spectra, under the assumption that at high redshift and high luminosity the continuum can be approximated as a single power law. In the range $1200 \text{ \AA} > \lambda > 600 \text{ \AA}$ we chose the average far UV spectral slope $\alpha_{\nu,FUV} = -1.70$ described in Lusso et al. (2015), who took into account a statistical correction on the neutral hydrogen column density of the intergalactic medium. On the X-ray side, we used the *pexrav* model (Magdziarz & Zdziarski 1995) in the XSPEC software (Arnaud 1996) to simulate a power-law spectrum with a photon index Γ equal to the average photometric photon index derived from *XMM-Newton* photometry (see Sect. 4.8) with a cutoff at 500 keV, and a reflection component with solid angle of 2π , inclination angle of $\cos i = 0.5$, and solar metal abundances⁴. Each X-ray spectrum was then scaled to the average 2 keV luminosity in each bin.

⁴ This is a simplistic description of the X-ray spectrum, as we do not take into account the so-called “soft excess” component, nor any additional spectral complexity. However, for the sake of simplicity and consistency with M04 we adhere to this description.

Lastly, we connected the 600 \AA and the 0.5 keV luminosities with a power law. We defined 8 template SEDs according to the 8 values of $L_{2500 \text{ \AA},\text{int}}$ and $\alpha_{\nu,UV}$, and used them to compute the hard X-ray (2–10 keV) bolometric corrections $k_{\text{bol},X}$. We estimated the uncertainty on $k_{\text{bol},X}$ by building 100 mock SEDs in each bin assuming Gaussian distributions for the adopted parameters with standard deviations $\sigma(\alpha_{\nu,IR}) = 0.1$, $\sigma(\alpha_{\nu,UV})$ equal to the uncertainty on the continuum slope evaluated from the fit of each de-reddened spectrum, $\sigma(L_{2500 \text{ \AA},\text{int}})$ equal to the standard deviation on the mean of the $L_{2500 \text{ \AA},\text{int}}$ values in each redshift bin, $\sigma(\alpha_{\nu,FUV}) = 0.6$, $\sigma(L_{2 \text{ keV}})$ and $\sigma(\Gamma)$ equal to the standard deviation on the mean of the $L_{2 \text{ keV}}$ and Γ distributions in each redshift bin.

The average SEDs resulting from this procedure are shown in Fig. 14, together with other literature samples as considered in Duras et al. 2020 (D20). Our $k_{\text{bol},X}$ values from the de-reddened SEDs clearly detach from the high-luminosity extrapolation of the M04 best-fit relation. At the same time, they are all located on the upper envelope of the D20 best fit for luminous type-1 sources. This implies that, if such sources existed, we would be selecting an exotic population made of the most luminous, yet dust-extincted quasars. These sources would exhibit the highest known bolometric conversion factors, while still showing the UV-optical colours and spectral properties of typical blue quasars. Moreover, if our SED templates are consistent by construction with M04, which then represents a proper comparison, the same does not hold for the empirical SEDs in D20. Indeed, we did not include any physically motivated emission in the spectral region between the far UV and the soft X-rays, as we only linked the $L_{600 \text{ \AA}}$ and $L_{0.5 \text{ keV}}$ with a straight power law, while in the latter work the authors could also rely on far-UV photometry from GALEX. In order to quantify the possible inconsistency of our average estimates with those from the other literature samples, we performed a basic fit of the $k_{\text{bol},X} - L_{\text{bol}}$ relation in our de-reddened sources. We adopted a simple power-law functional form, which is a good approximation for the high luminosity ($L_{\text{bol}} \geq 10^{46}$ erg s $^{-1}$) end of the D20 distribution:

$$\log(k_{\text{bol},X}) = a + b \left[\log \left(\frac{L_{\text{bol}}}{10^{47.5} \text{ erg s}^{-1}} \right) \right]. \quad (5)$$

The best fit parameters and their respective uncertainties are listed in Table 3. In Fig. 14 we show the $k_{\text{bol},X} - L_{\text{bol}}$ relation for our sample and the D20 compilation of type-1 objects. The bolometric corrections of our higher redshift bins are located outside the 95% confidence interval for the best fit of the high-luminosity tail of the D20 distribution (red pentagons). The offset a of the best fit relation for our sources is significantly higher (3.7σ) than for those from D20. We checked that this result is not biased by the choice of the luminosity threshold for the objects included in the fit, by performing again the fit increasing the threshold by steps of $\Delta \log(L_{\text{bol}}/\text{erg s}^{-1}) = 0.5$ from 45.5 to 47. In each case, the high-redshift ($z \geq 2.5$) bolometric correction factors remain located above the 95% confidence interval. In conclusion, the bolometric correction estimates derived from our average de-reddened SEDs are significantly offset with respect to the expectations for high-luminosity type-1 quasars, hinting at a seemingly exotic – and yet unknown – population.

As a result of all the tests performed in this Section, we conclude that a dust extinction scenario as an explanation for the tension with the Λ CDM model emerging in the quasar Hubble–Lemaître diagram seems disfavoured for several reasons. Firstly, significant amounts of extinction at high redshift, in combination with steeper than the average UV spectra, are required in

Table 3. Results of the fit of the $k_{\text{bol,X}} - L_{\text{bol}}$ relation for this work and the template SEDs of the most luminous sources in D20.

	a	b
This work	2.402 ± 0.042	0.484 ± 0.050
D20	2.184 ± 0.041	0.455 ± 0.040

order to reconcile the tension with the Λ CDM model, unless we assume a cosmological evolution from high-redshift extremely blue sources to low-redshift typical quasars. Secondly, the reddening should increase with redshift, in opposition to the current paradigm of increasing dust with time (see, e.g. Aoyama et al. 2018, and references therein). In addition, the dust content in the high-redshift quasars should increase by the exact amount needed to compensate for the steepening of the intrinsic spectrum, thus producing the typical spectrum of blue quasars at all redshifts. Lastly, the hard X-ray bolometric correction factors derived from the average de-reddened SEDs are in tension with the high luminosity extrapolation from known sources.

All the above considerations would imply the selection of exotic objects at high redshift whose emission properties are intrinsically different from those at lower redshift and luminosity, while their apparent broadband properties are consistent with those of typical blue quasars.

4.8. The X-ray stacks

The availability of X-ray data, which underlies the creation of the L20 sample, allowed us to address the question on the evolution of the shape of X-ray emission with respect to the accretion parameters and redshift. However, at this stage, a full stacking of the X-ray spectra, taking into account the largely different redshift intervals sampled, as well as the different response functions is still unfeasible. For this reason, we adopted a slightly different approach and built ‘photometric X-ray stacks’. The intrinsic X-ray spectrum of an unobscured quasar above 0.2 keV can be roughly described as a power law with spectral index $\alpha = 1 - \Gamma$, where Γ is the photon index. For our sample the best estimate for the Γ parameter is the photometric photon index (Γ_{ph}) derived following the procedure described in Section 4 of L20 (where also a more detailed description of photometric spectral indices and of their limitations can be found). Although the photometric Γ values are somewhat less reliable than those spectroscopically derived (Γ_{sp}), the lack of a systematic offset in the $\Delta\Gamma$ distribution (defined as $\Delta\Gamma = \Gamma_{\text{sp}} - \Gamma_{\text{ph}}$) and its non-evolution with redshift allow us to use these estimates in a statistical sense. For each source we built a synthetic spectrum between 0.2 and 10 keV using Γ_{ph} and the 2 keV flux reported in L20. Then, the procedure to build the X-ray stacks is analogous to that explained in Sect. 3.2.

We note, however, that the X-ray stacks are less informative than those produced in the UV-optical for at least two reasons: 1) they are not computed using the actual X-ray spectra, but rather photometric proxies, which are not sensitive to the peculiar features of each spectrum other than the steepness of the spectrum itself; 2) the quality cuts performed in the X-ray selection (especially that concerning the lower limit on the photon index, i.e. $\Gamma_{\text{ph}} - \delta\Gamma_{\text{ph}} > 1.7$, with $\delta\Gamma_{\text{ph}}$ being the uncertainty) are more conservative, given the generally lower quality of the X-ray data, and thus have a much greater impact on obliterating the spectral diversity than the optical cuts. Nonetheless, the anal-

ysis of the slope of the stacked X-ray spectra provides valuable information about the goodness of our selection as well as the possible evolution of the X-ray properties within our parameter space. In Fig. 15 we recover the well-known correlation between $\Gamma - \log \lambda_{\text{Edd}}$ (Shemmer et al. 2008; Risaliti et al. 2009; Jin et al. 2012; Brightman et al. 2013; Trakhtenbrot et al. 2017), with a correlation index $S = 0.38$. Interestingly, the correlation index is marginally larger for the $\Gamma - \log M_{\text{BH}}$ ($S = -0.41$), but this could be due to the way the photometric data were grouped to produce the stacks, rather than to an intrinsically tighter correlation.

In theory, the dependence of the 2 keV fluxes on M_{BH} or λ_{Edd} implies that adding one of these parameters, or both, to the $L_X - L_{\text{UV}}$ space, the dispersion could be further narrowed down. For instance, the dependence on M_{BH} , proportional to the Mg II FWHM, has been explored in Signorini et al. (2023), where it has been shown that the inclusion of this parameter has a negligible effect in improving the intrinsic dispersion of the relation. At the same time, the zero-covariance between the X-ray flux estimated at the pivot point (i.e., the effective-area-weighted average energy of any given band of interest) and Γ (see Section 4 in L20) assures that even a systematic steepening of the photon index would not affect our estimates of the X-ray luminosity.

5. Discussion

The suitability of quasars as ‘standardized candles’ relies upon the consistency of the method (i.e. the absence of selection/analysis biases, or correction thereof) and the intrinsic non-evolution of quasar spectra with redshift. Although increasingly larger samples are desirable, our group has been lately devoting a substantial effort to investigate any possible issues affecting the L20 sample with SDSS spectral coverage. In this context, this work has been chiefly focused on the analysis of the average spectral properties of 2316 of the sources described in L20. The main results obtained through the various analyses performed in the previous sections are the following: 1) The overall shape of the composite spectra (apart from emission-line properties) does not show any clear evolution across the explored parameter space. 2) Host galaxy contamination is not a major issue in our sources below $z \leq 0.8$, and we therefore expect it to be completely negligible for the bulk of our sample, at higher redshift and luminosity. 3) Our sample is compatible, on average, with a low degree of intrinsic extinction, $E(B - V) \lesssim 0.01$. 4) The deviation of the cosmographic best fit of the quasar–SN Ia Hubble–Lemaître diagram from the Λ CDM model reported by our group is unlikely to be caused by dust reddening, based on the spectral properties of our sample. These findings have far-reaching consequences on the physical properties of the sources described here, as well as on the solidity of the criteria adopted to select our cosmological sample.

The similarity of quasar spectra has led through the years to the production of several composite spectra (e.g. Francis et al. 1991, VB01, Zheng et al. 1997; Brotherton et al. 2001; Harris et al. 2016), which have found large employment as templates for the typical quasar emission, in terms of line and continuum properties. In our case, building high signal-to-noise composite spectra allowed us to explore the spectral diversity within the parameter space underlying SMBH accretion. We do not report any clear evidence for a redshift-dependent evolution of the average continua in our sample, as the agreement with either the global L20 or the independent VB01 composite spectra does not deteriorate with redshift. This appears to hold up to very high redshift, where luminous quasars have been observed to show spectral properties very akin to those seen at




Fig. 15. Values of Γ for our X-ray stacks across the parameter space. Bins not meeting the representativeness criterion are marked with a circle.

lower redshift (e.g. Fig. 1 in Mortlock et al. 2011). Under the paradigm that such emission is driven by SMBH accretion, these results would imply that the emission mechanism underlying the observed spectra does not vary significantly in the blue quasars spanning the last 12 Gyr. Although we do observe the expected correlations between line equivalent width and continuum luminosity (i.e. the Baldwin effect, Baldwin 1977) and between line width and BH mass (virial relations; e.g. Peterson et al. 2004), the non-evolution of the continuum shape of our composite with any of the quantities in our parameter space is remarkable (see Fig. G.1).

Contamination from the host galaxy light is a known issue when trying to isolate the AGN component (e.g. Baldwin et al. 1981; Kauffmann 2003). This effect is capable, in principle, of inducing the overestimate of the actual 2500 Å fluxes, by introducing a spurious UV contribution and eventually biasing the estimate of the cosmological parameters based on the quasar $L_X - L_{UV}$ relation. For this reason, our selection criteria require steep photometric continua to single out sources where the contribution from the galactic stellar component is negligible with respect to the AGN emission (Elvis et al. 1994; Richards et al. 2006; Elvis et al. 2012; Krawczyk et al. 2013; Hao et al. 2013). As shown in Section 4.5, host galaxy contamination does not play a major role in our sample, as we generally observe steep continua, coupled with the lack of strong stellar absorption features and weak narrow emission lines from stellar photoionisation. These findings are particularly evident on the UV side of the spectrum, where our composites are dominated by higher redshift, luminous objects, in which the host contribution is expected to be negligible above $L_{bol} \sim 10^{46}$ erg s⁻¹ (Shen et al. 2011). A detailed disentanglement of the host galaxy and quasar emission by means, for instance, of a principal component analysis (e.g. Yip et al. 2004) is beyond the scope of this work. Nonetheless, we can roughly estimate the average incidence of host galaxy light by comparing the absorption features between our full composite and VB01. The EW of the Ca II λ 3934 absorption in our composite (0.7 ± 0.4 Å) is similar to that observed in the VB01 template (0.9 ± 0.1 Å), yet a fair comparison is hampered by the differences in the resolution and the signal to noise between the two templates. This goes in the direction of our sample having an average host contribution comparable to that of VB01, deemed 7%–15% at such wavelengths, with an incidence decreasing bluewards.

Recently, Zajaček et al. (2023) have investigated the possibility that luminosity distance measurements based on the $L_X - L_{UV}$ relation are biased by extinction on a sample of 58 quasars at $z \leq 1.7$ with available reverberation mapping data (Khadka et al. 2023). In the latter work, the authors used Mg II time-delays to infer distances by means of the radius–luminosity relation (Mg II $R-L$ relation, with intrinsic dispersion $\sigma \sim 0.3$ dex; Khadka et al.

2022). Assuming that the difference between the distances based on the Mg II $R-L$ and $L_X - L_{UV}$ relations were attributable to local reddening, and adopting different cosmological models, they derive an intrinsic colour excess $E(B - V) \simeq 0.002-0.003$, an amount of extinction that would be intrinsic to all type-1 quasars. However, such an effect, even if present, would still not be enough to explain the reported high-redshift tension with the current cosmological model.

In this framework, the shape of the quasar continuum can be interpreted as ultimately resulting by the combination of two main parameters: the intrinsic spectral slope of the UV spectrum and the column density of gas and dust causing the extinction. The degeneracy between these two parameters could be capable, in principle, of producing similar spectra by combining increasingly larger amounts of dust in high redshift, luminous sources with intrinsically steeper spectra (e.g. Xie et al. 2015).

There are several reasons to believe instead that dust extinction, even if present in small amounts in our sample, is not systematically biasing our results. First, the photometric selection, besides excluding sources where the host galaxy significantly reddens the UV-optical SED, also singles out objects where the amount of dust extinction is expected to be negligible. Indeed, assuming that the photometric slopes of the average quasar SED described in Richards et al. (2006) are truly representative of blue quasars, by selection, the colour excess of our sources should not exceed $E(B - V) \simeq 0.1$. We also note that the incidence of significantly reddened objects with $E(B - V) > 0.1$, at least in early SDSS data (whose average properties are encapsulated in the VB01 template) was estimated to be modest, below <2% according to Hopkins et al. (2004). In this work, we took a step further, searching a posteriori for spectral signatures of dust reddening, but we found none. Moreover, the distribution of the broad-line Balmer decrements from the H α and H β reported in the WS22 catalogue for the low-redshift tail ($z < 0.57$) of our sample is in excellent agreement with other samples made of allegedly unobscured sources (e.g. Zhou et al. 2006; Dong et al. 2007). In Section 4.6 we showed that the spectral shape of our global template is consistent with the VB01 template reddened by minimal levels of local extinction, i.e. $E(B - V) \lesssim 0.01$. Similar conclusions have also been drawn in Reichard et al. (2003) who estimated the intrinsic reddening in non-BAL AGN from the early SDSS release to be $E(B - V) \sim 0.004$. This implies, under the assumption that the VB01 template is made of mostly blue unobscured quasars, a low incidence of reddening (dust/gas extinction) and galactic contamination. Finally, the similarity of the spectral shape between our sample and other quasar composite spectra from several surveys designed to select luminous blue quasars (Francis et al. 1991; Brotherton et al. 2001; VB01; Selsing et al. 2016; Harris et al. 2016) can be interpreted as a proof of the lack of significant reddening, and a clue of the

non-evolution through cosmic time of the intrinsic emission properties of quasars. The other possible explanations to interpret the overall agreement among composite spectra from different samples involve a certain degree of fine-tuning. Several samples of quasars on different scales of luminosity and redshift should have intrinsically steep spectral shapes and significant amounts of dust, but still producing similar spectral shapes. As a corollary to this consideration, we should also conclude that all the quasars used to build the average SEDs available in the literature (Richards et al. 2006; Krawczyk et al. 2013; Gu 2013), whose colours have been adopted to select our sample, are also affected by the same issue.

If we were to attribute the tension with the Λ CDM in the Hubble–Lemaître diagram to dust reddening, a constant extinction with redshift would not be enough. This would require a differential extinction (increasing with redshift) combined with increasingly steeper intrinsic spectra, resulting in observed spectral shapes virtually independent of redshift.

Yet, this introduces additional questions, as larger intrinsic luminosities would also imply larger SMBH masses, enlarging the discrepancy between the largest possible SMBH masses in the Eddington-limited regime and those observed, unless the accretion efficiency varies systematically at high redshift (see, for example, Inayoshi et al. 2020 and Lusso et al. 2022 for recent reviews on this topic). In addition, significant local extinction $E(B - V) \sim 0.1$ would imply an intrinsic UV flux larger by a factor whose exact value depends on extinction curve, but on average between 3 and 4 times at 1450 Å. In turn, assuming as a rough approximation a first order constant bolometric correction, the intrinsic bolometric luminosity should be accordingly larger than the observed one. Therefore, the total mass in black holes in the Universe as evaluated by Soltan (1982), would need a significant increase with increasing redshift, unless the M_{BH} values are already underestimated by similar amounts. The same argument applies when considering greybody-like extinction curves (e.g. Gaskell et al. 2004), which do not modify the continuum slope, but underestimate the bolometric luminosities.

The discrepancy between the observed DM and those derived in the Λ CDM framework cannot even be solved invoking for cosmic opacity on the 2500 Å photons which would cause the observed DM to diminish. In addition, an opposite effect on DM would arise when considering cosmic opacity acting on X-ray photons. We can thus consider the decrease in the DM caused by the dimming of the UV emission as an upper limit to the detachment from the Λ CDM model caused by free electrons extinction along the line of sight. In quantitative terms, the extrapolation of the cosmic dimming values derived for SNe Ia in B band in Zhang (2008), Hu et al. (2017) lie below 0.02 at $z < 3$, while our Hubble–Lemaître diagram (Fig. 13) shows much larger values, as ΔDM approaches 1 at $z = 3$. Alternatively, some ubiquitous – yet unknown – material in the IGM should be responsible for a grey-body like absorption, though this option seems unlikely.

On the X-ray side, gas absorption plays an opposite role with respect to the UV counterpart, as the correction for absorption would increase the departure from the Λ CDM model. Generally speaking, the incidence of gas absorption among our sources is minimised by the $\Gamma_{\text{ph}} > 1.7$ criterion. In addition, this effect is further mitigated at high redshift by the shift within the observed 0.5–2 keV band of much higher rest-frame energies, for which the effect of gas absorption becomes increasingly weaker. Indeed, in all the high-redshift objects for which a dedicated X-ray spectral analysis was performed (Vito et al. 2019; Nardini et al. 2019; Salvestrini et al. 2019; Sacchi et al. 2022),

column densities in excess of 10^{23} cm^{-2} would be required to modify the spectral slope and thus affect the extrapolated rest-frame 2 keV flux, but such values were generally excluded.

As a final consideration, the results of the spectroscopic analysis described in this work ultimately prove the reliability of the photometric selection criteria adopted in our works to define a sample of quasars with typical properties. Using the photometric cuts, rather than the more time-consuming spectroscopic analysis, will allow us to considerably enlarge the cosmological sample with the forthcoming surveys (e.g. LSST, DESI), while retaining a high degree of consistency with the spectroscopic properties.

6. Conclusions

In this sixth paper of the series, we presented a thorough analysis of the average spectral properties of 2316 sources with available SDSS data from the sample described in Lusso et al. (2020). Our main aims were to investigate the possible evolution of the UV-optical properties with redshift, and to study the possible systematic effects (e.g. host galaxy contamination, dust extinction) that could undermine the validity of quasars as standard candles. To this end, we built spectral composites in the parameter space defined by M_{BH} , L_{bol} , and redshift to investigate how the spectral properties of the cosmological sample vary as a function of the key accretion parameters and cosmic time. We also carried out specific tests to determine the presence of residual galactic contamination and dust reddening based on their spectral signatures.

Our main findings are listed below:

- Apart from the line emission, which varies according to well-known relations with L_{bol} and M_{BH} , we report a strong similarity in terms of continuum shapes among the composite spectra across the explored parameter space. The slopes of the continuum of our composite spectra cluster around $\langle\alpha_{\lambda}\rangle = -1.50 \pm 0.14$ without statistically significant correlations (i.e. $|S| \leq 0.3$) with the explored parameters (M_{BH} , L_{bol} , λ_{Edd} , z). This points to a non-evolution of the broadband continuum emission within our sample with respect to the considered parameters.
- Our global composite spectrum shows a good agreement with other templates of typical blue quasars as well as with the individual composite spectra in the various regions of the parameter space. This confirms that the L20 catalogue is, on average, made of typical objects.
- By comparing the composite spectrum of our low-redshift ($z \leq 0.8$) sources with SDSS objects affected by heavy host galaxy contamination, we were able to verify that host galaxy emission is negligible in our sample. We estimate the host galaxy fraction at optical wavelengths to be $\leq 10\%$, and likely even lower at UV wavelengths.
- After performing several tests, as listed in Section 4.6, we demonstrated that the objects constituting the L20 sample are consistent with a low degree of intrinsic extinction, with an average colour excess value $E(B - V) \leq 0.01$.
- We showed that the tension between our cosmographic best fit of the joint quasar–SN Ia Hubble–Lemaître diagram, and the Λ CDM model cannot be simply accounted for by allowing for cosmic opacity or differential dust extinction. Otherwise, a colour excess $E(B - V) \gtrsim 0.1$ increasing with z , accompanied by anomalously steep UV spectra, should ubiquitously affect high- z sources while delivering similar spectral shapes at all redshifts.
- The spectroscopic analysis carried out in this work corroborates the photometric selection criteria adopted to define a

sample of quasars with typical properties. This in turn allows for the continued use of the more efficient photometric cuts rather than a more time-consuming individual spectroscopic analysis.

The search for possible systematics performed in this work has broad implications for the cosmological applications of our compilation of quasars. First, the spectral confirmation that dust or galaxy contamination do not significantly affect, without invoking peculiar or anomalous explanations, our cosmological sample further corroborates the reliability of our methods and results. In this context, quasars emerge as consistent standardisable candles, and the high-redshift deviation from the current cosmological model remains outstanding. Secondly, as our method has proven fruitful in producing large samples of homogeneous sources, we can extend and refine its application to the forthcoming data from new facilities. As a matter of fact, the extended ROentgen Survey with an Imaging Telescope Array (eROSITA, Merloni et al. 2012) began populating the pivotal $z \sim 1$ region in the X-ray domain, allowing for a more robust cross-calibration between quasars and SNe Ia and a more precise determination of Ω_{Λ} (Lusso 2020). In the mid- to long-term future, the cross-match between *Euclid* (Laureijs et al. 2011) and LSST (Ivezić et al. 2019) observations in the optical–UV and *Athena* (and/or other proposed facilities) in the X-rays will significantly enlarge the sample, helping to fill up the region at intermediate to high redshift. New large and reliable samples will be key to testing the cosmological models in synergy with other well-established probes with the intent to precisely constrain the cosmological parameters.

Data availability

The spectrum in Fig. 3 and Table D.1 are available at the CDS via anonymous ftp to cdsarc.cds.unistra.fr (130.79.128.5) or via <https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/689/A109>

Acknowledgements. We gratefully acknowledge the anonymous referee for the thoughtful comments which resulted in a significantly improved paper. For catalogue cross analysis we have used the Virtual Observatory software (Taylor 2005), available online at <http://www.star.bris.ac.uk/mbt/topcat/>. The figures were generated with (Hunter 2007), a PYTHON library for publication-quality graphics. The reddening/unreddening of the spectral templates was carried out through the PYTHON package ‘dust_extinction’ (Gordon et al. 2021). For catalogue cross analysis we have used the Virtual Observatory software (Taylor 2005), available online at <http://www.star.bris.ac.uk/mbt/topcat/>. The figures were generated with (Hunter 2007), a PYTHON library for publication-quality graphics. The reddening/unreddening of the spectral templates was carried out through the PYTHON package ‘dust_extinction’ (Gordon et al. 2021).

References

Abdurro’uf, N., Accetta, K., Aerts, C., et al. 2022, *ApJS*, 259
 Antonucci, R., Kinney, A., & Ford, H. 1989, *ApJ*, 342, 64
 Aoyama, S., Hou, K.-C., Hirashita, H., Nagamine, K., & Shimizu, I. 2018, *MNRAS*, 478, 4905
 Arcodia, R., Merloni, A., Nandra, K., & Ponti, G. 2019, *A&A*, 628, A135
 Arnaud, K. A. 1996, *ASP Conf. Ser.*, 101, 17
 Assef, R., Kochanek, C., Brodwin, M., et al. 2010, *ApJ*, 713, 970
 Baldwin, J. A. 1977, *ApJ*, 214, 679
 Baldwin, J. A., Phillips, M. M., & Terlevich, R. 1981, *PASP*, 93, 5
 Bañados, E., Venemans, B. P., Mazzucchelli, C., et al. 2018, *Nature*, 553, 473
 Bargiacchi, G., Risaliti, G., Benetti, M., et al. 2021, *A&A*, 649, A65
 Bargiacchi, G., Benetti, M., Capozziello, S., et al. 2022, *MNRAS*, 515, 1795
 Bargiacchi, G., Dainotti, M., & Capozziello, S. 2023, *MNRAS*, 525, 3104
 Brightman, M., Silverman, J., Mainieri, V., et al. 2013, *MNRAS*, 433, 2485
 Brotherton, M., Tran, H. D., Becker, R., et al. 2001, *ApJ*, 546, 775

Czerny, B., Beaton, R., Bejger, M., et al. 2018, *Space Sci. Rev.*, 214, 1
 Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2012, *AJ*, 145, 10
 Dong, X.-B., Zhou, H.-Y., Wang, T.-G., et al. 2005, *ApJ*, 620, 629
 Dong, X., Wang, T., Wang, J., et al. 2007, *MNRAS*, 383, 581
 Duras, F., Bongiorno, A., Ricci, F., et al. 2020, *A&A*, 636, A73
 Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, *AJ*, 142, 72
 Elvis, M., & Karovska, M. 2002, *ApJ*, 581, L67
 Elvis, M. S., Wilkes, B. J., McDowell, J. C., et al. 1994, *ApJS*, 95, 1
 Elvis, M., Hao, H., Civano, F., et al. 2012, *ApJ*, 759, 6
 Evans, D. 1989, *A&AS*, 78, 249
 Evans, I. N., Primini, F. A., Glotfelty, K. J., et al. 2010, *ApJS*, 189, 37
 Faucher-Giguère, C.-A., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008, *ApJ*, 688, 85
 Fitzpatrick, E. L. 1999, *PASP*, 111, 63
 Francis, P., Hewett, P. C., Foltz, C. B., et al. 1991, *ApJ*, 373, 465
 Gallerani, S., Maiolino, R., Juarez, Y., et al. 2010, *A&A*, 523, A85
 Gaskell, C. M., & Ferland, G. J. 1984, *PASP*, 96, 393
 Gaskell, C. M., Goosmann, R. W., Antonucci, R. R., & Whysong, D. H. 2004, *ApJ*, 616, 147
 Goodrich, R. W. 1990, *ApJ*, 355, 88
 Gordon, K. D., Fouesneau, M., Arab, H., et al. 2016, *ApJ*, 826, 104
 Gordon, K. D., Misselt, K. A., Bouwman, J., et al. 2021, *ApJ*, 916, 33
 Greene, J. E., & Ho, L. C. 2005, *ApJ*, 630, 122
 Gu, M. 2013, *ApJ*, 773, 176
 Haardt, F., & Maraschi, L. 1991, *ApJ*, 380, L51
 Haardt, F., & Maraschi, L. 1993, *ApJ*, 413, 507
 Halpern, J., & Steiner, J. 1983, *ApJ*, 269, L37
 Hao, H., Elvis, M., Bongiorno, A., et al. 2013, *MNRAS*, 434, 3104
 Harris, D. W., Jensen, T. W., Suzuki, N., et al. 2016, *AJ*, 151, 155
 Hopkins, P. F., Strauss, M. A., Hall, P. B., et al. 2004, *AJ*, 128, 1112
 Hu, J., Yu, H., & Wang, F. 2017, *ApJ*, 836, 107
 Hunter, J. D. 2007, *Comput. Sci. Eng.*, 9, 90
 Inayoshi, K., Visbal, E., & Haiman, Z. 2020, *ARA&A*, 58, 27
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111
 Jalan, P., Rakshit, S., Woo, J.-J., Kotilainen, J., & Stalin, C. 2023, *MNRAS*, 521, L11
 Jin, C., Lusso, E., Ward, M., Done, C., & Middei, R. 2024, *MNRAS*, 527, 356
 Jin, C., Ward, M., & Done, C. 2012, *MNRAS*, 425, 907
 Kauffmann, G., et al. 2003, *MNRAS*, 346, 1055
 Kellermann, K. I., Sramek, R., Schmidt, M., Shaffer, D. B., & Green, R. 1989, *AJ*, 98, 1195
 Khadka, N., Zajaček, M., Panda, S., Martínez-Aldama, M. L., & Ratra, B. 2022, *MNRAS*, 515, 3729
 Khadka, N., Zajaček, M., Prince, R., et al. 2023, *MNRAS*, 522, 1247
 Kratzer, R. M., & Richards, G. T. 2015, *AJ*, 149, 61
 Kravczyk, C. M., Richards, G. T., Mehta, S. S., et al. 2013, *ApJS*, 206, 4
 La Franca, F., Bianchi, S., Ponti, G., Branchini, E., & Matt, G. 2014, *ApJ*, 787, L12
 Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, ArXiv e-prints [arXiv:1110.3193]
 Lenart, A. Ł., Bargiacchi, G., Dainotti, M. G., Nagataki, S., & Capozziello, S. 2023, *ApJS*, 264, 46
 Lusso, E. 2020, *Front. Astron. Space Sci.*, 7, 8
 Lusso, E., & Risaliti, G. 2017, *A&A*, 602, A79
 Lusso, E., Comastri, A., Vignali, C., et al. 2010, *A&A*, 512, A34
 Lusso, E., Worseck, G., Hennawi, J. F., et al. 2015, *MNRAS*, 449, 4204
 Lusso, E., Risaliti, G., Nardini, E., et al. 2020, *A&A*, 642, A150
 Lusso, E., Valiante, R., & Vito, F. 2022, The Dawn of Black Holes, in Bambi, C., Santangelo, A. (eds), *Handbook of X-ray and Gamma-ray Astrophysics* (Springer)
 Ma, Y.-S., Li, S.-J., Gu, C.-S., et al. 2023, *MNRAS*, 522, 5680
 Madau, P. 1995, *ApJ*, 441, 18
 Magdziarz, P., & Zdziarski, A. A. 1995, *MNRAS*, 273, 837
 Marconi, A., Risaliti, G., Gilli, R., et al. 2004, *MNRAS*, 351, 169
 Markwardt, C. B. 2009, *ASP Conf. Ser.*, 411, 251
 Marziani, P., & Sulentic, J. W. 2014, *Adv. Space Res.*, 54, 1331
 Mejía-Restrepo, J. E., Trakhtenbrot, B., Koss, M. J., et al. 2022, *ApJS*, 261, 5
 Menzel, M.-L., Merloni, A., Georgakakis, A., et al. 2016, *MNRAS*, 457, 110
 Merloni, A. 2003, *MNRAS*, 341, 1051
 Merloni, A., Predehl, P., Becker, W., et al. 2012, ArXiv e-prints [arXiv:1209.3114]
 Moré, J. J. 1978, *Numerical Analysis* (Springer), 105
 Mortlock, D. J., Warren, S. J., Venemans, B. P., et al. 2011, *Nature*, 474, 616
 Nardini, E., Lusso, E., Risaliti, G., et al. 2019, *A&A*, 632, A109
 Neugebauer, G., Oke, J., Becklin, E., & Matthews, K. 1979, *ApJ*, 230, 79
 Nicastro, F. 2000, *ApJ*, 530, L65
 Novikov, I. D., & Thorne, K. S. 1973, *Black holes (Les astres occlus)*, 1, 343
 Osterbrock, D. E. 1977, *ApJ*, 215, 733

- Pâris, I., Petitjean, P., Aubourg, É., et al. 2018, *A&A*, **613**, [A51](#)
- Peterson, B. M., Ferrarese, L., Gilbert, K. M., et al. 2004, *ApJ*, **613**, [682](#)
- Prevot, M. L., Lequeux, J., Prevot, L., Maurice, E., & Rocca-Volmerange, B. 1984, *A&A*, **132**, [389](#)
- Prochaska, J. X., Worseck, G., & O’Meara, J. M. 2009, *ApJ*, **705**, [L113](#)
- Rakshit, S., Stalin, C., & Kotilainen, J. 2020, *ApJS*, **249**, [17](#)
- Reichard, T. A., Richards, G. T., Hall, P. B., et al. 2003, *AJ*, **126**, [2594](#)
- Richards, G. T., Fan, X., Schneider, D. P., et al. 2001, *AJ*, **121**, [2308](#)
- Richards, G. T., Lacy, M., Storrle-Lombardi, L. J., et al. 2006, *ApJS*, **166**, [470](#)
- Risaliti, G., & Lusso, E. 2015, *ApJ*, **815**, [33](#)
- Risaliti, G., & Lusso, E. 2019, *Nat. Astron.*, **3**, [272](#)
- Risaliti, G., Young, M., & Elvis, M. 2009, *ApJ*, **700**, [L6](#)
- Sacchi, A., Risaliti, G., Signorini, M., et al. 2022, *A&A*, **663**, [L7](#)
- Salvestrini, F., Risaliti, G., Bisogni, S., Lusso, E., & Vignali, C. 2019, *A&A*, **631**, [A120](#)
- Schlafly, E. F., & Finkbeiner, D. P. 2011, *ApJ*, **737**, [103](#)
- Scolnic, D. M., Jones, D. O., Rest, A., et al. 2018, *ApJ*, **859**, [101](#)
- Selsing, J., Fynbo, J. P., Christensen, L., & Krogager, J.-K. 2016, *A&A*, **585**, [A87](#)
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, **24**, [337](#)
- Shemmer, O., Brandt, W., Netzer, H., Maiolino, R., & Kaspi, S. 2008, *ApJ*, **682**, [81](#)
- Shen, Y., Richards, G. T., Strauss, M. A., et al. 2011, *ApJS*, **194**, [45](#)
- Signorini, M., Risaliti, G., Lusso, E., et al. 2023, *A&A*, **676**, [A143](#)
- Soltan, A. 1982, *MNRAS*, **200**, [115](#)
- Tananbaum, H., Avni, Y., Branduardi, G., et al. 1979, *ApJ*, **234**, [L9](#)
- Taylor, M. B. 2005, *ASP Conf. Ser.*, **347**, [29](#)
- Trakhtenbrot, B., Ricci, C., Koss, M. J., et al. 2017, *MNRAS*, **470**, [800](#)
- Vanden Berk, D. E., et al. 2001, *AJ*, **122**, [549](#)
- Vignali, C., Brandt, W. N., & Schneider, D. P. 2003, *AJ*, **125**, [433](#)
- Vito, F., Brandt, W. N., Bauer, F. E., et al. 2019, *A&A*, **630**, [A118](#)
- Wang, T., Zhou, H., Grupe, D., et al. 2009, *AJ*, **137**, [4002](#)
- Wang, J.-M., Songsheng, Y.-Y., Li, Y.-R., Du, P., & Zhang, Z.-X. 2020, *Nat. Astron.*, **4**, [517](#)
- Wang, F., Yang, J., Fan, X., et al. 2021, *ApJ*, **907**, [L1](#)
- Webb, N. A., Coriat, M., Traulsen, I., et al. 2020, *A&A*, **641**, [A136](#)
- Wild, V., Groves, B., Heckman, T., et al. 2011, *MNRAS*, **410**, [1593](#)
- Wu, Q., & Shen, Y. 2022, *ApJS*, **263**, [42](#)
- Xie, X., Shen, S., Shao, Z., & Yin, J. 2015, *ApJ*, **802**, [L16](#)
- Yang, J., Wang, F., Fan, X., et al. 2020, *ApJ*, **897**, [L14](#)
- Yip, C., Connolly, A., Berk, D. V., et al. 2004, *AJ*, **128**, [2603](#)
- Yoshii, Y., Kobayashi, Y., Minezaki, T., Koshida, S., & Peterson, B. A. 2014, *ApJ*, **784**, [L11](#)
- Zafar, T., Möller, P., Watson, D., et al. 2015, *A&A*, **584**, [A100](#)
- Zajaček, M., Czerny, B., & Khadka, N. 2023, AAS, accepted [arXiv:2305.08179]
- Zappacosta, L., Piconcelli, E., Fiore, F., et al. 2023, *A&A*, **678**, [A201](#)
- Zhang, P. 2008, *ApJ*, **682**, [721](#)
- Zheng, W., Kriss, G. A., Telfer, R. C., Grimes, J. P., & Davidsen, A. F. 1997, *ApJ*, **475**, [469](#)
- Zhou, H., Wang, T., Yuan, W., et al. 2006, *ApJS*, **166**, [128](#)

Appendix A: Stacking methods

To build the stacks, we considered different techniques to resample the flux values, average them, and estimate the uncertainty in each spectral channel. The abbreviations in italics are used to denote these options in Table B.1.

A.1. Stacking algorithm

Here we list the stacking algorithms adopted to combine the single spectra into the composites.

Direct stacking (*avg*): each spectrum was rebinned on a fixed wavelength grid designed to contain all the possible wavelengths in the same redshift bin. Then, each spectrum was normalised (see below). In each spectral channel the final flux was computed as the mean, median, or geometric mean of the normalised non-zero fluxes. The uncertainty on the flux on a spectral channel was computed as the standard deviation or the inter-percentile deviation divided by the number of objects contributing to the spectral channel.

MC resampling (*mc*): a target number was decided before the stacking (generally $N_{\text{target}} = 10,000$). Then, we randomly extracted N_{target} spectra allowing for replacement and created a mock spectrum by randomizing the flux according to the uncertainty vector, in the same way as described in Section 3.3 to generate the mock spectra. The mock spectra are subject to the same procedure as in the case of direct stacking.

Bootstrap resampling (*bs*): a target number was decided before the stacking (generally $N_{\text{target}} = 200$). Then we randomly extracted N_{draw} spectra allowing for replacement and stacked them together to create the i -th composite spectrum. N_{draw} was chosen to be equal to the total number of good spectra in the bin, according to the quality criteria defined in Sec. 3.2. This procedure was performed N_{target} times, building the distribution of the mean fluxes in each spectral channel. The final spectrum in each spectral channel was obtained as the mean of the mean fluxes and the uncertainty as the standard deviation.

A.2. Normalisation

Here we list the normalisation options adopted to scale the single spectra during the combination process.

Monochromatic (*mcr*): the spectrum is normalised by its value at a certain wavelength λ , generally chosen in a relatively emission-free region. Actually, the normalisation is not performed using the exact monochromatic interpolated flux, but rather the mean flux values in a wavelength range (10-100 Å) around λ , in order to avoid anomalously high or low fluxes due to bad pixels.

Integral (*itg*): the spectrum is normalised by its integral over the whole wavelength range.

A.3. Average estimator

Here we list the several estimators adopted to get the average value in each spectral channel when building the spectral composites.

Mean: the direct mean is the most straightforward way to average the flux in each spectral channel. However, especially close to the wings of the spectra where the SNR lowers, composites are more prone to anomalously high fluxes, likely due to bad pixels.

Median: using the median value in each spectral channel is safer when there are strong outliers, such as bad pixels, which can significantly alter the average flux. Moreover, the median spectrum preserves the relative fluxes of the emission features (see Sec. 3 in VB01).

Geometric mean (*gmean*): the geometric mean spectrum preserves instead the global continuum shape. Indeed, it can be shown (Sec. 3 in VB01) that the geometric mean of a distribution of spectra describable as power laws is equivalent to a power law with a spectral index given by the arithmetic mean of the spectral indices $\langle f_{\lambda} \rangle_{\text{gmean}} \propto \lambda^{-\langle \alpha_{\lambda} \rangle}$.

A.4. Uncertainty estimator

Here we list the estimators adopted to evaluate the uncertainty on the average flux in each spectral channel.

Standard deviation (*std*): the uncertainty is estimated as the standard deviation of the non-zero normalised fluxes divided by the square root of the objects contributing to a spectral channel.

Semi-interpercentile range (*pct*): the uncertainty is estimated as half the difference between the 99th and 1st percentile values divided by the square root of the objects contributing to a spectral channel.

Appendix B: The test bin




Fig. B.1. Stack-to-stack variations of α_{λ} (green), 2500-Å monochromatic flux (red), and X-ray Γ (violet). The solid black line represents the mean value, and dashed lines enclose values within a standard deviation. Stacks are numbered according to Table B.1.

To test how the different stacking options could affect the final result we performed all the 36 possible combinations of the stacking options using the averagely populated ($N_{\text{obj}} = 77$) bin ‘cB4’. For the MC (bootstrap) stacks 10,000 (200) re-samples were performed.

We performed the spectral fits of all the test-bin stacks, in order to assess whether the slope of the continuum and the monochromatic fluxes were affected by the choices adopted for building the stack. In Fig. B.1 we show that α_{λ} , f_{λ} , and Γ are barely sensitive to the different stacking assumptions, as their relative variations are of the order of 2%.

Table B.1. Labelling of the test-bin stacks. Abbreviations correspond to those defined in the list of stacking options in Section A.

N_{stack}	stack features	N_{stack}	stack features	N_{stack}	stack features
1	avg gmean int pct	13	bs gmean int pct	25	mc gmean int pct
2	avg gmean int std	14	bs gmean int std	26	mc gmean int std
3	avg gmean mcr pct	15	bs mean itg std	27	mc gmean mcr pct
4	avg gmean mcr std	16	bs gmean mcr std	28	mc gmean mcr std
5	avg mean int pct	17	bs mean int pct	29	mc mean int pct
6	avg mean int std	18	bs mean int std	30	mc mean int std
7	avg mean mcr pct	19	bs mean mcr pct	31	mc mean mcr pct
8	avg mean mcr std	20	bs mean mcr std	32	mc mean mcr std
9	avg median int pct	21	bs median int pct	33	mc median int pct
10	avg median int std	22	bs median int std	34	mc median int std
11	avg median mcr pct	23	bs median mcr pct	35	mc median mcr pct
12	avg median mcr std	24	bs median mcr std	36	mc median mcr std

Appendix C: Impact of the extinction curve




Fig. C.1. Test-bin composite spectrum de-reddened according to different extinction curves reported in the inset and assuming $E(B - V) = 0.1$ and $R_V = 3.1$. The slopes produced from the spectral fits according to the different extinction curves are represented in the inset. The spectrum and the slope of the continuum before the de-reddening ($\alpha_{\lambda, obs}$) are reported in magenta as references. The slope derived assuming the extinction curve from [Gordon et al. \(2016\)](#) is the closest to the average (black line).

With the end to explore the effect of the choice of the extinction curve adopted to redden or de-redden the composites throughout this work, we tested different extinction curves on

the composite spectrum in the ‘cB4’ bin. We included SMC-like curves ([Prevot et al. 1984](#); [Gordon et al. 2016](#)), as well as AGN-derived extinction curves ([Gaskell et al. 2004](#); [Gallerani et al. 2010](#), excluding BALs; [Wild et al. 2011](#); [Zafar et al. 2015](#)). After de-reddening the composite we performed the spectral fit to infer the slope of the continuum α_{λ} and compare it to that of the actual composite. We show the results of this procedure in Fig. C.1. The slope of the de-reddened spectrum assuming the [Gordon et al. \(2016\)](#) SMC-like curve (adopted in this work) is close to the average of the distribution. The strongest outlier in the extinction distribution is the [Gaskell et al. \(2004\)](#) extinction curve, which is akin to a grey-body for wavelengths below $\lesssim 3300 \text{ \AA}$, and therefore does not alter the spectral shape at short wavelengths. However, their sample is based on 72 radio-selected quasars with the addition of mostly radio-quiet objects from the LBQS survey, and it is possible, as already suggested by [Hopkins et al. \(2004\)](#), that the flattening of their extinction curve is driven by the spectral properties of the radio-loud quasars in their sample. Our choice of a SMC-like curve as representative of the intrinsic quasar extinction curve seems also justified by the findings of the latter work. There, the authors found a red tail in the distribution of SDSS photometric colours according to different redshifts. This trend proved to be consistent with SMC-like reddening, but not with Large Magellanic Cloud- or Milky Way-like.

Appendix D: Table of results

In Table D.1 we report all the average accretion parameters of the bins in which the parameter space is divided, as well as the spectral properties of the composites.

Table D.1. Relevant quantities for each bin of the parameter space.

bin	N_{obj}	$\log(M_{\text{BH}}/M_{\odot})$	$\log(L_{\text{bol}}/\text{erg s}^{-1})$	$\log(\lambda_{\text{Edd}})$	α_{λ}	$f_{2500\text{\AA}}$	Γ	$\Delta\alpha$
aC1	6	7.63 ± 0.02	45.47 ± 0.02	-0.27 ± 0.02	-1.432 ± 0.005	1.298 ± 0.001	2.33 ± 0.09	0.014 ± 0.002
aD1	40	7.45 ± 0.03	45.07 ± 0.03	-0.50 ± 0.03	-1.541 ± 0.073	1.324 ± 0.018	2.34 ± 0.11	-0.060 ± 0.010
aC2	109	8.17 ± 0.02	45.66 ± 0.02	-0.62 ± 0.02	-1.484 ± 0.007	1.311 ± 0.002	2.25 ± 0.13	-0.022 ± 0.003
aD2	126	8.04 ± 0.02	45.16 ± 0.02	-0.99 ± 0.02	-1.437 ± 0.025	1.299 ± 0.006	2.16 ± 0.11	0.011 ± 0.002
aC3	141	8.67 ± 0.01	45.76 ± 0.02	-1.02 ± 0.02	-1.584 ± 0.007	1.335 ± 0.002	2.08 ± 0.08	-0.090 ± 0.014
aD3	39	8.58 ± 0.02	45.19 ± 0.02	-1.50 ± 0.03	-1.623 ± 0.021	1.344 ± 0.005	2.01 ± 0.07	-0.117 ± 0.018
aB4	6	9.35 ± 0.05	46.55 ± 0.05	-0.92 ± 0.08	-2.025 ± 0.028	1.446 ± 0.007	2.10 ± 0.11	-0.394 ± 0.060
aC4	13	9.32 ± 0.05	46.10 ± 0.08	-1.33 ± 0.09	-1.919 ± 0.013	1.419 ± 0.003	1.95 ± 0.05	-0.321 ± 0.049
bC2	60	8.21 ± 0.02	45.72 ± 0.03	-0.61 ± 0.03	-1.342 ± 0.026	1.277 ± 0.006	2.26 ± 0.10	0.077 ± 0.012
bD2	45	8.14 ± 0.03	45.18 ± 0.02	-1.07 ± 0.04	-1.374 ± 0.151	1.285 ± 0.035	2.15 ± 0.07	0.054 ± 0.010
bB3	20	8.85 ± 0.04	46.53 ± 0.03	-0.44 ± 0.04	-1.751 ± 0.003	1.376 ± 0.001	2.17 ± 0.13	-0.205 ± 0.031
bC3	253	8.75 ± 0.01	45.93 ± 0.02	-0.93 ± 0.02	-1.487 ± 0.024	1.311 ± 0.006	2.16 ± 0.10	-0.023 ± 0.004
bD3	43	8.61 ± 0.02	45.27 ± 0.02	-1.46 ± 0.03	-1.409 ± 0.023	1.293 ± 0.005	2.08 ± 0.07	0.030 ± 0.005
bB4	23	9.29 ± 0.03	46.66 ± 0.05	-0.75 ± 0.04	-1.664 ± 0.004	1.354 ± 0.001	2.13 ± 0.09	-0.145 ± 0.022
bC4	32	9.27 ± 0.03	46.12 ± 0.04	-1.26 ± 0.04	-1.521 ± 0.007	1.320 ± 0.002	2.06 ± 0.05	-0.047 ± 0.007
cC2	35	8.22 ± 0.02	45.80 ± 0.04	-0.54 ± 0.04	-1.445 ± 0.011	1.301 ± 0.003	2.22 ± 0.08	0.036 ± 0.000
cD2	9	8.09 ± 0.05	45.17 ± 0.05	-1.04 ± 0.06	-1.419 ± 0.020	1.295 ± 0.005	2.19 ± 0.07	0.053 ± 0.001
cB3	49	8.92 ± 0.02	46.57 ± 0.02	-0.47 ± 0.03	-1.565 ± 0.003	1.330 ± 0.001	2.22 ± 0.08	-0.045 ± 0.000
cC3	231	8.76 ± 0.01	45.97 ± 0.02	-0.90 ± 0.11	-1.451 ± 0.008	1.303 ± 0.002	2.22 ± 0.09	0.032 ± 0.000
cD3	6	8.79 ± 0.05	45.27 ± 0.08	-1.63 ± 0.02	-1.333 ± 0.011	1.275 ± 0.003	2.08 ± 0.06	0.110 ± 0.001
cB4	77	9.34 ± 0.02	46.69 ± 0.02	-0.76 ± 0.02	-1.568 ± 0.004	1.331 ± 0.001	2.19 ± 0.09	-0.046 ± 0.000
cC4	61	9.25 ± 0.01	46.19 ± 0.02	-1.17 ± 0.02	-1.495 ± 0.006	1.313 ± 0.002	2.16 ± 0.08	0.002 ± 0.000
dC2	9	8.20 ± 0.05	45.82 ± 0.07	-0.49 ± 0.08	-1.137 ± 0.192	1.230 ± 0.043	2.23 ± 0.12	0.201 ± 0.034
dB3	66	8.93 ± 0.02	46.62 ± 0.02	-0.43 ± 0.02	-1.453 ± 0.005	1.303 ± 0.001	2.25 ± 0.10	-0.021 ± 0.000
dC3	141	8.81 ± 0.02	46.06 ± 0.02	-0.87 ± 0.02	-1.577 ± 0.024	1.333 ± 0.006	2.22 ± 0.10	-0.108 ± 0.002
dA4	7	9.66 ± 0.05	47.55 ± 0.04	-0.22 ± 0.06	-1.557 ± 0.004	1.328 ± 0.001	2.16 ± 0.05	-0.094 ± 0.001
dB4	173	9.40 ± 0.01	46.82 ± 0.02	-0.70 ± 0.02	-1.479 ± 0.011	1.309 ± 0.003	2.20 ± 0.09	-0.039 ± 0.000
dC4	69	9.32 ± 0.02	46.16 ± 0.03	-1.27 ± 0.03	-1.374 ± 0.051	1.285 ± 0.012	2.17 ± 0.08	0.035 ± 0.001
dA5	7	9.98 ± 0.03	47.70 ± 0.06	-0.39 ± 0.06	-1.560 ± 0.123	1.329 ± 0.030	2.00 ± 0.04	-0.096 ± 0.008
dB5	6	9.92 ± 0.02	46.89 ± 0.13	-1.14 ± 0.14	-1.336 ± 0.020	1.276 ± 0.005	2.16 ± 0.08	0.061 ± 0.001
eC2	22	8.20 ± 0.05	46.10 ± 0.04	-0.21 ± 0.05	-1.308 ± 0.019	1.269 ± 0.004	2.17 ± 0.09	0.169 ± 0.004
eB3	96	8.86 ± 0.02	46.68 ± 0.02	-0.30 ± 0.03	-1.491 ± 0.016	1.312 ± 0.004	2.17 ± 0.09	0.052 ± 0.001
eC3	63	8.81 ± 0.03	46.19 ± 0.02	-0.73 ± 0.03	-1.558 ± 0.018	1.329 ± 0.004	2.18 ± 0.10	0.010 ± 0.000
eA4	26	9.53 ± 0.04	47.70 ± 0.04	-0.06 ± 0.06	-1.623 ± 0.154	1.344 ± 0.038	2.00 ± 0.03	-0.032 ± 0.003
eB4	174	9.39 ± 0.02	46.87 ± 0.02	-0.64 ± 0.02	-1.447 ± 0.017	1.302 ± 0.004	2.24 ± 0.09	0.080 ± 0.002
eC4	23	9.31 ± 0.04	46.27 ± 0.02	-1.15 ± 0.03	-1.152 ± 0.012	1.234 ± 0.003	2.17 ± 0.08	0.268 ± 0.006
eA5	44	10.06 ± 0.03	47.80 ± 0.04	-0.37 ± 0.04	-1.452 ± 0.004	1.303 ± 0.001	2.08 ± 0.08	0.077 ± 0.002
eB5	12	9.98 ± 0.06	47.15 ± 0.06	-0.95 ± 0.09	-1.732 ± 0.030	1.371 ± 0.008	2.14 ± 0.06	-0.101 ± 0.003

Appendix E: Correlation index tables

Table E.1. Correlation index and associated P -value (S, P) pairs for the UV and optical slopes against the accretion parameters. In the top row all the bins are included in the computation, in the bottom only those meeting the representativeness criterion.

subsample	$\alpha_\lambda - \log M_{\text{BH}}$	$\alpha_\lambda - \log L_{\text{bol}}$	$\alpha_\lambda - \log \lambda_{\text{Edd}}$
all	(-0.300, 0.067)	(-0.263, 0.110)	(-0.006, 0.972)
selected	(-0.281, 0.125)	(-0.229, 0.214)	(-0.017, 0.928)

In Table E.1 we report the correlation index S between the slope α_λ and the accretion parameters, including/excluding the bins ruled out by the representativeness criterion.

Appendix F: Reddening expression

In this section, we explicitly derive Eq. 4. We begin from Eq. 7 in L20 in order to express the distance modulus DM starting from the luminosity distance d_L computed from the $L_X - L_{\text{UV}}$ relation:

$$\text{DM} = 5 \left[\frac{\log f_X - \beta - \gamma (\log f_{2500\text{\AA}} + 27.5)}{2\gamma - 1} - \frac{1}{2} \log(4\pi) + 28.5 \right] - 5 \log(10 \text{ pc}).$$

Keeping only the dependence on $f_{2500\text{\AA}}$, DM can be written as

$$\text{DM} = 5 \frac{\gamma}{2(\gamma - 1)} \log f_{2500\text{\AA}} + \theta,$$

where θ encapsulates all the terms where $f_{2500\text{\AA}}$ does not appear. Next, we assume, by hypothesis, that the observed UV fluxes are

significantly extinguished to the point that the discrepancy from the ΛCDM model that we report is driven by such extinction. We can then denote the observed DM as

$$\text{DM}_{\text{ext}} = -5 \frac{\gamma}{2(\gamma - 1)} \log f_{2500\text{\AA},\text{ext}} + \theta,$$

and the intrinsic (unextinguished) DM as

$$\text{DM}_{\text{int}} = -5 \frac{\gamma}{2(\gamma - 1)} \log f_{2500\text{\AA},\text{int}} + \theta.$$

The intrinsic UV flux is related to the extinguished one by

$$f_{2500\text{\AA},\text{int}} = f_{2500\text{\AA},\text{ext}} 10^{0.4 e_\lambda R_V E(B-V)}$$

where e_λ is the assumed extinction curve, R_V is the total to selective extinction ratio, and $E(B - V)$ is the colour excess. We can then write DM_{int} in terms of the extinguished UV flux:

$$\text{DM}_{\text{int}} = -5 \frac{\gamma}{2(\gamma - 1)} [\log f_{2500\text{\AA},\text{ext}} + 0.4 e_\lambda R_V E(B - V)] + \theta.$$

The difference between the intrinsic and the observed distance moduli (ΔDM) thus becomes

$$\Delta\text{DM} = \text{DM}_{\text{int}} - \text{DM}_{\text{ext}} = -5 \frac{\gamma}{2(\gamma - 1)} 0.4 e_\lambda R_V E(B - V).$$

Supposing that the DM_{int} is the DM corresponding to the ΛCDM model, ΔDM can be evaluated by imposing $\text{DM}_{\text{int}} = \text{DM}_{\Lambda\text{CDM}}$. By isolating the $E(B - V)$ term, Eq. 4 is found:

$$E(B - V) = \frac{\Delta\text{DM}}{-\frac{5}{2} \frac{\gamma}{\gamma - 1} 0.4 R_V \frac{A_{2500\text{\AA}}}{A_V}}.$$

Appendix G: Synopsis of the UV stacks

In Fig. G.1, we show all the UV spectral stacks across the parameter space in different colours according to their redshift bin,

along with the average spectrum of our whole sample in black as a reference. Regions without any spectrum are those discarded in Sec. 3.1. Only the UV side for the spectra is shown, as the optical one is only available in the lower redshift bins.

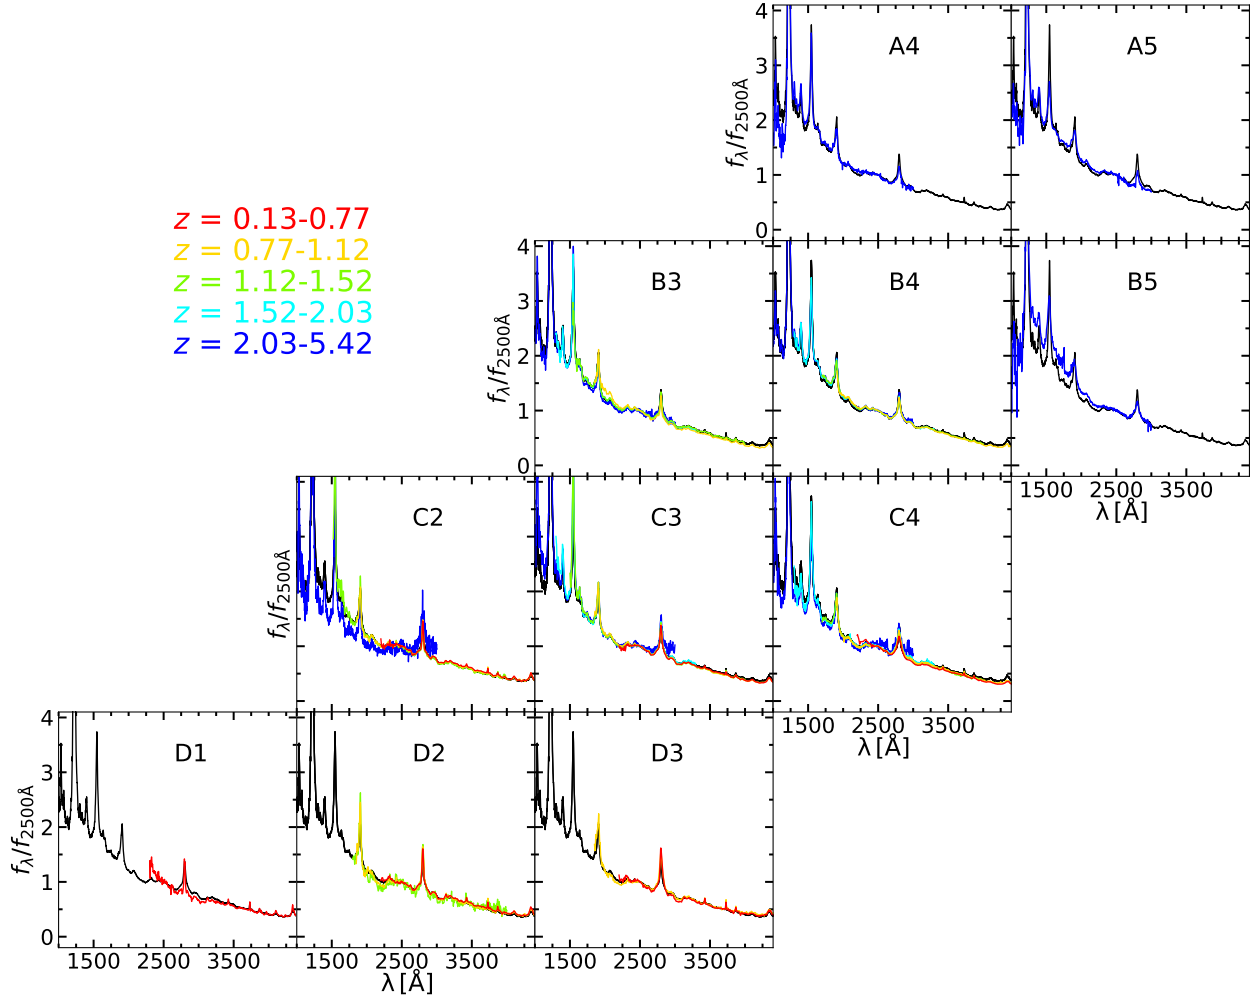


Fig. G.1. Similarity of the stacked spectra across the parameter space. All the spectra are normalised by their average emission between 2490–2510 Å. The black spectrum represents the average spectrum of our full sample. Only spectra of the bins meeting the representativeness criterion are included in the plot.