FACULTY OF SCIENCE
UNIVERSITY OF COPENHAGEN

Master's thesis

Dust reddening of quasars:
A study of redshift and luminosity dependence of extinction, and of dust as main explanation of quasar UV SED diversity.

Amanda Bredahl Larssen

Academic Supervisor:
Marianne Vestergaard
Dark Cosmology Centre
Niels Bohr Institute
University of Copenhagen

Submitted: October 28, 2016

Abstract

Establishing an accurate account of the properties of dust in quasar nuclei and host galaxies has historically been hampered by the fact that the intrinsically emitted spectrum is not unique and universal. For a large sample of quasar spectra obtained from the Sloan Digital Sky Survey I investigate the extinction of quasars to test if and how it changes as a function of redshift and source luminosity - an aspect that has not previously been investigated. I also test if the diversity of the observed quasars can be explained by dust extinction alone.

I select a subsample of nearly 2500 UV spectra in the redshift range $1.90 < z < 2.35$ with luminosities in the range $46.3 < \log(L_{bol}/erg\ s^{-1}) < 48.2$, for which I measure the amount of extinction through continuum fits for two models of extinction, namely dust from the Small Magellanic Cloud (SMC) and the Large Magellanic Cloud (LMC), respectively.

Dust in quasars is usually assumed to be SMC like, but I find that the extinction in many cases can be described equally well with the LMC dust model. This has interesting implications for our understanding of quasars. I see no dependence on source luminosity or redshift for the average amount of extinction of quasars for either the SMC or LMC dust models. I do however find indications of an evolution in the type of dust affecting quasars, with LMC like dust more common at higher redshifts. This is an intriguing result that deserves a closer investigation in the future.

Dansk resumé

Historisk set har det været problematisk at give en nøjagtig beskrivelse af støvs egenskaber i kvasarer, fordi det specktrum kvasarer udsender ikke er unikt og universelt. Jeg vil undersøge rødfarvningen af kvasarer for et stort sæt af spektre hentet fra Sloan Digital Sky Survey, for at teste om, og i så fald hvordan, rødfarvningen afhænger af et objekts rødforskydning og lysstyrke. Dette er ikke tidligere blevet undersøgt. Jeg vil også undersøge om variationer i kvasarernes spektre kan tilskrives rødfarvning pga. støv.

For en udvalgt gruppe af næsten 2500 UV spektre fra rødforskydning 1,90 til 2,35 med bolometrisk luminositet i området $46.3 < \log(L_{bol}/erg\ s^{-1}) < 48.2$ måler jeg mængden af rødfarvning ved at fitte to forskellige støvmodeller til kvasarernes kontinuum emission. Nemlig støv med egenskaber som observeret i henholdsvis den Lille og den Store Magellanske Sky (LMC og SMC).

Normalt antages det at støv i kvasarer er af SMC typen, men jeg kan i mange tilfælde beskrive rødfarvningen af spektrene lige så godt med LMC modellen. Dette har interessante implikationer for vores forståelse af kvasarer. For hverken SMC eller LMC støv finder jeg at mængden af rødfarvning afhænger af rødforskydning eller kvasarens lysstyrke. Til gengæld er der indikationer af at typen af støv ændrer sig, og LMC støv ser ud til at være mere almindeligt i kvasarer ved højere rødforskydning. Dette er et spændende resultat der fortjener at blive undersøgt nærmere i fremtiden.

Acknowledgements

This thesis is written as part of a master's degree in physics with specialization in astrophysics at the Niels Bohr Institute, University of Copenhagen.

For personal reasons, this work became a much bigger challenge than I expected, and it has taken a lot longer to finish than I first anticipated. I am therefore extremely grateful for all of the help and support I have been given along the way.

I want to say thank you to my family for encouragement and support through my education, and for helping me nourish my interests in science. To my boyfriend for care and comfort in troublesome times, and for being a silly, funny and awesome guy, and as much a nerd as I am. Thanks to my friends for all of the fun times when I needed to take my thoughts off of my studies.

And last but not least, thank you Marianne for finding such an interesting topic for me to write about, for all of your advice and guidance, and for endless patience and understanding.

Contents

1 Introduction

Quasars are extremely bright objects that can be observed all the way to redshifts greater than 6, at which the age of the universe is less than 1Gyr. Quasars are located at the centers of active galaxies from where they interact with the host galaxies through their powerful radiation. The evolution of quasar and host galaxy is linked (Hopkins et al. 2007). If we can understand quasars, we may therefore gain knowledge not only of the quasars, but also of the early universe and how the galaxies have formed and evolved.

A subclass of quasars are very strong sources of radio emission, and the first quasars were discovered in the late 1950s in the first radio surveys (Peterson 1997, p. 5). But quasars are still not fully understood today. Studies of quasars are complicated because the observed properties do not always reflect their intrinsic properties. For example, many quasars are hidden behind dust that absorbs and changes the light we receive from them. The amount of dust that modifies the emitted spectrum is different from one quasar to another, and it is therefore difficult to know how much each individual spectrum is altered, and what the intrinsic emission from the quasar really is.

1.1 What are quasars?

Quasars have been studies for more than half a century, and there exists an innumerable number of papers and books on the subject. Based on Peterson (1997), Sparke & Gallagher (2007) chapter 9 and Rosswog & Brüggen (2007) chapter 8, I will here give a short description of what kind of object a quasar is.

Quasars are a very luminous type of active galactic nuclei (AGN). They are some of the brightest objects in the Universe, and can be seen across enormous distances and lookback times. This property makes them very useful in cosmology. Their population peaked around redshift $z = 2$ when the universe was $\sim$ 3Gyr old, but they can be observed through most of the history of the universe. The highest redshift of a quasar registered so far is $z = 7.085$ (Momjian et al. 2014). AGNs are galaxies with an active nucleus. Their centers are relatively small and compact. Most of the radiation from an AGN is produced in the central engine, a region with a typical size similar to our solar system. Yet the brightest quasars have an energy output equivalent to a thousand galaxies, and completely outshine their host galaxies.

We cannot take an image of the central region of AGNs to see what happens in the core. It is simply too small to be spatially resolved, even with modern technology and equipment. But quasar spectra are both very energetic and contain emission lines that cannot be produced by stars. The only thing that can continuously emit such spectra is a supermassive black hole with an accretion disk.

Figure 1 provides a schematic view of an AGN on a logarithmic length scale. At the very center is a supermassive black hole. The typical mass of the black holes ranges from 10^6 to $10^9 M_{\odot}$, similar to the supermassive black holes we see in our Milky Way and other normal galaxies today. These are believed to be the remnants of an earlier active phase in the evolution of galaxies.

The central black hole is surrounded by an accretion disk, from which it accretes matter. The gas of the disk is pulled inward by the gravity of the black hole, but angular momentum keeps the gas from falling directly towards the center. Instead it forms a disk in which the gas slowly spirals closer, before it eventually crosses the event horizon.

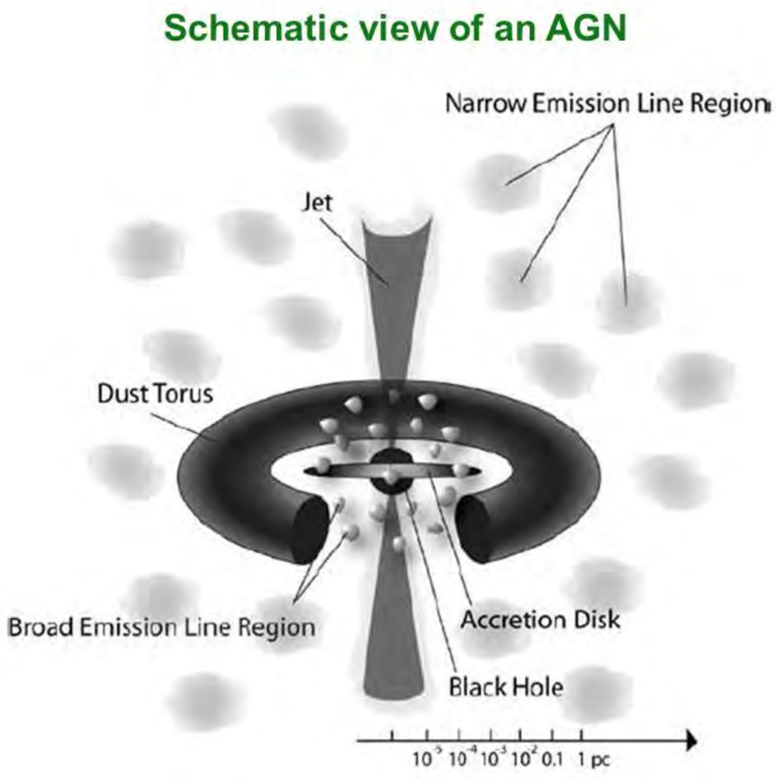

Figure 1: Schematic view of an AGN (Rosswog & Brüggen 2007, p. 301). The typical dimensions of each component are indicated on a logarithmic length scale.

Contrary to normal galaxies, quasars emit strongly over a wide range of frequencies. They produce radiation all the way from radio to γ-ray energies. The material of the accretion disk gives up potential energy as it moves closer to the black hole. The released energy is converted to an increase in velocity as the gas moves to lower orbits, and friction heats the gas to higher and higher temperatures as it gets closer to the center. The conversion of rest mass to energy is very efficient for accretion, and a supermassive black hole only needs to accrete a few solar masses of gas per year to power the most luminous of quasars.

Each annulus of the disk produces thermal radiation with a blackbody spectrum corresponding to the temperature of the gas. The blackbody spectra are combined across the disk to produce the continuum emission from optical light at the outer edge, ultraviolet (UV) light further in, and X-ray radiation from the hottest, innermost part of the disk. The continuum of the disk can be well approximated by a powerlaw function. The temperature of the disk depends on the mass and spin of the black hole and the accretion rate. The disk is hotter around smaller black holes. Though they are similar in many ways, each optical/UV spectrum is therefore uniqe, and there is no universal template for the intrinsic quasar spectrum.

Strong magnetic field lines in the disk gets pulled in with the accretion of matter, and can cause narrow jets of relativistic outflow to emerge from near the center. Electrons are accelerated to highly relativistic velocities in the jets, and they spiral around the magnetic field lines to produce synchrotron radiation, which we observe as radio emission. The relativistic electrons can also boost lower energy photons from the accretion disk to X-ray and γ-ray energies through inverse Compton scattering.

A prominent feature of optical and UV spectra of some quasars is their broad emission lines. The lines originate from dense clouds of gas near the accretion disk and black hole, in what is called the broad line region (BLR). The Doppler shift of the fast moving BLR clouds broadens the emission lines significantly, and the line widths indicate velocities between 2000 and 10.000km/s.

Beyond the accretion disk and BLR is a large torus of molecular gas and dust that lies in the plane of the disk. The dust in the torus absorbs the light it receives from the central engine, particularly optical and UV light. The dust is heated by the radiation, and the energy is re-emitted as thermal radiation in the infrared (IR).

Depending on the angle at which we see the quasar, the torus can hide the accretion disk and BLR from view. If the central engine is obscured the spectrum has no broad emission lines, and the quasar is classified as type 2 instead of type 1 in which the inner regions are visible.

The central engine is also surrounded by a region of diffuse gas clouds that produce narrow emission lines, the narrow line region (NLR). The NLR extends further out than the dust torus, and narrow emission lines are seen in spectra of both types of quasars. The clouds are further away from the black hole and move slower than in the BLR. The line widths correspond to velocities of the order 100 to 1000km/s.

The dust that affects the quasar spectra is hidden in the gas of the torus and the NLR as well as in the host galaxy. Dust evaporates at temperatures above $\gtrsim$ 2000K and cannot exist within the sublimation radius at which this occurs. For this reason there is no dust in the BLR and accretion disk where the temperature is simply too high for dust to survive (MacAlpine 1985, p. 274-278).

In most quasars the dust is located at the redshift of the quasar, implying that the dust is intrinsic to the object (Hopkins et al. 2004). But sometimes there can be absorption from intervening dust in the spectra as well, from e.g. galaxies that lie along our line of sight to the quasar. There can also be dust in the intergalactic medium between us and the quasar. However, in this thesis I am only concerned with dust associated with the quasar itself.

1.2 Dust and extinction models

Extinction is the combined effect of absorption and scattering of light caused by dust particles. Dust grains absorb and scatter light most efficiently at wavelengths comparable to their size (Li 2007). Small dust grains therefore preferentially absorb the light at short wavelengths (e.g. UV light), and larger grains absorb at longer wavelegths (e.g. optical light and NIR). The energy absorbed by the dust is re-emitted as thermal radiation in the IR.

Extinction is generally stronger in blue light and UV than at long wavelengths. There are more of the small dust grains which absorb at the short wavelengths, than of the larger dust grains. This causes a reddening of the affected spectra since primarily blue light is removed. Quasar spectra subject to strong extinction bend downward at short wavelengths in an 'n'-shape as a result of the stronger extinction at shorter wavelengths (Francis et al. 2000).

The only information we receive from a quasar is the spectrum that we observe.

But the observed spectrum and the emitted spectrum are not the same if the light is altered by dust extinction. Dust is therefore important in order to correct especially the UV/optical emission which probes the gas of the accretion disk, and in our understanding of the structure and geometry of quasars (Li 2007).

Known extinction models

The extinction as function of wavelength is described by extinction curves. Figure 10 shows three examples hereof. Only extinction curves for sight lines in the Milky Way and the nearest galaxies have been determined accurately. The shape of a dust extinction curve depends on the chemical composition and size distribution of the dust grains (Pei 1992). Interstellar dust in the Milky Way consists mainly of silicate and graphite dust grains. The extinction curve increases towards shorter wavelengths and has a bump centered at 2175Å caused by small graphite dust particles in the form of PAHs (Cardelli et al. 1988, 1989; Li 2007).

The average Milky Way extinction curve is well described by the CCM model (Cardelli et al. 1988, 1989) which is parameterized by the total-to-selective extinction $R_V = A(V)/E(B-V)$, where $A(V)$ is the exinction of the V filter, and $E(B-V) = A(B) - A(V)$ is the color excess between the B and V filters.

R_V is a rough measure of the average dust grain size, and different sight lines in the Milky Way can have different R_V values. Different processes influence the size of dust grains. Small dust grains can stick together (coagulate) to form larger dust grains, and hard radiation like extreme UV or X-rays can destroy the grains and make them smaller (Cardelli et al. 1989; Hopkins et al. 2004; Li 2007). For Milky Way extinction, small values of R_V (i.e. small dust grains) have a strong 2175Å bump, and the UV rise of the extinction is steep. For larger values (i.e. larger dust grains) the 2175Å bump is small and the UV rise is flatter (Cardelli et al. 1988, 1989; Li 2007).

Dust in the nearby dwarf galaxies, the Small Magellanic Cloud (SMC) and the Large Magellanic Cloud (LMC), is different from dust in the Milky Way (Gordon et al. 2003; Pei 1992). The SMC extintion curve has no 2175Å bump, and it is much steeper than the Milky Way curve, which means the dust grains are comparatively smaller. The LMC curve is intermediate between the two. It is steeper in UV than Milky Way extinction, but less so than the SMC extinction curve. The 2175Å bump is also present in LMC extinction, though weaker than in the Milky Way.

Together, the Milky Way, LMC and SMC extinction curves form a sequence that can be well described by a single dust model of graphite and silicate dust by adjusting the relative abundances of these elements (Gordon et al. 2003; Pei 1992). Dust in the Milky Way consists of roughly equal amounts of graphite and silicate, whereas the SMC dust is nearly pure silicate grains. The composition of LMC dust is somewhere in between. The Milky Way, LMC and SMC then form a series with diminishing strength of the 2175Å bump, and increasing strength and steepness of the far-UV extinction.

Measuring extinction curves

The Milky Way, LMC and SMC extinction curves are measured using the pair-method (Cardelli et al. 1988, 1989; Gordon et al. 2003; Li 2007; Pei 1992), which extracts an extinction curve by comparing reddened and unreddened stars of same spectral type, i.e. with the same intrinsic spectrum. This method works well for sightlines in the Milky Way and the Magellanic Clouds, where single stars can be observed and the spectral type identified. The intrinic spectrum is well determined for stars of a given spectral type and does not vary significantly. This provides a good reference point for measuring the amount of extinction.

Extinction curves of quasars are more difficult to obtain. The pair-method does not apply for quasars because their distance is too great for observations of single stars in their galaxies. We can only measure the spectrum of the whole quasar and host galaxy together, except for the very nearest of AGN in which the largest NLRs may be resolved. But even for these, stellar spectra are still unavailable.

Attempts of measuring quasar extinction curves have been done by comparing single reddened quasar spectra to unreddened objects representing the intrinsic quasar spectrum (Crenshaw et al. 2001). But one must be careful with such comparisons, since the intrinsic SED of quasars, contrary to stars, can vary from object to object. The resulting extinction curves are very steep in the UV and have no 2175Å bump. Studies of colors of reddened and unreddened quasars in large quasar surveys such as the Sloan Digital Sky Survey (SDSS) also support the concept of SMC like dust (Hopkins et al. 2004), and it is widely accepted that quasar extinction curves are very similar to the SMC curve.

Others have measured extinction curves of quasars by comparing composite spectra (Czerny et al. 2004; Gaskel et al. 2004). Contrary to extinction curves derived from single quasars, these curves are very flat in the UV. This is most likely a result of bias from the composites. The high redshift quasars, which contribute to the UV part of the composite, are on average less reddened than nearer objects, so the extinction becomes relatively smaller at UV wavelengths (Willott 2005).

1.3 Motivation and goal of my thesis

It has so far been difficult to account for the properties of dust in quasars because their intrinsic continuum emission is not unique.

The goal of my thesis is to investigate the extinction curves of quasars to test if and how they change as a function of redshift and source luminosity. A confirmation or refutation of variations in the extinction will be helpful to future studies of quasars.

I will also test if there is some connection between luminosity, redshift and the intrinsic spectrum. That is, for a given luminosity and redshift, can the diversity of the observed quasars be explained by dust extinction alone, so that the intrinsic spectra of these objects are the same? Other studies have shown that quasar spectra are more alike if they are selected for a limited range of luminosities only (Krawczyk et al. 2013). So throughout my thesis it will be a fundamental assumption that quasars have identical intrinsic spectra for a small range in redshift and luminosity.

To achieve these goals I need a large sample of quasars over a wide range of luminosities and redshifts. This is provided by the Sloan Digital Sky Survey which contains more than 100.000 spectroscopically confirmed quasars.

2 Data analysis

2.1 Outline of analysis method

The goal of this thesis is to investigate if and how extinction curves of quasars change with redshift and luminosity of the source. For ease of understanding, I start with a quick description of the strategy behind the analysis method. Each step will be explained in more detail in the following sections.

A large number of high quality spectra of quasars are required. These are provided by one of the largest quasar catalogs available, the Sloan Digital Sky Survey[1] (SDSS). Find more details about the survey in section 2.2. The SDSS quasar catalog contains a huge number of objects, and I can divide the quasars into a mesh of small bins where all objects have almost the same redshift and luminosity. For each bin I then assume that the quasars it contains all have the same intrinsic spectral energy distribution (SED), and any variations in the continuum of the spectra from object to object are caused by dust extinction alone.

By fitting a powerlaw function to the continuum of each spectrum I identify the bluest (i.e. steepest) spectrum of each bin. I then assume this spectrum is the most unaffected by dust extinction, and let the powerlaw continuum fit of this spectrum represent the intrinsic continuum of each spectrum in that same bin for a given redshift and luminosity. To determine the amount of dust that reddens each of the other spectra in a given bin, I make a new continuum fit with a combination of the blue unreddened powerlaw and a dust model, where the only free parameter is the amount of dust. I investigate two different models of dust extinction; the Small Magellanic Cloud model and the Large Magellanic Cloud model. These are described in section 2.7.

The reddening of an object can change its luminosity significantly, and some quasars may be extinguished enough, that the intrinsic continuum belongs in a more luminous bin. To account for this, each spectrum is fitted using the bluest powerlaw of each luminosity bin in the same redshift range, to see which of the bluest powerlaws give the best fit to the shape of the continuum after modeling the effects of dust. The fitted quasar is then moved to the bin of the most appropriate blue object.

Finally, I can compare the amount of reddening measured in the spectra for bins across both redshift and luminosity. I also compare the different impacts of the two dust models.

All programming involved in this work is done by me unless otherwise mentioned.

2.2 SDSS DR7

My thesis is based on data from the Sloan Digital Sky Survey (SDSS) Data Release 7 (DR7)[2]. SDSS DR7 is a catalog containing 105,783 spectroscopically confirmed quasars with great variety in redshift and luminosity. I use a subset of quasars consisting of nearly 2500 objects, which I select after criteria described in detail in section 2.4.

The SDSS identify quasars using imaging data taken in five broad optical bands (*ugriz*) with a CCD. They select quasar candidates on basis of their position in color space. For DR7 the spectra of the selected quasars cover the wavelength range 3800Å $<$

[1] www.sdss.org
[2] http://classic.sdss.org/dr7/

$\lambda < 9200$Å at a spectral resolution of $R \simeq 2000$. The spectral resolution R is defined as $R \equiv \lambda/\Delta\lambda$, where $\Delta\lambda$ is the smallest difference in wavelength that can be distinguished at the wavelength λ. The quasars have a luminosity of at least $M_i = -22.0$mag, and the redshifts lie in the range $0.065 < z < 5.46$ with a typical redshift error of $z_{err} \approx 0.004$. The catalog covers an area of approximately 9380deg^2 on the sky equivalent to 23% of the entire sky. The sky covered is located mainly in two areas. The largest area is centered on the North Galactic Pole, and the other is located near the Celestial Equator in three narrow strips (Schneider et al. 2010).

2.3 Extraction and preparation of spectra

The spectra in the SDSS database are given in vacuum wavelengths and are not corrected for Galactic extinction. Therefore, before I can use any of the spectra, I need to change the format to restframe wavelengths, so I can compare the spectra directly with each other. I also need to correct for Galactic extinction to get flux values that are not affected by the dust of the Milky Way, which would otherwise interfere with my measurements of intrinsic dust in the quasars.

Plate, fiber and mjd uniquely identifies the spectra. With these informations I can use the IDL routine `readspec`[3] to load a specific spectrum from the SDSS DR7 dataset in the form of two arrays. One array contains the vacuum wavelength values of the spectrum and the other contains flux density values.

To correct for Galactic extinction I use an IDL routine `qso_dered` provided by my supervisor. It dereddens the spectrum according to the amount of Galactic extinction $E(B-V)$ in the line of sight towards an object using the O'Donnell Milky Way extinction model (O'Donnell 1994). This IDL program also shifts the spectrum to rest frame wavelengths using the cataloged redshift of the quasar as the systemic redshift.

The DR7 dataset includes error and mask arrays of the spectra. The error array gives the uncertainty of each measured flux value in the spectrum. The mask array contains information on pixels that SDSS find to be unreliable measurements. These include for instance contaminated pixels and areas where the sky background is badly modelled or too strong compared to the measured source. Both the error array and the mask array are outputs of `readspec`. I wrote a program `mask_out` which uses the mask array to remove pixels of bright sky background (specified by mask bit 23 in the array)[4] and bad sky fit χ^2 (mask bits 27 and 28). I use the error array to weight the flux data points when I fit the spectra as described in section 2.6.

2.4 Sample selection

The SDSS quasar catalog contains a huge number of objects, and it would be far too time consuming to process all of their spectra. Instead, I have focused on a smaller sample of the full dataset. The following describes how and why I selected the specific subset of quasars and their spectra.

The distribution of SDSS quasars in bolometric luminosity and redshift is shown in figure 2. Each dot represents a quasar. The blue dots show my final sample, and

[3]The IDL routine readspec.pro was downloaded from http://spectro.princeton.edu/idlspec2d_doc.html#READSPEC
[4]For a detailed list of mask bits see http://classic.sdss.org/dr7/dm/flatFiles/spPlate.html

this sction explains how and why these objects were selected. In figure 2 only the more luminous objects are visible at high redshifts. This is because of the sensitivity of the instruments, and because the less luminous objects are too faint to detect. At small redshifts there are only few high luminosity objects because the sampled volume of the universe is much smaller and high luminosity objects are rare. But because of the smaller distance at low redshifts, we see a lot more faint objects. At higher redshifts around $z = 2.7$ and $z = 3.5$ there are areas with only few quasars detected. The colors of quasars at these redshifts are similar to those of stars in our galaxy, and the algorithm used by SDSS to identify quasar candidates can therefore not easily distinguish between stars and quasars (Schneider et al. 2010). Many quasars are misclassified as stars and are therefore overlooked by the survey.

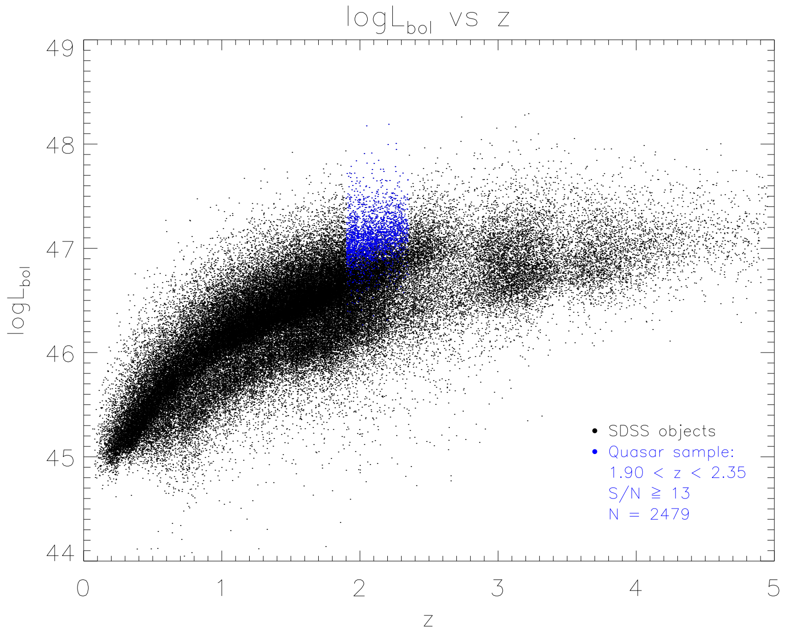

Figure 2: Distribution of SDSS objects in bolometric luminosity and redshift. Blue dots mark the 2479 quasars within the redshift and S/N range that I have selected for my sample.

The quasars I investigate only cover a relatively short redshift range of the SDSS dataset. I use quasars for which $1.90 < z < 2.35$. I choose the lower limit $z = 1.90$ for two different reasons. For objects with $z < 0.8$ light from the host galaxy contributes to the spectra in the infrared. This is an extra source of light that I would have to take into account when I model the continuum of the spectrum, and it would make my work a lot more complicated. For this reason I have excluded these objects from my sample. Most of the objects meassured by the SDSS lie in the range $1 \lesssim z \lesssim 2$ with a median redshift value of $z = 1.49$ (Schneider et al. 2010). Unfortunately, the spectra of these objects are dominated or strongly influenced by the iron bump feature around 3000Å, which makes it difficult to estimate the underlying continuum powerlaw slope. For this reason, these objects are also removed from the sample. I therefore choose the lower limit $z = 1.90$ because I evaluate this to be the lowest redshift at which neither the host galaxy nor the iron bump complicate my measurements of the continuum.

I choose the upper limit $z = 2.35$ in order not to get too close to the area around $z = 2.7$ where the completeness of the catalog is reduced due to the colors of the quasars and the selection algorithm used by the SDSS. I need a high density of quasars in terms of redshift and luminosity for this study, and at redshifts higher than $z = 2.35$ the quasars are too few and too thinly spread out. I also choose the redshift range $1.90 < z < 2.35$ because it has a relatively flat distribution of quasars in luminosity over redshift, which makes it easier to compare the results across redshift.

To get the best fits to the continua I want the spectra to have easily distinguishable features, and the continuum and emission or absorption lines should be easy to tell apart. To ensure this I need spectra with a high signal to noise ratio (S/N). The S/N is simply a number given by the signal (flux) of a spectrum divided by the noise of the spectrum. When S/N is low, there is a lot of noise in the spectra. But when S/N is high there is only a small amount of noise compared to the strength of the signal, and smaller details of the spectra emerge.

Figure 3 shows the same data as in figure 2, but colored so that the color of each object represents the median S/N value of its spectrum, (see the legend of figure 3). The figure shows that for a given redshift, the spectra with the highest median S/N value (e.g the blue dots; S/N $\geqq 20$) are also typically the most luminous quasars. More light is received from the most luminous objects, and they therefore have a stronger signal, while the noise from the instruments remains the same.

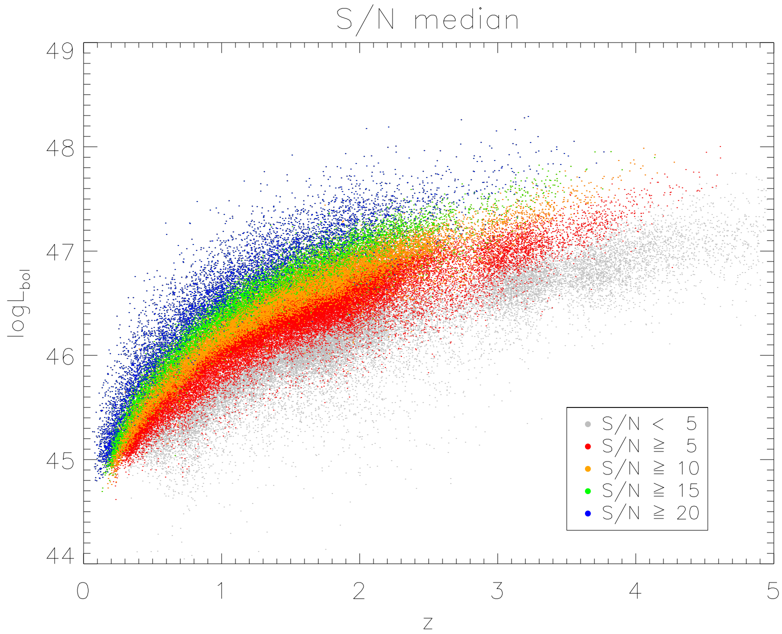

Figure 3: Distribution of median Signal to Noise ratio values for the spectra of SDSS quasars. Gray dots represent objects that have a median S/N value less than 5. For red dots S/N $\geqq 5$, for orange dots S/N $\geqq 10$, for green dots S/N $\geqq 15$ and for the blue dots S/N $\geqq 20$. Luminous quasars have the highest S/N at a given redshift.

I impose a lower limit on S/N to ensure a minimum quality of the spectra in my sample. According to the distribution of S/N values in figure 3, I primarily lose low luminosity objects in this process.

To decide on a value for the lower limit of S/N, I visually inspect spectra with different S/N median values. Figure 4 shows four spectra with a redshift close to 2, and

S/N median values close to 5, 10, 15 and 20 respectively, from the bottom and up. The spectra are plotted on log-log axes and using the keyword `nsum` set to 3, which averages the flux over every 3 data points and smooths out the spectra a little for clarity. From visual inspection of the spectra in the example of figure 4 and other plots like it, I decide the lower limit of S/N should lie between 10 and 15 for me to be able to make acceptable continuum fits. I then repeat this process looking at spectra with S/N values close to 10, 11, 12, 13, 14 and 15. For steps of only one in S/N the spectra look very alike in terms of noise level, so here the resulting number of quasars in my sample also influece my choice of S/N limit.

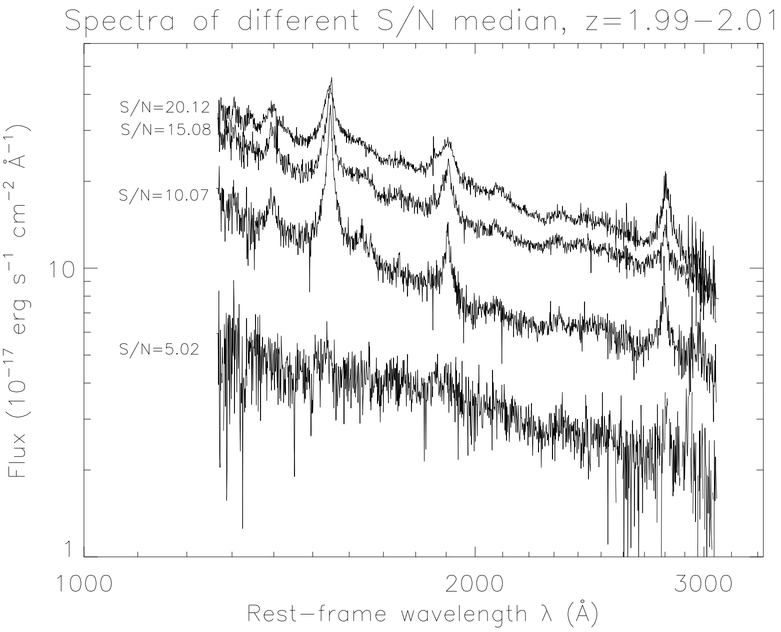

Figure 4: Four spectra at a redshift $z \approx 2$, representing spectra with a S/N value of 5, 10, 15 and 20. There is a lot of noise in the spectum with S/N $\approx$ 5. For spectra with higher S/N, smaller and smaller features of the spectrum are distinguishable. The spectra are plotted using the keyword `nsum=3` which averages over every three data points for clarity.

The redshift range inspected is $1.90 < z < 2.35$. Figure 5 shows a histogram of median S/N values for this redshift range. The majority of spectra have median S/N values below 10, but there is a tail of high quality spectra with larger S/N values and low noise. The total number of objects within the redshift range is 14763. I choose S/N = 13 as the lower limit. When I ignore all objects with S/N smaller than 13 (illustrated by the red dashed line), I am left with a total of 2479 objects in my sample, equivalent to 17% of the available quasars.

The blue dots in figure 2 illustrate my sample, and show the distribution of quasars in redshift and luminosity within the boundaries set. In summary, I select my sample of nearly 2500 quasars for this study to have redshift in the range $1.90 < z < 2.35$ and with S/N values of their spectra larger than 13.

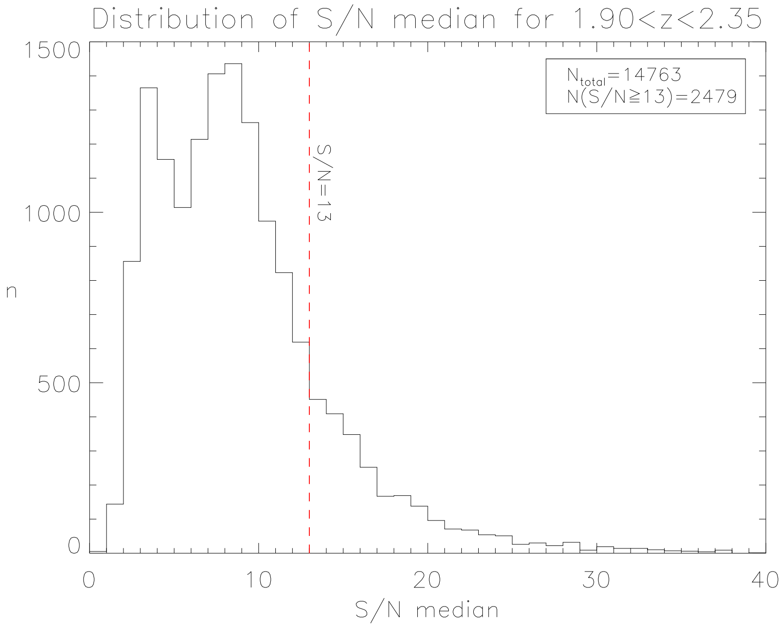

Figure 5: Distribution of median S/N values of quasar spectra in the redshift range $1.90 < z < 2.35$. The total number of quasars is 14763. I impose a limit S/N $\geqq$ 13, illustrated by the red dashed line, which leaves me a sample of 2479 quasars.

2.5 Defining the (z,L) bins

Redshift bins

I divide the quasars of my sample into smaller bins of redshift and luminosity to put them in smaller groups where I can assume these properties to be the same for all bin members. This means I can assume their intrinsic luminosities and spectral energy distributions, or SEDs, to be identical.

The bins cover the redshift range $1.90 < z < 2.35$ which is equivalent to a time span[5] of 645 Myr assuming the cosmological parameters $\Omega_m = 0.30$ and $H_0 = 70$km/s/Mpc. The width of each individual bin is set to $\Delta z = 0.05$. This is about ten times the typical error of the redshifts (Schneider et al. 2010, see also section 2.2), and it is therefore fairly certain that the quasars in a specific redshift bin do not belong in a neighbouring bin due to measurement errors on z. The bins are still small enough that there is no significant cosmological evolution (approximately 70 Myr) from one end of a bin to the other.

Different measures of luminosity

Shen et al. (2011) provide several spectral measurements of luminosity for the SDSS sources. These include the bolometric luminosity and a monochromatic luminosity at either 1350Å, 3000Å or 5100Å depending on the redshift of the source. The bolometric luminosity is computed from the monochromatic luminosity L_{5100} for $z < 0.7$, L_{3000} for $0.7 \leq z < 1.9$ and L_{1350} for $z \geq 1.9$, using the bolometric corrections $\mathrm{BC}_{5100} = 9.26$, $\mathrm{BC}_{3000} = 5.15$ and $\mathrm{BC}_{1350} = 3.81$ respectively (Shen et al. 2011).

[5]Using the Cosmological Calculator for a Flat Universe. http://home.fnal.gov/~gnedin/cc/

In figure 6 I compare bins divided by bolometric luminosity and by a monochromatic luminosity. The left panel shows seven spectra of quasars in a bin at $2.20 < z < 2.25$ covering a range of 0.1dex in $\log L_{bol}$. The spectra here are very different and show great variations in flux level and continuum slope, and it seems unreasonable to assume that the SEDs of these objects are intrinsically identical. The right panel shows spectra of a bin with six quasars at the same redshift, but over a range of 0.1dex in monochromatic luminosity, $\log L_{2600}$. I explain later how L_{2600} is calculated. The spectra of the right panel are more alike in terms of slope and flux level. The reason I divide the quasars into small bins over redshift and luminosity is to put intrinsically similar objects together. The bin content defined by monochromatic luminosity is better in this regard, so this is the method I will use for dividing bins of luminosity. There is also one other important advantage of this method that I will explain next.

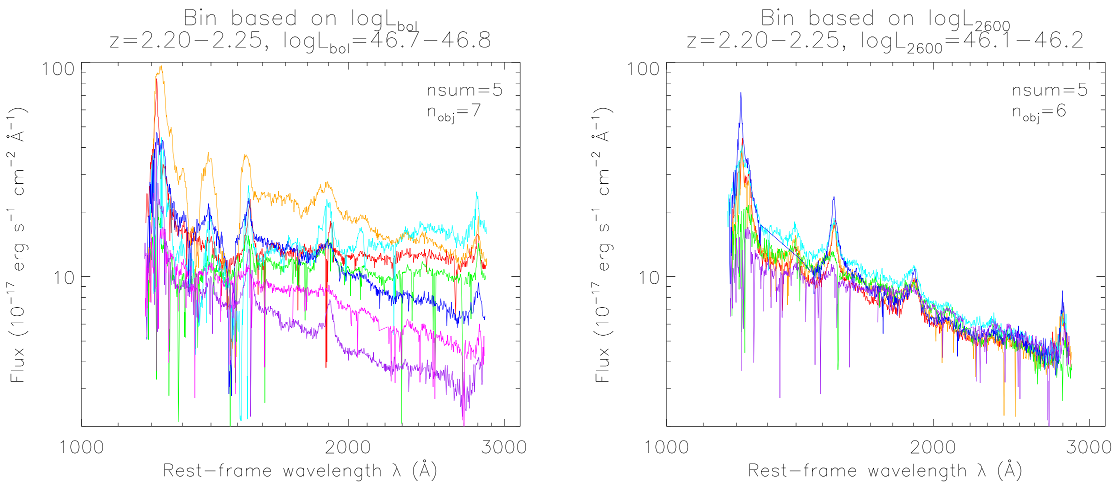

Figure 6: *Left panel:* A bin based on bolometric luminosity L_{bol} containing 7 quasars. *Right panel:* A bin based on monochromatic luminosity L_{2600} calculated at 2600Å containing 6 quasars.

For an observed wavelength λ_{obs} the corresponding restframe wavelength λ_{rest} is given by

$$\lambda_{rest} = \frac{\lambda_{obs}}{1+z}. \tag{1}$$

The spectra of SDSS are observed over the wavelength range 3800Å $< \lambda <$ 9200Å. For a spectrum at $z = 1.90$ this corresponds to the restframe wavelength range

$$\frac{3800\text{Å}}{1+1.90} = 1310\text{Å} \qquad \text{to} \qquad \frac{9200\text{Å}}{1+1.90} = 3172\text{Å},$$

and the range 1134Å $< \lambda <$ 2746Å for a spectrum at $z = 2.35$. My entire sample therefore falls into the $z \geq 1.9$ category for which the monochromatic luminosity at 1350Å is available (Shen et al. 2011). The wavelength range just calculated puts 1350Å in the blue end (i.e. at short wavelengths) of all the spectra in my sample. But a theoretical bin of quasars that are all intrinsically identical, but affected by different amounts of dust (assumed to be SMC like), would have the most similar flux values in the red end of the spectra, where dust has the smallest effect. A monochromatic luminosity at the long wavelength end of the spectra is therefore better for collecting similar spectra, when some of them are affected by dust. Also, at the redshifts of my sample L_{1350} is the luminosity used to calculate the bolometric luminosity (Shen et al. 2011), and as illustrated in figure 6 this is not the most useful quantity with respect to defining bins of intrinsically similar objects.

Instead I define my luminosity bins from the monochromatic luminosity L_{2600}. I calculate L_{2600} from the flux density F_λ in the red end of the spectra at a small range (±50Å) around $\lambda = 2600$Å. I take the median flux density of this range to even out any effects of photon noise in the spectra. L_{2600} reflects the flux in the red end of the spectra which is less affected by dust, and therefore provides a better representation of intrinsic properties than L_{1350} and L_{bol}.

The median flux density at 2600Å, F_{2600}, is affected by the broad 3000Å iron bump feature which may not be exactly the same in all quasars. But I do not expect the potential object to object variation of the UV FeII bump to affect my study for several reasons.

For one, strong UV FeII emission is expected in many quasars (Netzer & Wills 1993) as is also commonly observed. For example, the right panel of figure 6 illustrates that the strength of the UV FeII pseudocontinuum (this can be gauged as the intensity difference between the flux levels at 1950Å to 2300Å and at 2500Å to 2600Å) is rather similar between quasars at similar redshift and luminosity. If this was not the case, there would be a much larger dispersion of the flux levels at restframe $\sim$ 2100Å for these spectra that are normalized at restframe 2600Å.

Furthermore, the UV FeII emission is essentially insensitive to the element abundance in the broad line region compared to the optical FeII emission (Verner et al. 2003) which is also known to vary significantly from object to object, i.e. UV FeII emission varies much less than optical FeII emission. The UV FeII emission appears to vary by about 20% (Maoz et al. 1993) in a typical nearby AGN. Distant quasars vary with smaller amplitudes and on longer timescales because of time dilation (Kaspi et al. 2007; Watson et al. 2011), and therefore the UV FeII emission is not expected to vary much in time or from object to object for the quasars I study.

I also reorganize the quasars across luminosity bins independent of L_{2600} later. The small potential object to object Fe flux variations will therefore not have any impact on the later distribution of quasars across luminosity bins, and I therefore do not correct for this potential object to object Fe flux variation.

The luminosity distance d_L relates the luminosity L and flux F of an object by the following equation (Ryden 2003, p. 107)

$$d_L \equiv \left(\frac{L}{4\pi F}\right)^{1/2}. \tag{2}$$

I use the IDL program `lumdist`[6] to calculate d_L. It calculates the luminosity distance for an input redshift value, for a given cosmology. I use the redshifts given by SDSS and the default cosmological parameters of the program (flat universe with $\Omega_m = 0.30$ and $H_0 = 70$km/s/Mpc). I can then calculate L_{2600} by isolating L in eq. (2):

$$L_{2600} = 4\pi d_L^2 F_{2600}. \tag{3}$$

[6]Found at http://idlastro.gsfc.nasa.gov/ftp/pro/astro/lumdist.pro

I use the standard deviation of the flux density, δF, at 2600Å ±50Å to estimate the errors of L_{2600}. I calculate upper and lower limits of L_{2600} from $F_{2600} + \delta F$ and $F_{2600} - \delta F$ respectively, and take their average deviation from L_{2600} to be the uncertainty δL_{2600}. The average error on $\log L_{2600}$ of all quasars in my sample is then 0.039dex. The luminosity bins should be at least this size to have physical meaning.

Luminosity bins

I consider two methods for distributing the quasars into bins of luminosity. One option is to make bins of the same size in luminosity $\Delta \log L$. Another is to make bins with different spans in luminosity, but adjusted so that each bin contains the same number of objects. So the choice is between bins of equal luminosity range and bins of equal numbers. Both methods are illustrated in figure 7.

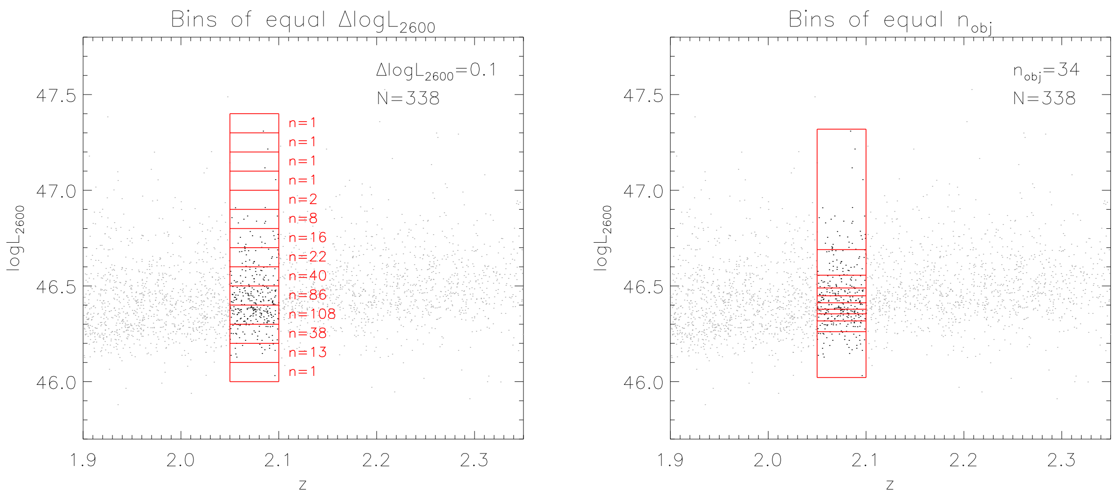

Figure 7: *Left panel:* Example of bins distributed according to fixed $\Delta \log L_{2600}$. Each dot represent an object from the sample. The red number next to each bin indicates the number of objects in that bin. *Right panel:* Example of bins distributed so that each bin contains the same number of objects. In this case 10 bins with $n_{obj} = 34$ objects in each, except the most luminous bin which contains only 32 objects.

The first method is shown in the left panel of figure 7 for a single redshift range. The bins are all of the same width in luminosity, $\Delta \log L = 0.1$dex, and the numbers to the right of each bin indicate the number of quasars included in that bin. This method provides bins that are easy to compare across redshift, because the bins at higher and lower redshifts will cover the same ranges of luminosity. The main complication with this method is that it puts most of the objects in the central luminosity bins and leave only a few in the outlying high and low luminosity bins. Several bins contain only one object, which will then automatically represent the intrinsic spectrum as the bluest spectrum of that bin, whether or not there is significant reddening of the spectrum. For a bin with many objects it is less likely for all of the quasars to be affected by dust, and the actual blueness of the bluest object is more reliable and genuine. But for a bin with only few quasars, my assumption that the bluest object is unaffected by dust is more easily broken. Also, some bins might be empty if I use this method, in which case that luminosity and redshift will not be represented by any spectrum.

The second method is illustrated in the right panel of figure 7 for a single redshift range. In this example there are 10 bins with 34 objects in each, except for the topmost bin which contains only 32 quasars. This method ensures that all bins are equally likely to contain an unreddened object. And unless rearranging the quasars between luminosity bins according to their dust fits changes the number of members significantly, there will be roughly the same number of quasars in each bin after the fits when I am ready to look at the statistics of my results. None of the bins are empty or wasted because there are too few quasars. This method is not without problems either however. For a given redshift range the luminosity distribution is concentrated towards the middle of the represented luminosity values. As a result, the middle bins are very narrow, and some are smaller than the average error of $\log L_{2600}$. The higher and lower luminosity bins on the other hand have to cover a large range of luminosities to include the same number of objects. The blue spectra found in the middle bins then represent almost the same intrinsic spectrum, while the blue spectra of the higher and lower luminosity bins represent a larger range of intrinsic luminosities. My assumption that spectra of the same bin have the same intrinsic SED is then more likely to break down at the outmost luminosity bins. Also, as the bins will vary in size for different redshifts, the bins will not be directly comparable across redshift.

I decide it is better to have fewer bins with many objects, but which I can compare easily across redshift, than to make sure there is a good chance of a reasonable blue object for all bins. It is easier to ignore bins with few objects or odd looking blue spectra than to compare bins across redshift that are not aligned in luminosity. As a result, the final distribution of bins follows the first method described above. The bins are illustrated in figure 8. I choose the width of the bins $\Delta\log L_{2600} = 0.1$dex so that the majority of objects are divided between 4-5 bins. This makes it possible to compare the results across luminosity for several bins, and $\Delta\log L_{2600}$ is sufficiently large compared to the errors I calculated ($\delta\log L_{2600} = 0.039$dex), that the luminosities are well represented by the values of the bins. Also, when $\Delta\log L$ is constant, whatever small variations there may be of the SEDs within a bin are of the same order for all of the bins.

As mentioned, the biggest drawback of bins of equal size is that some of the bins are more or less empty. To compensate for this I will ignore bins containing less than a specified number of objects, after the spectra have been fitted and the quasars rearranged, when I compare the bins. The number of the cut off depends on what I want to illustrate and whether I consider a single redshift or luminosity range or the whole sample, and therefore how many available objects I have in total. The specific value of the cut will be mentioned when used.

2.6 Powerlaw fits

With the bins in place I can now write the program I need to make powerlaw fits to the continuum of each spectrum. From the powerlaw fits I will identify the bluest spectrum of each bin. To fit the continuum I need to specify which parts of the spectrum contain no significant absorption or emission lines, but only continuum emission. I have identified these continuum windows by visual inspection of a handfull of high S/N spectra, and with inspiration from Natali et al. (1998); Vanden Berk et al. (2001) and Vestergaard & Wilkes (2001). I have found five continuum-windows in the relevant wavelength range. These are listed in table 1.

These are the continuum windows I use for all of the fits, both the pure powerlaw

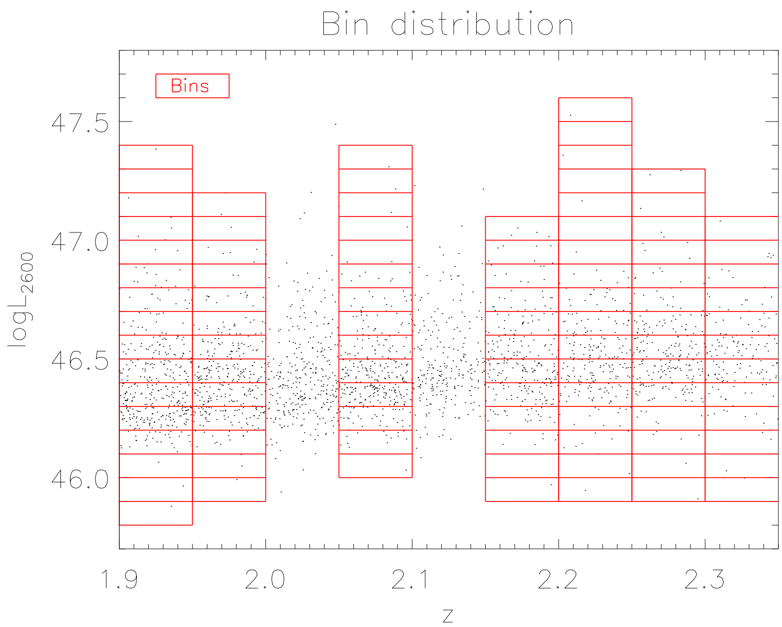

Figure 8: Distribution of all the bins used. Dots show the distribution of objects in the sample. A total of 98 bins are lined out. Two redshift ranges have not been divided into bins because I did not have enough time to include them in my studies.

Table 1: Continuum windows.

1350Å - 1370Å
1450Å - 1490Å
1690Å - 1710Å
2020Å - 2050Å
2170Å - 2250Å

continuum fits and the dust model fits. There are not always flux measurements available in all of the windows, in which case the empty windows are excluded and do not contribute to the fits.

The intrinsic continuum emission of quasars at UV wavelengths is powerlaw shaped. I assume that all objects in a single bin have the same intrinsic continuum emission (section 2.1). The spectrum least affected by dust extinction must then also be the spectrum with the steepest and bluest powerlaw continuum.

To identify the bluest and least dust affected spectrum in each bin, I fit a powerlaw function to all of the spectra. The powerlaw function is

$$F_\lambda = b\lambda^{-a}, \tag{4}$$

where F_λ is the flux density, λ is the wavelength, b is a normalization constant and a is the slope of the powerlaw. A larger value of a means a steeper slope, and thereby a bluer continuum.

To perform the fits I wrote a program in IDL. The spectra are loaded one at the time (using `readspec`), corrected for Galactic extinction (with `qso_dered`) and masked for bad pixels (`mask_out`). I then select the flux data points within the continuum windows of table 1, and do a least squares fit with a powerlaw function to the flux densities in these continuum windows. The data points are weighted by $1/\sigma_{flux}^2$ where σ_{flux} is the error of the flux density given in the error spectrum. The errors of the flux densities, or standard deviation σ_{flux}, is simply the square root of the variance of the flux values, or equivalently: $variance = \sigma^2$.

After an initial continuum fit I "sigma clip" the flux values to minimize the effects of narrow absorption lines, hot and cold pixels, cosmic rays etc., that might pull the fit away from the continuum. I have written the sigma clipping procedure `sigmaclip` inspired by a program provided by my supervisor. The program calculates the standard deviation of data points around the fit and then identifies data points more than a specified number of standard deviations away from the fit. It identifies and removes data points more than the specified number of standard deviations from the fit, and then recalculates the standard deviation for the remaining points. This process repeats until the number of remaining points converge or until a maximum specified number of iterations is reached. Typically 1 to 6 iterations are needed.

I exclude data points that are more than 2.5 standard deviations below the fit and run a new fit to the remaining points. I only sigma clip below the fit (i.e. negative deviating flux points) since narrow absorption lines are the main problem. The new fit is then used in another sigma clipping to the original set of data points (to include any earlier excluded data points that may be valid to the new fit). I repeat sigma clipping and fitting of the spectrum until the χ^2 of the fit converges, or the loop reaches a maximum specified number of iterations. The maximum number of iterations is set to 20, but typically no more than 2 to 5 iterations are needed.

After I fit all of the spectra with a powerlaw function to the continuum, I identify the bluest spectrum of each bin as the spectrum with the largest fitted powerlaw slope, a.

Figure 9 shows two examples of powerlaw fits. The left panel shows a spectrum with flux in all five continuum windows, and only a few narrow absorption lines. The powerlaw fit matches well with the continuum in this case, and the sigma-clip based fit (blue line) is not much different from the initial fit (gray dashed line).

The right panel of figure 9 shows an example of a more difficult fit. In this spectrum the flux is missing from approximately 1850Å to 2100Å, including one of the continuum windows, and there are many absorption lines around the two continuum windows at the shortest wavelengths. The powerlaw fit in which I use sigma clipping is clearly an improvement to the initial fit, as it ignores a lot of the data points in the absorption lines, but it is still not a very good fit, because it overestimates the flux at the longest wavelengths. However, this is due to the curved nature of this spectrum which clearly is affected by dust. The dust absorbs more light at shorter wavelengths and gives the spectrum a slightly arched appearance. Hence, the flatter powerlaw spectrum does in this case indicate the presence of dust in this quasar system, which is exactly what I hope to quantify in this study.

The curved nature of a dust affected spectrum makes it difficult to make a good powerlaw fit to the continuum. Fortunately, this is not important because I only use these fits to identify the bluest spectrum, and only use this one powerlaw fit of each bin in the following analysis. If the bluest spectrum is unreddened, as it should be, the powerlaw will be at good model for the continuum. I have looked through all of the bluest spectra and their continuum fits to identify any bad fits. In total, I identified

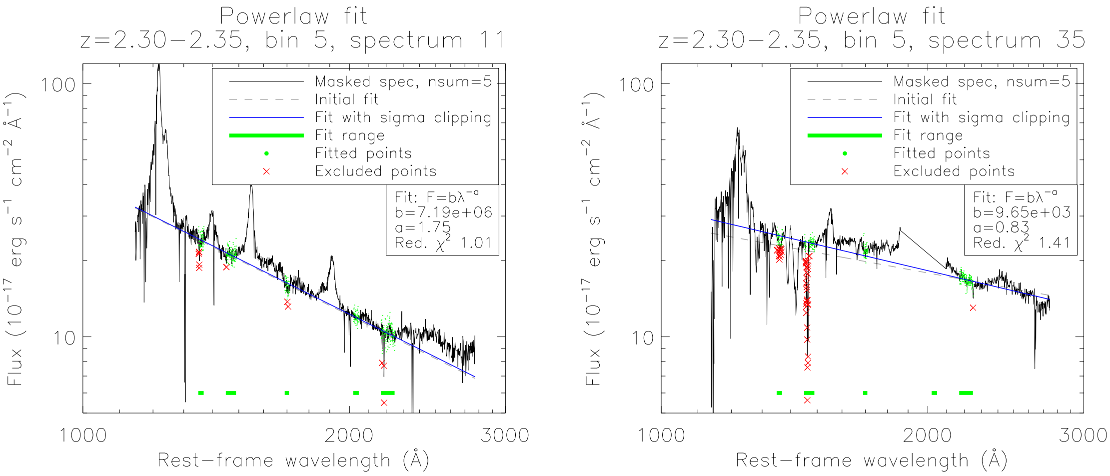

Figure 9: *Left panel:* Example of a good powerlaw continuum fit to an uncomplicated spectrum. The gray dashed line (difficult to see in this panel) shows the initial fit without sigma clipping, the blue line shows the fit after sigma clipping. The green areas at the bottom indicate the continuum windows used in the fit. Green dots are data points included in the fit and red x's are data points excluded by sigma clipping. *Right panel:* A fit to a complicated spectrum with many absorption lines and part of the spectrum missing. The powerlaw fit is improved by sigma clipping of the flux, but because of dust extinction, the spectrum is not quite the right shape to be fitted by a powerlaw continuum.

5 such spectra. I remove these from the sample, and identify a new bluest powerlaw for each of the implicated bins. All other spectra will be fitted again with a different function, as I explain in the next section.

2.7 Dust models

I assume the bluest spectrum of each (z, L) bin to be unaffected by dust extinction, or only with insignificant amounts of dust reddening, and let it represent the intrinsic spectrum of all other spectra in the same bin. This provides a reference point to which I can compare the dust extinction of other spectra. To measure the reddening of the rest of the spectra I fit a model of a continuum reddened by dust. The powerlaw fit to the bluest object of each bin gives me the intrinsic continuum of the other spectra, which I then redden by adding an extinction curve.

The dust models I use are the extinction curves of the SMC and LMC. The models are shown in figure 10. The figure shows linear interpolations to points of the average SMC and LMC extinction listed in table 4 of Gordon et al. (2003). For $\lambda < 3000$Å the points of the extinction curves are closely spaced, so even though a linear interpolation is a crude approximation, it is sufficiently detailed that the 2175Å bump of the LMC model is well resolved, and detailed enough for these studies. The linear interpolation is done with the IDL task `linterp`[7].

Also shown in figure 10 is the average Milky Way extinction curve for $R_V = 3.2$ (Cardelli et al. 1989). The Milky Way and the LMC extinction curves are very similar, so I have adopted the LMC curve as representative of both. My results from fits using

[7]Found at http://idlastro.gsfc.nasa.gov/ftp/pro/math/linterp.pro

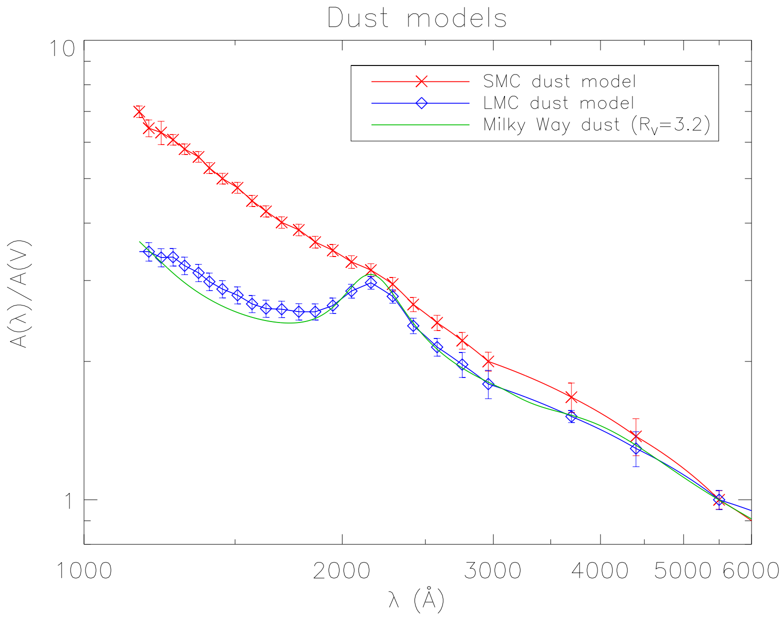

Figure 10: Average absolute extinction $A(\lambda)/A(V)$ for SMC (red), LMC (blue) and Milky Way dust (green). The SMC and LMC extinction curves are linear interpolations to table 4 of Gordon et al. (2003). Both curves are shown with error bars and normalized to unity at 5500Å, equivalent to the middle of the V-band filter.

SMC and LMC like extinction show that an extinction curve without the 2175Å bump (SMC) can be fitted to a larger number of the spectra than an extinction curve with the 2175Å bump (LMC) (see also section 4.1). The Milky Way extinction curve has a slightly more pronounced 2175Å bump than the LMC curve, and would therefore most likely give even fewer fits than the LMC model.

2.8 Fitting the dust-model

If $F_{\lambda,0}$ is the flux of an unextinguished source at wavelength λ, then the flux recieved from a source reddened by dust with optical depth τ_λ is (Sparke & Gallagher 2007, p. 33):

$$F_\lambda = F_{\lambda,0} e^{-\tau_\lambda}. \tag{5}$$

The extinction A_λ at a wavelength λ, can be measured as a difference in magnitude. The difference in apparent magnitude m for two objects given their flux is (Sparke & Gallagher 2007, p. 18):

$$m_1 - m_2 = -2.5\log(F_1/F_2). \tag{6}$$

If m_1 and m_2 are the apparent magnitudes of a reddened source and the same unextinguished source respectively, equation 6 can be used to describe the extinction of the source caused by dust. Combining equations 5 and 6 gives the extinction at the wavelength λ expressed in terms of optical depth τ_λ:

$$\begin{aligned}
A_\lambda &= m - m_0 \\
&= -2.5\log(F_\lambda/F_{\lambda,0}) \\
&= -2.5\log\left(\frac{F_{\lambda,0}e^{-\tau_\lambda}}{F_{\lambda,0}}\right) \\
&= -2.5\log(e^{-\tau_\lambda}) \\
&= -2.5\frac{\ln(e^{-\tau_\lambda})}{\ln(10)} \\
&= 1.086\tau_\lambda. \qquad (7)
\end{aligned}$$

Here m and m_0 are the apparent magnitudes of the reddened and unextinguished sources respectively. If the wavelength dependent optical depth τ_λ is known, the extinction A_λ towards the source is given by equation 7. But τ_λ is a function of λ, and it is not easy to determine through a fit. This is why the dust models are needed. With a dust model τ_λ can be separated into a wavelength dependent component containing the shape of the dust, and a wavelength independent component that describes the amount of dust. I can rewrite equation 7 using the total extinction of the V band $A(V)$ to get

$$\begin{aligned}
\tau(\lambda) &= \frac{A(\lambda)}{1.086} \\
&= \frac{A(V)}{1.086}\cdot\frac{A(\lambda)}{A(V)} \\
&= \tau(V)k(\lambda), \qquad (8)
\end{aligned}$$

where τ_V is the optical depth in the V band, and $k_\lambda \equiv A(\lambda)/A(V)$ is the absolute extinction, which is shown for the SMC and LMC dust models in figure 10. By using k_λ I only need to fit the constant τ_V, not the whole dust model.

To find the best fit values of τ_V I fit the spectra with a combination of a dust model and a powerlaw representing the intrinsic continuum. By merging equations 4 (powerlaw funtion) with 5 and 8 (extinction) I get a function that describes the continuum of a dust reddened quasar:

$$F(\lambda) = b\lambda^{-a}e^{-\tau_V k_\lambda}. \qquad (9)$$

I can fit this function to the continuum of a reddened spectrum. As I assume the bluest spectrum of each bin to be unextinguished and to represent the intrinsic continuum of all quasars in the same bin, I can fix the powerlaw constants a and b to the values found in the pure powerlaw continuum fit of this spectrum. I get k_λ for each dust model from a linear interpolation of the table values of Gordon et al. (2003), also plotted in figure 10. This makes the constant τ_V the only free parameter.

I interpolate the extinction models to the wavelengths of each individual spectrum, and use the powerlaw function fitted to the continuum of the bluest objects to get a model which is directly comparable to the data points of each individual spectrum. I fit the model to the continuum using the same continuum windows as earlier (listed in table 1). This part of the procedure is the same as for the first powerlaw fit (section 2.6). I weight the flux data points of the spectrum with their flux errors, and use sigma clipping iterations to exclude any extreme flux values. Each spectrum is fitted twice, once for each of the SMC and LMC dust models, so that I end up with a value of τ_V for each extinction model.

Figure 11 shows a spectrum fitted with the SMC dust model. The spectrum is fitted with the function of equation 9 using the powerlaw (blue line) of the bluest spectrum of the bin (light gray spectrum). When dust is added, it dims the blue powerlaw continuum until it fits the shape of the continuum of the reddened spectrum (red line). The figure also shows that the pure powerlaw continuum is a straight line, while the dust reddened continuum is slightly curved.

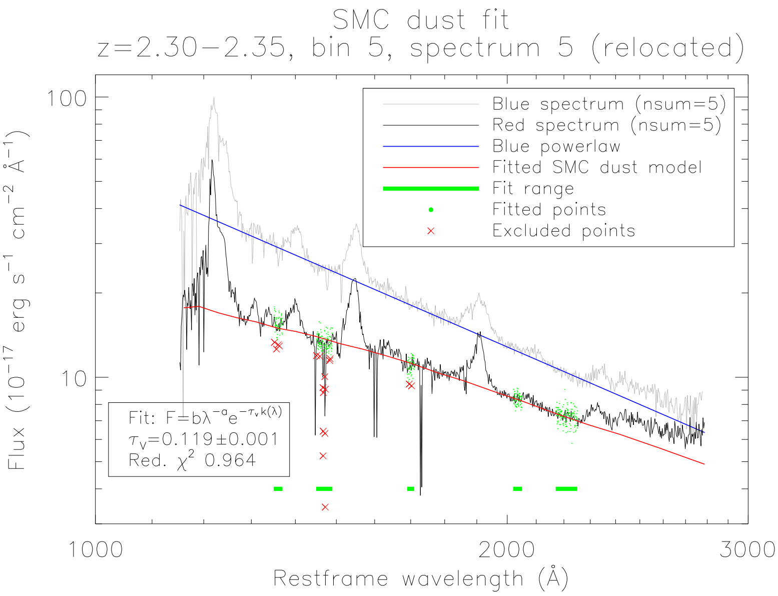

Figure 11: A spectrum fitted using the SMC dust model. The light gray spectrum is the bluest spectrum of the bin, shown with its powerlaw continuum fit (blue line). SMC dust is added to the blue powerlaw to make it match the continuum of the reddened spectrum (red line).

Depending on the amount of dust, a spectrum may be reddened so that the quasar is located in a less luminous bin of my sample, than if it had been unreddened. In that case, the blue spectrum that should represent the intrinsic continuum of the reddened spectrum is to be found in a more luminous bin. It is therefore not enough to treat each bin across luminosity as an isolated case.

To account for this shift to lower luminosity bins by extinction I fit the spectra with the function in equation 9 for each of the bluest spectra within the same redshift range. If a fit with an unreddened powerlaw from a more luminous bin is a better match, I relocate the fitted object to that bin. The program I have written to do these fits automatically moves the quasar to the luminosity bin that gives the lowest χ^2 value to the fit, (on the condition that it is a more luminous bin, since a reddened spectrum cannot have a less luminous intrinsic spectrum). Afterwards, I do a visual inspection of the fits to make sure they all look reasonable. If a fit has a high χ^2 or the shape is not quite right I take an extra look at the spectrum and compare the fits of all the luminosity bins to that spectrum. By eye and from χ^2 values I then decide whether to keep the fit as it is, move the quasar to another bin, or discard the fits if none of them are acceptable. The spectrum shown in figure 11 was originaly placed in a luminosity bin one step lower. In most cases where I discard a fit it is because the dust model has a different shape than the continuum and does not provide a good fit to the spectrum, or if there is strong absorption or flux measurements are missing (as in e.g. the right panel of figure 9) in the continuum windows and hinders an acceptable fit. A fit is e.g. unacceptable if the model is above or below one standard deviation (more or less, I judge by eye) of the average flux level of the spectrum within two or more of the fitted continuum windows. A fit is also discarded if the model clearly overshoots the spectrum in other parts of the spectrum outside the continuum windows. By these standards the fit on figure 11 is quite acceptable. Only a few percent of the SMC fits were discarded, mainly because of broad absorption lines in the continuum windows affecting the model fits. On the other hand, LMC extinction gives the continuum a markedly different shape, and I am only able to fit a little less than half of the spectra with this model. Table 4 in section 4.1 shows how many spectra I have fittet with each dust model for individual redshift ranges.

I redistribute the objects for the SMC and LMC dust models separately. A particular spectrum may therefore not appear in both sets, or in the same bin of both sets. For the same reason, a specific (z, L) bin does not necessarily contain the same number of SMC and LMC fitted objects.

As mentioned, I model the reddened spectra with a combination of an unreddened continuum and a dust model to determine the amount of dust.

An alternative approach could be to simply add dust to the bluest spectrum of a bin until it matches the shape of the fitted spectrum. Assuming the fitted spectum is intrinsically identical to the bluest one, the amount of dust added to the bluest spectrum is then the same amount that reddens the fitted spectrum.

But the properties of emission and absorption lines differ between quasars, and cannot be accounted for in a simple χ^2 optimization routine. To fit the spectra using this method, I would therefore have to process them one at the time, by adding a random amount of dust to a spectrum and inspect it visually until I find the best fitting τ_V value by trial and error. I have a large number of spectra which would make this method very time consuming. Alternatively I would have to write a very complicated program to do pattern matching of the spectra for me. Neither seems a very practical solution. The method I use here is much simpler and thus preferred.

3 Uncertainties

I use the IDL fitting procedure `curvefit` for the continuum fits. The uncertainties of τ_V obtained from the curvefitting routine are quite small. They turn out to be less than 1% of the fitted τ_V values for both the SMC and LMC dust models. The average of all measured values of τ_V are listed in table 2 for both dust models along with the average uncertainties of the fits, $\langle\sigma_{\tau,fit}\rangle$.

The errors from `curvefit` are small because I fix all parameters but τ_V in equation 9 without taking their uncertainty into account. But both the uncertainty of the dust model and the uncertainty in determining the powerlaw continuum of the spectra contribute to the errors of τ_V.

In section 3.1 I calculate the error contributions from the dust models, σ_{τ,k_λ}, and from the powerlaw uncertainty, $\sigma_{\tau,slope}$. I also show that σ_{τ,k_λ} is negligible. The last column of table 2 shows the average uncertainty of τ_V from the powerlaw slopes, $\langle\sigma_{\tau,slope}\rangle$ (calculated according to equation 14 for $\lambda = 2200$Å). These errors are much larger than $\sigma_{\tau,fit}$, and I can therefore ignore the errors from the `curvefit` procedure, and base my error estimates of τ_V solely on $\sigma_{\tau,slope}$.

In section 3.4 I use $\sigma_{\tau,slope}$ to determine which of my τ_V measurements are significant detections of dust, and which measurements can be explained by uncertainties of the continuum powerlaw slopes.

Table 2: Average values and uncertainties of τ_V.

Dust model	$\langle\tau_V\rangle$	$\langle\sigma_{\tau,fit}\rangle$	$\langle\sigma_{\tau,slope}\rangle$
SMC	0.091	0.78%	26%
LMC	0.127	0.72%	19%

Average values of τ_V and average errors from the fitting procedure $\langle\sigma_{\tau,fit}\rangle$ and powerlaw slopes $\langle\sigma_{\tau,slope}\rangle$ in percent of average τ_V. The error $\sigma_{\tau,slope}$ is much greater than $\sigma_{\tau,fit}$.

3.1 Uncertainty from powerlaw slopes

To find an estimate of the uncertainty of τ_V from powerlaw slopes and k_λ, I take a look at equation 5 describing dust extinction: $F_\lambda = F_{\lambda,0}e^{-\tau_\lambda}$. The unextinguished continuum $F_{\lambda,0}$ is a powerlaw function. If the amount of dust τ_V added to a powerlaw continuum is small, then $e^{-\tau_\lambda}$ is close to 1 and the continuum will still look very much like a powerlaw function. So when τ_V is small the dust affected continuum F_λ is also a powerlaw function. Combining equations 4 and 5 under this assumption I can set up the following equation:

$$e^{-\tau_\lambda} = \frac{F_\lambda}{F_{\lambda,0}} = \frac{b\lambda^{-a}}{b_0\lambda^{-a_0}}.$$

Here a, b and a_0, b_0 are the powerlaw parameters of the dust affected and unaffected continuum powerlaws respectively. I take the natural logarithm to both sides of the above equation and get:

$$\begin{aligned}\ln(e^{-\tau_\lambda}) &= \ln\left(\frac{b\lambda^{-a}}{b_0\lambda^{-a_0}}\right)\\ -\tau_\lambda &= \ln(b/b_0)+\ln(\lambda^{-a})-\ln(\lambda^{-a_0})\\ &= \ln(b/b_0)-a\ln(\lambda)+a_0\ln(\lambda)\\ &= \ln(b/b_0)-(a-a_0)\ln(\lambda). \end{aligned} \tag{10}$$

I know that $\tau_\lambda = \tau_V k_\lambda$ (equation 8). I can then reorganize equation 10 to get:

$$\tau_V = \frac{(a-a_0)\ln(\lambda)}{k_\lambda} - \frac{\ln(b/b_0)}{k_\lambda}. \tag{11}$$

In principle, τ_V can be calculated from equation 11 under the assumptions made, (i.e. for objects containing small amounts of dust), if the dust model and powerlaw parameters are known. But I don't know in advance the amount of dust in any of my sample objects and cannot assume τ_V to be small. Equation 11 is therefore not applicable in this context. But as long as I deal with errors I will assume that τ_V is small. I do not find several magnitudes of extinction in the spectra (see equations 21 and 22 and my calculations in chapter 5), so this is a reasonable assumption.

To determine the uncertainty of τ_V from equation 11 I can write σ_{τ_V} using the rule for error propagation for independent and random errors (Taylor 1997, p. 75):

$$\sigma_{\tau_V} = \sqrt{\left(\frac{\partial\tau_V}{\partial a}\sigma_a\right)^2 + \left(\frac{\partial\tau_V}{\partial a_0}\sigma_{a_0}\right)^2 + \left(\frac{\partial\tau_V}{\partial b}\sigma_b\right)^2 + \left(\frac{\partial\tau_V}{\partial b_0}\sigma_{b_0}\right)^2 + \left(\frac{\partial\tau_V}{\partial k_\lambda}\sigma_{k_\lambda}\right)^2}. \tag{12}$$

I can shorten down equation 12 and make it simpler. For this I need to figure out how the uncertainties of the powerlaw parameters are constructed by the fitting procedure `curvefit`. Figure 12 illustrates the uncertainties hereof for a powerlaw fit. The upper panel of the figure shows a spectrum with a powerlaw function fitted to the continuum (black line). Above and below are colored lines representing the powerlaw functions produced by combinations of adding and/or subtracting the powerlaw parameter errors σ_a and σ_b. The green dots show the fitted flux points. The lower panel shows the same dots and lines but with the powerlaw fit subtracted. The green bars show the spread of fitted points for each of the fit windows.

I compare the colored lines with the spread of the flux points and conclude that the error powerlaws with only one of either σ_a or σ_b (red and blue lines respectively) best represent the spread of the flux in the spectrum. If both errors are added/subtracted (orange lines) they double the error, and if one error is added and the other subtracted (magenta lines) they nearly cancel the effect of each other. This means that the errors from `curvefit` are meassured with all of the uncertainty on one parameter at the time, while other parameters are fixed.

If I put all of the uncertainty on the powerlaw slopes σ_a and σ_{a_0} I have $\sigma_b = \sigma_{b_0} = 0$. Then the third and fourth terms under the square root in equation 12 disappears and the uncertainty of τ_V is solely due to uncertainty in the powerlaw slopes and k_λ. Equation 12 then becomes:

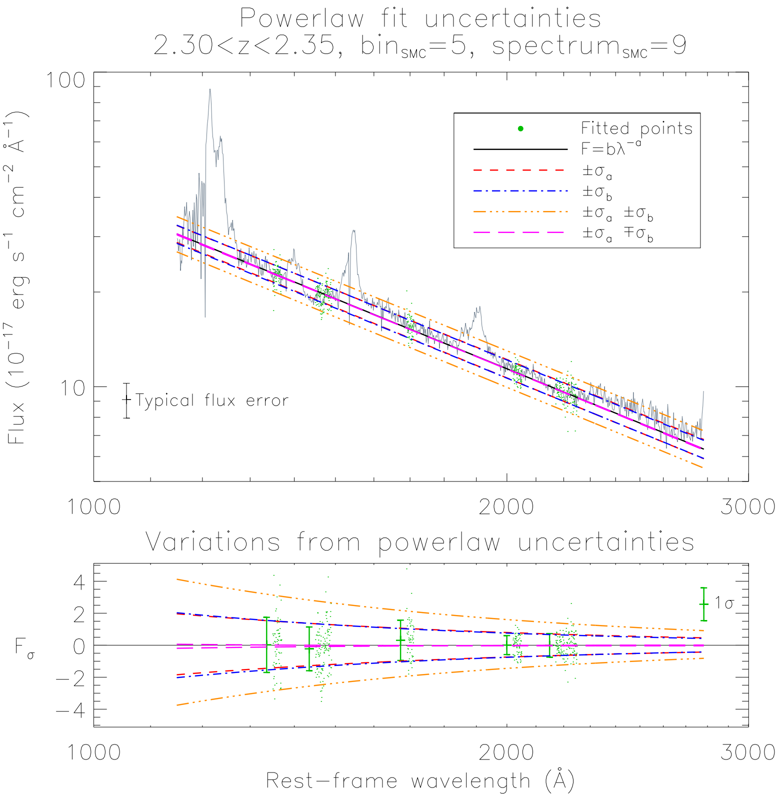

Figure 12: Illustration of the errors calculated by `curvefit`. *Upper panel:* A spectrum (gray) and the powerlaw fitted to the continuum (black). Colored lines represent the powerlaw functions produced by combinations of adding and/or subtracting the powerlaw parameter errors σ_a and σ_b. *Lower panel:* Error powerlaws with the continuum powerlaw subtracted. Green dots and error bars show the spread of the fitted points. The error of a single parameter represents the spread of the fitted points best.

$$\begin{aligned}
\sigma_{\tau_V} &= \sqrt{\left(\frac{\partial \tau_V}{\partial a}\sigma_a\right)^2 + \left(\frac{\partial \tau_V}{\partial a_0}\sigma_{a0}\right)^2 + \left(\frac{\partial \tau_V}{\partial k_\lambda}\sigma_{k_\lambda}\right)^2} \\
&= \sqrt{(\sigma_a^2 + \sigma_{a0}^2)\left(\frac{\ln(\lambda)}{k_\lambda}\right)^2 + \left(\tau_V \frac{\sigma_{k_\lambda}}{k_\lambda}\right)^2} \\
&= \sqrt{\sigma_{\tau,slope}^2 + \sigma_{\tau,k_\lambda}^2}, \qquad (13)
\end{aligned}$$

where $\sigma_{\tau,slope}$ and σ_{τ,k_λ} are the uncertainty contributions from powerlaw slope and dust model respectively, given by:

$$\sigma_{\tau,slope} = \sqrt{\sigma_a^2 + \sigma_{a0}^2}\,\frac{\ln(\lambda)}{k_\lambda}, \qquad (14)$$

$$\sigma_{\tau,k_\lambda} = \tau_V \frac{\sigma_{k_\lambda}}{k_\lambda}. \qquad (15)$$

I want to show that the contribution to σ_{τ_V} from the dust model is much smaller than the contribution from the powerlaw fit, and that I can therefore safely ignore the uncertainty from the dust models. The errorbars of the SMC and LMC dust models in figure 10 are similar, so my estimate of σ_{τ,k_λ} does not depend strongly on which model I use or on wavelength. At $\lambda = 2160$Å (as close to 2200Å as possible) the relative uncertainty of both dust models is $\sigma_{k_\lambda}/k_\lambda = 0.032$ (Gordon et al. 2003). The largest values of τ_V I measure are around $\tau_V = 0.35$ for the SMC model, and smaller for LMC (see figure 17). With these numbers $\sigma_{\tau,k_\lambda} = 0.011$ according to equation 15. In equation 17 of the next section I estimate the uncertainty from the powerlaw at $\lambda = 2200$Å to be $\sigma_{\tau,slope} = 0.03$. At first glance the difference between σ_{τ,k_λ} and $\sigma_{\tau,slope}$ is not enough to ignore the former, but in equation 13 the errors are added in quadrature, which makes the contribution from σ_{τ,k_λ} quite small. It only increases σ_{τ_V} by 7% when I include σ_{τ,k_λ}, which is only a small change. This is a maximum estimate. For a more average amount of dust, e.g. $\tau_V = 0.1$, the contribution from the dust models are even smaller, and increases σ_{τ_V} by less than 1%. The error from k_λ is therefore negligible in equation 13, and I can write the error of τ_V as:

$$\sigma_{\tau_V} = \sigma_{\tau,slope} = \sqrt{\sigma_a^2 + \sigma_{a_0}^2}\,\frac{\ln(\lambda)}{k_\lambda}. \tag{16}$$

The errors of both powerlaw slopes contribute to the uncertainty of τ_V, and are added in quadrature. Not all of the dust reddened quasars have a very powerlaw like continuum, but the pure powerlaw fits are still the best estimates I have of a powerlaw shaped continuum for the reddened objects. I therefore use the slope uncertainties of the pure powerlaw fits as σ_a for dust affected spectra.

Equation 16 tells me that $\sigma_{\tau,slope}$ also depends on wavelength, both directly through $\ln(\lambda)$ and from k_λ of the dust models. I have to consider this dependence when I determine the significance of the fitted τ_V values.

3.2 Properties of powerlaw slope uncertainties

The powerlaw slope uncertainties have a constant distribution across redshift. I illustrate this in figure 13, which shows σ_a for the quasars fitted with the SMC dust model (black dots), and for the blue objects used as intrinsic continuum in the fits (blue dots). The distribution is the same for the LMC dust model, except fewer of the available redshift ranges were fitted. (Section 4.1 includes a table that shows which redshift ranges were fitted for the SMC and LMC models.)

The distribution of σ_a in figure 13 is very uniform across redshift. There seems to be an upper limit cut off of the distribution, and the red line shows a constant σ_a value to mark where 90% of the distribution has a smaller σ_a. Only few points lie thinly spread above this line. The cut off originates from the S/N limit I set for the sample. The S/N-limit ensures an upper limit to the noise of the spectra. If there is severe absorption or other irregularities in the continuum windows of the spectra however, it could influence the powerlaw fits and result in a higher uncertainty. This may explain why some of the σ_a values are above the cut off.

I can use the 90% limit value $\sigma_{a,lim} = 0.0089$ as an estimate for the maximum general powerlaw slope uncertainty. At e.g. $\lambda = 2200$Å I have $\ln(\lambda) = 7.7$, and $k_\lambda(2200\text{Å}) \approx 3$ for both SMC and LMC (see figure 10), so the estimate will be valid for both dust models. Then from equation 16, with both σ_a and σ_{a_0} substituted by $\sigma_{a,lim}$, I find:

$$\sigma_{\tau,slope}(2200\text{Å}) \approx 0.03. \tag{17}$$

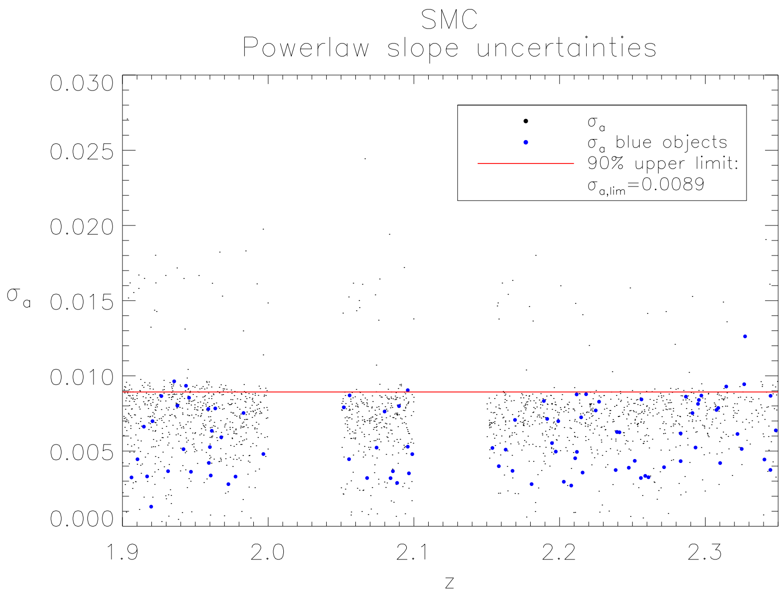

Figure 13: Powerlaw slope uncertainty σ_a for the objects fitted with the SMC dust model shown across redshift. Blue dots indicate unreddened quasars and a red line shows the upper limit for 90% of σ_a. The dustribution is very even across redshift.

The blue dots of figure 13 indicate σ_a values of the quasars I have used as the unreddened objects of each bin. (So these are the σ_{a_0}'s of equation 16). Almost all of these stay below or very close to the 90% limit, and only one of them stands out significantly from the main distribution. This tells me that the blue quasars, which influence the dust fits of multiple other quasar spectra, are generally among the powerlaw fits of relatively high quality, which is good for the quality of my fits. There are 6 of the bluest quasars above $\sigma_{a,lim}$ in figure 13. My choice of $\sigma_{a,lim}$ at the 90% upper limit is arbitrary, and I could just as well have placed it at $\sigma_a = 0.1$. This is a small adjustment and would only change the error estimate of equation 17 by 12%. But then all of the bluest objects, except for one, would be below this limit.

Figure 14 shows a histogram of the distribution of σ_a. The black histogram shows the distribution of all σ_a, including those of the bluest objects, while the blue histogram at the bottom is for the bluest objects alone. As in figure 13, the 90% upper limit is illustrated by the red line.

In this figure it is easy to see the sharp decline in number of objects above the 90% upper limit. Interestingly, the σ_a values above the cut off seem to cluster at approximately twice the value of the main distribution. I currently have no explanation for this, but it would be interesting to look into at a later time.

The distribution peaks at values just below the cut off. This reflects the S/N distribution of the quasars (figure 5) where most of the sampled quasars have S/N values just above the S/N limit. A smaller S/N value corresponds to larger scattering of the flux and leads to a higher uncertainty of the powerlaw fits. The tail of high S/N objects then form the low uncertainty objects.

The bluest objects have a broader distribution that goes to lower uncertainties. So these include relatively more high quality spectra than the main distribution.

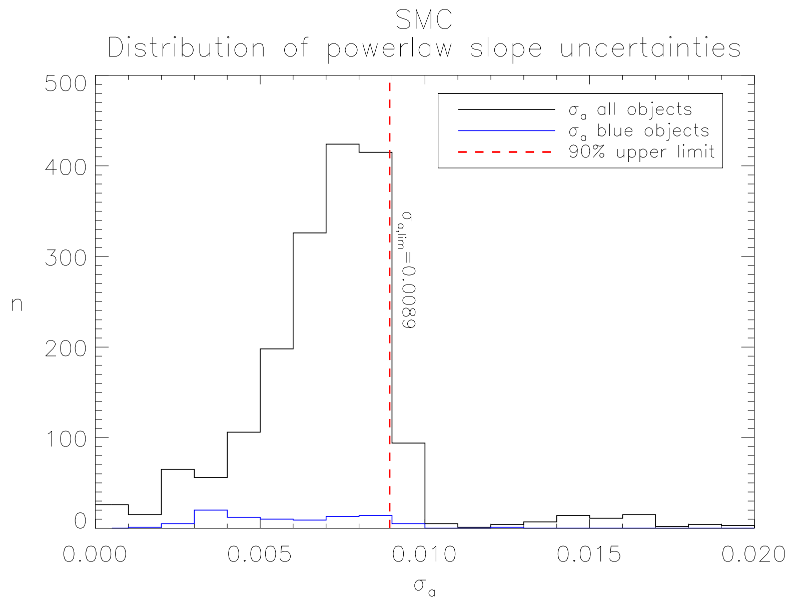

Figure 14: Histogram of powerlaw slope uncertainties, σ_a, for the objects fitted with the SMC dust model. The smaller blue histogram indicates the distribution of σ_a for the bluest objects. The 90% upper limit of the distribution is marked by a red dashed line, at which there is a sharp decline in number of objects.

3.3 Effects of powerlaw continuum uncertainty

Equation 16 shows how $\sigma_{\tau,slope}$ depends on wavelength. Figure 15 shows two examples of τ_V and $\sigma_{\tau,slope}$ over wavelength, one for a spectrum fitted with the SMC dust model (top), and one with the LMC dust model (bottom). The black lines illustrate how much τ_V the flux difference between the powerlaw of the blue spectrum of the fit and the powerlaw of the fitted spectrum corresponds to, when calculated by equation 11. The blue lines show $\sigma_{\tau,slope}$, the amount of τ_V allowed by the powerlaw slope errors. The orange line is the fitted τ_V value, which is independent of wavelength due to the nature of the fit. To the right is an errorbar to show the uncertainty $\sigma_{\tau,fit}$ found through `curvefit` doing the fit, which is quite small compared to $\sigma_{\tau,slope}$. This is generally the case. For this reason I ignore $\sigma_{\tau,fit}$ and use only $\sigma_{\tau,slope}$ to determine which dust fits give a significant τ_V.

The fitted τ_V and the slope difference τ_V always lie close to each other, which is only to be expected since the slope difference τ_V illustrates the same flux difference I attempt to fit with the dust models. The shapes of the slope difference τ_V and the slope error $\sigma_{\tau,slope}$ come from k_λ of the dust model used (see figure 10).

For the SMC model, $\sigma_{\tau,slope}$ is monotonically increasing, and the smallest $\sigma_{\tau,slope}$ is found at short wavelengths where the dust has the largest effect on the spectrum. Correspondingly, the largest $\sigma_{\tau,slope}$ value is at the long wavelength end of the spectrum where it takes more dust to produce the same effect on the flux. For the LMC model, $\sigma_{\tau,slope}$ has a more complicated shape because of a dip caused by the 2175Å bump of the dust model. It is still mostly an increasing function, but the value of $\sigma_{\tau,slope}$ changes more slowly with λ because the LMC dust model is flatter as a function of wavelength than the SMC model.

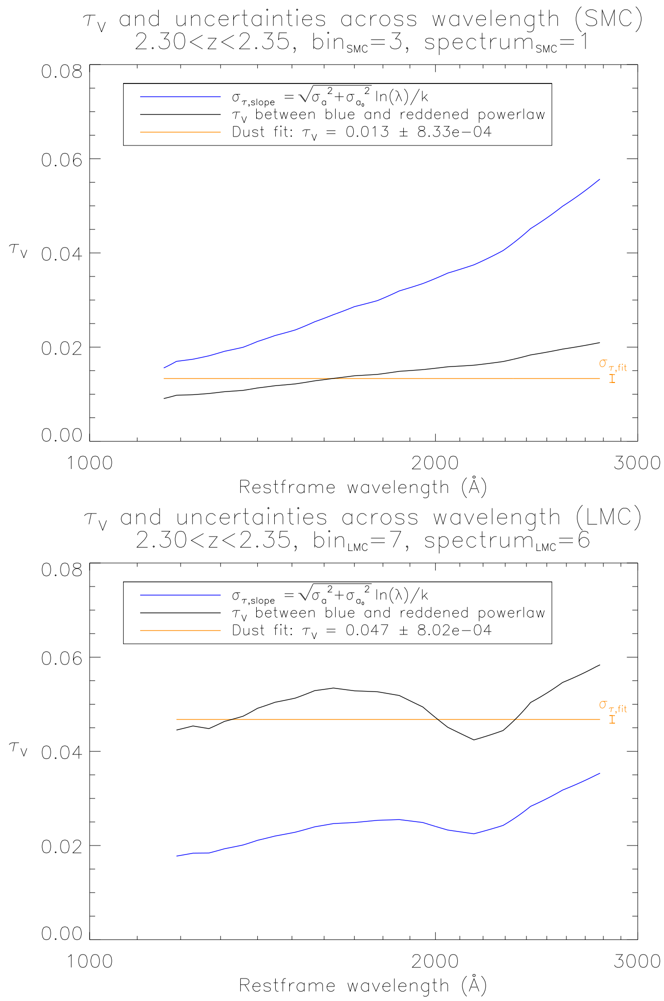

Figure 15: Example of fitted τ_V (orange), τ_V corresponding to the difference between powerlaw slopes (black) calculated from equation 11, and the uncertainty from powerlaw slopes $\sigma_{\tau,slope}$ (blue). *Top panel:* An example using the SMC dust model. The fitted τ_V is smaller than the error due to powerlaw uncertainty, and τ_V is insignificant. *Bottom panel:* An example using the LMC dust model. The fitted τ_V is larger than the powerlaw slope uncertainty, and τ_V is significant.

Depending on $\sigma_{\tau,slope}$ the powerlaw continuum of the dust affected quasars can be more or less indistinguishable from the powerlaw of the unreddened spectrum. Figure 16 gives three examples hereof. The topmost panel shows an example of two spectra (black and gray lines) and their very similar powerlaws (blue and purple lines) with slope errors (illustrated by dashed lines). Both powerlaws are contained within the errors of the other. The difference between the continua can be explained by the small amount of dust fitted, but may as well originate from the uncertainty in determining the powerlaw slopes.

The powerlaws of the middle panel are more separated, and their errors do not overlap at short wavelengths, though still at long wavelengths. In this case it is more likely dust that causes the continuum difference, at least at short wavelengths. But at long wavelengths the difference can still be explained by slope uncertainty alone. That is, the flux uncertainties in the spectra are not small enough to establish an effect from

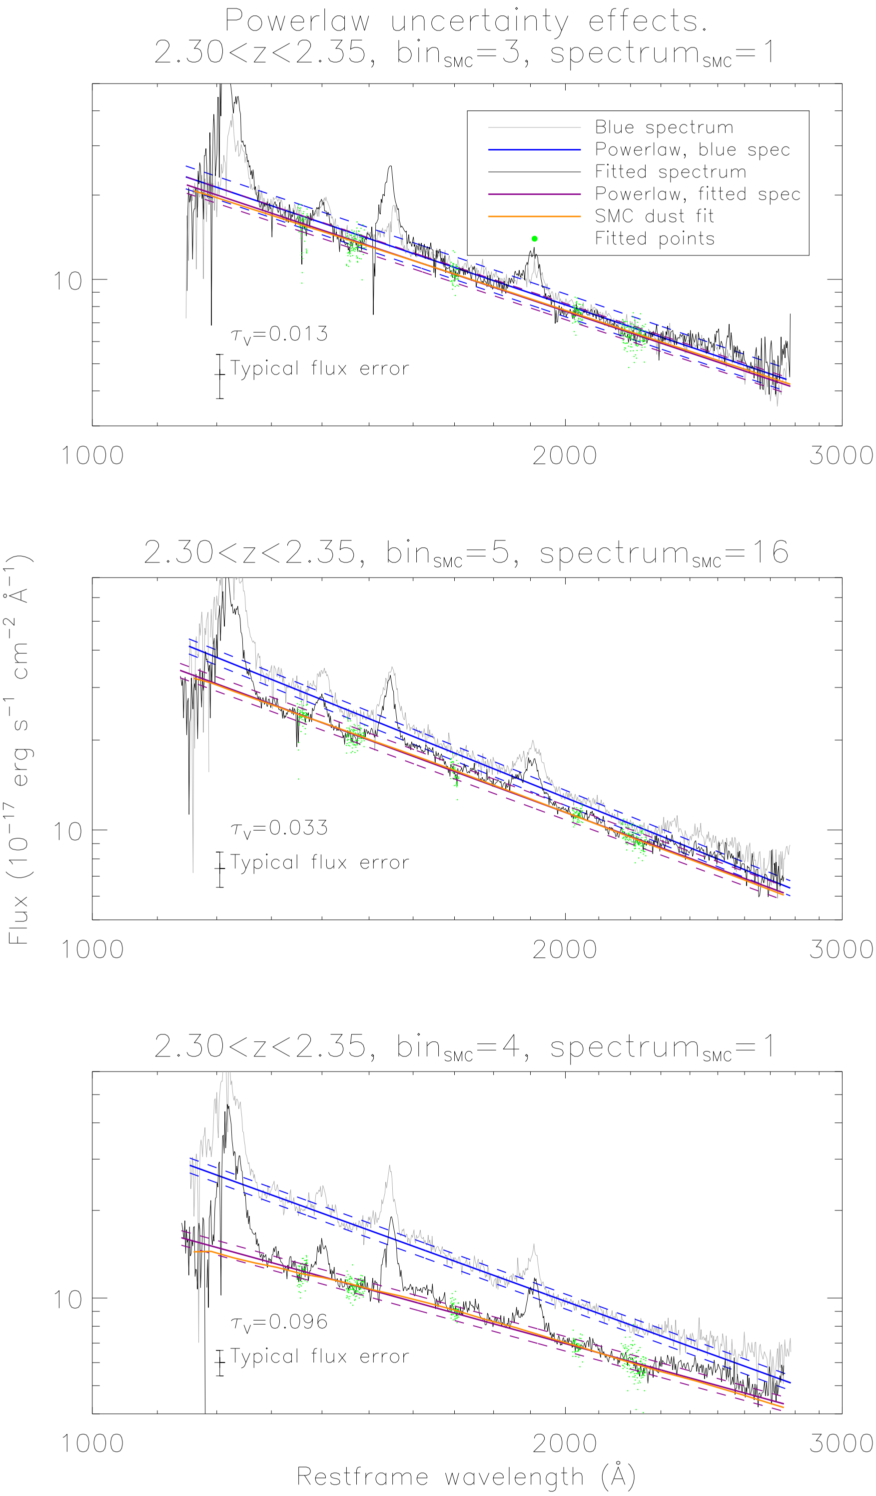

Figure 16: Three dust fits (orange) with different amounts of dust τ_V, including powerlaws of the bluest quasar (blue) and fitted quasar spectrum (purple). Dashed lines indicate the powerlaw slope uncertainty of each, to illustrate whether the continuum difference can be explained by powerlaw slope errors or if dust reddening is more likely.

dust with certainty. Given the significant effect seen at shorter wavelengths, the fitted spectrum is most likely affected by dust, but the quality of the spectrum is not high enough that the amount of reddening can be established across the entire wavelength range (as can be done for the quasar in the bottom panel).

The bottom panel of figure 16 shows two powerlaws that are completety separated for the entire wavelength range shown, even considering the slope errors of both. Here, the continuum difference cannot be caused by slope uncertainty alone, there has to be dust involved. The dust fit (orange line) gives a much larger τ_V than for the other examples and even starts to deviate from the shape of a powerlaw.

In summary, these are three very different cases. The last example seems to contain a lot of dust, the first probably none at all, and in the third case (middel panel) it is a bit ambiguous as to whether or not dust is required to explain the slope difference. So how do I determine if τ_V is significant? I have to decide in which part of the spectrum the effect of τ_V should exceed $\sigma_{\tau,slope}$, and if τ_V has to be significant for the whole wavelength range, or if part of the spectrum is enough to establish the presence of dust.

3.4 Selection of significant τ_V memsurements

The most interesting objects of my sample for further study are the subset with significant measurements of τ_V. But as explained in the previous section, some of the fits do not indicate a significant amount of dust, or the results are unambiguously with respect to the significance of the dust detection.

The powerlaw slope uncertainties, $\sigma_{\tau,slope}$, have by far the largest contribution to the uncertainty of τ_V, and I can disregard the contributions from the curvefit routine, $\sigma_{\tau,fit}$, and from the uncertainty of the dust models, σ_{τ,k_λ}, when I consider the significance of the τ_V measurements. Since $\sigma_{\tau,slope}$ is much larger than σ_{τ,k_λ} and $\sigma_{\tau,fit}$, this will be the uncertainty that decides if τ_V is significant, and τ_V needs to be larger than $\sigma_{\tau,slope}$ in order to be significant. Otherwise the dust can be explained by uncertainty in determining the continuum slopes.

The number of significant τ_V detections depends on the wavelength at which I test significance, because $\sigma_{\tau,slope}$ is wavelength dependent. The powerlaw errors give room for a larger flux difference between the intrinsic powerlaw continuum and the reddened spectrum at short wavelengths (see figure 12, lower panel), but dust also has a larger effect at short wavelengths (figure 10). So which of these is the dominant factor? And where does the uncertainty allow for the largest and smallest amounts of τ_V?

The easiest way to find out is to simply test the significance of τ_V in different wavelengths of the spectra. I have sampled $\sigma_{\tau,slope}$ at 1350Å, 1700Å and 2200Å for the SMC and LMC dust model fits. These wavelengths are all inside my continuum windows (table 1), two at opposite ends of the spectra, and one in the middle. With tests at 1350Å and 2200Å I get an idea of the properties of the fits at both extremes, but avoid extrapolating my results to shorter or longer wavelengths than I have investigated. I also include a test of significance at 1700Å, because I can see in figure 15 that for LMC the slope uncertainty is a bit larger here than at 1350Å and at 2200Å where the 2175Å bump of the dust model is. I therefore expect a larger difference in the number of significant fits between 1350Å and 1700Å than between 1350Å and 2200Å for the LMC dust model. While for the SMC model I expect to see the largest difference between 1350Å and 2200Å.

The errors are based on a normal distribution for which $\pm 1\sigma$ includes 68% of the data points, $\pm 2\sigma$ includes 95% and $\pm 3\sigma$ includes 99.7%. I require τ_V to be larger than $3\sigma_{\tau,slope}$ to make absolutely sure the dust is significant.

Table 3: Number of significant τ_V detections.

Dust model	N_{sample}	sign(1350Å)	sign(1700Å)	sign(2200Å)	sign(1350Å,2200Å)
SMC	1703	1436 (84%)	1240 (73%)	1004 (59%)	1004
LMC	276	225 (81%)	201 (73%)	205 (74%)	205

Number of significant τ_V detections for the SMC and LMC dust model fits. The significant percentage is indicated for each wavelength. More τ_V are significant at 1350Å than at 1700Å and 2200Å, and all τ_V significant at 2200Å are also significant at 1350Å.

Table 3 lists the number of significant τ_V fits for the SMC and LMC dust models. The second column shows the total number of fits (not counting unreddened quasars and discarded fits). The following columns show the number of τ_V significant at 1350Å, 1700Å, 2200Å and those significant at both 1350Å and 2200Å. The significance of τ_V clearly depends on wavelength for the SMC model, but for LMC it is almost constant and varies only by 8%. Fewer τ_V are significant at the long wavelength end of the spectra (1700Å and 2200Å). When I compare how many objects are significantly reddened at 2200Å and at both 1350Å and 2200Å I see that all significant τ_V at 2200Å are also significant at 1350Å. From this I conclude that if τ_V is significant at the long wavelength end of the spectrum, it is significant for the whole spectrum (at least for the SMC model). Then I only need to test the significance at 2200Å.

I prefer to determine the significance of τ_V at the same wavelength for both dust models. The LMC fits are slightly more restrictive at 1700Å than at 2200Å, but the difference is so small (only 4 fits of 276) that I will ignore it, and test the significance of dust for the LMC model at 2200Å as well, to be consistent with the SMC model.

One could argue that if a fit is significant only at 1350Å, it *is* significant in part of the spectrum, and some dust *is* required to explain the change of the continuum flux. But I choose to test significance at 2200Å to be conservative. Any significant measurements are then unambiguously significant.

I find significant dust extinction in most of the quasars. The LMC dust model gives a larger fraction of significant fits than the SMC dust model. Based on the numbers of table 3, 74% of the LMC dust fits are significant at 2200Å , while the same number for the SMC dust model is only 59%. So although the SMC dust model generally is easier to fit and yields slightly better fits (as described in greater detail in sections 4.1 and 4.4), the LMC dust model gives a larger fraction of significant fits. This is maybe a bit surprising since the LMC dust model has a more complicated shape because of the 2175Å bump. The bump is on top of the two longest wavelength continuum windows used for the fits, so I expected the SMC dust model with a more uniform shape to be the most adaptable. On the other hand, the SMC model is steeper, which makes the dust at short wavelengths more 'aggressive'. This may not agree well with some of the strongly reddened quasars and in these cases the more moderate LMC model may be more compatible.

The histograms of figure 17 show the distribution of significant τ_V measurements for the SMC and LMC dust models. Black histograms show the distribution of all fitted τ_V values. Blue histograms show the values significant at 1350Å and red show the values significant at 2200Å. The latter is more restictive than the former. The figure also lists mean values and standard deviations of the distributions. Primarily low values of τ_V are insignificant and missing from the blue and red histograms, as these are the ones that may be explained by the uncertainties.

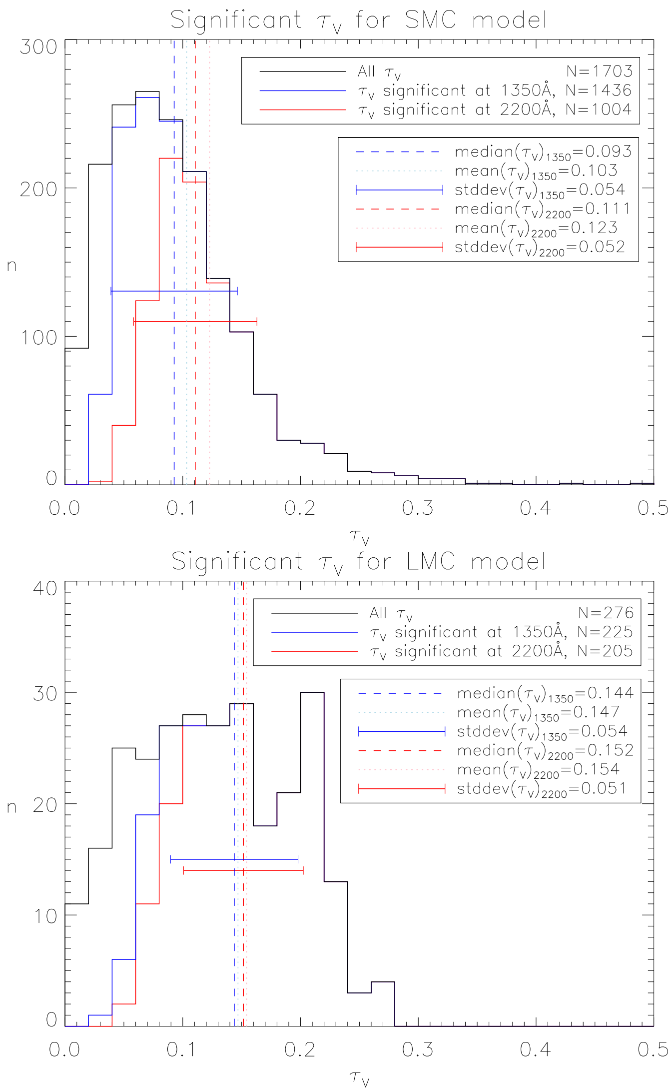

Figure 17: Histogram of τ_V for the SMC (top) and LMC dust models (bottom). The red and blue histograms show the distribution of τ_V meassurements that are significant at 2200Å and 1350Å respectively. The dashed and dotted lines mark median and mean values, and the error bars show the standard deviation of the distributions.

The distribution of τ_V fitted using the LMC dust model is flatter and broader than for the SMC, and the LMC model fits give relatively more large τ_V values. The difference between significant τ_V at 1350Å and 2200Å is largest for the SMC model, while the two are almost the same for LMC. This reflects the differences of the dust models; the SMC model is steep and the effect of the dust is different at the two wavelengths, while the

LMC model is comparatively flat in this wavelength range, and the effect of dust more uniform.

The distribution of significant τ_V values at 2200Å (red histogram) versus insignificant τ_V (black above red) compared with the estimate made for the uncertainty, $\sigma_{\tau,slope}(2200\text{Å}) \approx 0.03$ (equation 17), are in good agreement for both dust models. If I use the estimate, the cut for significant τ_V is at $3\sigma_{\tau,slope}(2200\text{Å}) \approx 0.09$, which is also approximately where the distribution of significant objects meets the overall distribution.

4 Results

4.1 Distribution of quasars in the bins

I have explained how τ_V is obtained, and how it represents the amount of dust extinction. Now I present the results and compare τ_V for the SMC and LMC dust models.

Table 4 provides an overview of the redshift bins (column 1) in my sample. It lists the number of objects in each redshift bin for the full sample (column 3) and for the spectra I successfully fitted with the SMC and LMC dust models (columns 5 and 6 respectively). The bluest objects are not included in the sample, and their numbers are shown in column 4 by themselves. The bold face numbers of columns 5 and 6 are the number of fits that reveal significant dust measurements. Also listed in column 2 are the timespans Δt for the given redshift ranges[8].

Table 4: Redshift bin overview.

z range	Δt (Myr)	N_{sample}	N_{blue}	N_{spec} SMC	N_{spec} LMC
1.90 - 1.95	83.3	365	14	332 / **213**	-
1.95 - 2.00	80.0	317	11	303 / **234**	-
2.00 - 2.05	76.9	305	15	-	-
2.05 - 2.10	73.9	324	14	312 / **170**	104 / **80**
2.10 - 2.15	71.1	234	12	-	-
2.15 - 2.20	68.5	249	11	239 / **115**	-
2.20 - 2.25	66.0	212	14	190 / **103**	95 / **61**
2.25 - 2.30	63.6	216	14	197 / **105**	-
2.30 - 2.35	61.3	135	12	130 / **64**	77 / **64**
$1.90 < z < 2.35$	645	2357	117	1703 / **1004** (94% / **55%**)	276 / **205** (41% / **31%**)

Column 1 list the redshift range of each bin. Column 2 gives the corresponding timespan in units of Myr. Column 3 lists the number of objects in the sample. Column 4 lists the number of bluest spectra (one for each non-empty luminosity bin), these are not included in N_{sample}. Columns 5 and 6 shows the number of fitted objects for SMC and LMC dust respectively. For the last two columns, a '-' means the bin was not processed for the corresponding model. The bold numbers show the number of significant dust measurements. The bottom row shows the total of each column. For columns 5 and 6 it also show the numbers in percentage of the processed sample.

The hyphens in columns 5 and 6 of table 4 indicate that the spectra of these redshift bins were not fitted with the corresponding model. I did not have time to process all of the redshift ranges for both dust models. I gave SMC dust fits higher priority than LMC dust fits, simply because SMC extinction is the generally accepted extinction model for quasars (Hopkins et al. 2004).

The redshift bins all cover the same range in redshift, but the second column of table 4 shows that the time span of the bins decreases with increasing redshift. The change

[8] Δt calculated using the Cosmological Calculator for a Flat Universe using the cosmological parameters $\Omega_m = 0.30$ and $H_0 = 70$km/s/Mpc. http://home.fnal.gov/~gnedin/cc/

is not significant to this work however, since the time spans are all of the same order of magnitude, and Δt only decrease by 26% from the lowerst redshift bin to the one at highest redshift.

Also, the number of quasars in the sample (column 3) generally decreases with increasing redshift. Both due to the larger distance combined with the sensitivity of the detectors used by SDSS, and a more significant decrease in number of quasars at the high redshift end of the sample as the redshift bins get closer to the void in the SDSS catalog around $z = 2.7$ (see figure 2).

The LMC dust model is applicable to somewhat fewer objects than the SMC model. For the 7 out of 9 redshift bins fitted with the SMC dust model, I successfully fitted 94% of the quasars. For LMC I fitted only 3 of 9 available redshift bins, and for these only 41% of the spectra could be fitted with the model. I am able to fit the SMC dust model to nearly all of the quasars. There is some degeneracy in the SMC dust model fits because the dust model is powerlaw shaped, just like the intrinsic continua of the spectra. This makes it easier to fit the SMC model than the LMC model. The differences of the observed quasar continua can therefore be explained by SMC dust reddening for most objects.

The dust models are a little more equal in terms of the fraction of significant dust measurements they give. See also figure 18. About half, 55% on average, of the processed sample show significant amounts of dust for the SMC model. For the LMC model I find significant dust in 31% of the quasars. More interesting perhaps, is that out of the successfully fitted objects, 74% are significant for the LMC model, and only 59% for the SMC model, though these numbers depend on where in the spectra I determine the significance (see also table 3). The SMC dust model may be easier to fit, but out of the objects that can be fitted with the LMC dust model, more of them have significant amounts of dust.

There is an interesting trend in the fraction of fitted spectra across redshift. In figure 18 I plot the fitted percentage of the sample for both dust models. SMC is shown in red and LMC in blue. Points connected by dotted lines show the total fraction of successfully fitted spectra. One can see that essentially all spectra can be fitted with an SMC dust model with some degree of extinction. Points connected by dashed lines show the fraction of spectra in the sample with significantly fitted dust.

The fraction of successful LMC fits, both those with significant detections of dust and otherwise, seems to increase with redshift. Between the two highest redshift points of significant LMC fits there is an increase of 4.5 or 2.6 standard deviations, depending on whether I compare to the uncertainty of the lower or higher redshift point. The increase is less evident from the first LMC point, just below $z = 2.1$, to the middle one. But the line above showing the fitted fraction of spectra for all LMC fits has the same trend, so there may be a real evolution. This implies that LMC dust is more common at higher redshifts.

For the SMC dust model there is a weaker trend in the opposite direction, with fewer significant τ_V at higher redshifts. The fraction of successful fits for all SMC fits is almost constant, so there must be something that affects the fits differently for different redshifts, and therefore yields fewer significant measurements at higher redshifts.

The space between the red lines represent the insignificant dust measurements, and therefore the accuracy with which τ_V is measured. The weak decline of SMC dust could

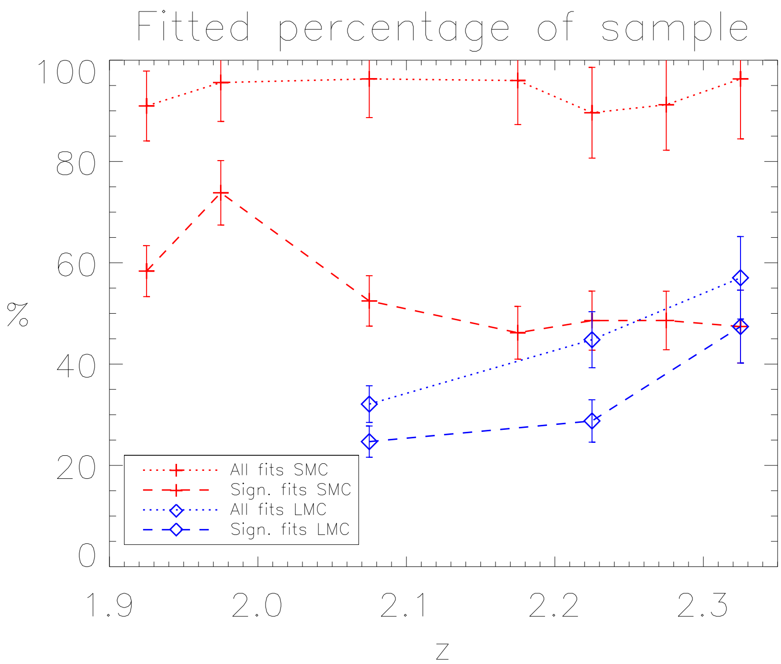

Figure 18: The fraction of successfully fitted spectra for different redshift ranges. The points are placed at the center of their redshift bin. The fraction of fits for the SMC model is shown with red colors and the LMC dust model with blue. Points connected by dotted lines show the total fraction of spectra that could be fitted by the corresponding model. Points connected by dashed lines show fits with significant amounts of dust. The error bars are calculated using the standard deviation of a Poisson distribution and error propagation.

therefore in principle arise from an overall shift in the S/N distribution. It is more difficult to observe high S/N spectra at larger redshifts (see figure 3), because a high redshift quasar would have to be comparatively brighter, and brighter objects are more rare. If the quality decreases, the uncertainty of τ_V will increase, and there will be fewer significant measurements.

A change in S/N distribution is not the cause in this case however. I have measured the average S/N for each of the redshift ranges of my sample, and it is constant within the standard deviation of the distribution. I currently have no alternative explanation either. I would have to take a closer look at all of the fits to determine exactly why there is this weak trend. Perhaps there really is a small decrease in the amount of significant SMC dust at higher redshifts, but it is within the errors and therefore not a significant change.

There is one redshift bin in the $1.95 < z < 2.00$ redshift range with a clearly larger fraction of significant τ_V for the SMC model that is puzzling. The most realistic explaination for this deviation is that for some reason the continuum slopes are more accurately determined for the spectra of this bin. It might be enough that the bluest spectra have larger S/N and are of a higher quality, since they also contribute to the uncertainty of τ_V. The powerlaw slope uncertainties of the blue spectra in this redshift range are actually slightly smaller than for most other redshift ranges in figure 13. But the same could be said for $2.15 < z < 2.20$, which does not deviate from the other points of figure 18. Again, the S/N distribution, and hence uncertainty distribution, seems an unlikely cause.

The trend of the LMC fits is more evident than for SMC. But the increase of successful LMC fits towards higher redshifts might still be artificial. For higher redshifts the emitted spectrum is pushed to longer wavelengths. The SDSS observes over a fixed range of wavelengths. For my sample it means that the fitted 2175Å bump moves closer to the long wavelength edge of the spectra. I have no continuum windows longward of the bump and do not fit directly to this area, but I use the flux level as a guide in visual inspections the fits. The LMC fits have a tendency to overshoot the spectra longward of the bump (see also section 4.4), and at higher redshifts there is a shorter wavelength range to constrain the fits at this end. That could lead to larger values of τ_V and an increase in significant fits.

However, if the trend of LMC dust is real there is an interesting evolution in the type of dust in quasars. To investigate further would require data at higher redshifts. Spectra of a longer wavelength range may also be necessary to better constrain the LMC fits at longer wavelengths. A possible source of such spectra is the X-shooter instrument installed at the VLT.[9].

I have moved each spectrum to the luminosity bin that gives the best dust model fit. This changes the distribution of spectra over luminosity to represent the intrinsic luminosities instead of the observed ones. This is illustrated in figure 19, which shows how the spectra are distributed across luminosity bins. The bars show what percentage of the spectra are in each luminosity bin.

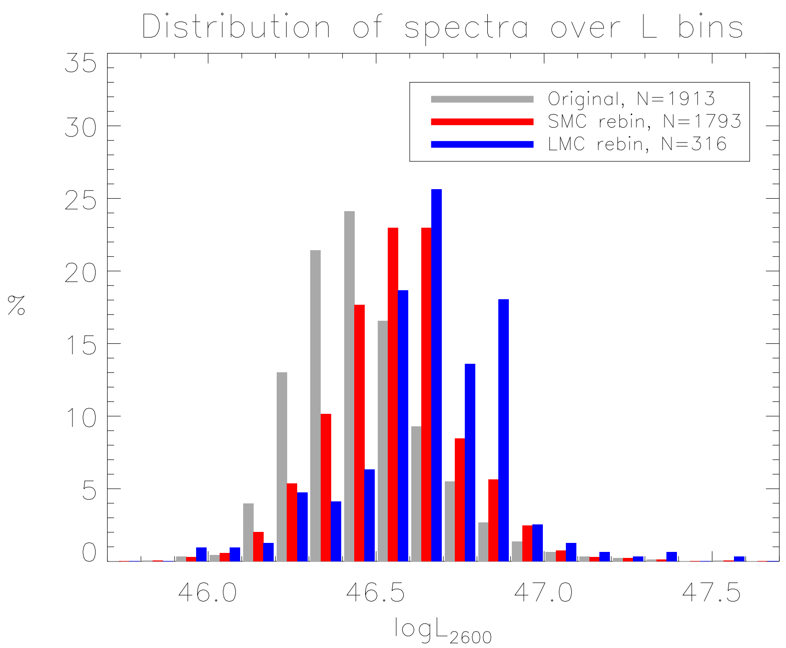

Figure 19: Distribution of spectra across luminosity bins. The bars indicate the percentage of spectra in a bin. Gray bars show the original distribution of the spectra fitted with the SMC model. Red bars show the distribution of spectra according to the SMC dust fits. And blue bars show the distribution according to the LMC dust fits for the three relevant bins (see table 4). The distributions after dust fits are shifted to higher luminosities.

[9]https://www.eso.org/sci/facilities/paranal/instruments/xshooter.html

The gray bars show the original distribution of spectra in the sample, though not including the two redshift bins that were not fitted with the SMC dust model. These were excluded to make the original distribution comparable with the SMC fitted distribution (red bars). The SMC fitted distribution is clearly pushed towards higher luminosity, as to be expected since this would be the unextinguished distribution of the spectra.

The distribution according to the LMC dust fits (blue bars) is even more shifted towards higher luminosity than for the SMC model. This indicates that more dust can be fitted with the LMC model than with SMC. The LMC fitted distribution does not include all of the same redshift ranges as for SMC (see table 4), but should still be comparable to the other distributions, as there are enough objects to make it representable for the sample.

4.2 Smaller bins in luminosity reveal more similar SEDs

An interesting result that I can reproduce from my data, but on an even smaller scale, was originally found by Krawczyk et al. (2013). They divide the quasars of the SDSS DR7 into three equally populated bins of luminosity and then measure and compare the SEDs. Krawczyk et al. (2013) find that the SEDs of the sub-samples are luminosity dependent, and differ in spectral slopes at e.g. far-UV and optical wavelengths, but also that there is less dispersion between SEDs for smaller luminosity ranges than for the full sample.

At UV wavelengths the powerlaw function (equation 4) is representative of the SED of quasars. Figure 20 shows the distribution of the fitted continuum powerlaw slope a for all quasars in a single redshift range. This covers all luminosities in my sample. The mean value of a is $\langle a(z) \rangle = 1.51$ and the standard deviation is $\sigma_{a(z)} = 0.50$. The mean values differ with only a few percent for the other redshift ranges, but the standard deviation is consistently smaller, down to $\sigma_{a(z)} = 0.35$, for other z. The same end conclusion still applies for all redshift ranges though.

Figure 21 shows the distribution of fitted powerlaw slope values for quasars in the same redshift range as for figure 20, but for each luminosity bin separately (for the original distribution of quasars). I have ignored (z, L) bins with less than 20 objects, because a minimum number of quasars are required to give an idea of the distribution and spread. The mean values of the single bins vary more than from one full redshift range to another and are overall decreasing with increasing luminosity. But the mean values are still consistent with the overall mean of the redshift range they belong to. More interesting, is that the standard deviations are significantly smaller for all of the single bins compared to the full collection of luminosity bins. Combined with the small systematic changes in mean slope, this implies a luminosity dependence of the spectral slopes. This is the same result as in Krawczyk et al. (2013), but for much smaller luminosity ranges.

A few objects in figure 20 have negative powerlaw slopes. I have not inspected every powerlaw fit made, because I had to fit all except the bluest quasar of each bin again with another model. For this reason there are probably a small number of bad powerlaw fits in the sample. A cause for bad powerlaw fits could for instance be cases of strong absorption of the spectra within the fitted continuum windows. The sigma clipping procedure is not always able to remove all of these flux points if most of the continuum window is affected. Leftover flux points with absorption can then influence the fit and

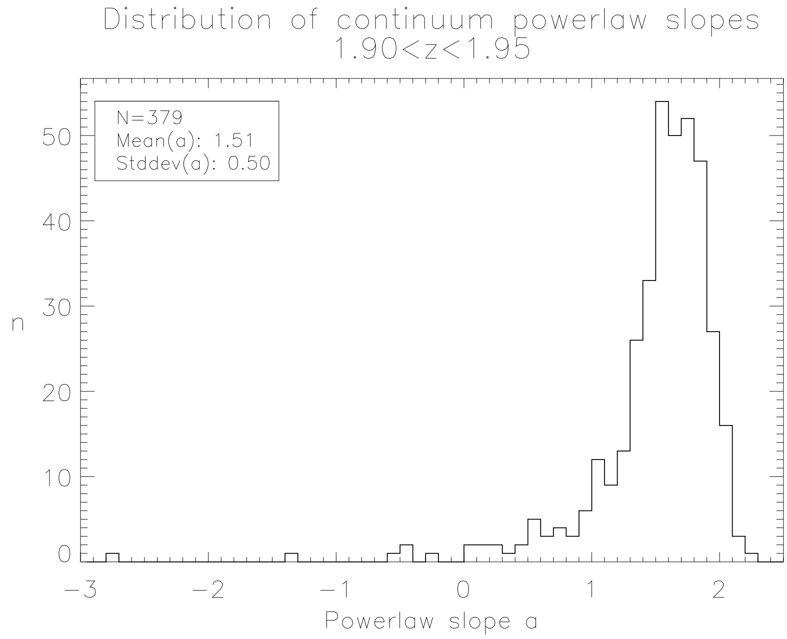

Figure 20: Distribution of fitted powerlaw slopes for all sample objects in the redshift range $1.90 < z < 1.95$. The mean value is 1.51, with a standard deviation of 0.50, measured for 379 quasars. A few objects have very low values which indicate strong reddening.

for instance pull it down at short wavelengths to give the powerlaw a negative slope.

Another possible cause for negative fitted slopes is, that there is so much dust in a quasar and the extinction is so strong, that the spectrum bends downward at short wavelengths and becomes 'n'-shaped (Francis et al. 2000). In that case the best powerlaw fit can also have a negative slope, though a powerlaw is then no longer very representative of the continuum. There are a total of 27 objects fitted with a negative powerlaw slope in my sample. Of these, 11 are simply bad fits caused by strong absorption in the fitted continuum windows. But 16 objects, corresponding to only 0.8% of the fitted sample, show signs of strong reddening.

The general picture is that the standard deviation of the powerlaw slope a within a single (z, L) bin is smaller than when including all of the luminosities. But whenever there is a negative or many small values of a in a (z, L) bin, the standard deviation is much greater, and in some cases become larger than for the full luminosity range. If I remove the negative a's (i.e. bad fits and strongly reddened objects) from the bins, the new standard deviation is smaller and fits in with the overall picture.

The middle panel of figure 21 has a single object with negative a. It is also the bin with the largest standard deviation, $\sigma_{a(z,L)} = 0.43$. If I exclude the one strongly reddened object from the bin I would get $\sigma_{a(z,L)} = 0.29$ instead. The reddened object did not make the standard deviation larger than for the full luminosity range in this case, but they do for a few of the other bins with less quasars. And it still made quite an improvement to the standard deviation to remove the negative value from the bin in this case.

Krawczyk et al. (2013) also exclude objects showing signs of strong reddening from their sample. If I do the same, it will most likely strengthen the same conclusion.

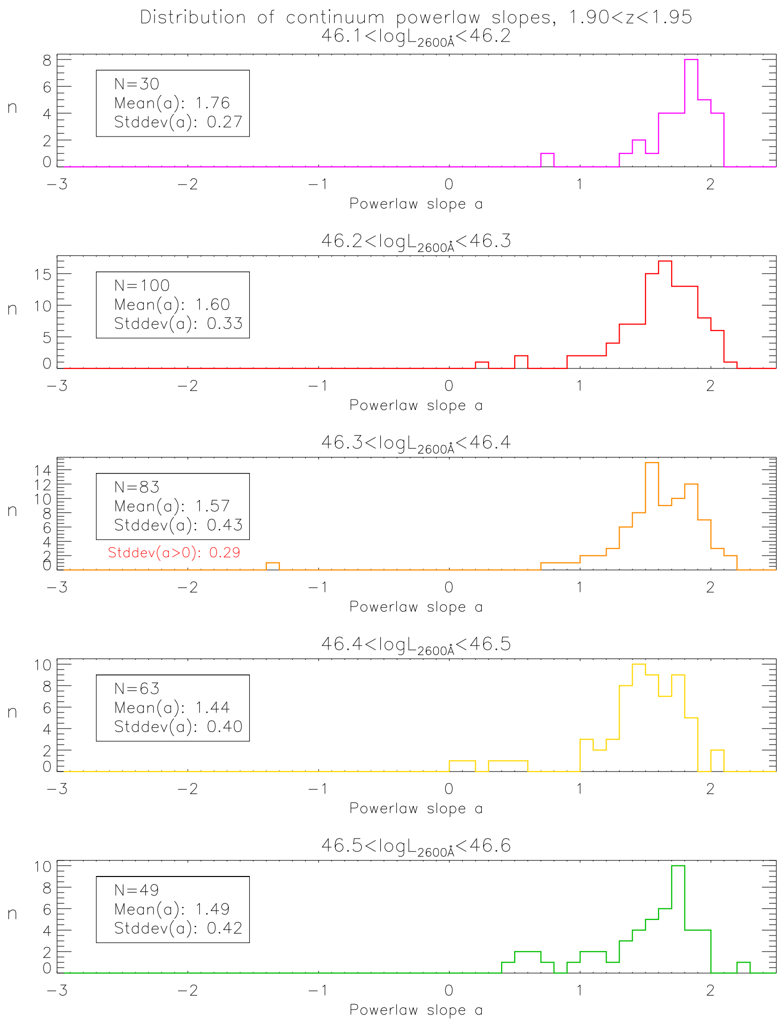

Figure 21: Distribution of fitted powerlaw slopes for luminosity bins (the original luminosity distribution of quasars) containing at least 20 objects in the redshift range $1.90 < z < 1.95$. The mean values vary, but are consistet with the total distribution of figure 20. The standard deviations are smaller than for the total distribution.

One of the fundamental assumptions I make in order to measure τ_V, is that the intrinsic SEDs of quasars within a (z, L) bin of small size are the same. Krawczyk et al. (2013) and especially the results here support this assumption.

4.3 The effect of dust for different redshift and luminosity bins

The main goal of my thesis is to investigate if and how the extinction curves of quasars depend on redshift and source luminosity. Figure 22 shows the significant measurements of dust and their distribution across redshift for the SMC and LMC dust models in top and bottom panels respectively. The squares indicate the mean values of τ_V for each full redshift range, and the error bars show one standard deviation. Above each error bar is written the number of quasars in the bin.

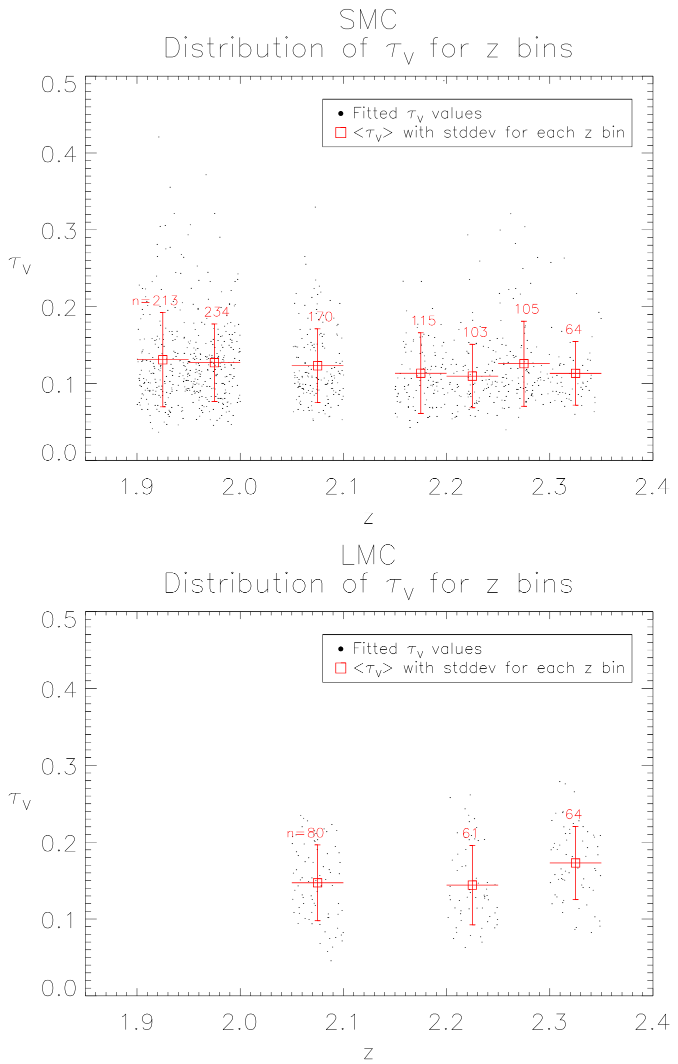

Figure 22: Distribution of significant measurements of τ_V across redshift for the SMC (*top*) and LMC (*bottom*) dust models. Each dot represents a measured τ_V value. The mean value of τ_V for each redshift bin is marked by a red square, and the standard deviation is shown by an error bar. The horizontal red bar shows the width of the redshift bin. The numbers above each error bar indicate the number of significant dust measurements of their respective redshift ranges. The mean values are consistent with no evolution of τ_V across redshift.

The mean τ_V values are consistent with no evolution across z within the standard deviations of each redshift bin for both the SMC and LMC dust models. The average amount of extinction does not change with redshift. The mean values $\langle \tau_V \rangle_{SMC}$ and $\langle \tau_V \rangle_{LMC}$ and standard deviations of significant measurements across all of the tested redshift bins are:

$$\langle \tau_V \rangle_{SMC} = 0.12 \pm 0.05, \tag{18}$$

$$\langle \tau_V \rangle_{LMC} = 0.15 \pm 0.05. \tag{19}$$

The values of equations 18 and 19 agree within the standard deviations, but the average τ_V measured with the LMC model is a bit larger than for the SMC dust model.

If I split the sample up into individual (z, L) bins, and only look at one luminosity range across all z at the time, they show the same behavior as the full sample. Different luminosity ranges do not behave significantly different from each other. See figures 27 and 26 in appendix A for more details.

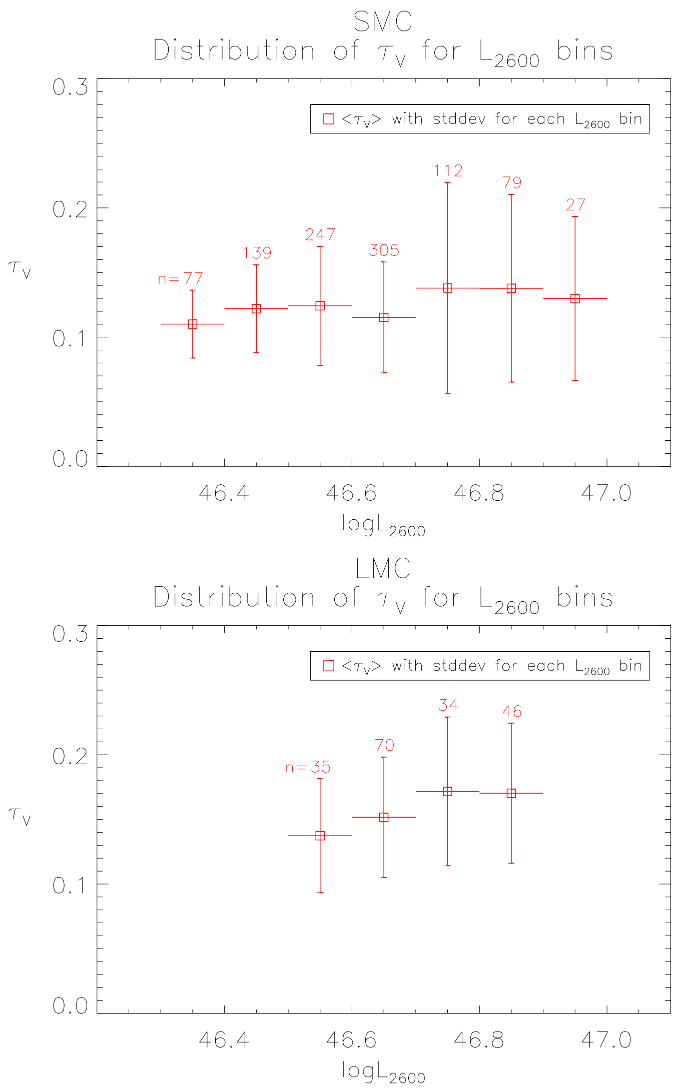

Figure 23: Distribution of significant measurements of τ_V across luminosity for the SMC (*top*) and LMC (*bottom*) dust models. The quasars are reorganized in luminosity according to their best dust fit. The mean values of τ_V for luminosity bins are marked by red squares, and standard deviations are shown with error bars. The horizontal red bar shows the width of the luminosity bin. The numbers above each error bar indicate the number of significant dust measurements of their respective luminosity ranges. I do not shown luminosity ranges with less than ten quasars. The mean values are consistent with no evolution of τ_V across luminosity, though there is a weak tendensy for larger amounts of dust at higher luminosities.

Figure 23 shows the average τ_V of significant measurements across monochromatic luminosity L_{2600}, when the results are collected for all investigated redshifts. The top and bottom panels show results for the SMC and LMC dust models respectively. The squares indicate the mean values of τ_V for each luminosity range and the error bars show one standard deviation. Above each error bar is written the number of quasars in the bin. The figures do not show luminosity ranges with less than ten objects.

As is the case with redshift, the mean τ_V values of figure 23 for individual luminosity ranges are consistent with no evolution across luminosity within the standard deviations of each luminosity range, and the average amount of extinction does not change with luminosity. This is valid for both the SMC and LMC dust models. There may however be a weak tendency for more dust in higher luminosity objects. This is most pronounced for the LMC model results.

Individual redshift ranges behave the same as the full sample for both dust models, and different redshift ranges do not behave significantly different from each other. See figures 29 and 28 in appendix A for more details.

The quasars are reorganized in luminosity bins according to their best dust fit. The luminosity ranges shown therefore represent the intrinsic luminosities of the quasars. The results of the LMC model are spread out over fewer luminosity bins than for the SMC model. This suggests that the intrinsic luminosities of the quasars are more alike if the LMC dust model is adopted instead of the SMC model. But this is possibly just an effect of not having the same amount of data for the LMC model. In figure 19 of the intrinsic luminosity distributions of SMC and LMC fitted spectra, the widths of the distributions are similar and do not show this tendency.

4.4 How are SMC and LMC different?

In figure 24 I compare τ_V values obtained for the SMC and LMC models for objects where both have been measured. Each colored dot shows τ_V for one of the 269 quasars that have been successfully fitted with both dust models. The x-axis shows $\tau_{V,LMC}$ and the y-axis shows $\tau_{V,SMC}$ values. If a measurement is insignificant for the SMC dust model it is marked with a + sign, and measurements that are insignificant for the LMC dust model are marked with an x.

The points are distributed on slanting lines or belts. To make this more visible I have colored the dots according to the luminosity bin offset between the dust model fits. For instance, an orange dot means that an object was placed in a luminosity bin two steps brighter using the LMC dust model, than by using the SMC dust model in the continuum fit.

The distribution of the blue dots, where the models agree with respect to the best luminosity bin, is very narrow. Whereas the belts are broader when there is an offset in luminosity bin between the dust models. The blue dots form a line with a slope of almost 1:1, but it is a little flatter because the LMC dust model is flatter than the SMC model in the wavelength area of the spectra (see figure 10). LMC extinction is weaker than SMC extinction, so it takes a little more LMC dust to produce the same effect on a spectrum or powerlaw. The other colored belts seem to have the same slope, and are roughly ordered by the offset in luminosity bin. The dots mix with a color that is plus or minus one in luminosity bin offset, but not more.

Figure 24 shows more clearly what I also mention in section 4.1; that the LMC dust model gives a larger number of significant τ_V measurements than the SMC model, and more dust can be fitted to some of the quasars if I use the LMC dust model. The colors show that many of the quasars have been fitted with the powerlaw continuum of

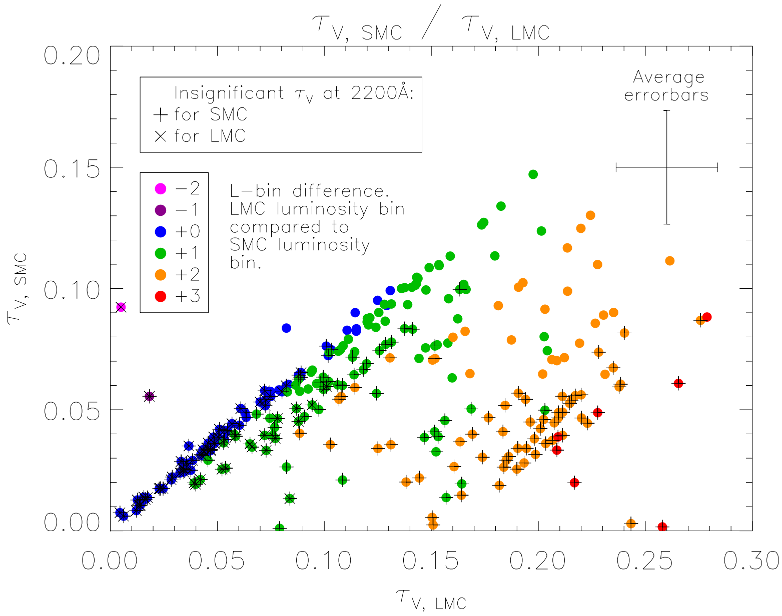

Figure 24: Measurements of τ_V made on quasars for both the SMC and LMC dust models. Insignificant measurements are marked by '+' and 'x' for the SMC and LMC model respectively. The color of each dot indicates the difference in luminosity bin after the dust fits. E.g. for a green dot, the LMC fit moves the quasar to a bin one step brighter than the SMC fit does. For a lot of object it is possible to fit more dust using the LMC dust model than with the SMC model. Many objects fitted with a significant amount of LMC dust are insignificant for the SMC model.

a brighter luminosity bin using the LMC dust model, than for the SMC model fits. The amount of dust fitted by the two models can be very different. There are quasars where I find no dust at all using the SMC model (i.e. when $\tau_{V,SMC}$ is close to 0), while I find plenty of dust if I fit the LMC dust model (e.g. at $\tau_{V,LMC} \approx 0.25$).

Many of the quasars have either significant or insignificant dust whether I use LMC or SMC dust in the fits. But there is an overweight of objects that favor the LMC dust. Many of the fits that give significant measurements of τ_V for the LMC model are insignificant for the SMC model (objects marked only with +). In figure 24 there are 109 LMC-significant and SMC-non-significant measurememnts of τ_V out of a total of 269. For comparison, the pink dot shows the only quasar that is significant in SMC dust and not in LMC dust (the quasar marked only with an x). As I mention in section 4.1 there is a degeneracy in the SMC fits. The much larger number of significant dust measurements with the LMC model, along with the more complex shape of this dust model, therefore also shows that the LMC dust measurements are probably more trustworthy than the SMC dust measurements.

I do a visual inspection of the fits represented in figure 24 to see if the quality of the fits are equal for the two dust models. I simply plot the spectra of each quasar one at the time along with both continuum fits, and compare the shapes, fit residuals and reduced χ^2 values. For each quasar I choose which model gives the best fit. The SMC dust model gives the best fit for 177 ($\sim 2/3$) of the quasars and the LMC model fits best for 92 ($\sim 1/3$) of the quasars. The SMC dust model fits are then slightly better, but in many cases one fit is not obviously better than the other, so the SMC and LMC models are close to equal in quality.

Figure 25 shows two examples of quasar spectra fitted with both dust models to illustrate how the different dust models shape the continua differently, and how the LMC dust model can sometimes be a better alternative to the commonly used SMC dust model. The unreddened powerlaw functions are shown with red dashed and blue dotted lines for the SMC and LMC model fits respectively. The dust model continuum fits themselves are shown with solid lines, red for the SMC model and blue for the LMC model. I also show the reduced χ^2 value of each fit. The fitted τ_V values are written in the legends. For both spectra τ_V is significant for the LMC model fit, but not for the SMC fit.

The upper panel of figure 25 shows a spectrum fitted using the same unreddened powerlaw continuum for both dust models, (i.e. one of the blue points of figure 24). The influence of dust is too strong at short wavelengths for the SMC model, and the SMC fit is clearly too flat to match the continuum of the spectrum properly. The LMC fit follows the spectrum through all of the fitted areas, though it overshoots the spectrum a little around 1450Å in the second continuum window. This may be due to a small amount of absorption in this area however. For this spectrum the LMC dust model gives a much better continuum fit than the SMC model. This shows that SMC is not necessarily the best model for dust in quasars. There are of course other spectra where the SMC model gives the best fit, so it all depends on the intrinsic powerlaw continuum of the fits. I have limited choices in this regard because the intrinsic continuum of each bin is represented by one specific powerlaw.

In the lower panel of figure 25 the LMC model is fitted using a much more luminous powerlaw, and τ_V is more than four times larger than for the SMC fit. The SMC fit is not bad, though it overshoots the spectrum at short wavelengths near 1300Å. This is a typical problem for the SMC fits. Many of the fits have only a little amount of dust and therefore only curve a little. When the continuum of a spectrum is bent by extinction as in this figure, the SMC fits often overshoots the spectrum at one end because it is too straight. This, I believe, supports that SMC extinction is not always most representative of quasar extinction.

The LMC fits have the opposite problem. The model curves too much, and overshoots the spectra in the middle around 1700Å and at the longest wavelengths after the 2175Å bump of the dust model. The LMC fits have a tendency to add too much dust to the spectra. As mentioned earlier (in section 2.8), I visually inspect any problematic fits to see if a fit using the unreddened powerlaw of another luminosity bin is a better match, and ignore the fit if not. This way I make sure the LMC fits do not overestimate the amount of dust.

The lower panel also illustrates well how the LMC fits behave differently because of the 2175Å bump in the dust model, which the SMC model does not have. The SMC fit is close to a powerlaw and indicates a lot of iron emission from around 2200Å and towards longer wavelengths. But the 2175Å bump causes the LMC fit to dip right before the iron emission begins, and then rises above the SMC fit to include much of the iron emission at longer wavelengths as continuum emission. Therefore, depending on which dust model I use to fit the continuum, I find a spectrum with either a lot or very little iron emission. It would be interesting to see how the LMC model fits behave at longer

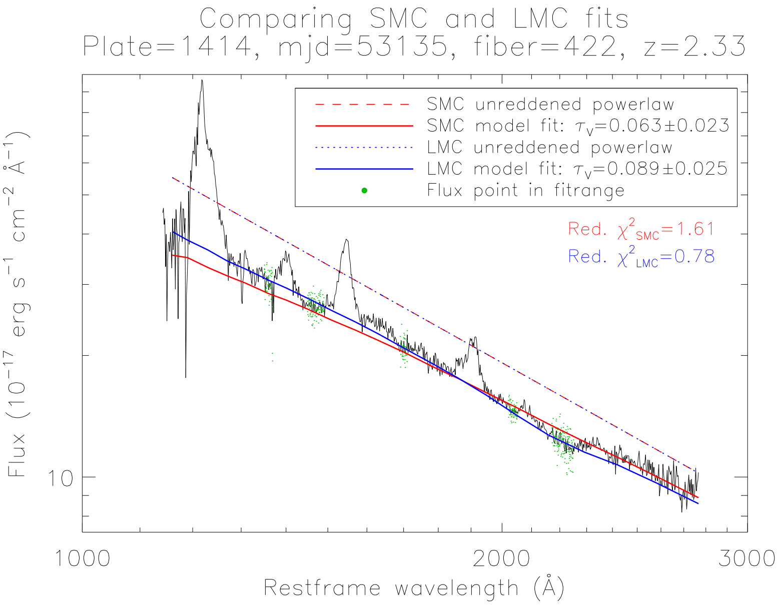

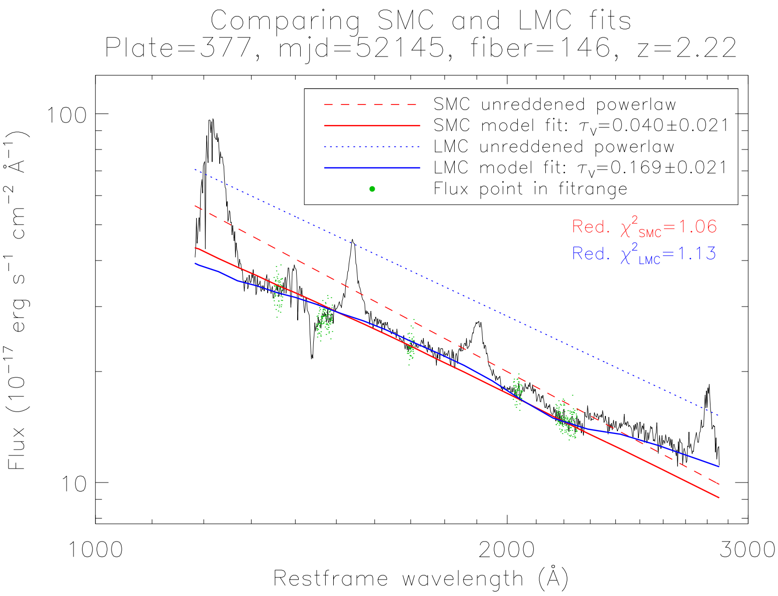

Figure 25: Two examples of quasar spectra fittet with both the SMC (red solid line) and LMC (blue solid line) dust models. In both cases τ_V is significant for the LMC model, but not for the SMC dust model. The unreddened powerlaws used in the fits are shown with red dashed and blue dotted lines for the SMC and LMC models respectively. *Upper panel:* Both dust models are fitted using the same unreddened powerlaw. The SMC model fit does not quite match the slope of the spectrum, while the LMC fit follows the shape of the continuum closely. *Lower panel:* The LMC dust model is fitted using the powerlaw of a much more luminous bin than for SMC, and shows a lot more dust. The most distinct difference between the SMC and LMC model fits is how they fit the iron bump feature starting around 2200Å and towards higher wavelegths. The SMC model continuum fit indicates strong emisson in this area, while the LMC continuum fit includes most of the emission as part of the continuum.

wavelengths than represented in my sample, but I did not have time to investigate this myself.

5 Discussion

I set out to investigate if and how dust in quasars depend on redshift and luminosity, but I have also reached other interesting conclusions during my studies. In short summary my results are the following:

- I find no dependence of τ_V on either luminosity or redshift in the range $1.90 < z < 2.35$ (section 4.3 and figures herein).
- For a specific bin the intrinsic SED is assumed to be the same for all objects. With this I am able to fit nearly all of the spectra with a dust model (section 4.1). I can therefore explain the diversity of observed quasar SEDs within a (z, L) bin with dust extinction.
- The SMC model is the commonly adopted extinction curve for quasars, but it is not the only extinction model that can be matched with the spectra. It is possible to fit the LMC model instead, and in some cases this extinction model gives better fits (section 4.4).
- When LMC dust can be fitted, there is a tendency for larger values of τ_V. The SMC model can be successfully fitted to more objects than the LMC model, but gives a smaller fraction of significant dust measurements (section 4.4).
- There is no redshift or luminosity dependent evolution of the average amount of extinction, but there are however indications of an evolution in the type of dust, with LMC dust being relatively more common at higher redshifts.
- Krawczyk et al. (2013) show that quasar SEDs have a smaller spread when sampled in smaller ranges of luminosity. I have reproduced these results and shown that this is also valid in even smaller luminosity bins yet (section 4.2).

Comments on dust measurements

My results are somewhat limited by the quasars selected as the bluest objects of each bin, as these provide the available intrinsic powerlaw continua of the dust model fits. These quasars do not all have the same continuum slope, and because they are each selected from the quasars within a specific luminosity range, they are not equally separated in luminosity. The bluest spectra of two bins adjacent in luminosity can therefore be either nearly identical, or have both different powerlaw slopes and be separated by almost two bins of luminosity, or anything in between these two extremes. In a more optimal configuration the bluest spectra would be equidistant in luminosity.

All values of τ_V that I have measured are lower limits. I assume the bluest quasar of each (z, L) bin to be unaffected by dust, but I don't know if it is. If one of the bluest spectra is reddened by a small amount of dust, the intrinsic powerlaw continuum is bluer than the powerlaw I fit. It would then require more dust to dim the powerlaw when I use it in the following dust model fits. This is mostly a concern for the bins containing few objects, since there is a better chance of finding an unreddened quasar in a bin containing many objects.

The SDSS is biased against very dusty quasars
Strongly reddened quasars are important if we want to achieve a better understanding of the nature of dust in quasars, and could help clarify if quasar dust is mainly of the SMC type as usually assumed, or if some of them in fact have LMC type dust. For instance, a bump at 2175Å in the extinction curve would be easier to detect in a quasar with a lot of dust than in a weakly reddened quasar. The strongly reddened quasars propably represent only a small fraction of the whole quasar population, but they are to a high degree missed by large quasar surveys such as the SDSS that rely on optical color selection criteria alone (Francis et al. 2000; Fynbo et al. 2013; Krogager et al. 2015, 2016). Quasars with a lot of dust appear redder than most quasars. They are often classified as stars instead and therefore excluded from these surveys.

I also see in my sample that the SDSS do not contain many strongly reddened quasars (section 4.2). Figure 20 shows the slopes of all powerlaw continuum fits to the quasars in a single redshift range. There are a total of 379 objects in this redshift range, and only 6 of these have been fitted with a negative powerlaw slope, indicating strong reddening with the characteristic 'n'-shaped continua (Francis et al. 2000). Two of them are simply bad fits because of strong absorption in the spectra, and do in fact have a positive continuum slopes. That leaves only 4 out of 379 quasars ($\sim 1\%$) that show strong reddening by dust.

The 2175Å bump of the LMC dust model
Some features of a quasar spectrum are constrained, e.g. the ratio of the strength of some emission lines. But the strength of the iron bump emission in the area $2000\text{Å} \lesssim \lambda \lesssim 2400\text{Å}$ seems to vary at random from spectrum to spectrum, unrelated to any other features.

The LMC continuum fits behave different than the SMC fits because of the bump at 2175Å in the dust model, which partially hides in the iron emission (see section 4.4 and figure 25). This raises the question if the variations of the iron bump are caused by LMC dust, which would contribute with different amounts of extinction from one spectrum to another, and the iron bump itself is actually a constant feature. Maybe LMC dust in quasars has gone unnoticed because the variations of the Fe emission have been thought of as a difference in emission rather than a difference in absorption.

To investigate this further I would need spectra of a wider wavelength range. A spectrum that extends to longer wavelengths gives better constrains to the fit around the iron bump, and makes it possible to investigate the influence of LMC dust and the 2175Å bump on the iron emission more closely. For instance X-shooter[10] spectra could be useful, as it observes all the way from UV at 3000Å to the near infrared at close to 2.5μm. For a quasar at redshift $z = 2$ this corresponds to approximately $1000\text{Å} < \lambda < 8000\text{Å}$ in restframe wavelengths, and goes well past the iron bump.

X-shooter spectra could also be used to further investigate the trend that LMC dust is more common at higher redshifts (as also mentioned in section 4.1).

LMC dust as alternative to SMC dust
It is customary to assume SMC like dust when modeling the extinction of quasars. Studies such as Hopkins et al. (2004) show that SMC like dust is the most common type of dust in the SDSS quasars and for quasars in general. But there are also examples of

[10]https://www.eso.org/sci/facilities/paranal/instruments/xshooter.html

extinction curves that are either steeper or flatter than the SMC curve. For instance Fynbo et al. (2013) and Zafar et al. (2015) find extinction curves of quasars that are much steeper than the SMC curve in the UV. Other studies (Czerny et al. 2004; Gallerani et al. 2010; Gaskel et al. 2004; Maiolino et al. 2004) find extinction curves that are flat in the UV. The method used by Czerny et al. (2004) and Gaskel et al. (2004) is probably biased because their extinction curves are based on composite spectra, but Gallerani et al. (2010) and Maiolino et al. (2004) find UV flat extinction curves in high redshift quasars ($z \gtrsim 4$). The dust is thought to be produced by a completely different mechanism than for lower redshift galaxies though. These flatter extinction curves are therefore not necessarily comparable with the LMC curve, although LMC is also flatter than SMC in the UV, and I see indications that LMC may be more common at high redshifts in agreement with these objects. My point however, is that there are plenty of cases indicating different kinds of extinction than SMC. If we limit our studies of quasar extinction to SMC dust alone, we may get a distorted picture of the properties of quasars and of dust in our universe.

It is interesting that I can fit the LMC dust model to so many of the SDSS spectra. I will not attempt to list all of the implications LMC dust in quasars will have for our understanding of these objects, but to mention one, we would have to reconsider the dust production mechanism. The many significant detections of LMC dust in these quasars open the possibility that a lot of quasars can have a flatter extinction curve, and therefore a different dust composition than if SMC dust is assumed. The processes that produce the dust and influence the size distribution of the dust grains must then also be different.

The first part of section 4.1 and the numbers of table 4 show that the SMC dust model is easier to match with the spectra than the LMC model. This suggests that SMC dust is more common in quasars than LMC dust, as also concluded by Hopkins et al. (2004). But the SMC model is nearly a powerlaw function, and has no features like the 2175Å bump of the LMC model to put constraints on the fits. To add a small amount of SMC dust to a powerlaw is almost like adjusting the powerlaw slope, and as the spectra do not necessarily have the exact same intrinsic powerlaw continuum slope, this does not automatically mean that SMC dust is present. One of the reasons there are fewer matches with the LMC model could be that it has fewer false positives, and that some of the spectra fitted with the SMC dust model simply reflect a small difference in the slope of the intrinsic powerlaw continuum.

There is a tendency for more dust inferred by the LMC model fits than the SMC model fits. See equations 18 and 19. This too could indicate that the LMC fits find actual dust, while the SMC model also measures dust where there is merely a small difference in the continuum slopes and therefore give false measurements of dust.

My studies imply that the generally used SMC dust model is not the only dust model applicable for quasars. LMC type dust is also a possibility, and probably Milky Way type dust as well, since it is very similar to LMC dust. Alternative extinction models should therefore not be ruled out, just because the SMC model is the usually applied one. Maybe a mix of the SMC and LMC dust or an entirely different dust model is even better at describing quasar extinction. A quasar is a very different and more extreme environment than the Milky way or the SMC and LMC galaxies (Peterson 1997, p. 36-38), and there is no reason the properties of quasar dust cannot differ from those cases. There are indications of a steeper extinction curve towards the galactic center of the Milky Way (Fynbo et al. 2013; Sumi 2004). In fact the Milky Way, SMC and

LMC extinction curves usually refered to are only average curves, and all three of these galaxies have variations in their extinction curves for different lines of sight. So even within a single object, extinction curves can vary. As an alternative to the well known extinction curves of the Milky Way, SMC and LMC galaxies, Zafar et al. (2015) suggest a continuum of steeper than SMC extinction curves could apply for quasars. This can be extended to include flatter extinction curves as well.

Impact of dust on luminosity measurements and black hole mass estimates
Extinction can be measured as a difference in magnitude. I use this in section 2.8 to find a model for the UV continuum. If I combine equations 7 and 8, I have the following relation:

$$A_\lambda = m - m_0 = 1.086\tau_\lambda = 1.086\tau_V k_\lambda. \tag{20}$$

The average values of significant τ_V I have fitted for the SMC and LMC dust models are given in equations 18 and 19 respectively. The extinction curves are normalized to $k_\lambda = 1$ in the V band ($\lambda = 5500$Å). Using equation 20 the average fitted τ_V values then corresponds to a difference in magnitude in the V band of:

$$\text{SMC:} \qquad A_V = 1.086\langle\tau_V\rangle_{SMC} \cdot 1 = 0.13\text{mag} \tag{21}$$

$$\text{LMC:} \qquad A_V = 1.086\langle\tau_V\rangle_{LMC} \cdot 1 = 0.16\text{mag} \tag{22}$$

for the SMC and LMC models respectively.

When I combine equations 7 and 3, I can translate this magnitude difference into a difference in luminosity:

$$\begin{aligned} A_\lambda &= -2.5\log\left(\frac{L_\lambda/4\pi d_L^2}{L_{\lambda,0}/4\pi d_L^2}\right) \\ -\frac{1}{2.5}A_\lambda &= \log(L_\lambda/L_{\lambda,0}) \;=\; \log L_\lambda - \log L_{\lambda,0} \;=\; \Delta\log L_\lambda. \end{aligned} \tag{23}$$

Here L_λ is the luminosity at λ for the reddened quasar, and $L_{\lambda,0}$ is the luminosity of the same unextinguished source. Next, I plug in the extinctions of equations 21 and 22 to find the luminosity difference in the V band:

$$\text{SMC:} \qquad \Delta\log L_V = -0.052\text{dex} \tag{24}$$

$$\text{LMC:} \qquad \Delta\log L_V = -0.065\text{dex} \tag{25}$$

For the average fitted amounts of dust, the luminosities in the V band are underestimated with 11% and 14% for the SMC and LMC dust models respectively. For comparison, the average error of $L_{5100\text{Å}}$ (which is closest to the V band) in the SDSS DR7 catalogue as given by Shen et al. (2011) is only 0.044dex, and even smaller when low S/N objects are excluded. If we do not correct for extinction, we will underestimate the luminosity of many quasars. Dust is therefore an important factor in luminosity measurements, and the better our understanding of dust is, the better our capability to correct for it.

When the luminosities are underestimated it causes a systematic error when used to estimate other quasar properties, like for instance the mass of the central supermassive black hole. Vestergaard & Peterson (2006) present an equation that can be used to estimate black hole masses of quasars using optical or UV continuum luminosities of the sources. For the CIV emission line and UV continuum luminosity at $\lambda = 1350\text{Å}$ they give the following relation:

$$\log M_{BH}(\mathrm{C}_{IV}) = \log\left[\left(\frac{\mathrm{FWHM}(\mathrm{C}_{IV})}{1000\mathrm{km\ s}^{-1}}\right)^2\left(\frac{\lambda L_\lambda(1350\text{Å})}{10^{44}\mathrm{ergs\ s}^{-1}}\right)^{0.53}\right] + (6.66 \pm 0.01), \qquad (26)$$

where M_{BH} is the mass of the central black hole. With equations 23 and 26 I can calculate how large an impact dust reddening has on the mass estimates.

$$\begin{aligned}
\Delta\log M_{BH} &= \log M_{BH} - \log M_{BH,0} \\
&= \log\left[\left(\frac{\mathrm{FWHM}(\mathrm{C}_{IV})}{\mathrm{FWHM}(\mathrm{C}_{IV,0})}\right)^2\left(\frac{\lambda L_\lambda}{\lambda L_{\lambda,0}}\right)^{0.53}\right] + (6.66 \pm 0.01) - (6.66 \pm 0.01) \\
&= \log[(L_\lambda/L_{\lambda,0})^{0.53}] \\
&= 0.53 \cdot \log(L_\lambda/L_{\lambda,0}) \\
&= -\frac{0.53}{2.5}A_\lambda \qquad (27)
\end{aligned}$$

To go from line 2 to 3 in the above equations I assume that the extinction does not change the width of the CIV line. Equation 27 is valid for monochromatic luminosities measured at $\lambda = 1350\text{Å}$, so to use the equation I need to calculate A_λ at this wavelength. I use the table values of k_λ at the nearest wavelength, 1360Å (Gordon et al. 2003), and the average measured τ_V values to calculate A_λ from equation 20. The difference the extinction causes in the mass estimates is then:

$$\text{SMC:} \qquad \Delta\log M_{BH} = -0.15\mathrm{dex} \qquad (28)$$

$$\text{LMC:} \qquad \Delta\log M_{BH} = -0.11\mathrm{dex} \qquad (29)$$

The masses are underestimated a little. The above values correspond to an underestimation of 30% and 22% for SMC and LMC extinction respectively. But for comparison, Vestergaard & Peterson (2006) list the accuracy of the mass estimates of their mass scaling law (equation 26) to be at least 0.56dex (a factor of 3.6). This value is based on the uncertainty in the zeropoint of the mass scale itself (as determined by the reverberation mapping-based masses), which is of a factor 2.9, and on the scatter of the single-epoch mass measurements around the reverberation mapping masses to which the mass scaling law is calibrated. The average measured amount of dust is therefore not enough to make a noticable impact with the current accuracy of black hole mass estimates and measurements. It requires strong reddening with 4-5 times the average measured amount of dust to make the change in mass comparable to the error of the estimate. For that level of dust the quasar spectra will show clear signs of extinction. If the mass measurements of black holes should improve to an uncertainty of e.g. 50% or lower, it might however be a good idea to start considering the systematic underestimation due to dust extinction.

6 Conclusion

I have studied the extinction of intrinsic dust in quasars for a large sample of SDSS spectra, with the specific objective to find out if the extinction depends on redshift or luminosity of the source. To do so, I have divided my sample into a mesh of small bins for which the change in redshift and luminosity is very small, so that these properties are essentially identical for all objects in a bin. I then assume that the intrinsic SED of quasars is the same for objects in a bin, i.e. for quasars of the same luminosity and redshift.

I have fitted two different extinction models to the quasar continua to determine the degree of extinction, the SMC dust model and the LMC dust model. For the SMC model I am able to fit all of the spectra except for a few percent. This confirms that my assumption of intrinsically similar SEDs for quasars of same redshift and luminosity is realistic, and the diversity of observed quasar spectra can be ascribed to dust extinction.

Many authors argue that a SMC dust model is more appropriate for quasars. But I have produced acceptable continuum model fits with the LMC dust model as well for a substantial part of my sample, and with a higher rate of significant dust detections. This shows that SMC dust is not exclusively representative of quasar dust. It is possible to model the extinction of some quasars with LMC dust as well.

My sample of quasars covers the redshift range $1.90 < z < 2.35$. In this range I do not find any evolution in the amount of extinction as function of source luminosity or redshift for either dust model. But I do find indications that LMC dust is more common in higher redshift quasars, and it is therefore possible there is an evolution of the type of dust affecting quasars. This would necessarily be the subject of another study.

My results with LMC dust are intriguing but not without reservations. The wavelength range of the spectra in my sample makes it difficult to constrain the fit on the long wavelength side of the 2175Å bump of the LMC dust model. It would therefore be relevant to follow up on my results with spectra including longer wavelengths, for instance X-shooter spectra.

A Distribution of dust over redshift and luminosity bins in detail

Figures 26 and 27 show the average values and standard deviations of τ_V for individual (z, L) bins across redshift, for the LMC and SMC dust models respectively. The bins show greater variation than when added up across luminosity, but they are all still consistent with no evolution of dust for the redshift ranges investigated. The colors used correspond to the same luminosity ranges for both dust models. Bins containing less than five quasars are not shown in any of figures 26 to 29.

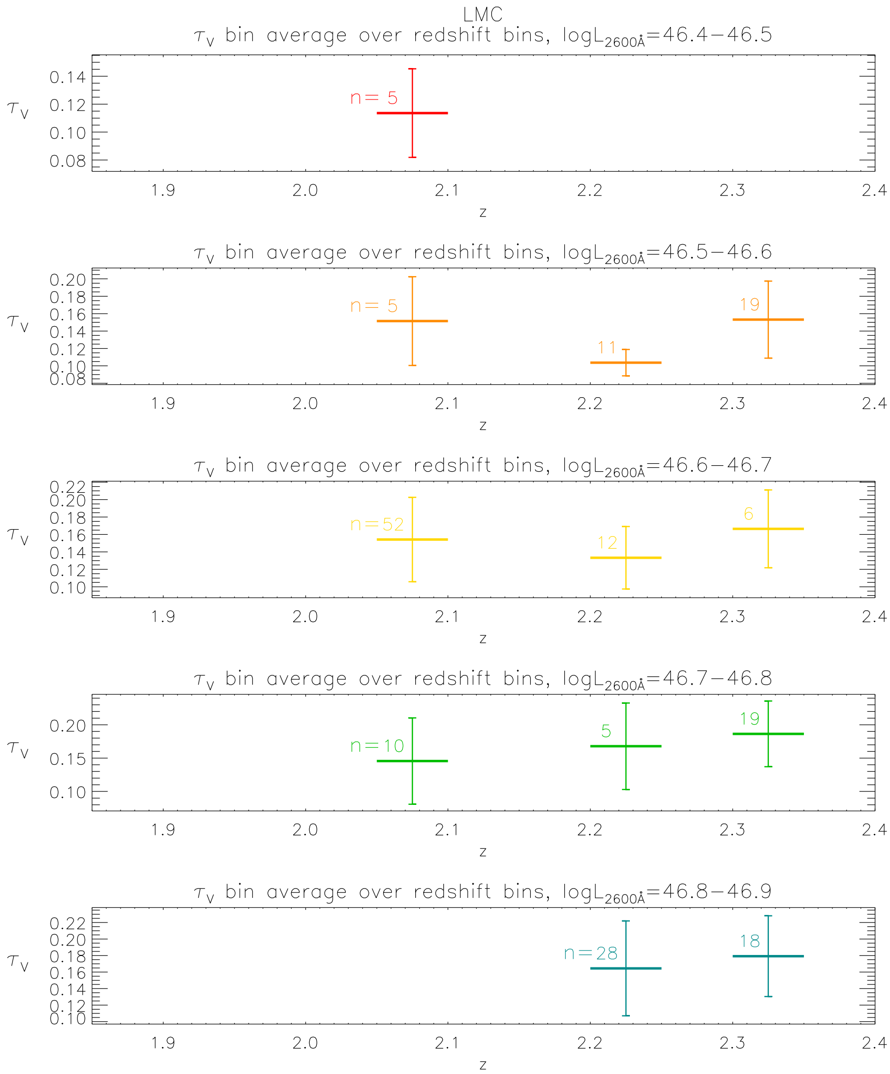

Figure 26: Mean values and standard deviations of significant τ_V measurements for individual (z, L) bins across redshift for the LMC dust model. The numbers indicate the number of significant dust measurements of each bin. Only bins with more than five objects are included. The measurements vary more than in figure 22, but are still consistent with no evolution across redshift.

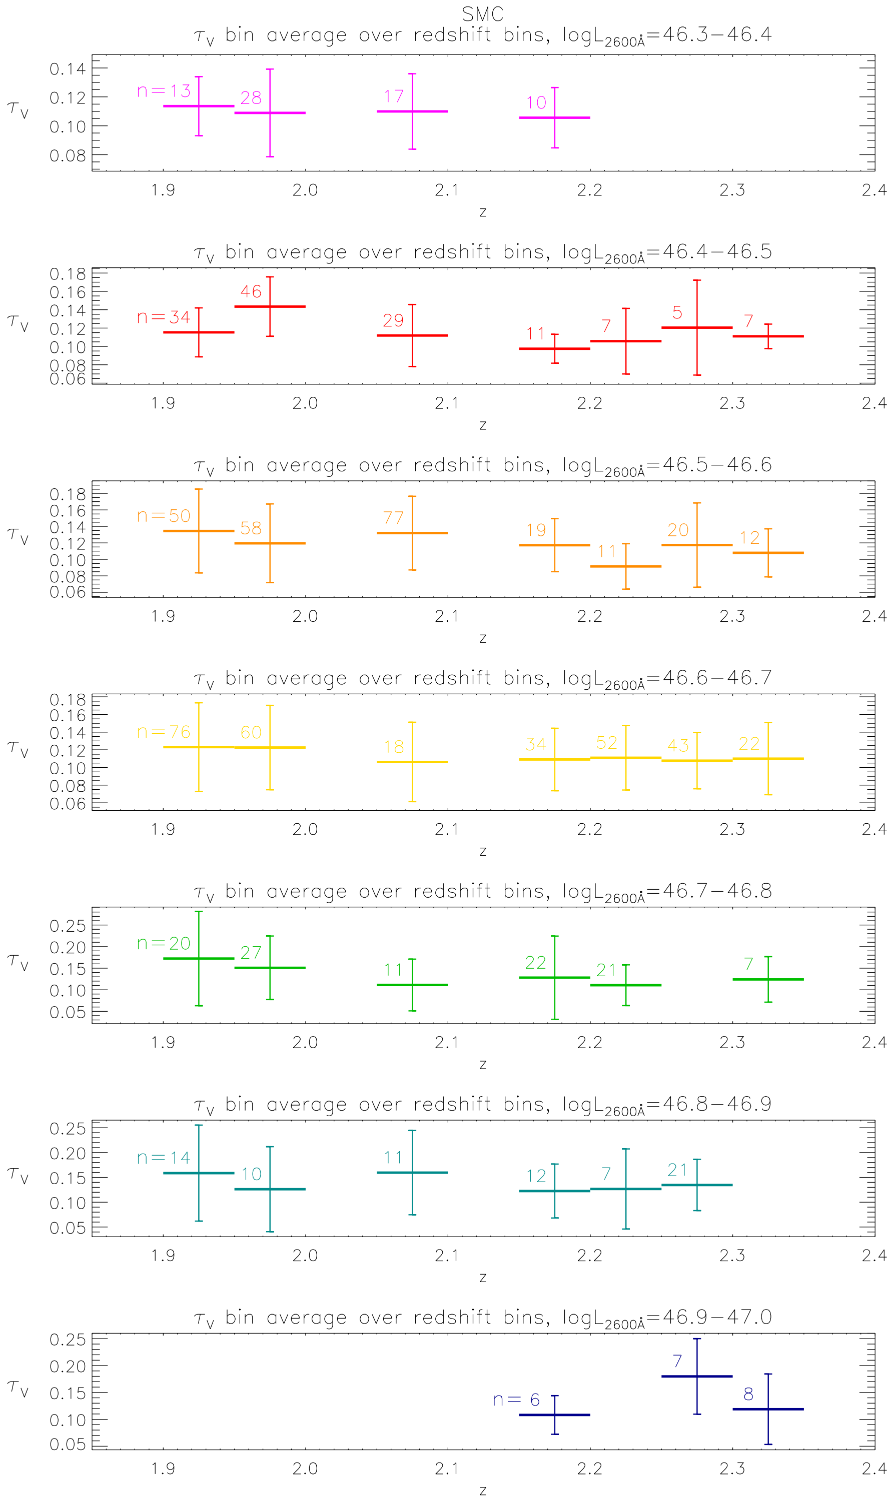

Figure 27: Mean values and standard deviations of significant τ_V measurements for individual (z, L) bins across redshift for the SMC dust model. The numbers indicate the number of significant dust measurements of each bin. Only bins with more than five objects are included. The measurements vary more than in figure 22, but are still consistent with no evolution across redshift.

Figures 28 and 29 also show average values and standard deviations of τ_V for the LMC and SMC dust models respectively, but now for individual (z, L) bins across luminosity. This distribution of bins also show greater variation than when added up across redshift, but is still consistent with no dependence on luminosity. The colors used correspond to the same redshift ranges in both figures.

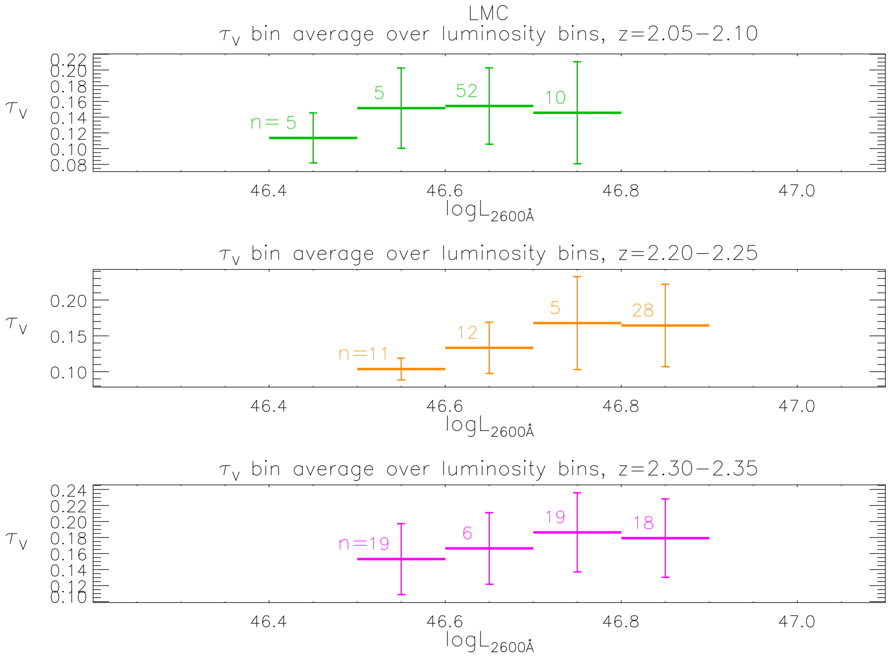

Figure 28: The mean values and standard deviations of significant τ_V measurements for individual (z, L) bins for the LMC dust model. Each panel shows the luminosity bins of a separate redshift range. The colors match with the redshift ranges shown for the corresponding figure for the SMC model. The numbers indicate the number of significant dust measurements of the bins. Only bins with more than five objects are included. The mean values are consistent with no evolution of τ_V across luminosity, though there is a weak tendensy for slightly larger amounts of dust at higher luminosities.

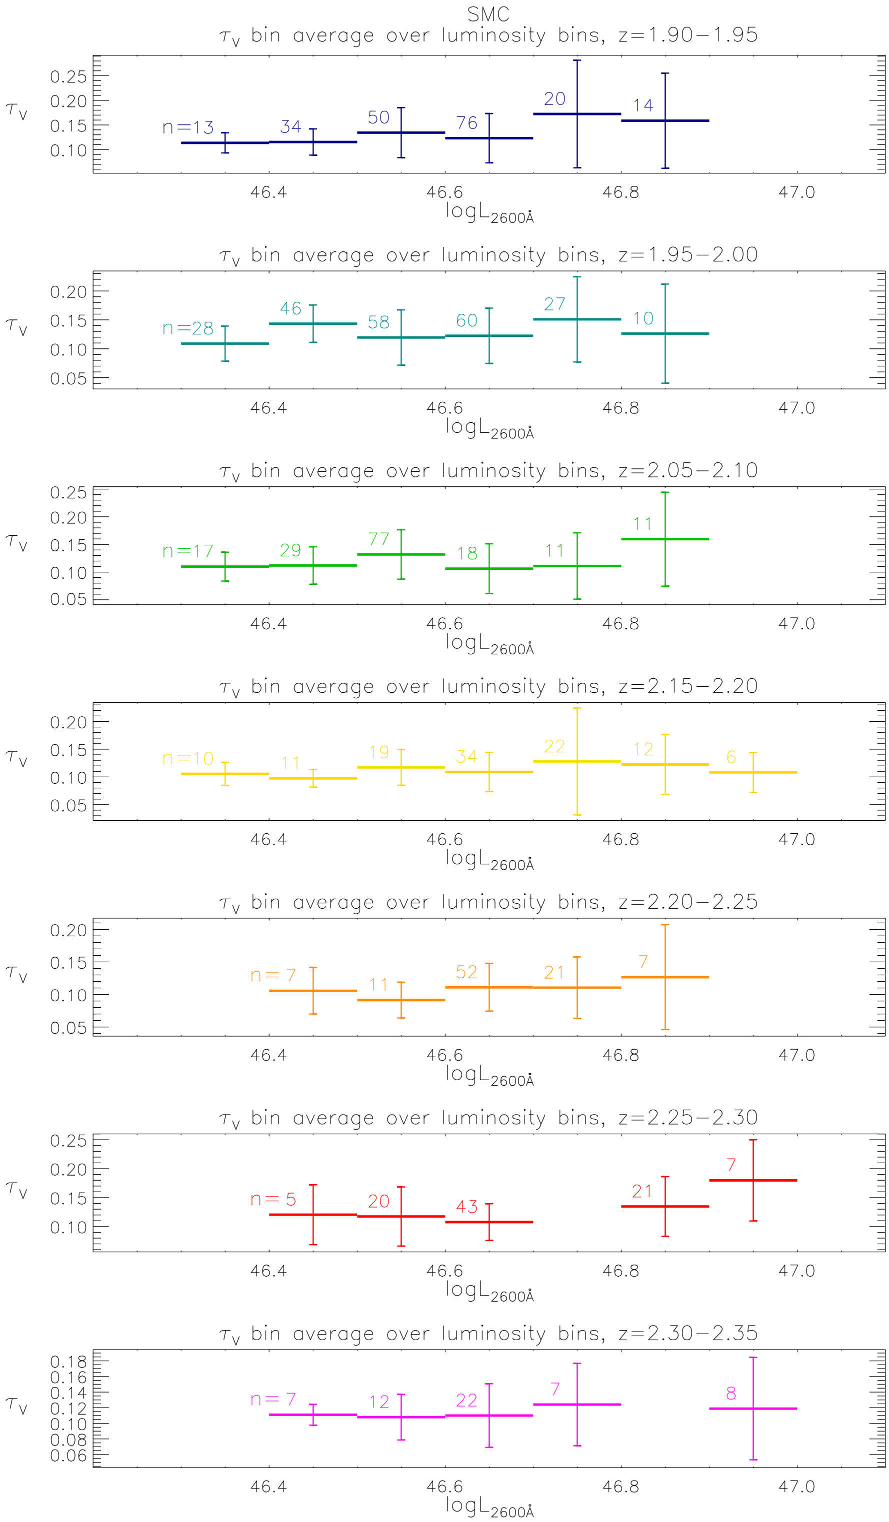

Figure 29: The mean values and standard deviations of significant τ_V measurements for individual (z, L) bins for the SMC dust model. Each panel shows the luminosity bins of a separate redshift range. The numbers indicate the number of significant dust measurements of the bins. Only bins with more than five objects are included. The mean values are consistent with no evolution of τ_V across luminosity, though there is a weak tendensy in most redshift ranges for slightly larger amounts of dust at higher luminosities.

References

Cardelli, J. A., Clayton, G. C., & Mathis, J. S., 1988, ApJ, 329, L33

Cardelli, J. A., Clayton, G. C., & Mathis, J. S., 1989, ApJ, 345, 245

Crenshaw, D. M., Kraemer, S. B., Bruhweiler, F. C., & Ruiz, J. R., 2001, ApJ, 555, 633

Czerny, B., Li, J., Loska, Z., & Szczerba, R., 2004, MNRAS, 348, L54

Francis, P. J., Whiting, M. T., & Webster, R. L., 2000, PASA, 17, 56

Fynbo, J. P. U., Krogager, J.-K., Venemans, B., Noterdaeme, P., Vestergaard, M., Møller, P., Ledoux, C., & Geier, S., 2013, ApJS, 204, 6

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

Hopkins, P. F., et al. 2004, AJ, 128, 1112

Hopkins, P. F., Bundy, K., Hernquist, L., & Ellis, R. S., 2007, ApJ, 659, 976

Kaspi, S., Brandt, W. N., Maoz, D., Netzer, H., Schneider, D. P., & Shemmer, O., 2007, ApJ 659, 997

Krawczyk, C. M., Richards, G. T., Mehta, S. S., Vogeley, M. S., Gallagher, S. C., Leighly, K. M., Ross, N. P., & Schneider, D. P., 2013, ApJS, 206, 4

Krogager, J.-K., et al. 2015, ApJS, 217, 5

Krogager, J.-K., Fynbo, J. P. U., Noterdaeme, P., Zafar, T., Møller, P., Ledoux, C., Krühler, T., & Stockton, A., 2016, MNRAS, 455, 2698

Li, A., 2007, in The Central Engine of Active Galactic Nuclei, eds. L. C. Ho, & J.-W. Wang, ASP Conf. Ser., 373, 561

MacAlpine, G. M., 1985, in Astrophysics of Active Galaxies and Quasi-Stellar Objects, ed. J. S. Miller (University Science Books), 259

Maiolino, R., Schneider, R., Oliva, E., Bianchi, S., Ferrara, A., Mannucci, F., Pedani, M., & Roca Sogorb, M., 2004, Nature, 431, 533

Maoz, D., et al., 1993, ApJ, 404, 576

Momjian, E., Carilli, C. L., Walter, F., & Venemans, B., 2014, AJ, 147, 6

Natali, F., Giallongo, E., Christiani, S., & La Franca, F., 1998, AJ, 115, 397

Netzer, H., & Wills, B. J., 1983, ApJ, 275, 445

O'Donnell, J. E., 1994, ApJ, 422, 158

Pei, Y. C., 1992, ApJ, 395, 130

Peterson, B. M., 1997, *An introduction to active galactic nuclei*, Cambridge University Press

Rosswog, S., & Brüggen, M., 2007, *Introduction to High-Energy Astrophysics*, Cambridge University Press

Ryden, B., 2003, *Introduction to cosmology*, Pearson Addison-Wesley

Schneider, D. P., et al. 2010, AJ, 139, 2360

Shen, Y., et al. 2011, ApJS, 194, 45

Sparke, L. S., & Gallagher, J. S., 2007, *Galaxies in the Universe: An Introduction*, Second Edition, Cambridge University Press

Sumi, T., 2004, MNRAS, 349, 193

Taylor, John R., 1997, *An Introduction to Error Analysis*, Second Edition, University Science Books

Vanden Berk, D. E., et al. 2001, AJ, 122, 549

Verner, E., Bruhweiler, F., Verner, D., Johansson, S., & Gull, T., 2003, ApJ, 592, L59

Vestergaard, M., & Peterson, B. M., 2006, ApJ, 641, 689

Vestergaard, M., & Wilkes, B. J., 2001, ApJS, 134, 1

Watson, D., Denney, K. D., Vestergaard, M., & Davis, T. M., 2011, ApJ, 740, L49

Willott, C. J., 2005, ApJ, 627, L101

Zafar, T., et al. 2015, A&A, 584, A100