


Dust and the intrinsic spectral index of quasar variations: hints of finite stress at the innermost stable circular orbit

John R. Weaver ^{1,2,3}★ and Keith Horne³

¹*Cosmic Dawn Center (DAWN)*

²*Niels Bohr Institute, University of Copenhagen, Jagtvej 128, DK-2200 Copenhagen, Denmark*

³*SUPA School of Physics & Astronomy, University of St Andrews, North Haugh, St Andrews KY16 9SS, UK*

Accepted 2022 January 24. Received 2022 January 24; in original form 2021 November 26

ABSTRACT

We present a study of 9 242 spectroscopically confirmed quasars with multiepoch *ugriz* photometry from the SDSS Southern Survey. By fitting a separable linear model to each quasar’s spectral variations, we decompose their five-band spectral energy distributions into variable (disc) and non-variable (host galaxy) components. In modelling the disc spectra, we include attenuation by dust on the line of sight through the host galaxy to its nucleus. We consider five commonly used attenuation laws, and find that the best description is by dust similar to that of the Small Magellanic Cloud, inferring a lack of carbonaceous grains from the relatively weak 2175-Å absorption feature. We go on to construct a composite spectrum for the quasar variations spanning 700–8000 Å. By varying the assumed power-law $L_\nu \propto \nu^\alpha$ spectral slope, we find a best-fitting value $\alpha = 0.71 \pm 0.02$, excluding at high confidence the canonical $L_\nu \propto \nu^{1/3}$ prediction for a steady-state accretion disc with a $T \propto r^{-3/4}$ temperature profile. The bluer spectral index of the observed quasar variations instead supports the model of Agol & Krolik, and Mummery & Balbus, in which a steeper temperature profile, $T \propto r^{-7/8}$, develops as a result of finite magnetically induced stress at the innermost stable circular orbit extracting energy and angular momentum from the black hole spin.

Key words: accretion, accretion discs – methods: statistical.

1 INTRODUCTION

The optical identification of quasi-stellar objects (quasars hereafter) by Matthews & Sandage (1963) enabled, for the first time, studies of the distant universe at $z > 0.1$. Quasars are now recognized as high-luminosity examples of active galactic nuclei (AGN), powered by accretion on to a supermassive black hole (SMBH) (Lynden-Bell 1969; Shakura & Sunyaev 1973). The continuum variability of quasars, known soon after their discovery, allows us to peer directly into their central engines. Varying by 10–20 per cent over time-scales of months to years, the intrinsic variability of quasar continuum emission has long been theorized to be caused by changes in the environment close to the SMBH.

Quasar spectral energy distributions (SEDs) provide insight into their underlying physics. Spanning the full range from gamma-rays to radio, quasar SEDs exhibit both thermal (accretion disc, dust) and non-thermal (corona, jet) components. In the rest-frame UV–optical, thermal emission from the accretion disc is thought to manifest as the ‘Big Blue Bump’ (Shields 1978; Malkan & Sargent 1982), described by a sum of blackbody spectra over a range of temperatures from $\sim 10^3$ K for the cool outer edges of the disc to perhaps $\sim 10^5$ K near the innermost stable circular orbit (ISCO). A related feature is the ‘Small Blue Bump’, caused by closely packed FeII emission and the Balmer recombination continuum (Elvis, Wilkes & Tananbaum 1985; Wills, Netzer & Wills 1985)

For a geometrically thin steady-state accretion disc (Shakura & Sunyaev 1973), the effective temperature profile is $T_{\text{eff}} \propto (M \dot{M})^{1/4} r^{-3/4}$, where M is the black hole mass, \dot{M} is the accretion rate, and r is the radial distance from the black hole. The corresponding spectrum, obtained by summing blackbody spectra weighted by solid angle, is $L_\nu \propto (M \dot{M})^{2/3} \nu^{1/3}$. This power-law spectrum applies in the spectral range corresponding to the minimum and maximum disc temperatures, $k T_{\text{min}} \ll h \nu \ll k T_{\text{max}}$, where h and k are the Planck and Boltzmann constants. For a more general power-law temperature profile, $T \propto r^{-b}$, the disc spectrum is $L_\nu \propto \nu^\alpha$ with $\alpha = (3b - 2)/b$. Thus, measuring the disc’s spectral slope α determines the power-law slope b of its temperature profile and tests the accretion disc theory. If the theoretical power-law slope, $\alpha = 1/3$ is confirmed, the results measure the product $M \dot{M}$. Moreover, since the disc spectrum scales with inclination angle i and luminosity distance D_L via $\cos i / D_L^2$, we may potentially be able to measure quasar luminosity distances.¹

Several obstacles stand in the way of realizing these motivating goals. First, there may be significant extinction and reddening due to dust along the line of sight. Correcting for dust in our Milky Way galaxy is relatively straightforward (e.g. Schlegel, Finkbeiner & Davis 1998; Schlafly & Finkbeiner 2011). More difficult is to correct for the adverse effects of reddening caused by scattering and absorption by dust grains within the host galaxy, a complication shared by

¹For Type 1 AGN ($i < 60^\circ$), the mean \pm rms of $\cos i$ is $3/4 \pm \sqrt{1/48} = 0.75 \pm 0.14$.

* E-mail: john.weaver.astro@gmail.com

the use of Type Ia supernovae as standard candles. Reddening along the line of sight presents a degeneracy since dust grains can redden quasar spectra with a wavelength dependence similar to the power law form expected for the disc spectrum. However, carbonaceous grains produce an absorption feature prominent in the dust extinction law observed in the Milky Way Galaxy (Nandy et al. 1975; Allen & Glass 1976; Seaton 1979) and the Large Magellanic Cloud (Fitzpatrick 1986). This absorption feature, described by a ‘Drude’ profile centred at 2175 Å (Fitzpatrick & Massa 1986; Draine & Malhotra 1993), can largely resolve the degeneracy, given sufficient spectral coverage. The 2175-Å absorption is weak or absent in other notable systems including the Small Magellanic Cloud (Gordon et al. 2003) and local starburst galaxies (Calzetti et al. 2000), which suffer from the full strength of this degeneracy. Moreover, it has been postulated that quasar dust may differ from the varieties studied closely in the local universe, for example, by lacking small grains that are evaporated by the quasar luminosity (e.g. Gaskell et al. 2004).

Secondly, the observed spectra of quasars are generally redder than the predicted disc spectrum, hinted at already by Sandage (1965). This is due in part to the comparably red starlight of the host galaxy. Quasars are often too distant to directly resolve the host galaxy, which means that their measurements are contaminated by host galaxy starlight captured within the aperture from which the photometry is performed. While this problem can be mitigated in small samples of nearby, resolved quasars where the host galaxy’s light profile can be modelled, or extrapolated inward and subtracted from images, this approach fails for larger samples of more distant, unresolved quasars. The advent of large multiwavelength monitoring campaigns provides a viable workaround. Instead of attempting to subtract host galaxy contributions from imaging data, a time-series of images or spectra can be used to isolate the spectrum of the variable light arising from the central engine and the accretion disc. By this method, one can extract separate spectra for the variable accretion disc and the non-varying host galaxy components for a large number of unresolved quasars, provided multiwavelength photometric monitoring with sub-year cadence over a sufficiently long baseline to probe the variations.

Several successful and innovative campaigns have marked the previous twenty years of the study of AGN. Most recently, assemblages of multiepoch, multiwavelength photometric data sets have been observed by the Sloan Digital Sky Survey (SDSS; York et al. 2000) and more recently the Zwicky Transient Facility (ZTF; Bellm et al. 2018; Graham et al. 2019), which have enabled fundamentally new comparisons with theoretical models of accretion disc structure and behaviour with statistically significant samples. Additional observations from the Rubin Observatory’s Legacy Survey of Space and Time (LSST; Ivezić et al. 2019) will greatly increase both sample sizes and epoch baselines, expected to be underway in 2024. In addition, precise, spectroscopic monitoring campaigns such as the Sloan Reverberation Mapping Project (SDSS-RM; e.g. Shen et al. 2015) are providing valuable details on the variability of continuum and line emission, although with smaller samples.

Several techniques have been applied to interpret these photometric data sets. For example, MacLeod et al. (2012) employed a damped random walk model to describe the stochastic variations for an ensemble of $\sim 10^4$ quasars from SDSS, finding good agreement as a viable description of the optical continuum variability. For the same dataset, Kokubo et al. (2014) employed a ‘flux-flux correlation’ technique to derive the colour of the flux difference spectrum, which was used to infer an accretion disc spectral slope of $L_\nu \propto \nu^{1/3}$, consistent with standard steady-state accretion models. Parallel work has been undertaken with this and similar data sets to determine the extinction

law most appropriate for quasars (e.g. Hopkins et al. 2004; Krawczyk et al. 2015), the results of which influence work on variability.

The objective of this work is to directly probe the accretion disc light and test theories of accretion physics. This will be accomplished as follows. In Section 2, we develop our method that leverages the source variability to isolate the accretion disc light. In Section 3, we apply this to a sample of 9 242 quasars observed with multiepoch multiwavelength photometry during the Sloan Digital Sky Survey Stripe 82 quasar campaign, including a de-reddening of the isolated accretion spectrum with five commonly used dust extinction laws. Section 4 presents the composite spectra. Our results are then discussed in Section 5 and our conclusions made in Section 6.

We adopt in this paper a concordance cosmological model with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_M = 0.3$, and $\Omega_\Lambda = 0.7$. All magnitudes are in the AB_ν system (Oke 1974), for which a flux f_ν in mJy ($10^{-26} \text{ erg cm}^{-2} \text{ s}^{-1} \text{ Hz}^{-1}$) corresponds to $AB_\nu = 16.4 - 2.5 \log_{10}(f_\nu/\text{mJy})$.

2 ISOLATING THE ACCRETION DISC LIGHT

In this section, we describe our method using a separable linear model to fit the photometric variations of a quasar observed with multiwavelength photometry. The method is illustrated for a particular SDSS quasar in Fig. 1, which we discuss below as we outline the steps of the analysis.

A quasar is observed at N_t times t in N_λ photometric bands, each labelled by its pivot wavelength² λ . The observations at time t are considered simultaneous if measured within a time interval so short that changes in the state of the accretion disc can be neglected. For UV and optical observations of quasars, this typically means measurements on the same night, or even over a few nights.

We fit the observed spectral flux variations with the following separable linear model:

$$F(\lambda, t) = A(\lambda) X(t) + B(\lambda). \quad (1)$$

Here, $F(\lambda, t)$ can be F_ν or F_λ , or indeed any suitable flux unit. The dimensionless light-curve shape $X(t)$ is shifted to zero mean and scaled to unit root mean square (rms):

$$\langle X \rangle_t = 0, \quad \langle X^2 \rangle_t = 1, \quad (2)$$

where $\langle \cdot \rangle_t$ denotes a suitably weighted time average. With this normalization, the model’s amplitude spectrum $A(\lambda)$ is the rms of the flux variations about the mean background spectrum $B(\lambda)$. This model has $2N_\lambda + N_t$ parameters, which are $A(\lambda)$, $B(\lambda)$, and $X(t)$. These are constrained by $N_t \times N_\lambda$ flux measurements plus two normalization constraints. For observations at a single wavelength, $N_\lambda = 1$, the model fits the N_t flux measurements exactly. For multiband observations, the model parameters are over constrained by the data, which permits optimizing the model parameters by fitting the data, and testing the validity of the model assumptions.

This model fitting is illustrated for a particular SDSS quasar in Fig. 1. The light curve of this quasar is sampled at $N_t = 62$ epochs with SEDs measured by $N_\lambda = 5$ bands, as shown in Fig. 1(b). Note here that the data follow a light-curve shape that is similar for all five bands. Fitting the model to these data, by minimizing χ^2 , the corresponding set of linear equations is solved to determine the model parameters $A(\lambda)$, $B(\lambda)$, and $X(t)$, with corresponding uncertainties. This is done in practice by using iterated linear regression fits. Start by constructing an initial guess for $X(t)$, for example using one

²See A2.1 of Bessell & Murphy (2012) for details.

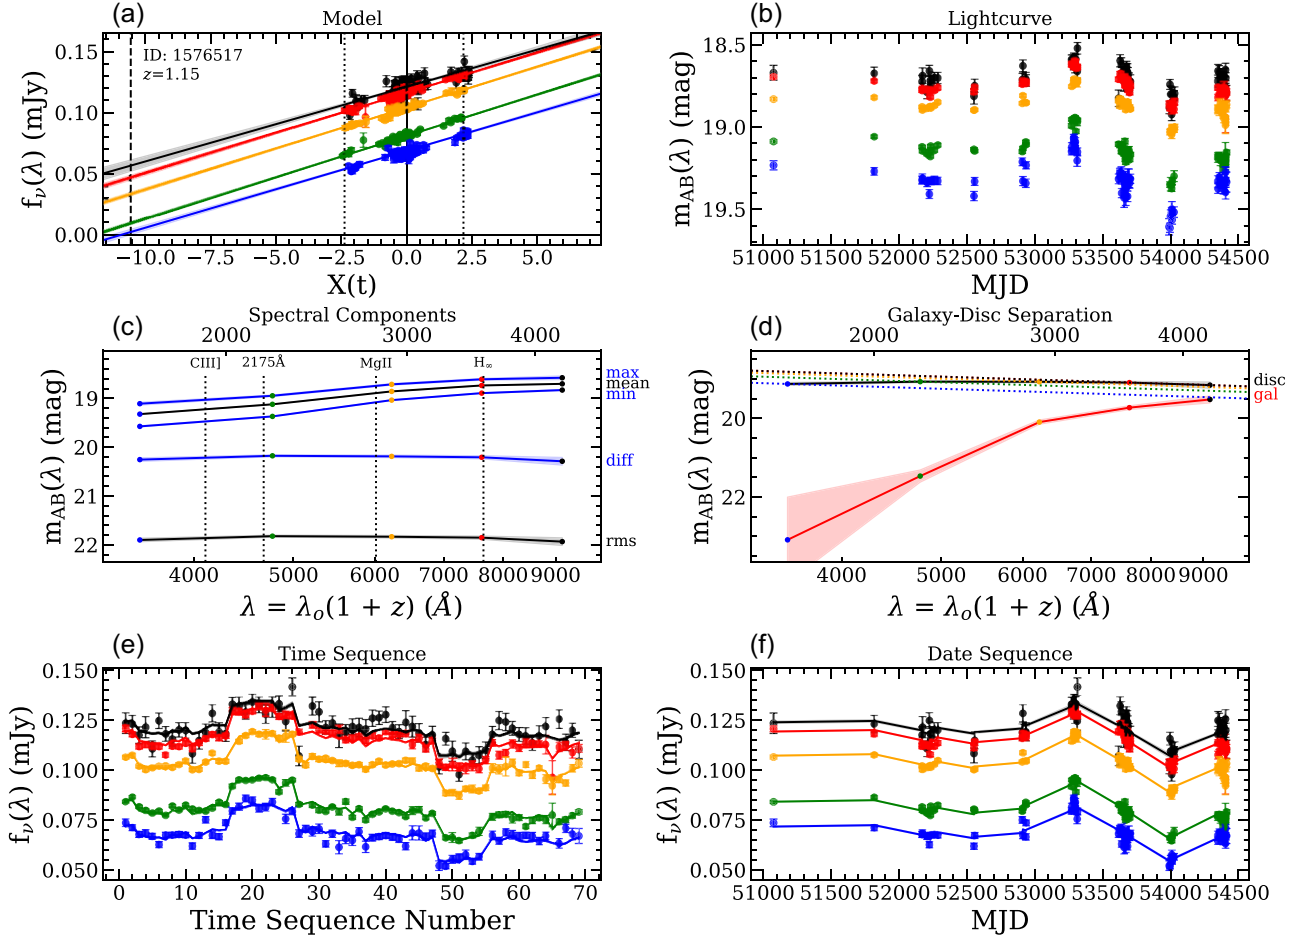


Figure 1. An example illustrating our light-curve decomposition method. Panel (b) shows the $N_\lambda = 5$ -band *ugriz* light-curve data at $Nt = 62$ epochs for the SDSS quasar ID 1576517 at redshift $z = 1.15$. Panel (a) shows that the flux variations are well fitted by a linear model, $F(\lambda, t) = B(\lambda) + A(\lambda)X(t)$, where $B(\lambda)$ is the mean flux, $A(\lambda)$ is the rms amplitude of the flux variations, and $X(t)$ is the light-curve shape, assumed to be the same for all bands, normalized to $\langle X \rangle = 0$ and $\langle X^2 \rangle = 1$. The maximum and minimum brightness states are indicated in Panel (a) by vertical dotted lines on either side of the mean state at $X = 0$. Extrapolating to fainter levels, the *u*-band flux becomes negative just below $X \approx -10.5$. Having thus turned off the disc, we attribute the extrapolated fluxes at the vertical dashed line, where *u* is 1σ above 0, to the non-variable host galaxy. Panel (c) shows the resulting five-band SEDs extracted at the maximum, minimum, rms, and mean states, as well as the bright – faint difference spectrum, with vertical dotted lines marking the wavelengths of relevant spectral features. Panel (d) shows the SEDs extracted for the disc (black) and host galaxy (red). Dotted lines show $L_\nu \propto \nu^{1/3}$ fixed to each band and coloured accordingly. Finally, the light-curve data and fitted model are shown versus time sequence number in Panel (e) and by date of observation in Panel (f). All error envelopes are shown at $\pm 1\sigma$.

of the observed light curves, suitably normalized. Then use two-parameter linear regression to find $A(\lambda)$ and $B(\lambda)$ assuming $X(t)$ is known. Next, revise $X(t)$ assuming $A(\lambda)$ and $B(\lambda)$ are known. Impose the normalization constraints on $X(t)$. Finally, iterate to convergence. Some care may be needed to identify and down-weight or reject significant outliers, using a robust procedure such as sigma clipping.

The two-parameter linear regression fits that determine $A(\lambda)$ and $B(\lambda)$, with $X(t)$ assumed to be known, are presented in Fig. 1(a). The fitted linear models are shown as solid coloured lines with $\pm 1\sigma$ envelopes. The flux data with error bars are plotted versus the dimensionless $X(t)$. This tracks the changing brightness of the quasar above and below the mean flux level. For each band, the slope in this diagram is the rms amplitude $A(\lambda)$ of the flux variations above and below the mean spectrum $B(\lambda)$ at $X(t) = 0$. Note here that the quasar variations are well described by linear flux variations. In particular, there is no evident curvature that could indicate a change in the disc spectrum between the faint and bright states. The shape of the light-curve $X(t)$ and comparison of the fitted model with the light-curve

data, are examined in Fig. 1(e), where the flux data are plotted versus time sequence number, and in Fig. 1(f) versus observation date. Here, the light-curve shape is determined as a weighted average of the variations seen in all bands.

The fitted model can now be used to predict fluxes expected at different variability states $X(t)$. Given N_λ observed bands, the fitted model predicts the SED for any variability state $X(t)$. Meaningful extrapolation is possible, above and below the range sampled by the monitoring data, with the usual caveat that the extrapolated model becomes progressively uncertain. Fig. 1(c) presents the SED obtained for several indicative states. The mean spectrum, $B(\lambda)$, is the SED derived for the mean state of the system, at $X(t) = 0$. Above and below the mean SED are SEDs for the faintest and brightest observed states, at X_{\min} and X_{\max} , respectively. These SEDs are relatively red, F_ν rising to longer wavelengths. However, the difference SED between the brightest and faintest states, evaluated for $\Delta X = X_{\max} - X_{\min}$, and the SED of the rms variations, for $\Delta X = 1$, are comparably blue. Such quasar variations are often described as ‘bluer when brighter’. However, the linearity seen in Fig. 1(a) shows that this is not due to

the disc spectrum becoming bluer when brighter, but rather to the relatively red (host galaxy) spectrum becoming dominant as light from the relatively blue disc dims.

Our linear model fit to the spectral variations determines the disc and host galaxy SEDs shown in Fig. 1(d). The host galaxy SED is obtained by extrapolating the linear model to fainter states until the disc is effectively turned off. Here we define this static point as the variable state X_{gal} at which the lower 1σ uncertainty envelope of any band is predicted to lie at zero flux. The model would be non-physical at any fainter state. With short-term variability assumed to arise by modulating the disc's SED, the variable disc's SED is the flux emitted at each band in excess of the galaxy's SED. Note in Fig. 1(d) that the disc SED is close to, but slightly redder than, the power-law $L_\nu \propto \nu^{1/3}$ spectra, indicated by dotted lines.

3 APPLICATION TO SDSS DATA

3.1 Sample selection: SDSS stripe 82 quasars

As described in the previous section and illustrated by Fig. 1, our disc + galaxy decomposition procedure requires multiwavelength coverage with a suitably long time baseline to adequately probe the variability of a given source. More importantly, the procedure is well-posed mathematically if and only if the multiwavelength coverage is near simultaneous ($\lesssim 1$ night) as to constrain all relevant regimes of the SED at any one time. Thankfully, multiwavelength photometric coverage for transient surveys are usually performed on nightly basis, thereby providing multiple samplings in wavelength per source, per night.

A suitable survey satisfying these requirements is the Southern Sample of the Sloan Digital Sky Survey. The Southern Sample catalogue (MacLeod et al. 2012) contains re-calibrated *ugriz* light curves for all of the spectroscopically confirmed quasars in SDSS DR7 Stripe 82. Summarily, the catalogue includes 9258 quasars over $\sim 290 \text{ deg}^2$ with an observational baseline of ~ 10 yr, observing each source for 2–3 consecutive months a year. The total number of epochs per source is ~ 60 with photometric accuracy between 0.02 and 0.04 mag.

The original photometry was adopted from the official SDSS quasar catalogue (Schneider et al. 2010) using PSF magnitudes which were re-calibrated (see MacLeod et al. (2012) for details). According to Schneider et al. (2010), 97 per cent of these objects are registered as having point-like morphology, with the remaining 3 per cent limited to $z \lesssim 0.7$; ~ 80 per cent of $z < 0.7$ sources are registered as point like. Future surveys such as LSST will be deeper and have higher angular resolution, relative to SDSS. As such, the task of accurately disentangling the nuclear quasar light from that of resolved host galaxies will require more detailed image modeling and/or aperture photometry.

Fig. 2 shows the photometric properties, redshift distribution, and sky density of sources within the catalogue. This sample provides a broad range in redshift, $0.1 < z < 6.0$, which extends the rest-frame spectral range deep into the ultraviolet and enables us to probe quasar variability and thus accretion disc structure out to remarkably early times. The photometry here is corrected for Galactic extinction using the coefficients provided in the catalogue, to thus be consistent with MacLeod et al. (2012).

As highlighted in Fig. 2(b) and (c), the typical observed-frame $g - i$ colour index of the SDSS quasars, taken from the initial catalogue prior to our decomposition analysis, is ~ 0.5 mag. redder than a power-law $L_\nu \propto \nu^{1/3}$ spectrum. This may be expected due to contamination of the disc spectrum by light from the host galaxy

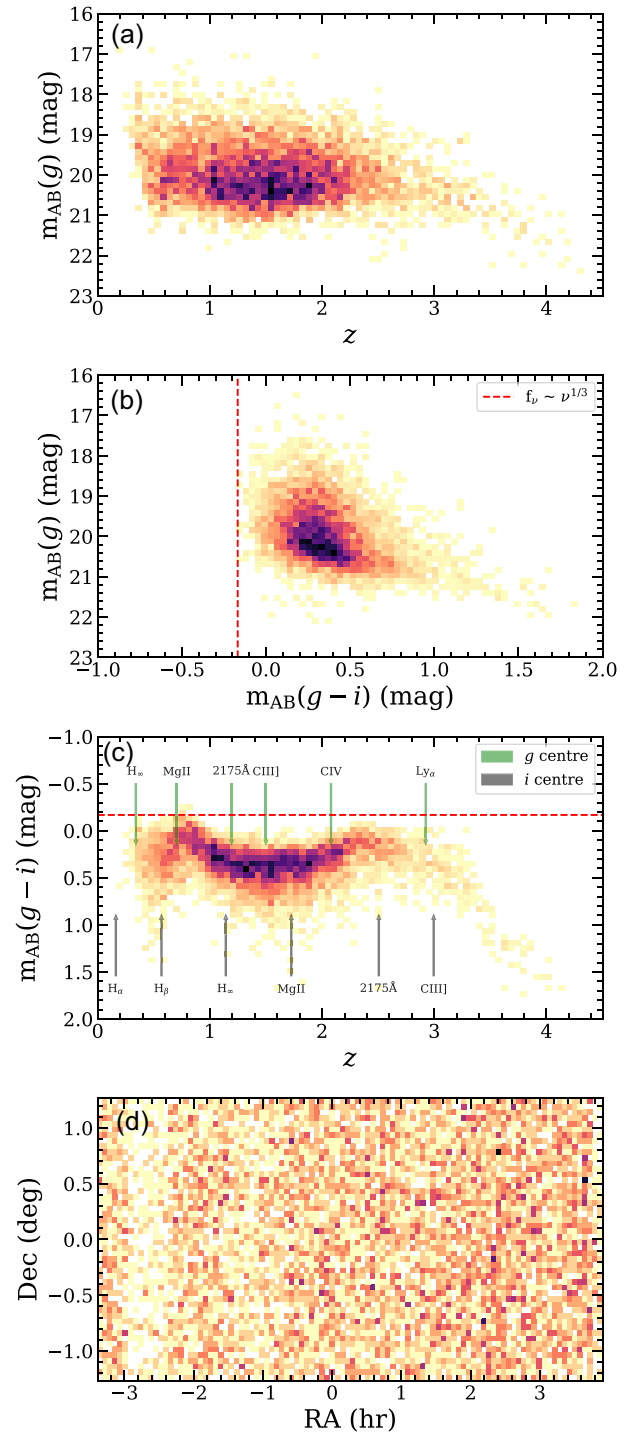


Figure 2. Summary of the raw input photometry from the SDSS Southern Sample. The g -band magnitude and $g - i$ colour distributions with redshift are shown in Panels (a) and (c). The g -band magnitude distribution is also shown against the $g - i$ colour distribution. The red dotted line indicates the expected $g - i$ colour from an $f_\nu \sim \nu^{1/3}$ spectrum. Panel (d) shows the density of sources on the sky.

and/or reddening due to dust on the line of sight to the quasar. In addition, the undulating redshift dependence of $g - i$ shown in Fig. 2(c) likely arises from quasar emission-line features redshifting into and out of the g and i passbands.

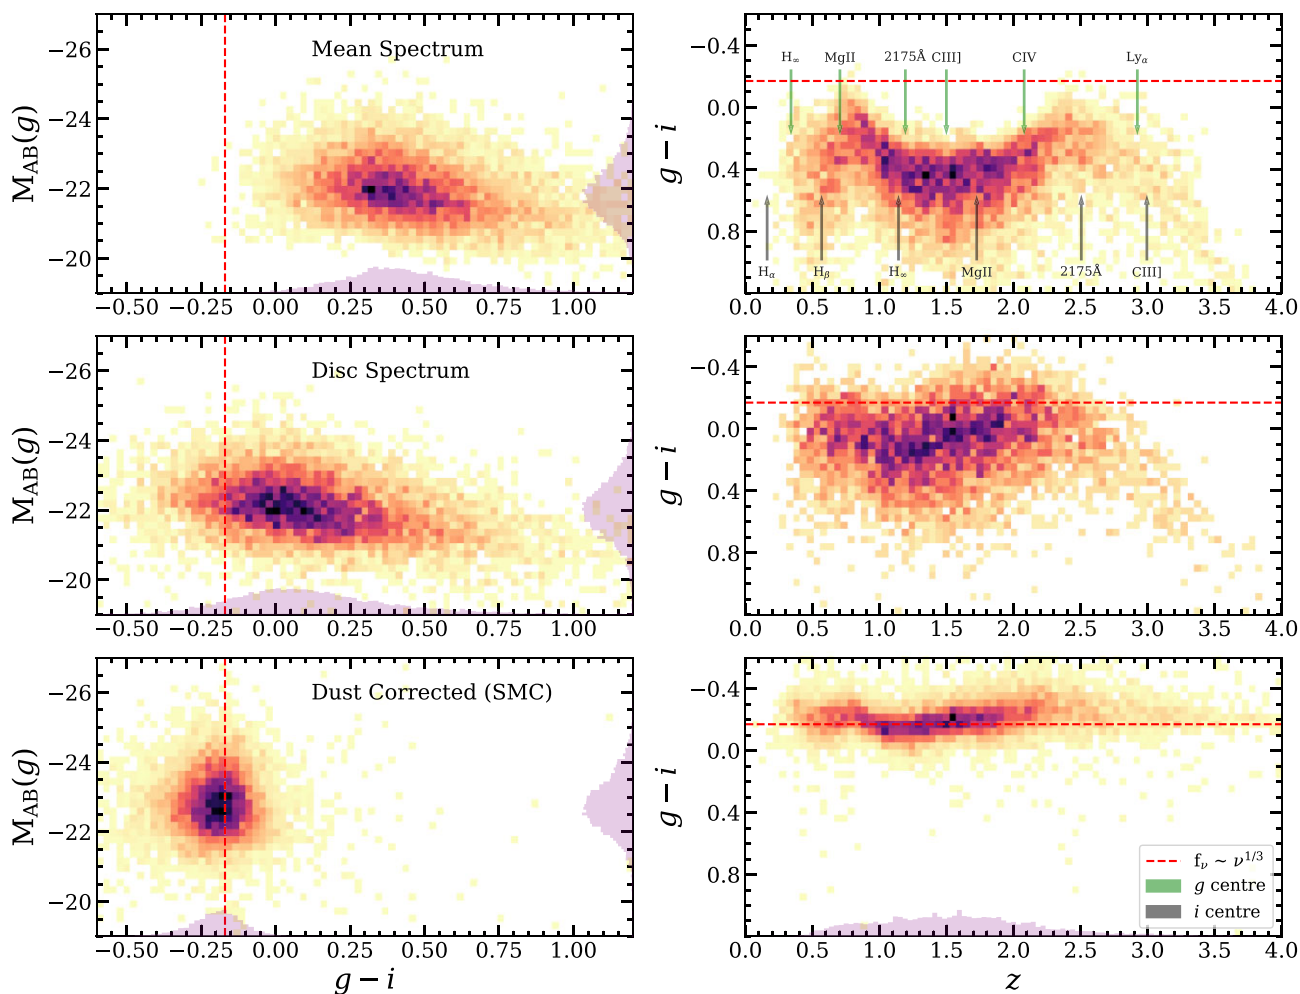


Figure 3. Luminosity (absolute g -band AB magnitude) versus $g-i$ colour index (left-hand panel) and $g-i$ versus redshift (right-hand panel) for the mean SED (top panel), disc SED (middle panel), and SMC de-reddened disc SED (bottom panel). Distributions are coloured darker with increasing density and shown as histograms projected on to each axis. The $g-i$ colour index for a power-law $L_\nu \sim \nu^{1/3}$ spectrum is indicated by the dashed red lines. Arrows in the top right-hand panel indicate the redshifts at which prominent spectral features are centred in the g - or i -band in green and grey, respectively.

3.2 Isolating disc SEDs using variations

For 99.8 per cent of the southern sample, the procedure explained in Section 2 succeeded in isolating the variable (accretion disc) and non-variable (host galaxy) SEDs. The analysis failed for just 0.2 per cent (16) of the sources, owing to either too sparsely sampled data (either in wavelength or epoch), insufficient variability signal, or a combination of the two.

The main effect of this disc + galaxy decomposition is evident for Object 1576517 in Fig. 1(c) and (d). Much of the red light is ascribed to the non-variable background galaxy component, shown in red in Fig. 1(d), thus isolating the relatively blue SED of the variable disc light, as shown in black in the same panel. In this case, the disc SED is slightly redder than the expected $L_\nu \propto \nu^{1/3}$ power-law SED, shown by dotted curves fixed to the observed magnitudes per band.

In Fig. 3, comparing the top and middle panels shows the effect of our disc + galaxy decomposition on $g-i$ colour indices over the full sample of SDSS quasars. The distribution of $g-i$ colours, for the mean and disc SEDs, are shown here as a function of magnitude and redshift. A red-dashed line marks the $g-i$ colour for the $f_\nu \propto \nu^{1/3}$ power law. The g -band absolute magnitude distribution is very similar for the disc and mean SEDs, indicating that these

quasars are typically brighter at g than their host galaxies. The $g-i$ distributions differ significantly – the disc SEDs are generally bluer than the mean SEDs. While none of the SDSS quasars has a mean spectrum as blue as the $L_\nu \propto \nu^{1/3}$ power law, most of the disc SEDs have bluer $g-i$ colours, moving toward and in some cases beyond the $L_\nu \propto \nu^{1/3}$ power law colour. But the $g-i$ distribution is not simply translated, rather it appears to be stretched towards bluer colors, leaving behind a long red tail of somewhat fainter quasars with $g-i$ similar in their disc and mean SEDs. One possible interpretation of these redder and fainter SEDs is dust along the line of sight to the quasar disc.

Note also that the stark effect of the emission lines causing $g-i$ to undulate with redshift is stronger for the mean than for the disc SEDs. This is consistent with broad UV emission lines being less variable than the disc continuum.

3.3 Accounting for dust extinction and reddening

We now investigate the possibility of dust along the line of sight to the quasar discs. This dust could be absent or differ significantly from the dust along lines of sight to other parts of the host galaxy,

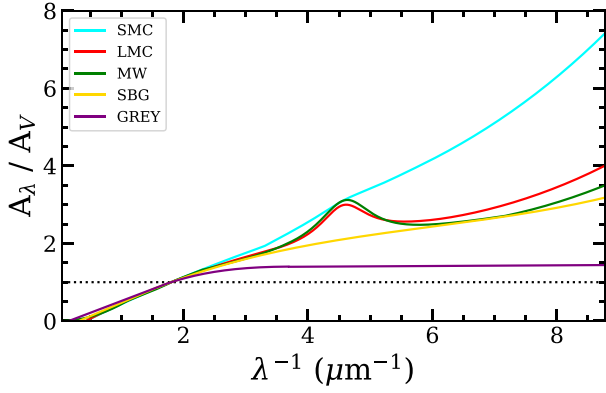


Figure 4. Attenuation curves of five commonly assumed dust attenuation laws used in this work. See Section 3.3 for details.

since the quasar luminosity can heat and evaporate dust in its vicinity. Nevertheless, there is evidence from infrared interferometry (Hönig et al. 2013; Asmus 2019) for both polar dust and equatorial dust. While equatorial dust is thought to obscure the disc and associated broad emission-line regions in Type 2 AGN, polar dust may attenuate and redden the observed disc spectra even for more face-on discs.

Although there is extensive discussion in the literature (e.g. Gallerani et al. 2010; Krawczyk et al. 2015; Zafar et al. 2015), there is as yet no definitive evidence and certainly no consensus as to the correct or possibly universal dust extinction law for quasars. Given this uncertainty, we consider five possible dust laws. Their attenuation curves, A_λ/A_V , are shown in Fig. 4. The five dust laws are as follows:

- (i) SMC – The Small Magellanic Cloud – a nearly smooth power-law-like curve with relatively high UV extinction due to small grains. Adopted from Gordon et al. (2003).
- (ii) LMC – The Large Magellanic Cloud – a flatter UV extinction curve with a strong 2175-Å graphite absorption feature. Adopted from Gordon et al. (2003).
- (iii) MW – The Milky Way – similar to the LMC dust law with a strong 2175-Å feature. Adopted from Seaton (1979) fitted by Fitzpatrick (1986).
- (iv) SBG – The Calzetti Starburst Law – a monotonic extinction curve similar to MW and LMC dust but lacking the graphite feature. Adopted from Calzetti et al. (2000).
- (v) GREY – The Gaskell AGN Law – flattens in the UV due to absence of small grains. Adopted from Gaskell et al. (2004).

The dust-attenuated power-law spectrum model, expressed in absolute AB magnitude versus rest wavelength λ , is

$$M_{\text{AB}}(\lambda) = M_0 + 2.5 \alpha \log_{10} \left(\frac{\lambda}{\lambda_0} \right) + R(\lambda) \times E(B - V). \quad (3)$$

Here, $R(\lambda) \equiv -2.5 \log_{10}(A_\lambda)/E(B - V)$ is the dust attenuation in magnitudes per colour excess $E(B - V)$. The intrinsic power-law spectrum is $L_\nu = L_0 (\lambda_0/\lambda)^\alpha$, with disc theory predicting a power-law index $\alpha = 1/3$. With no dust, $E(B - V) = 0$, the model’s absolute AB magnitude is M_0 at the fiducial rest wavelength $\lambda_0 = 2400 \text{ \AA}$, chosen because the vast majority of the SDSS quasars have rest-frame coverage at 2400 Å thus minimizing cases that pivot at wavelengths outside the observed *ugriz* range.

For all five dust laws, and for each SDSS quasar, we fit the observed five-band disc SED, holding $\alpha = 1/3$ fixed and minimizing χ^2 to estimate the two model parameters, M_0 and $E(B - V)$ in equation (3). Fig. 5 illustrates this fit and de-reddening procedure for the disc

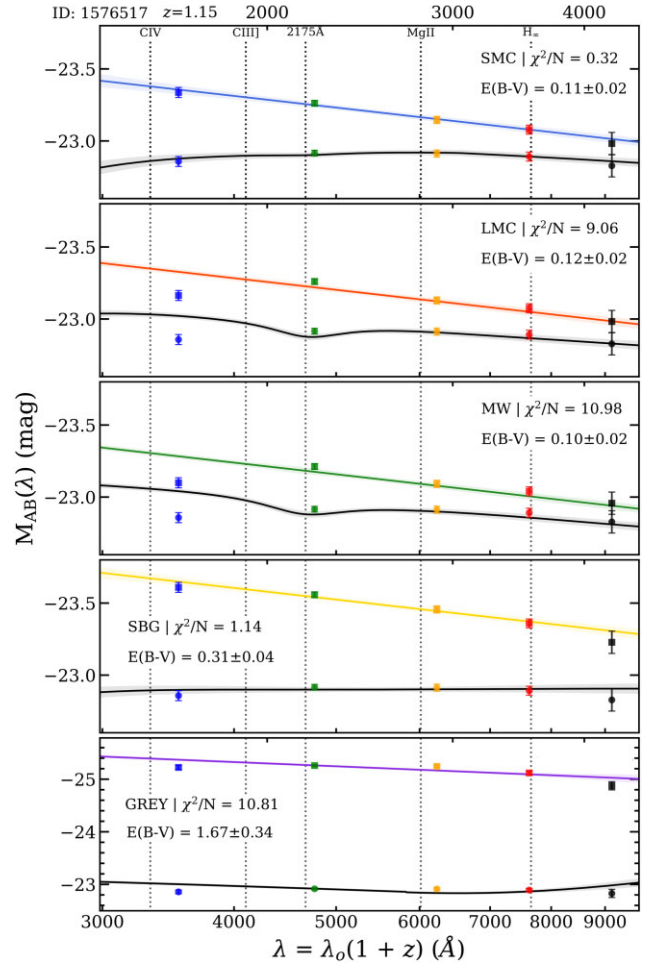


Figure 5. Dust-correcting the quasar disc spectrum (absolute AB magnitude versus observed- and rest-frame wavelength) to illustrate the procedure with each of the five dust laws in Fig. 4, for the same object as in Fig. 1. In each panel, the best-fitting dust-attenuated $L_\nu \propto \nu^{1/3}$ power-law spectrum (grey curve) is fitted to the observed disc fluxes (filled circles with error bars, coloured to correspond with the *ugriz* filters). The coloured lines show the same models after dust-correcting by setting $E(B - V) = 0$. The corresponding coloured square points are similarly dust-corrected data. The best-fitting $E(B - V)$ and the reduced χ^2/N is shown in each panel. With two parameters fitting five data, there are $N = 5 - 2$ degrees of freedom.

spectrum of SDSS ID 1576517 ($z = 1.15$) determined in Fig. 1. For each of the five dust laws, the de-reddened model spectrum, setting $E(B - V) = 0$, gives the intrinsic power-law $L_\nu \propto \nu^{1/3}$ fixed at the best-fitting value of M_0 . This also allows the photometric data to be dust-corrected by compensating for the dust extinction at each wavelength. This analysis delivers best-fitting estimates for $E(B - V)$ and M_0 , and a five-band dust-corrected SED, for each of the 9242 quasar discs.

With 2 parameters fitted to five data, there are $N = 3$ residual degrees of freedom. If the data and model are reliable, the reduced χ^2/N should be $1 \pm \sqrt{2/3}$, helping to discriminate among the five dust laws. For the $z = 1.15$ quasar in Fig. 5, the *g*-band happens to sample the redshifted 2175 Å feature that arises from graphite grains and is prominent in the MW and LMC dust laws. The observed disc SED is relatively smooth about the *g*-band. This strongly disfavours the LMC and MW dust laws, $\chi^2/N = 9.06$ and 10.98, respectively, for which the *g*-band datum is above and the *u*-band datum is well

below the best-fitting model. For the GREY dust law, the best fit requires a larger dust correction compared with the other dust laws. Also, the GREY dust law leaves relatively large residuals, and so is also strongly disfavoured, $\chi^2/N = 10.81$. For this particular quasar, and for the assumed power-law index $\alpha = 1/3$, the SMC and SBG dust laws remain viable, with $\chi^2/N = 0.32$ and 1.14, respectively.

A secondary metric to consider is the best-fit colour excess $E(B - V)$, which quantifies the line of sight dust column density. A prior on the dust reservoir of the quasar host galaxy may be set by the relatively small values observed in most extragalactic systems, with the notable exception of dusty starbursts (Casey, Narayanan & Cooray 2014; Talia et al. 2021). In our analysis, with $E(B - V)$ a free parameter, a fit requiring a much higher $E(B - V)$ should be rightly disfavoured. In Fig. 5, the best fit with the GREY dust law gives $E(B - V) = 1.67 \pm 0.34$ mag, the SBG dust law gives 0.31 ± 0.04 mag, and the SMC, LMC, and MW dust laws are consistent with $E(B - V) \approx 0.11 \pm 0.02$ mag. Thus, importantly, the SMC is not only the best-fitting dust-law as measured by χ^2/N , it also requires a significantly smaller $E(B - V)$ when compared to the next-best-fitting SBG dust law.

The similarities and differences among the dust law fits discussed above for SDSS ID 1576517 are found to hold statistically in the aggregate sample. For the SMC dust law, the lower panels of Fig. 3 demonstrate the dramatic tightening of the $g - i$ colour distribution effected by dust-correcting the quasar disc SEDs. For all five dust laws, Fig. 6 compares their dust-corrected colour–magnitude and colour–redshift distributions, reporting for each case the colour dispersion $\sigma(g - i)$ and the total χ^2 summed over all objects. The dust-corrected disc SEDs cluster around the assumed intrinsic $L_\nu \propto \nu^{1/3}$ power-law disc spectrum, with relatively mild dependencies on redshift. The tightest dispersions, $\sigma(g - i) \sim 0.23$ mag, are achieved similarly by the SMC, SBG, and GREY dust laws. This is closely followed by the LMC and MW dust laws, at 0.25 and 0.30 mag, respectively.

Despite their similar success in reducing the $g - i$ dispersion, we note several differences among the five dust laws. First, the GREY dust law, flat in the UV, is problematic as it spreads the dust-corrected disc SEDs over a wide range of implausibly large luminosities. In our view, this strongly disfavours the GREY dust law unless the intrinsic SEDs of quasar discs differ very substantially from a power-law spectrum.

For the LMC and MW, dust laws featuring graphite absorption at 2175 Å, the $g - i$ distribution has a tight core arising from the redshift range $0.9 < z < 1.6$, and broader wings from outside this range. This redshift structure stems from the 2175-Å feature redshifting across the g , r , and i bands, at $z \sim 1.2$, 1.9, and 2.5, respectively. At these redshifts, the evidence for absence of graphite absorption keeps $E(B - V)$ relatively small and better constrained than at intermediate redshifts where the feature falls between bands. This highly structured redshift dependence reduces the viability of our fits with these dust laws (see Appendix A).

In comparison, the SMC and SBG dust laws produce dust-corrected disc SEDs with tight distributions in both luminosity and colour, with small undulations in redshift that may plausibly be associated with emission-line features redshifting across the g and i bands, as indicated in the top right-hand panel of Fig. 6. Our fits with these dust laws also achieve the lowest total χ^2 , 2.5×10^5 for SMC and 3.3×10^5 for SBG, compared with $(5.0, 7.3, \text{ and } 8.8) \times 10^5$ for the (LMC, MW, and GREY) dust laws.

In conclusion, the SMC is our preferred dust law. It appears to be both reasonable and the best-fitting dust law overall, with a tight $g - i$ colour distribution centred about colour of an expected $L_\nu \propto \nu^{1/3}$

power law which is well constrained nearly equally at all redshifts. The SBG dust law is a close second choice, but with a somewhat higher χ^2 . For individual sources, the SMC provides the best fit in 43 per cent of cases, followed by the LMC, MW, and SBG at around ~ 17 per cent each, and lastly by GREY at < 7 per cent (see Fig. 12). We continue with all five dust laws, but consider the SMC dust law to be the most appropriate for our subsequent analysis.

3.4 Host galaxy SEDs

As a check on our SED decomposition procedure, using variability to separate the variable disc and non-variable galaxy SEDs, we examine the resulting galaxy SEDs. If our linear extrapolation to (sometimes much) lower fluxes than observed is a poor approximation, the resulting galaxy SEDs could be distorted.

Fig. 7 shows the SEDs inferred for the quasar host galaxies, sorted by redshift and by dust extinction. The red curves show the galaxy SEDs. The blue curves show the corresponding dust-corrected disc SEDs. Higher redshift host galaxies appear to be more luminous than those at lower redshifts. This is a natural consequence of the SDSS quasar sample being approximately magnitude limited, with fainter objects being detectable only at lower redshifts. Note that at the highest redshifts, $3 < z < 6$, the galaxy SEDs are strongly affected by the Lyman break moving into and thus suppressing the luminosity in the u band.

The host galaxy SEDs may be expected to be fainter and redder in quasars for which a large $E(B - V)$ is inferred to produce a $L_\nu \propto \nu^{1/3}$ intrinsic disc spectrum. However, comparing the right two columns in Fig. 7, we see no strong trend in this direction. The host galaxies of more attenuated discs are perhaps a bit fainter, but not much redder. This implies that dust along the line of sight to the quasar disc is not strongly correlated with dust on lines of sight to stars in the host galaxy.

Turning to trends with redshift, at $z < 1$, the quasar host galaxy SEDs covering 2000–6000 Å all look similar. They are fainter than and intermediate in spectral shape between the red SED of NGC 7585 and the blue SED of Mrk 930, typical red sequence and blue cloud galaxies, shown for comparison in each panel of Fig. 7. At $z > 1$, the quasar host galaxy SEDs are brighter, and an increasing fraction of them exhibit a UV component producing a V-shaped ugr dip, with g fainter than u or r . This can be interpreted as a young stellar population as in star-forming (blue cloud) or intermediate (green valley) galaxies. They constitute a minority at $1 < z < 2$, and a majority at $2 < z < 3$, compatible with maximum star formation at cosmic noon, and decreasing thereafter. At $z > 3$, virtually all of the quasar host galaxies are blue-cloud starbursts, with strong UV emission and brighter than the SED of the compact blue starburst galaxy Mrk 930. The Lyman break appears to depress the galaxy SEDs on the blue end.

These trends with redshift accord with our current understanding of the star formation history of galaxies over cosmic time (Madau & Dickinson 2014; Förster Schreiber & Wuyts 2020). As summarized in Fig. 8, star-forming hosts become increasingly faint with time. We find no significant change if we remove the 3 per cent of sources registered with resolved morphologies, as they constitute $\lesssim 20$ per cent of sources at $z \lesssim 0.7$, and $\ll 1$ per cent at higher redshifts. While these trends could be affected by unknown selection biases, they are broadly consistent with the well-known fading of star formation between $z \sim 2$ and the present epoch. In contrast to the quasar host galaxies, the dust-corrected disc luminosities are remarkably stable across all epochs, $M_{AB} \sim -22.6$ at $\lambda_0 = 2400$ Å, with a dispersion of ~ 0.4 mag, becoming less certain at z

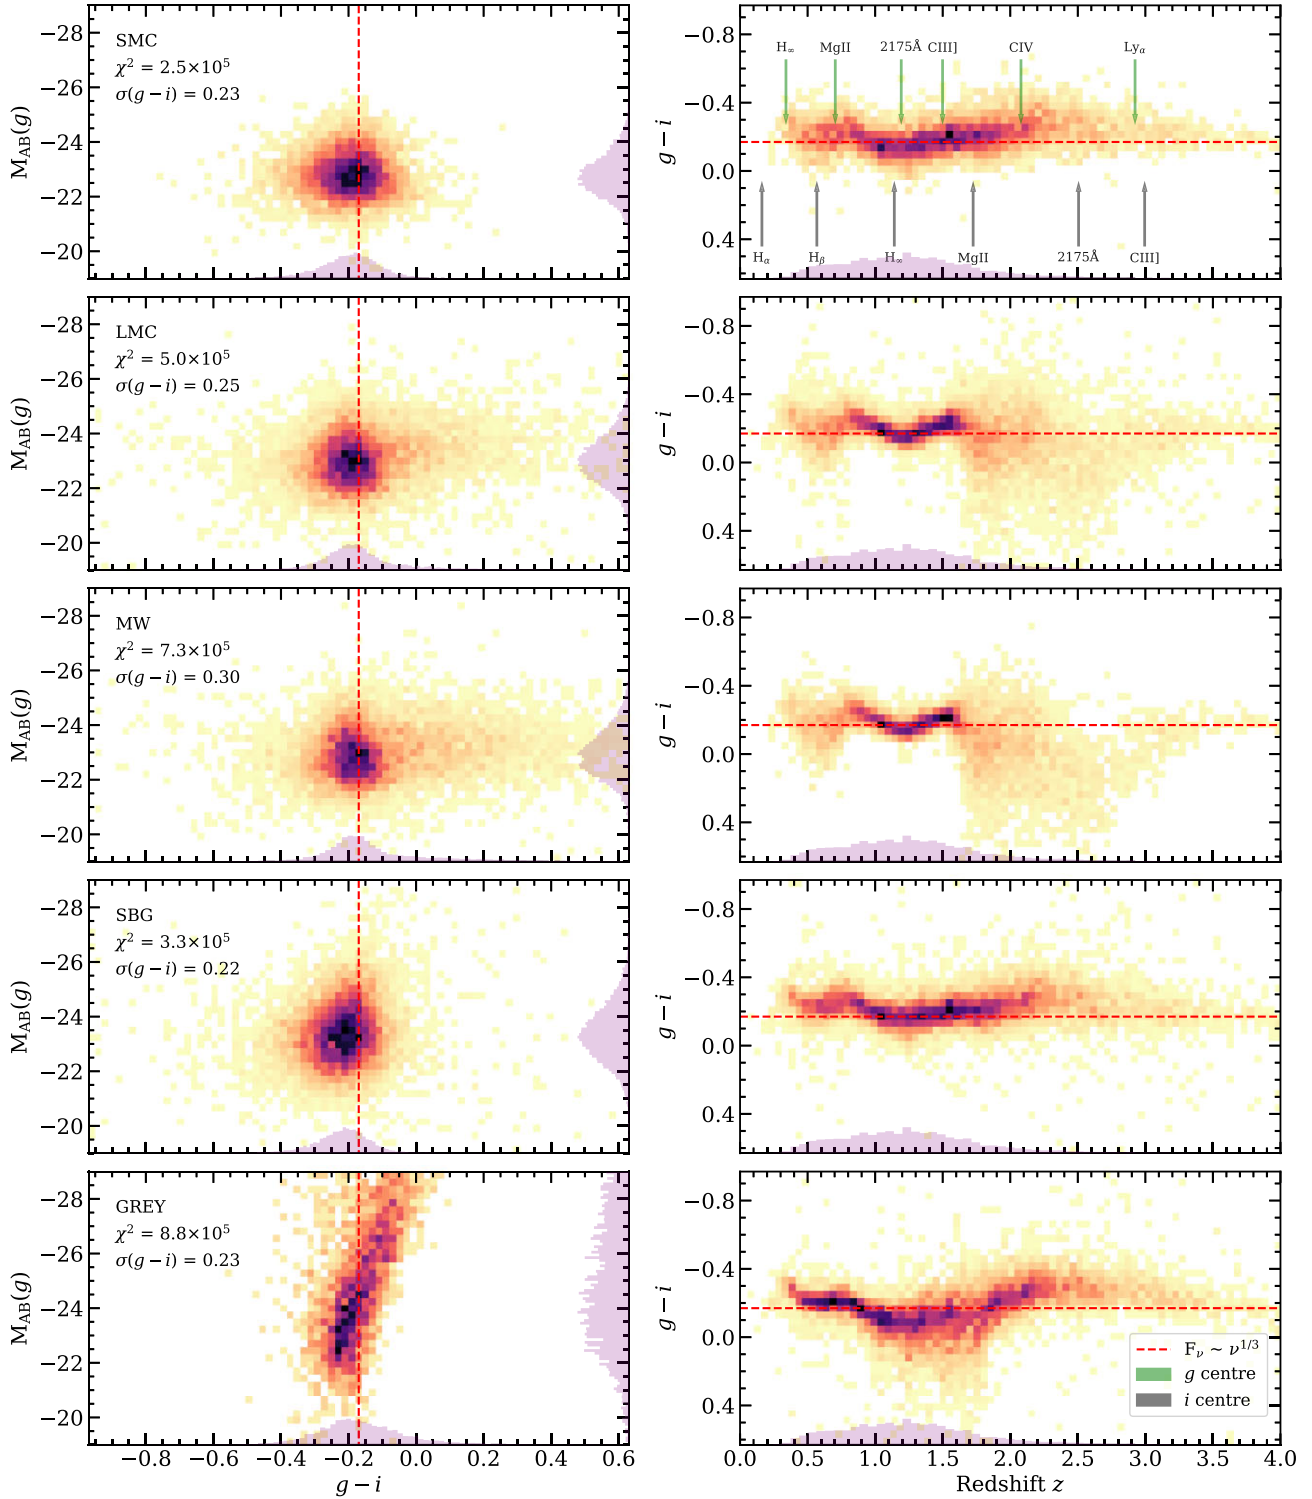


Figure 6. $g-i$ colour–magnitude (left-hand panel) and redshift–colour (right-hand panel) diagrams for each assumed dust law used to de-redden the 9156 *ugiz* disc SED. Distributions are coloured darker with increasing density and shown as a histogram projected on to each axis. The expected $F_\nu \sim \nu^{1/3}$ is indicated by the dotted red lines. Arrows indicate the presence of an emission feature in the center of the g - or i -band in green and grey, respectively. The reported values for χ^2 are total combined χ^2 statistics over the sample.

> 3 where extrapolation redward of the observed SED is required. Nevertheless, these encouraging results serve to validate our procedure using variability to separate the quasar disc and host galaxy light.

Some of our galaxy SEDs have u brighter than g , in fact rising more rapidly into the UV compared with a blue stellar population, perhaps even more rapidly than a Rayleigh–Jeans slope. This effect is likely a small flaw in our decomposition procedure. We currently

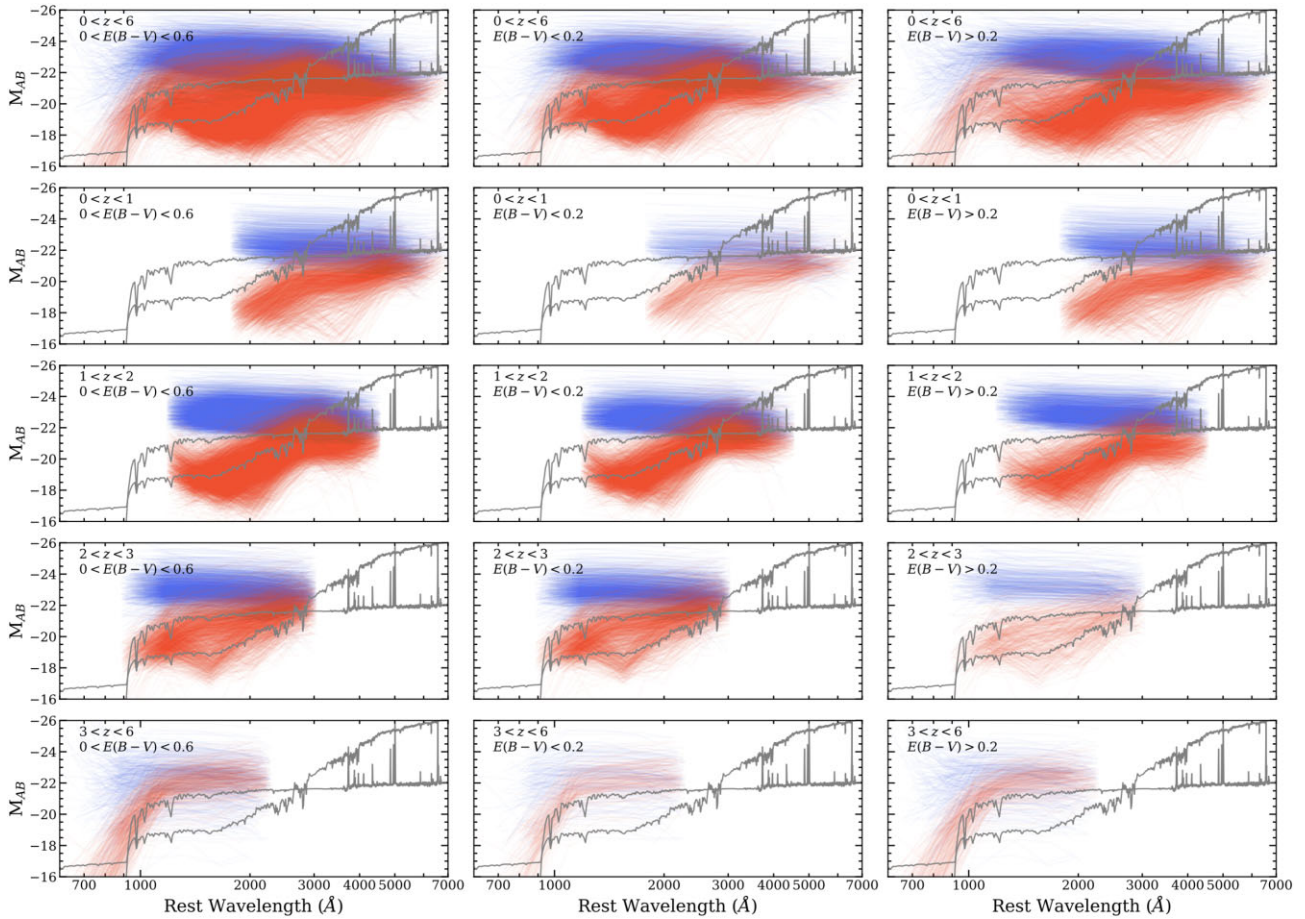


Figure 7. Quasar host galaxy SEDs (red) and dust-corrected disc SEDs (blue) sorted by redshift (columns) and $E(B - V)$ (rows). SEDs for typical galaxies in the red sequence (e.g. the lenticular shell galaxy NGC 7585) and blue cloud (e.g. the compact starburst Mrk 930), adopted from Brown et al. (2014), are shown in each panel for reference.

set X_{gal} at the lowest possible level, so that the extrapolated flux in one band, usually u or g , is 1σ above 0. A slightly higher level for X_{gal} could be used, thus moving a small fraction of the disc SED to the galaxy SED. The effect would be to make the V-shaped ugr dip in the galaxy SEDs less prominent in those cases where g is fainter than u , elevating the galaxy flux at g and moving the galaxy SEDs closer to the SED of a blue cloud galaxy. The disc SED would then have a correspondingly lower flux at g . We have not yet implemented this procedural tweak. We expect it to have a relatively small effect on the disc SEDs, which are much brighter than the galaxy.

These results follow the trends found in the analysis of Matsuoka et al. (2014), who performed a spatial decomposition to extract point-like quasar signals from their host galaxies, based on the same SDSS observations of Stripe 82. Limited to resolved sources at $z < 0.6$, they find that quasars are bluer than their host galaxies, with a quasar-to-host ratio of ~ 8 in u and ~ 1 in i . For our sources, at $z < 0.6$, host light is also typically fainter than our de-reddened discs, by a factor of ~ 90 in u and ~ 3 in i . However, a more equivalent comparison using our reddened (i.e. uncorrected) disc components produces a less extreme ratio of ~ 30 in u and ~ 2 in i , in better agreement with Matsuoka et al. The remaining discrepancy could be driven by a combination of selection effects, PSF-modelling biases in the analysis from Matsuoka et al. and that our definition of host galaxy may underestimate the host contribution at u . Regardless, it seems that the quasar discs are bluer than their hosts.

4 COMPOSITE SPECTRA OF VARIABLE QUASAR DISCS

The SDSS Stripe 82 quasar sample provides an unprecedented multiyear record of multiband quasar variations, but it is limited to just five optical bands ($ugriz$). Despite this drawback, we can leverage the cosmological redshift range to construct a composite quasar disc spectrum at somewhat finer spectral resolution and extending to much bluer rest-frame ultraviolet wavelengths. Our approach implicitly assumes that the spectral features of the accretion disc are universal, an assumption we made in the dust-correction procedure by assuming an $L_{\nu} \propto \nu^{1/3}$ power law for the intrinsic disc spectrum. As justified in Section 3.3, we assume that local extinction follows the SMC law for all sources when constructing our composite spectrum. It is also important to note that spectral features seen here will be smoothed out by the resolution of the filter profile of each band.

For each SDSS quasar, we have removed the host galaxy contribution by using the spectrum of the variable component, and corrected for possible dust extinction and SMC-like reddening in the host galaxy, assuming an $L_{\nu} \propto \nu^{1/3}$ power law for the intrinsic disc spectrum. We construct a composite disc spectrum by combining the resulting disc SEDs for quasars sampled across a continuous range of redshifts $0.1 < z < 6$. Our dust-correction procedure fits the power-law model, $L_{\nu}(\lambda) = L_0 (\lambda_0/\lambda)^{\alpha}$, assuming $\alpha = 1/3$, to determine for each quasar the luminosity L_0 at reference rest wavelength $\lambda_0 =$

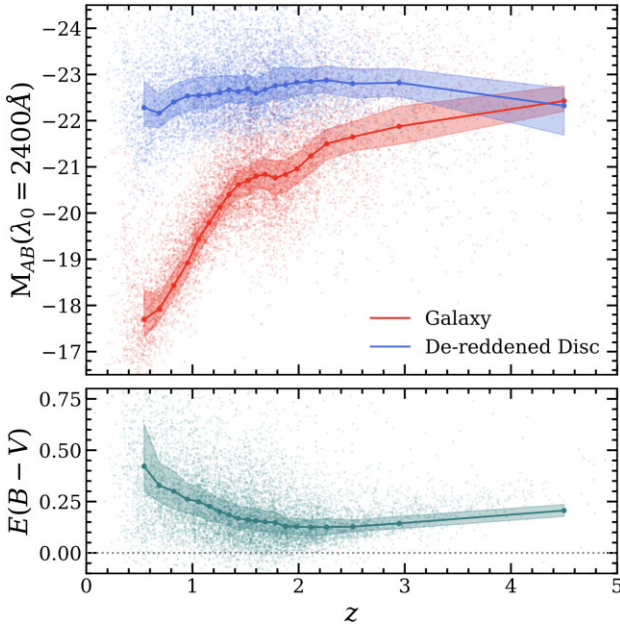


Figure 8. Upper panel: rest-frame 2400-Å absolute AB magnitudes as a function of z for the galaxy SED component (red) and disc component (blue), dust-corrected assuming an SMC-like attenuation, and $L_\nu \propto \nu^{1/3}$ power law. Lower panel: Dust attenuation $E(B - V)$ estimates assuming SMC. Bins have equal occupation at all redshifts and are used to determine the medians and 68 per cent ranges as shown.

2400 Å, and the required $E(B - V)$. The power-law spectrum provides a backbone for our composite disc model. We simply scale the L_0 for each quasar to a common value (-22.6 AB, as evidenced by Fig. 8), and scale the dust-corrected $ugriz$ fluxes by the same scale factor. This provides five measurements on the composite spectrum, at the rest wavelengths of the $ugriz$ bands at the redshift of that quasar. Doing that for the full sample gives $\sim 45\,000$ points, which we then average with a binned median to reduce the scatter.

In Section 4.1, we construct the composite disc spectrum assuming the canonical $L_\nu \propto \nu^{1/3}$ power law. In Section 4.2, we generalize this analysis by assuming $L_\nu \propto \nu^\alpha$ and solving for the best-fitting power-law slope α , thus testing the disc theory prediction $\alpha = 1/3$.

4.1 Composite disc spectrum for $L_\nu \propto \nu^{1/3}$

Fig. 9 presents the composite disc spectrum assuming a canonical $L_\nu \propto \nu^{1/3}$ power law. We remove 1 per cent of sources with the worst χ^2 (typically in excess of 100) to keep these poor fits from dominating the overall statistics. The top panel stacks the five-band dust-corrected disc SEDs for over the remaining 9 150 SDSS quasars, sorted by redshift, and coloured according to relative brightness, interpolating linearly between the pivot wavelengths of the $ugriz$ bands. Despite somewhat larger noise in the u and z bands, there is clearly a general increase in brightness toward bluer rest-frame wavelengths. This reflects the assumed $L_\nu \propto \nu^{1/3}$ power-law adopted for the dust corrections.

The middle panel presents the composite disc spectrum, which undulates above and below the assumed power-law spectrum (red dashed line). Here, the $5 \times 9\,150$ individual photometric measurements are summarized using a median with ~ 300 points in each bin (blue curve with a 68 per cent envelope). We additionally show 1 in 5 (i.e. 20 per cent) of sources to illustrate the object-to-object variations. The panel below this uses a similar format (in green) to

show residuals relative to the power law. The bottom panel shows the number of quasars contributing at each wavelength, nearly all 9150 in the middle at 2400 Å and dropping below 100 on the ends below 600 Å and above 7000 Å.

The power-law model provides a reasonable match to the data, with a reduced $\chi^2/N = 11.75$. Undulations around the power law are significant and plausibly attributed to variable spectral features such as the Balmer continuum emission around 3500 Å. Only a handful of low-redshift quasars contribute to the H α peak at 6563 Å. The downturn blueward of 1200 Å is expected due to intervening Lyman α forest absorption in the u band at $z > 2$. The more dramatic drop blueward of 900 Å is the Lyman break arising from Lyman continuum absorption depressing the u -band flux in the highest redshift quasars at $z > 3$.

4.2 Consideration of alternative spectral slopes

The common assumption that the underlying power law should be $L_\nu \propto \nu^{1/3}$ originates from classical theory (Shakura & Sunyaev 1973) and has not been conclusively demonstrated to be the true underlying spectral power-law. This work so far has made the same assumption, and so now we question it. While the disc-decomposition procedure is entirely independent of the assumed underlying power law, the de-reddening procedure is not. Hence, we re-derive all dust reddening solutions for each of the five assumed dust laws assuming a range of underlying power-law slopes.

As before, we compute the aggregate χ^2 statistic for each assumed dust law *and* underlying power-law index α such that $L_\nu \propto \nu^\alpha$, for a coarse grid of α values. The results are shown in Fig. 10, measured both using the aggregate χ^2 of the sample (left-hand panel) and only the scatter in $g - i$ colours (right-hand panel). Each coloured curve corresponds to an assumed dust law and each point in α is coloured by the median $E(B - V)$ achieved under those two assumptions. Whereas red colours indicate high median dust extinction, greyscale colours denote the median $E(B - V) < 0$ is non-physical, although some objects in a given collection may still have $E(B - V) > 0$. The best-fitting α for each curve is calculated by fitting a seventh-order polynomial to the samples in α and computing the minimum χ^2 . Estimates for the best-fitting values of α are shown also in Fig. 10 and are indicated by colour-corresponding vertical dotted lines with their uncertainty envelopes calculated where $\Delta\chi^2 = \chi_{\min}^2/N$, which are unexpectedly small for such a large data set. For reference, grey vertical dashed and dash-dotted lines are included to indicate α corresponding to the canonical $L_\nu \propto \nu^{1/3}$ as well as an alternative $L_\nu \propto \nu^{5/7}$, respectively. Table 1 summarizes the best-fitting parameters for each dust law, for fits with $\alpha = 1/3, 5/7$, and α optimized for each dust law.

From the left-hand panel of Fig. 10, showing the χ^2 landscape as a function of α , we find that the graphite-heavy dust laws for LMC and MW are strongly disfavoured. Their minimum χ^2 occurs for a red spectral slope ($\alpha < 0$) and under the assumption of $L_\nu \propto \nu^{1/3}$ they are disfavoured at high confidence in clear excess of 5σ . The best-fitting solution for the UV-flat GREY dust law occurs for an even redder spectrum, $\alpha = -0.39$, and with a non-physical median $E(B - V) = -0.67$ mag. LMC, MW, and GREY also have relatively high χ_{\min}^2 values, corresponding to reduced χ^2/N values 11.19, 13.19, 13.80, respectively.

The SMC and SBG laws fare rather better, achieving lower best-fit χ_{\min}^2 values of 5.99 and 7.37, respectively. The SMC dust law achieves its lowest overall χ_{\min}^2 at $\alpha = 0.71 \pm 0.02$ and is tightly constrained relative to the SBG, which achieves its relatively higher χ_{\min}^2 at $\alpha = 1.39 \pm 0.04$. While the SMC enjoys a smooth progression

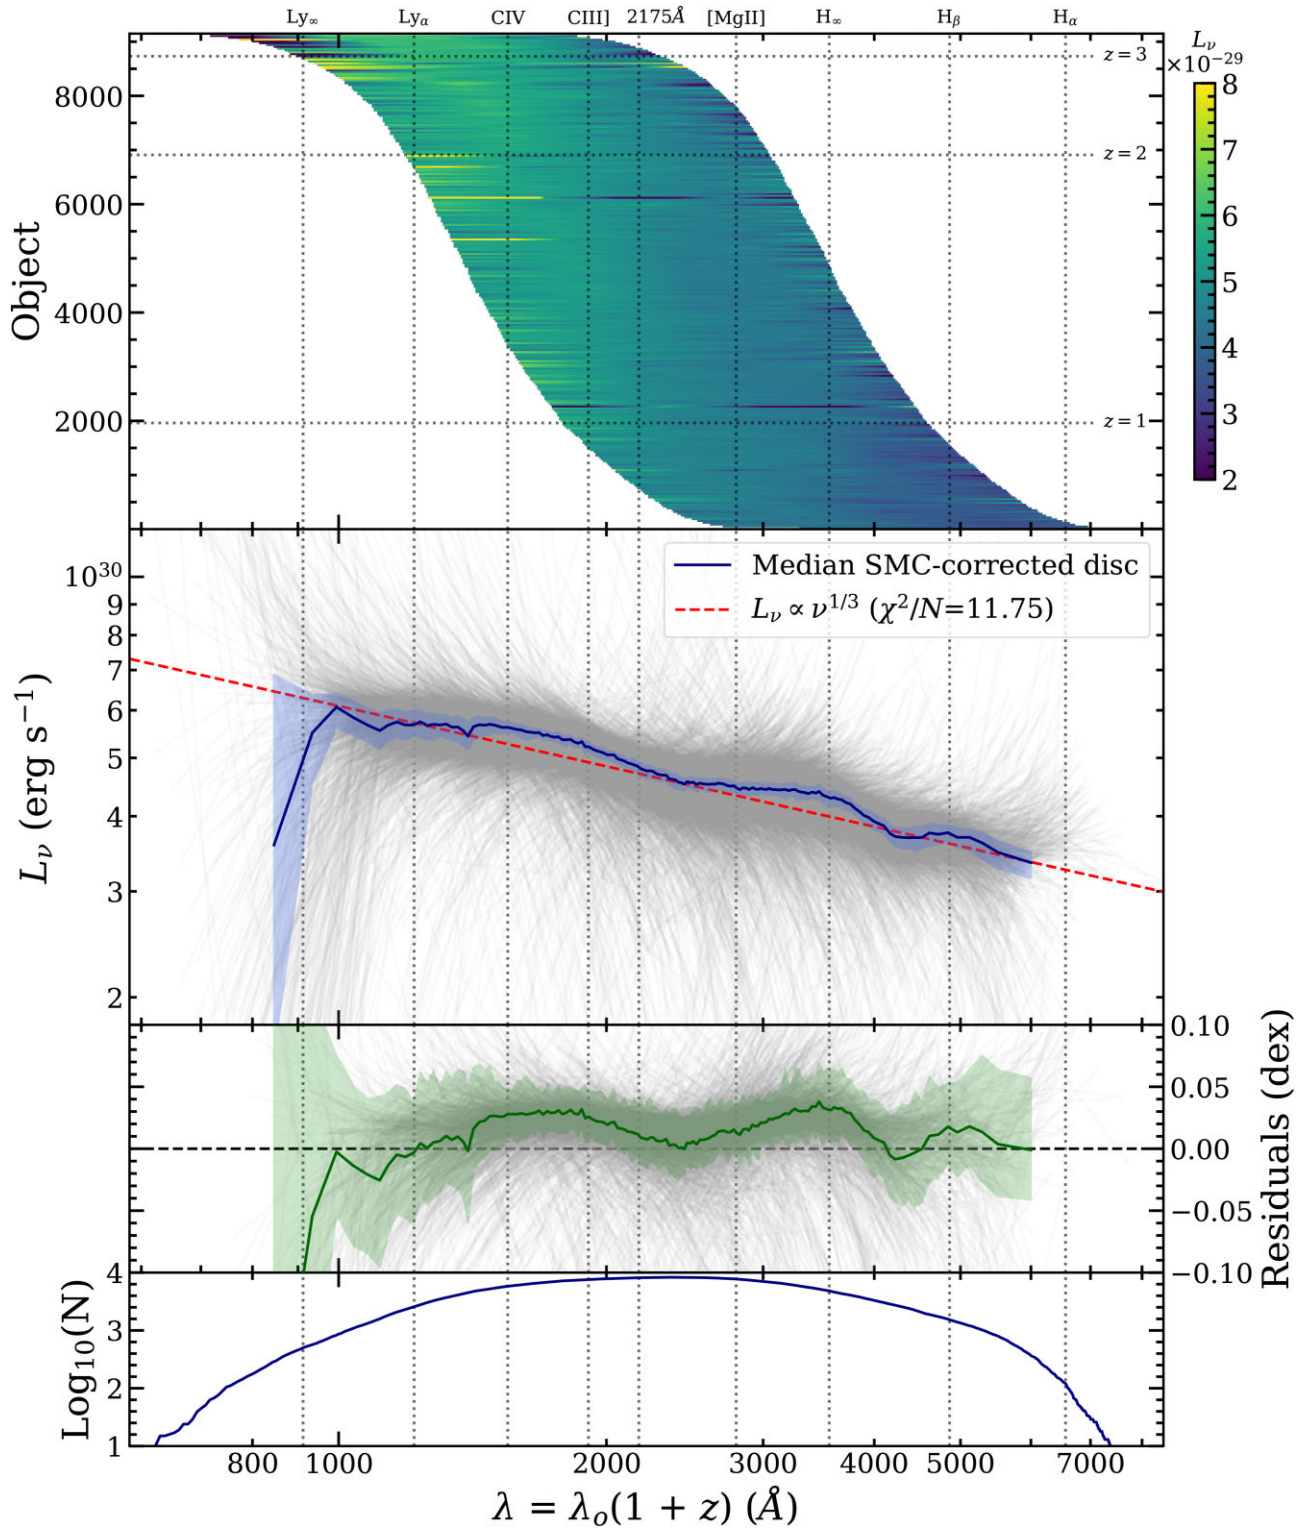


Figure 9. SEDs for the sample are shown at the top, shifted into the rest frame and ordered by z . Relevant emission lines are indicated by the vertical dotted lines. Shown in the middle is the composite spectrum of the de-reddened disc component computed with a binned median (blue, ~ 300 points per bin) with a 68 percentile envelope indicating the width of the distribution at that point, with scatter indicated behind by 1-in-5 SEDs. We assume an SMC-like reddening and $L_\nu \propto \nu^{1/3}$ power law, overlaid in red. The residuals are shown in the lower panel in green, with an envelope likewise from above. The lower panel shows the distribution of sources contributing to any give rest-frame wavelength.

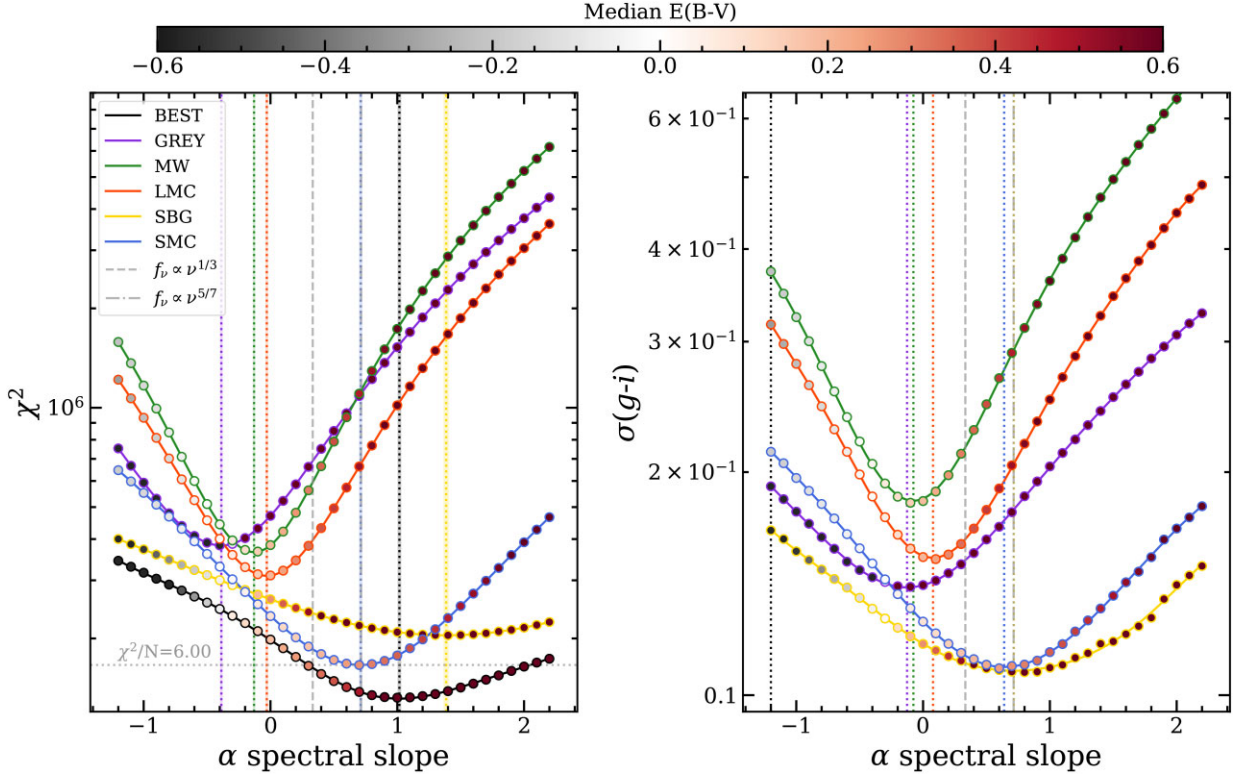


Figure 10. Two aggregate badness-of-fit (BoF) metrics, the total χ^2 (left-hand panel) and $\sigma(g-i)$ (right-hand panel), as functions of the accretion disc spectral index α . Points show the BoF values for each of the five candidate dust laws, fitting the α grid with seventh-order polynomials. For each dust law the median $E(B-V)$ is indicated by the filler of each point, red for positive values, and grey for (non-physical) negative attenuation. The best-fitting α values, at BoF minima, are marked on the left-hand panel by vertical dotted coloured lines within coloured bands denoting the uncertainty in α based on the $\Delta\chi^2 = \chi^2_{\min}/N$ criterion. Fiducial power-law models corresponding to $\nu^{1/3}$ and $\nu^{5/7}$ are marked by vertical grey lines. The best fit is achieved for SMC-like dust, median $E(B-V) = 0.28$ and $\alpha = 0.71 \pm 0.02$, close to $5/7$, as detailed in Table 1. Lastly, we compute the total χ^2 at each α using value corresponding to the best-fitting dust law for each source, shown in black circles.

Table 1. Summary of best-fitting χ^2_N for the total sample of 9 242 sources, reduced by the number of degrees of freedom ($N = 3 \times 9242$), including corresponding median $E(B-V)$ and $\sigma(g-i)$ for each of the five assumed dust laws, in the case, where $L_\nu \propto \nu^{1/3}$, $\propto \nu^{5/7}$, and the best-fitting spectral slope α whose uncertainty is boosted by $\sqrt{\chi^2/N}$.

	α	χ^2/N	$E(B-V)$ Median	$E(B-V)$ σ_{MAD}	$\sigma(g-i)$
$L_\nu \propto \nu^{1/3}$					
SMC	0.33...	6.70	0.18	0.13	0.11
SBG	0.33...	8.61	0.55	0.38	0.11
LMC	0.33...	14.26	0.34	0.22	0.16
MW	0.33...	21.30	0.32	0.23	0.22
GREY	0.33...	24.80	6.18	7.10	0.15
$L_\nu \propto \nu^{5/7}$					
SMC	0.71...	5.99	0.28	0.15	0.11
SBG	0.71...	7.88	0.86	0.37	0.11
LMC	0.71...	24.30	0.52	0.23	0.21
MW	0.71...	40.44	0.48	0.27	0.29
GREY	0.71...	39.62	11.63	11.61	0.18
Best-fitting α					
SMC	0.71 ± 0.02	5.99	0.28	0.15	0.11
SBG	1.39 ± 0.04	7.37	1.43	0.39	0.12
LMC	-0.03 ± 0.01	11.19	0.18	0.21	0.15
MW	-0.13 ± 0.01	13.19	0.13	0.21	0.18
GREY	-0.39 ± 0.02	13.80	-0.67	4.47	0.14

of median $E(B-V)$ values, the SBG requires even greater reddening for the same α despite turning over in $E(B-V)$ at the same α . Thus, the aggregate sample measured with χ^2 corroborates the aforementioned findings that the SMC provides the best-fitting dust solution to describe the de-reddened disc SED of the sample considered. However, we also measure the total χ^2 at each α by assigning each source a χ^2 corresponding to the best-fitting dust law, and recompute the best-fitting α (called ‘BEST’). Unsurprisingly, this combined sample finds a minimum $\alpha = 1.01 \pm 0.02$, mid way between the SMC and SBG laws which are the two best-fitting dust laws for any given α . However, this best-fitting α is suffers from a high median $E(B-V)$ compared to the SMC law at $\alpha \sim 1$, and may suffer from noisy measurements (as they are not weighted here), preferential χ^2 arising from unreasonably attenuated solutions, and other effects which make its interpretation non-trivial.

The constraining power of χ^2 in the left-hand panel of Fig. 10 is distinctly superior to that of $\sigma(g-i)$ in the right-hand panel. This makes sense, as χ^2 utilizes all five bands compared to only two bands in $\sigma(g-i)$. Nevertheless, comparing these may increase confidence in the results and deepen our understanding of the models. Note that the order of the best-fitting α values for different dust laws is the same for minima of χ^2 and minima of $\sigma(g-i)$. The best-fitting α values have a smaller range for minima of $\sigma(g-i)$. The MW, LMC, and GREY dust laws cluster around $\alpha = 0$, with relatively high $\sigma(g-i)$. The SMC and SBG are both close to $\alpha = 0.7$, with SBG achieving

a slightly lower $\sigma(g - i)$ than that for SMC. The $E(B - V)$ values at a given α are generally similar between the estimators.

Owing to the large sample size of this investigation, the constraint on the best-fitting slope α is remarkably tight, uncertain by of order 1 per cent for the χ^2 BoF. Both χ^2 and $\sigma(g - i)$ prefer a power-law spectral index α significantly bluer than the canonical $L_\nu \propto \nu^{1/3}$ accretion disc spectrum. The $L_\nu \propto \nu^{1/3}$ power law is statistically inconsistent with our results.

The merit of this result is further explored in Fig. 11 where we perform the same spectral composition procedure as shown by Fig. 9 but assume an $L_\nu \propto \nu^{5/7}$ while maintaining the assumption of an SMC-like power law. By doing so, we find an even more consistent picture with the $L_\nu \propto \nu^{5/7}$ power law, achieving a $\chi^2/N = 9.00$. This is lower than that achieved with the expected $L_\nu \propto \nu^{1/3}$, suggesting that $\alpha = 5/7$ is a more appropriate model. In addition, it is apparent that the bluest residuals have lessened, with the de-reddened disc spectrum now being consistent with a smooth $L_\nu \propto \nu^{5/7}$ power law within a 68 percentile range for wavelengths bluer than H_α and redder than the Lyman continuum break. We discuss the implications of this serendipitous finding in the following section.

5 DISCUSSION

5.1 Assumptions and caveats

Advantageous properties of the SDSS data set analysed here are its unprecedented number of quasars and the long time-span over which the five-band *ugriz* photometry has been obtained. For each quasar, we leverage its variable nature to separate the variable disc component from its static host galaxy. Our decomposition method treats each observation as an independent measurement of the galaxy+disc flux at some time-dependent dimensionless brightness level $X(t)$, where $X = 0$ is the mean level and $\Delta X = 1$ is the rms of the light-curve variations. This is illustrated in Fig. 1, where each flux measurement provides an independent constraint on the linear model, $F(\lambda, t) = B(\lambda) + A(\lambda)X(t)$. Here, the intercept $B(\lambda)$ is the mean galaxy + disc spectrum at $X = 0$ and the slope $A(\lambda)$ is the rms spectrum of the disc variations. Extrapolating the fit to fainter levels is assumed to effectively turn-off the variable disc light leaving just the galaxy spectrum at some minimum value of X . This point is somewhat arbitrary, particularly when the variations are small so that the extrapolation is a long one. In order to have a well-defined decomposition, we adopt the point at which the extrapolated flux is consistent with zero flux at 1σ , which we interpret as the limit below which the model is no longer physically meaningful.

Our estimates of $E(B - V)$ for five different extinction laws are computed for each source to quantify and compensate for extinction and reddening of the disc spectrum by dust along the line of sight to the accretion disc. This assumes that the observed disc spectrum is fainter and redder than the intrinsic disc spectrum due to line-of-sight reddening and extinction by dust, although the converse (requiring non-physical negative attenuation) is an allowed solution. For the intrinsic disc spectrum, we assume a power-law, $L_\nu \propto \nu^\alpha$. The estimate of $E(B - V)$ depends on the assumed power-law spectral index α , a bluer slope requires a larger $E(B - V)$. Our power-law disc model neglects possible contributions of emission lines and bound-free continua. The variability of this approximation is supported by Fig. 3, which shows that the variable disc spectrum has weaker emission features than the mean spectrum. However, as shown by Fig. 9, the final de-reddened disc spectrum, assuming an SMC-like extinction and $L_\nu \propto \nu^{1/3}$, shows that some emission features remain. These are of course smoothed by the broad bandwidths of the *ugriz*

filters, leaving wide and weak rather than narrow and strong emission features in the residuals. Although visually the residual features corresponding to the $\alpha = 1/3$ composite (Fig. 9) appear similar to those of the $\alpha = 5/7$ (Fig. 11), the latter achieves a significantly better fit, $\Delta\chi^2 \sim 20\,000$ or $\chi^2/N = 11.75 \rightarrow 9.00$. However, we caution that overfitting and unseen systematics may contribute to this effect.

Despite the straight-forward interpretation that a combination of SMC-like dust and a bluer spectral slope describes the variable accretion disc spectra of quasars, we note several caveats. First, for the MW and LMC laws the $E(B - V)$ distribution of the SDSS quasars has an implausible redshift dependence caused by the strong rest-frame 2175-Å absorption feature moving across the centre of a band. A ‘beating’ pattern is observed where the estimates of $E(B - V)$ have a large scatter when the 2175-Å feature is not directly observed (see Appendix A). This highlights a shortcoming in the modeling of the dust when adopting the MW and LMC models. We considered addressing this by using a prior favouring models that make $E(B - V)$ a smoother function of redshift, but decided in the interest of simplicity to omit this complication in our modelling. Our model also places no limitation on the extent to which the variable component can be reddened, which may permit extreme reddening requiring extraordinary dust column densities. We note that, with the exception of the GREY law, there are no instances of problematically dusty attenuation estimates.

5.2 On attenuation laws

The dust laws investigated in this work broadly fall into two categories. Either they are well described by a smooth power-law-like curve (e.g. SMC, SBG), or a power-law-like curve with a strong graphite absorption feature at 2175 Å (LMC, MW). The outlying case is that of the Gaskell’s dust law derived from a sample of AGN (GREY). Their differences are highlighted in Fig. 4.

As shown for a specific case in Fig. 5, the likelihood that a particular dust law is well suited for a particular source is assessed here using χ^2 to quantify the badness-of-fit and $E(B - V)$ to indicate cases where exceptional dust columns would be required. Similarly, to quantify the success in modelling the SDSS quasar sample as a whole, we employ two badness-of-fit metrics, χ^2 and $\sigma(g - i)$, along with the median $E(B - V)$. Given that the information presented by $\sigma(g - i)$ is contained in χ^2 and has generally less constraining power (see Fig. 10), we adopt χ^2 as the primary criterion, using $\sigma(g - i)$ and $E(B - V)$ for secondary considerations.

As presented in Section 3, assuming a $L_\nu \propto \nu^{1/3}$ power law, the least likely dust law is GREY which is statistically excluded at high confidence for the vast majority of the sources. However, 7 per cent of the sample finds a best-fitting solution with the GREY extinction law, but with an exceptionally large median $E(B - V) = 6.86$ mag for this sub-sample. Further, we find evidence to exclude the graphite absorption laws of the LMC and MW as a general best-fit, finding the best-fitting solution for only 17 and 16 per cent of the total sample, respectively.

We find the greatest success with the smooth power-law extinction laws, SMC and SBG. As measured by χ^2 in Fig. 10, the SMC provides the best fit to the sample as a whole and is consistent with a best-fitting $\alpha \sim 0.7$ from both the χ^2 and $\sigma(g - i)$ estimators. Assuming an $L_\nu \propto \nu^{1/3}$ power law, Fig. 12 shows that the SMC provides the best-fitting solution for 43 per cent of the sample while the SBG provides 17 per cent. Interestingly, although constrained with less information, the $\sigma(g - i)$ estimator finds that the SBG provides a solution similar to the SMC consistent with a $L_\nu \propto \nu^{5/7}$ power law. Thus, while we cannot exclude the SBG law outright, we

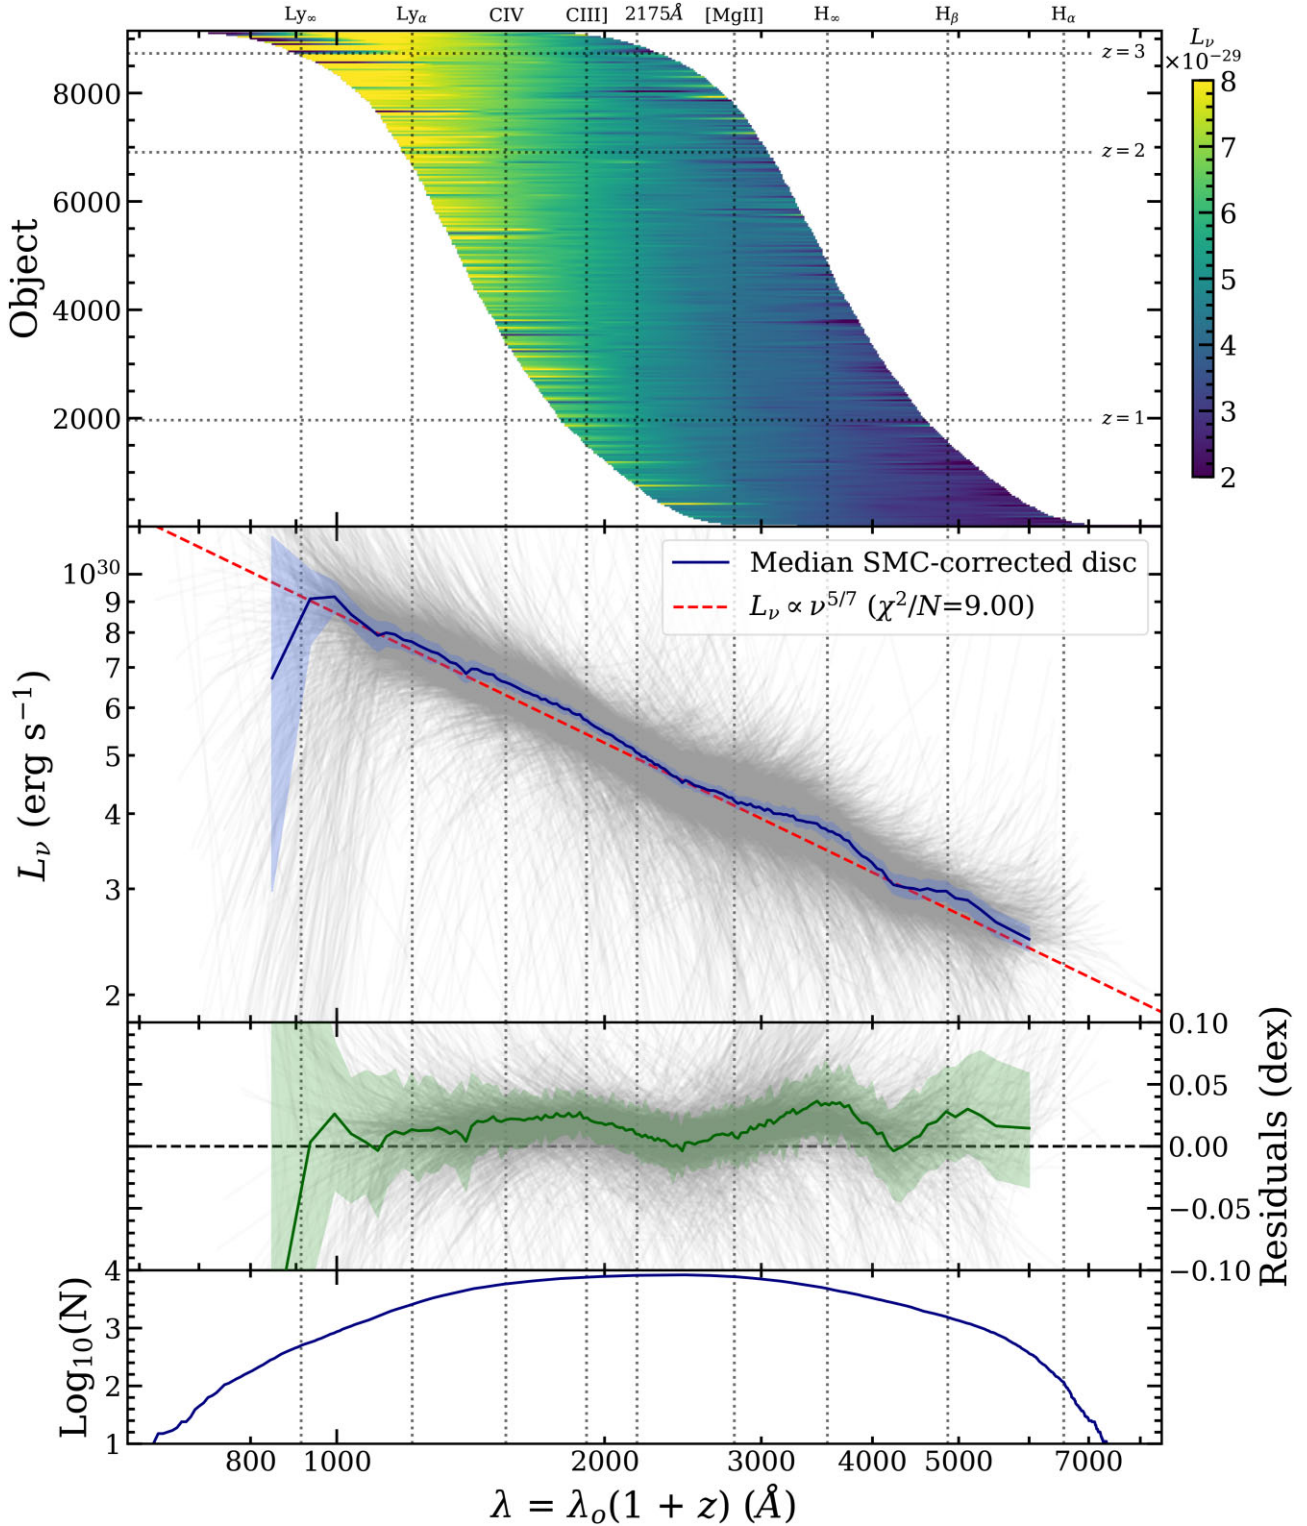


Figure 11. Same as Fig. 9, but assuming a $L_\nu \propto \nu^{5/7}$ power law.

none the less find the most likely dust law for this sample of quasars is the SMC. This result holds also in the case of an assumed $L_\nu \propto \nu^{5/7}$ power law.

Considering now the derived composite SEDs shown in Figs 9 and 11, under the assumption of an SMC-like dust law derived assuming either $L_\nu \propto \nu^{1/3}$ or $\propto \nu^{5/7}$, there are no discernible features consistent

with strong Balmer absorption or emission at the 10 per cent level. However, given that the emission is smeared across the broad-band filter, only the highest equivalent width lines could be detected. Continuum consistent with thermal emission from an optically thick accretion disc is clear (i.e. the ‘Big Blue Bump’; Malkan & Sargent 1982).

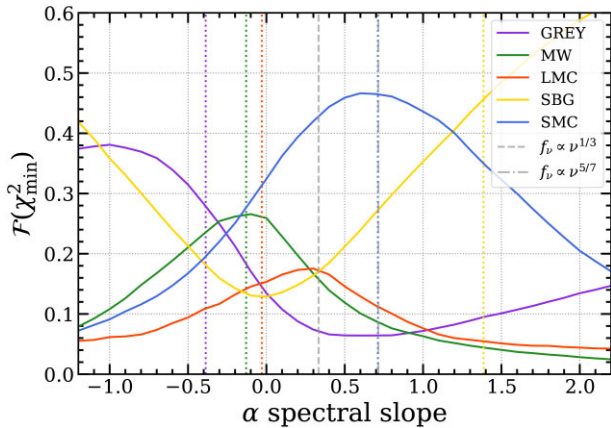


Figure 12. Fraction of sources per assumed spectral slope α best-fit by the given attenuation law, as measured by minimum χ^2 . Vertical dotted lines indicate the best-fitting α for each dust law.

The southern sample data set contains several sources observed at $3 < z < 6$. At these redshifts, the rest-frame u band intersects the expected rest-frame UV turnover of the accretion disc spectrum at $\sim 1000 \text{ \AA}$. Although this subset constitutes only a small fraction of the total sample, the effect of the turnover is evident from Fig. 9. The handful of these sources approaching $z \sim 6$ may also be affected by the neutral intergalactic medium which may be contributing to this observed turnover with resonant absorption by hydrogen gas clouds along the line of sight (i.e. the well-known Lyman Forest). Regardless of the physical mechanism driving this highly significant turnover, it is not accounted for in the continuous power law form assumed for the disc spectrum. Consequently, the continuous power-law model is not appropriate for the bluest bands for $z > 3$ objects, whose $E(B - V)$ may be overestimated due to the turnover acting as an extreme reddening of the u band.

5.3 Best-fitting power-law exponent

Initially, we assumed the canonical $L_\nu \propto \nu^{1/3}$ power law before expanding to a range of power law exponents to determine the most suitable power-law index and extinction law combination as assessed by the badness-of-fit χ^2 statistic. The result is shown in Fig. 10. The SMC law shows remarkable agreement from both χ^2 and $\sigma(g - i)$ estimators finding a best-fitting blue slope $\alpha \sim 0.7$ in both cases. For the latter estimator, SBG finds its minimum also at $\alpha \sim 0.7$. Taking χ^2 to be the more robust and more precise estimator, the corresponding α for the best-fitting SMC is 0.71 ± 0.02 . This is highly inconsistent with $L_\nu \propto \nu^{1/3}$ predicted for a geometrically thin steady-state disc, within these strict uncertainties.

Given this unexpected result, we verified the robustness of the de-reddening procedure by measuring 9000 simulated disc SEDs with $\alpha = 1/3$ perturbed with random noise corresponding to that of the observed photometry. We successfully recovered a best-fitting power-law index of $1/3$, confirming that the procedure is unbiased and that the result is indeed genuine.

This result contrasts with work from Kokubo et al. (2014), who used the same SDSS data set and employed a ‘flux-flux’ decomposition method most similar to ours in order to explore the variations in the photometric light curves. They found that the composite spectrum *before* decomposition features a red, $\alpha = -0.5$ slope relative to the much bluer $\alpha = 1/3$ ‘difference’ composite spectrum, in agreement with (Shakura & Sunyaev 1973). However,

Kokubo et al. caution that their method does not attempt to estimate or account for non-variable components of the host galaxy. Despite this work being the closest analogue to the present study available in the literature, the still different methodologies and assumptions make a concrete comparison of derived spectral slopes hazardous. It is possible though that the bluer slope derived in this work is found because we removed the static host galaxy component, which would otherwise produce a redder best-fit spectral slope.

Entirely serendipitously, we find our best-fit spectral slope $\alpha = 0.71 \pm 0.02$ to be statistically consistent with the recent theoretical prediction $\alpha = 5/7 \simeq 0.71$ for a thick disc heated predominantly by stresses conveyed by magnetic links between the inner disc and black hole (Agol & Krolik 2000; Mummery & Balbus 2020). We caution however that our result (1) assumes that all quasars have SMC-like dust and (2) is sensitive to the tail end of the χ^2 distribution used to compute the total χ^2 from Fig. 10. At face value, however, these stress-heated $\alpha = 5/7$ thick disc models cannot be ruled out by our results.

Our finding is intriguing, as it hints at additional accretion physics not considered in many previous studies. The possible importance of magnetic links connecting the inner disc with the black hole is noted by Thorne (1974). Consequences including a steeper disc temperature profile, $\tau \propto \lambda^{8/7}$, and bluer spectral shape, $L_\nu \propto \nu^{5/7}$, are worked out in detail for steady accretion by Agol & Krolik (2000), and for time-dependent accretion, as in tidal disruption events, by Mummery & Balbus (2020). These studies employ a fully relativistic framework of accretion discs, finding a temperature structure driven by energy liberated by both local viscous dissipation and by the torques arising from magnetic links to the spinning black hole and/or to gas inside the ISCO plunging toward the event horizon. This non-vanishing ISCO stress term augments the gravitational energy released by the inspiraling disc material, thus steepening the temperature structure. The steeper $T \propto r^{-7/8}$ temperature profile results in a bluer spectral slope, $L_\nu \propto \nu^{5/7}$, than in the traditional steady-state accretion disc model with $T \propto r^{-3/4}$ and $L_\nu \propto \nu^{1/3}$.

5.4 Comparison with intensive disc reverberation mapping

An independent method being used to probe accretion disc temperature profiles is to obtain intensive (sub-day) monitoring and then measure interband time-delays (Cackett, Horne & Winkler 2007; Edelson et al. 2019). This intensive disc reverberation mapping method (IDRM) assumes light traveltime delays $\tau \approx r/c$, and blackbody emission peaking near $\lambda \approx (hc/kT)$. A reverberating disc with $T \propto r^{-b}$ gives time-delays $\tau \propto \lambda^{1/b}$, and disc spectra $L_\nu \propto \nu^{(2b-3)/b}$. The standard disc model with $b = 3/4$ predicts $\tau \propto \lambda^{4/3}$ and $L_\nu \propto \nu^{1/3}$, while a steeper temperature profile with $b = 7/8$ predicts $\tau \propto \lambda^{8/7}$ and $L_\nu \propto \nu^{5/7}$.

The IDRM results to date are typically described as consistent with the standard disc prediction, $b = 3/4$, but with uncertainties large enough to admit $b = 1$. The most accurate IDRM results to date, from the Space Telescope and Optical Reverberation Mapping campaign monitoring of NGC 5548 in 2014, give a best-fitting power-law slope $b = 1.03 \pm 0.12$ from cross-correlation lags (Fausnaugh et al. 2016) or $b = 0.99 \pm 0.03$ from detailed fitting of a reverberating disc model to the light curves (Starkey et al. 2017). This corresponds to a steeper temperature profile, closer to $b = 7/8$ ($\alpha = 5/7$) than to $3/4$ ($1/3$).

A caveat, however, is that the disc sizes inferred from IDRM are typically larger than expected, by a factor of ~ 3 , and an excess lag in the Balmer continuum region suggests that bound-free emission from the (larger) broad emission-line region (BLR) may be contributing

significantly to the cross-correlation lags (Korista & Goad 2001; Lawther et al. 2018). Perhaps the clearest example is the lag spectrum from *HST* monitoring of NGC 4593 (Cackett et al. 2018). Work is underway to understand how best to disentangle the disc and BLR contributions to the measured lags.

6 SUMMARY AND CONCLUSIONS

In this work, we have separated the variable accretion disc light from their static host galaxies using broad-band photometric light curves, de-reddened their SEDs to account for dust, and leveraged the results to test accretion disc physics.

(i) We developed a method for decomposing quasar light curves by separating the contribution of the variable light from the static, background emission. Each disc SED was then de-reddened to provide a dust-free estimation of the underlying accretion disc light.

(ii) Of the five dust laws examined in this work, we find that the featureless laws of the Small Magellanic Clouds (SMC; Gordon et al. 2003) and starburst galaxies (SBG; Calzetti et al. 2000) are the most reasonable attenuation models as measured by their χ^2/N , with the SMC being slightly preferable as it typically requires less attenuation than the SBG.

(iii) Assuming an SMC-like dust attenuation, the best-fitting spectral slope α is found to be inconsistent with a standard $L_\nu \propto \nu^{1/3}$ corresponding the steady-state accretion model (Shakura & Sunyaev 1973). Instead, we find significant evidence for a $L_\nu \propto \nu^{5/7}$ accretion slope based on our best-fitting $\alpha = 0.71 \pm 0.02$, in agreement with the proposed disc model of Mummery & Balbus (2020).

If it is indeed the case that these ISCO-stress models better reflect the reality of accretion physics compared to previous models, then these observational findings challenge commonly made assumptions about the thermal structure of quasar accretion discs. Moreover, they have implications for deriving key properties of black holes and their accretion discs, including the Eddington luminosity and black hole mass. The observed variability of quasar discs might be plausibly attributed to changes in the magnetic links modulating ISCO stress and black hole spin. Future work is needed to confirm this model, requiring independent observations and continued monitoring of key sources in a way which these results can be confidently replicated.

Finally, we note that the methodology developed here should be useful for analysis of quasar variability data from the LSST, and could also be applied to data from multiobject spectroscopic monitoring surveys such as SDSS-RM.

ACKNOWLEDGEMENTS

The authors would like to thank Chelsea Macleod, Marianne Vestergaard, Michael Goad, and Julian Krolik for helpful discussions. We would also like to thank the organisers of the 2019 Quasars in Crisis meeting for facilitating so many useful conversations. Lastly, we would like to thank the anonymous referee for their generous insight and helpful comments which improved this work.

The Cosmic Dawn Center (DAWN) is funded by the Danish National Research Foundation under grant No. 140. JRW acknowledges support from the European Research Council (ERC) Consolidator Grant funding scheme (project ConTEXT, grant No. 648179), and from the University of St Andrews Undergraduate Research Assistant Scheme. KH acknowledges support from STFC grant ST/R000824/1.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National

Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is <http://www.sdss.org/>. The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

DATA AVAILABILITY

The data underlying this paper were accessed from Southern Survey of Stripe 82 Quasars as part of the Sloan Digital Sky Survey. The derived data generated in this research will be shared on a reasonable request to the corresponding author.

REFERENCES

- Agol E., Krolik J., 2000, *ApJ*, 528, 161
 Allen D. A., Glass I. S., 1976, *ApJ*, 210, 666
 Asmus D., 2019, *MNRAS*, 489, 2177
 Bellm E. C. et al., 2018, *Publ. Astron. Soc. Pac.*, 131, 018002
 Bessell M., Murphy S., 2012, *Publ. Astron. Soc. Pac.*, 124, 140
 Brown M. J. I. et al., 2014, *ApJS*, 212, 18
 Cackett E. M., Horne K., Winkler H., 2007, *MNRAS*, 380, 669
 Cackett E. M., Chiang C.-Y., McHardy I., Edelson R., Goad M. R., Horne K., Korista K. T., 2018, *ApJ*, 857, 53
 Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, *ApJ*, 533, 682
 Casey C. M., Narayanan D., Cooray A., 2014, *Phys. Rep.*, 541, 45
 Draine B. T., Malhotra S., 1993, *ApJ*, 414, 632
 Edelson R. et al., 2019, *ApJ*, 870, 123
 Elvis M., Wilkes B. J., Tananbaum H., 1985, *ApJ*, 292, 357
 Fausnaugh M. M. et al., 2016, *ApJ*, 821, 56
 Fitzpatrick E. L., 1986, *AJ*, 92, 1068
 Fitzpatrick E. L., Massa D., 1986, *ApJ*, 307, 286
 Förster Schreiber N. M., Wuyts S., 2020, *ARA&A*, 58, 661
 Gallerani S. et al., 2010, *A&A*, 523, A85
 Gaskell C. M., Goosmann R. W., Antonucci R. R. J., Whysong D. H., 2004, *ApJ*, 616, 147
 Gordon K. D., Clayton G. C., Misselt K. A., Landolt A. U., Wolff M. J., 2003, *ApJ*, 594, 279
 Graham M. J. et al., 2019, *Publ. Astron. Soc. Pac.*, 131, 078001
 Hönig S. F. et al., 2013, *ApJ*, 771, 87
 Hopkins P. F. et al., 2004, *AJ*, 128, 1112
 Ivezić e. et al., 2019, *ApJ*, 873, 111
 Kokubo M., Morokuma T., Minezaki T., Doi M., Kawaguchi T., Sameshima H., Koshida S., 2014, *ApJ*, 783, 46
 Korista K. T., Goad M. R., 2001, in Peterson B. M., Pogge R. W., Polidan R. S., eds, *ASP Conf. Ser. Vol. 224, Probing the Physics of Active Galactic Nuclei*. Astron. Soc. Pac., San Francisco, p. 411
 Krawczyk C. M., Richards G. T., Gallagher S. C., Leighly K. M., Hewett P. C., Ross N. P., Hall P. B., 2015, *AJ*, 149, 203

Lawther D., Goad M. R., Korista K. T., Ulrich O., Vestergaard M., 2018, *MNRAS*, 481, 533
 Lynden-Bell D., 1969, *Nature*, 223, 690
 MacLeod C. L. et al., 2012, *ApJ*, 753, 106
 Madau P., Dickinson M., 2014, *ARA&A*, 52, 415
 Malkan M. A., Sargent W. L. W., 1982, *ApJ*, 254, 22
 Matsuoka Y., Strauss M. A., Price Ted N. I., DiDonato M. S., 2014, *ApJ*, 780, 162
 Matthews T. A., Sandage A. R., 1963, *ApJ*, 138, 30
 Mummery A., Balbus S. A., 2020, *MNRAS*, 492, 5655
 Nandy K., Thompson G. I., Jamar C., Monfils A., Wilson R., 1975, *A&A*, 44, 195
 Oke J. B., 1974, *ApJS*, 27, 21
 Sandage A., 1965, *ApJ*, 141, 1560
 Schlafly E. F., Finkbeiner D. P., 2011, *ApJ*, 737, 103
 Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, 500, 525
 Schneider D. P. et al., 2010, *AJ*, 139, 2360
 Seaton M. J., 1979, *MNRAS*, 187, 73
 Shakura N. I., Sunyaev R. A., 1973, *A&A*, 500, 33
 Shen Y. et al., 2015, *ApJS*, 216, 4
 Shields G. A., 1978, *Nature*, 272, 706
 Starkey D. et al., 2017, *ApJ*, 835, 65
 Talia M., Cimatti A., Giuliatti M., Zamorani G., Bethermin M., Faisst A., Le Fèvre O., Smolčić V., 2021, *ApJ*, 909, 23

Thorne K. S., 1974, *ApJ*, 191, 507
 Wills B. J., Netzer H., Wills D., 1985, *ApJ*, 288, 94
 York D. G. et al., 2000, *AJ*, 120, 1579
 Zafar T. et al., 2015, *A&A*, 584, A100

APPENDIX A: PHYSICAL INTERPRETATION OF ATTENUATION ESTIMATES

As discussed briefly in Section 5, we do not place any priors on the allowed ranges of $E(B - V)$ parameter in our dereddening procedure. As a consequence of the prominent 2175-Å feature in the LMC and MW dust laws, estimates of $E(B - V)$ are more similar for sources where the bump is directly constrained by one of the five bands, which is shown to undulate with redshift in Fig. A1 for the LMC. This undulation is not seen, however, in the smooth dust laws of the SMC and SBG. We also find $E(B - V)$ flares up for sources at $z > 3$ where the u -band falls blueward of the Lyman continuum break. This discontinuity is not included in our continuous power-law model, and so mimics an extreme reddening. We do not interpret either feature as a genuine physical phenomena of accretion discs, but merely a limitation of our model. The consequences of this are described in Section 5.

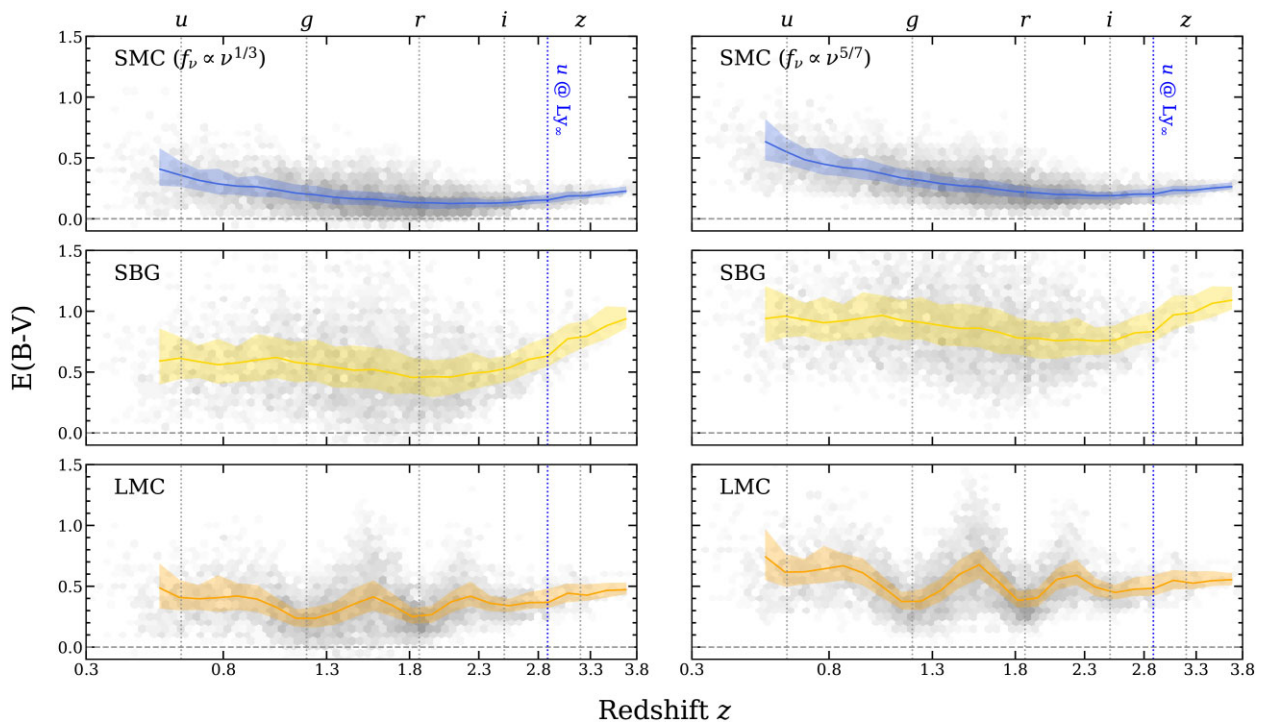


Figure A1. Binned median estimates of $E(B - V)$ over a range of redshift, assuming a $F_\nu \sim \nu^{1/3}$ (left-hand panel) or a $F_\nu \sim \nu^{5/7}$ (right-hand panel), for the two best-performing extinction laws (SMC, SBG) as well as the LMC which features significant absorption at 2175 \AA . The redshift at which the 2175-\AA feature is centred in each band is shown by the grey dotted lines. The coloured envelopes contain 68 per cent of sources in each bin.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.