Characterization of Newly Discovered Gravitational Lens Systems

Felipe Ávila Vera

Supervisor: Dr. Verónica Motta (UV)
Co-Supervisor: Dr. Paula Sánchez Sáez (ESO)
Internal reviewer: Dr. Patricia Arévalo (UV)
External reviewer: Dr. Evencio Mediavilla (IAC)



Tesis para optar al grado de Doctorado en Astrofísica
Instituto de Física y Astronomía
Facultad de Ciencias
Universidad de Valparaíso

21 de Agosto 2026
Valparaíso, Chile.

Dedico este trabajo a todas las personas que me han ayudado a llegar hasta este punto y a quienes ya no están, por sus consejos, por estar y simplemente por todos aquellos momentos que han contribuido a construir lo que soy hoy.
En especial, a mi madre, Silvana, sin quien nada de esto sería posible, a mi hermano, Sergio, a mi pareja Alejandra, a mis tíos, Sylvia y Lautaro y a mi abuela, Silvia, quien, de una u otra forma, nos mantiene a todos en pie. También quiero hacer una mención especial a mi perrito Rex, que me ha acompañado desde el comienzo de este camino.

This thesis is solely my own composition,
except where specifically indicated in the text.

Total or partial reproduction, for scientific or academic purposes,
is authorised including a bibliographic reference to this documment.



Felipe Ávila Vera
21 de Agosto 2026
Valparaíso, Chile.

Acknowledgements

First, I would like to thank my supervisor, Dr. Verónica Motta, for her constant support and guidance, and for always helping me remain focused throughout this journey, from the beginning of my master's degree to this point. I would also like to thank my co-supervisor, Dr. Paula Sánchez-Sáez, for encouraging me to pursue my own ideas.

I am also grateful to Dr. Evencio Mediavilla and Dr. Patricia Arévalo for their thoughtful comments throughout the development of this thesis.

I gratefully acknowledge the financial support provided by the FIB-UV Doctoral Fellowship Program, the ESO Early-Career Visitor Programme, the ESO Science Support Discretionary Fund through project No. 4/25-G, and ANID FONDECYT Regular Grant No. 1231418.

Finally, I would like to thank Edgar Sanders for his invaluable computational assistance.

Abstract

Gravitationally lensed quasars have proven to be one of the most promising tools in astrophysics. They provide a unique opportunity to study the mass distribution of lens galaxies and to probe the inner structure of quasars through microlensing events. However, the frequency of such events has not yet been well constrained. Furthermore, quasar optical spectra can themselves be affected by microlensing, making it necessary to develop new tools capable of analyzing their multiple spectral components and disentangling the different contributions present in the observed spectra. Both problems become increasingly relevant in the era of next-generation surveys, as the large volume of incoming data will put current methodologies to the test.

In this work, we investigate how frequently quasars are affected by gravitational microlensing events using magnification maps, which allow us to account for the granularity of the stellar population in the lens galaxy. We also incorporate corrected accretion-disk sizes in order to address the discrepancy between theoretical expectations and observational results. In addition, we study the probability of high-magnification events, which can provide insight into the peculiar velocities of the objects involved in these phenomena.

Finally, we present a new methodology for analyzing the thousands of quasar spectra that will be delivered by upcoming instruments, capable of separating their different spectral components. To achieve this, we take advantage of modern graphics processing unit (GPU) techniques to obtain fast and reliable results. We validate our results through comparisons with well-established methodologies, demonstrating that GPUs represent a natural next step in astronomical computation and in the analysis of lensed quasars.

Contents

1 Introduction **11**

1.1 Gravitational lensing 12

1.1.1 The deflection angle 16

1.1.2 The lens equation 18

1.1.3 Cosmological distances 20

1.1.4 Magnification and distortion 20

1.1.5 Critical curves and caustics 23

1.1.6 Simple mass distributions 23

1.1.7 Necessary and sufficient conditions for multiple imaging 28

1.2 Active Galactic Nuclei 30

1.2.1 The AGN structure 31

1.2.2 AGN classification 34

1.3 Quasar microlensing 39

1.3.1 Magnification maps 40

1.3.2 Microlensing as a probe of AGN structure 41

1.3.3 High-magnification events 44

1.4 Motivation 48

2 Microlensing time-scales and flux magnification probabilities of a sample of 204 lensed quasars **53**

2.1 Data 55

2.2 Methods 56

2.2.1 Macro lens modeling 56

2.2.2 Microlensing timescales 59

2.2.3 Microlensing magnification maps 63

2.3 Results 66

2.3.1 Homogeneous estimate of the Einstein radius 66

2.3.2 Microlensing timescales . . . 67
2.3.3 Microlensing magnification statistics . . . 68
2.4 Discussion . . . 75
2.4.1 Average size of the quasar source and properties of the microlenses population. . . . 75
2.4.2 Magnification statistics and flux-ratio anomalies. . . . 78
2.4.3 Microlensing timescales and frequency of HME. . . . 78
2.5 Conclusions . . . 80
2.6 Data availability . . . 82

3 Spectral Handling and Estimation of AGN Parameters: The first AGN fitting GPU-based code **83**

3.1 The optical ultraviolet spectra of AGN . . . 84
3.2 Modeling the optical spectra . . . 86
3.3 Methods . . . 87
3.3.1 Data pre-processing . . . 88
3.3.2 Modeling with `SHEAP` . . . 89
3.3.3 Optimization Strategy . . . 93
3.3.4 Uncertainty Estimation . . . 95
3.3.5 Post-fit derived quantities . . . 96
3.4 Performance validation of `SHEAP` . . . 97
3.4.1 The Wu&Shen22 sample . . . 98
3.4.2 The Pan+25 sample . . . 99
3.4.3 The Sánchez-Sáez+18 sample . . . 103
3.4.4 The Bernal+25 sample . . . 110
3.5 Computational time . . . 111
3.6 Discussion & Conclusions . . . 114
3.7 Data availability . . . 116

4 Conclusions **117**

5 Future work **123**

A Kinematic region limits **127**

B Broad Component Combination Methods **129**

C Emission lines 131

D Emission lines fitting examples 133

List of Figures

1.1 Composite image of QSO 0957+561 constructed from observations in the V and I filters. The image shows the field of view of a single HST WFPC2 wide-field CCD centered on the lensed quasar images. 14

1.2 Sketch of a typical gravitational lens system. A background source is deflected by the gravitational potential of a foreground lens, producing one or more images on the sky. The angles $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ denote the angular positions of the unlensed source and the observed image, respectively, as seen by the observer. The angle $\boldsymbol{\alpha}$ is the physical deflection angle produced at the lens plane, while the reduced deflection angle is $\hat{\boldsymbol{\alpha}} = (D_S/D_{LS})\boldsymbol{\alpha}$. The corresponding physical positions in the source and lens planes are $\boldsymbol{\eta}$ and $\boldsymbol{\xi}$, and D_L, D_S, and D_{LS} are the angular-diameter distances between observer and lens, observer and source, and lens and source, respectively. 15

1.3 Euclid imaging of Altieri's lens, an example of an Einstein ring (O'Riordan et al., 2025). The main panel shows a composite false-colour image constructed from Euclid VIS and NISP observations. The VIS instrument provides high-resolution optical imaging, while NISP provides lower-resolution near-infrared photometry in the Y_E, J_E, and H_E bands. In this composite, the broadband VIS I_E image defines the brightness structure, and the NISP bands provide the colour information. The central light of the lens galaxy has been suppressed to enhance the visibility of the lensed arc. The inset shows the central 8 arcsec using only the higher-resolution VIS data, corresponding to the square marked in the main panel. The angular scale and the physical scale at the lens redshift are shown in each frame. 17

1.4 Distortions produced by the convergence κ and the two components of the shear, γ_1 and γ_2. The original circular source is shown in green, while the corresponding lensed image is shown in black. The convergence term produces an isotropic change in size, with $\kappa > 0$ magnifying the image and $\kappa < 0$ demagnifying it. The shear components produce anisotropic distortions: γ_1 stretches the image along the horizontal or vertical axis, while γ_2 stretches it along axes rotated by 45°. Adapted from Wikipedia; see also Schneider et al. (1992); Meylan et al. (2006). . . 21

1.5 Examples of image configurations produced by simple lens mass models. The upper panels show a singular isothermal sphere (SIS), while the bottom panels show a singular isothermal ellipsoid (SIE). In each case, different source positions are shown relative to the caustics in the source plane, together with the corresponding image positions in the image plane. The comparison shows how changing the symmetry of the mass distribution changes the caustic structure, and the resulting number and arrangement of images. 24

1.6 Classical schematic representation of the main structural components of an AGN in the center of the host galaxy (figure not at scale). The central SMBH and accretion disk are surrounded by the broad-line region, dusty circumlunar torus, the narrow-line region, and, in some systems, a relativistic jet. Obtain from (Thorne et al., 2022) 29

1.7 Top: Schematic view of the main emitting regions of a quasar, with approximate characteristic sizes in cm. Bottom: Optical AGN spectrum showing how the continuum, broad emission lines, and narrow emission lines trace the accretion disk, BLR, and NLR, respectively. Compact regions, such as the X-ray corona and the optical/UV accretion disk, are the most sensitive to microlensing, whereas extended components, such as the dusty torus and the narrow-line region, are much less affected and can be used as reference emission. From Vernardos et al. (2024). 32

1.8 Example optical SDSS spectra of type-I and type-II AGN. The type-I spectrum, shown in blue, displays broad permitted emission lines produced in the broad-line region, while the type-II spectrum, shown in red, is dominated by narrow emission lines because the direct view of the central continuum source and broad-line region is obscured. 34

1.9 Example of an optical SDSS spectrum host-dominated AGN in the rest-frame. In this regime, stellar continuum and absorption features from the host galaxy can dominate part of the spectrum, making the nuclear AGN contribution more difficult to isolate. 36

1.10 Example of a microlensing magnification map computed for $\kappa = \gamma = 0.3$, assuming a mean microlens mass of $\langle M_\star \rangle = 0.3\,M_\odot$. The map covers a region of $15\,R_E$ on each side and has a size of 2000×2000 pixels with smooth matter fraction = 0 (100% stars). It shows the caustic network produced by compact objects in the lens galaxy and represents the magnification pattern experienced by a background source in the source plane. Compact sources are sensitive to individual caustics and can therefore show large magnification fluctuations, whereas extended sources average over the caustic pattern and produce smoother microlensing signals. Colors indicate the microlensing magnification, with red regions corresponding to the highest magnifications. 37

1.11 HST/NICMOS F160W image of the doubly imaged lensed quasar QJ0158−4325 (Faure et al., 2009), together with the corresponding microlensing magnification maps for images A and B. Each magnification map has a size of 1300×1300 pixels and spans $30\,R_E$ in the source plane. The maps were generated assuming a smooth-matter fraction of $s = 0.2$. For image A, the local convergence and shear are $\kappa = 0.38$ and $\gamma = 0.38$, while for image B they are $\kappa = 0.76$ and $\gamma = 0.76$, respectively. 38

1.12 Comparison between the original and convolved microlensing magnification maps. Left: original microlensing magnification map, computed for $\kappa = \gamma = 0.3$ and a mean microlens mass of $\langle M_\star \rangle = 0.3\,M_\odot$. The map is the same as in Figure 1.10, spanning $15\,R_E$ with 2000×2000 pixels. Center: probability distribution of the magnification values before and after convolution, showing how finite-source smoothing reduces the contrast of the magnification map. Right: map convolved with a finite source profile of characteristic size $R_S = 0.2\,R_E$, corresponding to approximately 0.8 light-days for $R_E \simeq 10^{16}$ cm. This size is comparable to the expected scale of the optical accretion disk. The convolution smooths the small-scale caustic structure because the finite source averages over the magnification pattern. 39

1.13 Same microlensing magnification map as in Figure 1.10, computed for $\kappa = \gamma = 0.3$ and a mean microlens mass of $\langle M_\star \rangle = 0.3, M_\odot$, with an example source trajectory overlaid. The right panel shows a zoomed-in view of the caustic-crossing event highlighted in grey, where the source passes through a high-magnification region and produces a clear increase in the microlensing signal. 45

1.14 Same microlensing magnification map as in Figure 1.10, computed for $\kappa = \gamma = 0.3$ and a mean microlens mass of $\langle M_\star \rangle = 0.3\, M_\odot$, with an example source trajectory overlaid. The track illustrates how a source moving across the magnification pattern samples different regions of the caustic network, producing a microlensing light curve whose variability depends on the source size and on the caustics crossed along the path. Compact emitting regions, such as the accretion disk, are more strongly affected by individual caustics and can experience larger magnification fluctuations, whereas more extended regions average over the caustic pattern and produce smoother, lower-amplitude microlensing signals. 46

1.15 Global comparison of the three HMEs and analytical fit to a single caustic crossing in Q2237+0305. The inset is a zoom in the region of the maximum. Dots correspond to the experimental light curves of each of the three events. From Mediavilla et al. (2015) 47

2.1 Astrometric error distribution for the sample of images modeled. 54

2.2 Microlensing magnification patterns produced by stars in the lensing galaxy corresponding to image A of the system DES2158−5812. *Left:* Original magnification map. The color scale encodes different magnification levels, while the diamond-shaped curves trace the caustics. *Right:* Same map after convolution with a Gaussian profile with a sigma size of 3.72 [pix], showing the effective magnification pattern. Both panels have a size of 1200 × 1200 pixels with a pixel size of 1.1 [pix/light-days] and 60 R_E. The maps were produced using the code developed in Jiménez-Vicente & Mediavilla (2022). 56

2.3 Comparison of θ_E in this work vs. values in the literature for modeled lens quasars. The systems that are clearly far from a 1:1 relation are SDSS 1155+6346 and HS 0818+1227. 62

2.4 Values of κ and γ for our sample (red filled diamonds), compared with literature samples: unfilled black diamonds Neira et al. (2025), unfilled circles Ertl et al. (2023), and unfilled black squares Schmidt et al. (2023). . 64

2.5 Histogram of the Einstein radii at the source plane for the systems. Blue shows the distribution of all the objects. Grey shows the distribution for the systems with the redshift of the lens known 65

2.6 Histogram of the S&S based sizes of the source, R_S, before and after applying a microlensing size correction (204 systems). 68

2.7 Histogram of the distribution of the effective velocities (Eq. 2.10). 69

2.8 *Top*: Distribution of Einstein timescales. *Bottom*: Distribution of source timescales. 70

2.9 Probability distribution of unsigned microlensing magnification for three different scaling factor values for 0.3 $\times$ $\{r_s\}$, 1.0 $\times$ $\{r_s\}$, and 2.0 $\times$ $\{r_s\}$, with two different stellar densities (α). 72

2.10 Histogram of unsigned microlensing magnification for three different scaling factors values for r_s, with two different stellar densities (α), comparing the results from Mediavilla et al. (2024) and this work, with all histograms normalized to the same total count. 73

2.11 Relative probability distribution ($\propto e^{-\chi^2/2}$) in the (half light radius [light days], fraction of total mass in stars) parameter space. The red contour marks the 68% confidence region, enclosing the highest-probability values. 73

2.12 Cumulative distribution functions of magnitude changes during microlensing events. *Left column*: Probability of observing a demagnification event with $\Delta\text{mag} > 0$. *Right column*: Magnification event with $\Delta\text{mag} < 0$. *Top row*: Stellar fraction of $\alpha = 0.1$. *Bottom row*: $\alpha = 0.2$. . . . 74

2.13 *Top*: Histogram of the probabilities of obtaining a magnification measurement of less than $\Delta\text{mag} = -0.32$ for different stellar densities and values of r_s. *Bottom*: Density histogram of the probabilities of obtaining a demagnification greater than $\Delta\text{mag} = 0.32$ for different stellar densities and values of r_s. 74

2.14 Same convolved microlensing magnification map as in Fig. 2.2 (*right*), With a zoom-in between pixel 0 and 400 in the Y axis and full length (1200) in the X axis. Here, the magnification map is masked for values over -0.32, highlighting the regions of strongest magnification within the caustic network. 75

2.15 *Top*: Probability of $\Delta\text{mag} \leq -0.32$ versus Einstein radius crossing time (t_E [years]), with blue (orange) representing values obtained for a stellar density of 0.1 (0.2) and a $1.0 \times \{r_s\}$, the red stars represent the object that shows the smaller Einstein radius crossing time and high magnification probability that correspond to DESJ2038-4008 and Q2237+030. *Bottom*: Probability or $\Delta\text{mag} \leq -0.32$ versus source crossing time (t_S [years]), with blue (orange) representing values obtained for a stellar density of 0.1 (0.2) and a $R_S \times 1$, in black stars the objects that show the smallest source crossing time and hight probability that are CY2201-3201, WG2100-4452, SDSSJ1138+0314, SDSSJ0821+4542, J0343-2828, WG0214-2105, and Q2237+030. 76

2.16 Correlation between the values of κ and γ in our sample, with the color bar representing the HME probability. The panels are separated by α and by multiples of r_s. 79

3.1 Multiple spectra observed with SDSS in rest frame 84

3.2 Left: Redshift distribution by sample. Right: Mean signal-to-noise ratio (S/N) distribution by sample and the distribution of S/N from DR16 in black. Sample names are shown at the top of the figure, with the number of objects in parentheses, ... 95

3.3 Right: Distribution of reduced χ^2 by sample, normalized to the fraction of objects in each sample. The vertical step line marks $\chi^2_{\text{red}} = 1$. Left: Relation between signal-to-noise ratio (S/N) and the reduced chi-square, χ^2_{red}. Each color represent a different sample. 97

3.4 Comparison between our measurements and those of Wu&Shen22. Top: C IV emission-line FWHM in logarithmic scale. Bottom: Monochromatic continuum luminosity at 1350 Å in logarithmic scale. In both panels, the black dashed line indicates the 1:1 relation, while the gray dashed lines mark the ± 0.3 dex region. 100

3.5 Distributions of the logarithmic differences, $\log_{10}(X_{\rm SHEAP}) - \log_{10}(X_{\rm lit})$, for the set of spectral parameters compared in each sample. Panel A shows against Wu&Shen22, while Panel B shows against Pan et al. (2025). In each panel, the legend reports the compared parameters. The grey dashed line indicates the 1:1 relation, while the black dashed-lines demark the $\pm$0.3 dex. 101

3.6 Differences between our measurements and those from the literature as a function of signal-to-noise ratio (S/N; x-axis). The y-axis shows $\log_{10}(X_{\rm SHEAP}) - \log_{10}(X_{\rm lit})$. Panel A against Wu&Shen22, and Panel B against Pan+25. The legend identifies each parameter, and the horizontal dashed lines mark the $\pm$0.3 dex band. 102

3.7 Comparison between our measurements and those of Pan+25. Top: Mg II emission-line FWHM in logarithmic scale. Bottom: Monochromatic continuum luminosity at 3000 Å in logarithmic scale. In both panels, the black dashed line indicates the 1:1 relation, while the gray lines mark the $\pm$0.3 dex region. Red circles represent values computed with `FSF`, while black squares correspond to `DAS`. 104

3.8 Comparison between our results and Pan et al. (2025) for the $R_{\rm FeII}$ in logarithmic scale. The black dashed line indicates the 1:1 relation, while the grey lines demark the $\pm$0.3 dex. 105

3.9 Comparison between our measurements and those of Sánchez-Sáez+18. Top: comparison of the emission-line FWHM in logarithmic scale. Bottom: comparison of the monochromatic continuum luminosity in logarithmic scale. In both panels, the black dashed line indicates the 1:1 relation, while the gray lines mark the $\pm$0.3 dex region. In the top panel, black circles correspond to Hβ measurements, while red squares represent Mg II. In the bottom panel, black circles correspond to L_{5100} measurements, while red squares represent L_{3000}. 107

3.10 Same as Fig. 3.5 but this time Panel A shows the comparison against Sánchez-Sáez+18, meanwhile Panel B shows the comparison against Bernal+25. 108

3.11 Same as Fig. 3.6 but this time Panel A compares against Sánchez-Sáez+18, and Panel B compares against Bernal+25. 109

3.12 Comparison between the stellar-to-continuum flux ratio at 5100 Å derived by Bernal+25 and the values obtained in this work. Red points show all objects in the sample, while blue points highlight sources with a confirmed broad emission-line component. The dashed line marks the 1:1 relation, and the shaded region indicates the ±0.3 dex agreement band. 111
3.13 Comparison between our measurements and those of Bernal+25. Top: Hα emission-line FWHM in logarithmic scale. Bottom: Hα emission-line luminosity in logarithmic scale. In both panels, the black dashed line indicates the 1:1 relation, while the gray lines mark the ±0.3 dex region. Point colors represent the EW reported by Bernal+25. 112
3.14 Comparison between our results and Bernal+25 for the monochromatic continuum luminosity at 5100 Å in log scale.The black dashed line indicates the 1:1 relation, while the grey lines demark the ±0.3 dex. 113

D.1 Example of fited spectra for the Wu&Shen sample for the object with median SN 133
D.2 Example of fited spectra for the Wu&Shen sample for the object with median χ^2_{reduced} 134
D.3 Example of fited spectra for the Pan+25 sample for the object with median SN 134
D.4 Example of fited spectra for the Pan+25 sample for the object with median χ^2_{reduced} 134
D.5 Example of fited spectra for the Sánchez-Sáez+18 sample for the object with median SN 135
D.6 Example of fited spectra for the Sánchez-Sáez+18 sample for the object with median χ^2_{reduced} 135
D.7 Example of fited spectra for the Bernal+25 sample for the object with median SN 135
D.8 Example of fited spectra for the Bernal+25 sample for the object with median χ^2_{reduced} 136

List of Tables

2.1 Mean microlensing statistics for different stellar mass fractions and source-size scaling factors. 71

3.1 Summary of fit quality, model complexity, and computational performance for the literature comparison samples. 113

3.2 $\widetilde{\chi}^2_{\rm red}$ denotes the median reduced chi-square, reported together with the 16th and 84th percentiles as asymmetric uncertainties. For Wu&Shen+22, we report the unmasked value, with the value obtained after masking $\lambda < 1300\,\text{Å}$ given in parentheses. $\widetilde{\rm S/N}$ is the median signal-to-noise ratio of the sample. $N_{\rm par}$ is the number of free parameters used in the fit, corresponding to the length of the optimized parameter vector. N_λ denotes the number of pixels per spectrum. $t_{\rm fit}$ is the runtime for a single best-fit optimization per spectrum, and $t_{\rm MC}$ is the runtime for the Monte Carlo posterior stage per spectrum. 113

A.1 `SHEAP` kinematic region limits. 128

A.2 Columns list the region name, the allowed FWHM range, the maximum absolute velocity offset, and the references used to define or motivate these limits. 128

C.1 Emission lines used by the sample. Counts: Wu&Shen22 (broad=3); Pan+25 (broad=1); Sánchez-Sáez+18 (broad=2, narrow=1, outflow=1); Bernal+25 (broad=2, narrow=1). 132

Chapter 1

Introduction

Gravitational lensing is a direct consequence of general relativity (Einstein, 1936) and has become one of the most powerful tools in modern astrophysics and cosmology (e.g. Schneider et al., 1992; Treu, 2010; Birrer et al., 2024). It enables the detection and study of exoplanets through microlensing (Bond et al., 2004; Gaudi, 2010; Rektsini & Batista, 2024); the observation of some of the most distant objects in the Universe through natural magnification by massive foreground structures (Mahler et al., 2023; Bergamini et al., 2023); the investigation of the distribution and nature of dark matter (Massey et al., 2010; Vegetti et al., 2012, 2024); the study of the mass distribution, formation, and evolution of galaxies (Treu & Koopmans, 2004; Koopmans et al., 2006; Treu, 2010; Knabel et al., 2025); the placement of constraints on the stellar initial mass function (Ferreras et al., 2010; Leier et al., 2016; Jiménez-Vicente & Mediavilla, 2019); and the independent determination of cosmological parameters, especially the Hubble constant, through time-delay cosmography (Wong et al., 2020; Birrer et al., 2020, 2024; Tdcosmo Collaboration et al., 2025). The latter is particularly relevant in the context of the current tension between early- and late-Universe measurements of the Hubble constant (Hu & Wang, 2023; Birrer et al., 2024; Di Valentino et al., 2025).

In addition, gravitational lensing provides a unique way to study the inner regions of active galactic nuclei (AGN), which are powered by accretion onto supermassive black holes located at the centers of galaxies. Among the AGN population, quasars (QSO) represent the most luminous and powerful class, making them ideal targets for

studying accretion physics and the structure of compact emitting regions at cosmological distances (Khachikian & Weedman, 1974; Bianchi et al., 2012; Padovani, 2017). In particular, microlensing analyses of lensed quasars make it possible to probe spatial scales that cannot be directly resolved with current observational techniques (Guerras et al., 2013a; Jiménez-Vicente et al., 2014; Fian et al., 2021; Vernardos et al., 2024; Fian et al., 2024; Mediavilla et al., 2026). These unresolved regions, including the accretion disk and the broad-line region, remain under active discussion (Popović, 2020; Vernardos et al., 2024). By combining strong lensing, microlensing, and spectral analysis, one can access spatial scales that are otherwise unreachable with present-day instruments. This combination is the central focus of this work.

This chapter is organized as follows. In Section 1.1, we introduce the basic formalism of gravitational lensing, from the deflection of light to the main observables produced in the strong-lensing regime. In Section 1.2, we introduce AGN and quasars, describing their main spectral components and the physical regions responsible for their emission. Finally, in Section 1.3, we discuss microlensing in lensed quasars and how it can be used to study the structure of unresolved AGN emitting regions.

1.1 Gravitational lensing

After centuries during which gravitational theory was dominated by the Newtonian description, the idea that gravity could affect the propagation of light was proposed within a classical framework. In his 1804 paper, "On the Deflection of a Light Ray from Its Rectilinear Motion by the Attraction of a Celestial Body Which It Nearly Passes," Johann Georg von Soldner used Newtonian mechanics to calculate the deflection of a light ray caused by a massive object. In this approximation, a light ray (that possess mater properties) passing a point mass (M) with impact parameter (ξ) is deflected by an angle

$$\alpha_{\rm N} = \frac{2GM}{c^2\xi}. \tag{1.1}$$

This result, however, is only half of the value predicted by General Relativity. The development of Einstein's theory revolutionized the understanding of gravity by replacing the concept of a force acting at a distance with the curvature of space–time. In General Relativity, matter and energy determine the geometry of space–time, and light propagates along null geodesics of this curved metric. Therefore, the trajectory

of a light ray is modified when it passes near a massive object. For a point mass, the relativistic deflection angle is

$$\alpha_{\rm GR} = \frac{4GM}{c^2\xi}, \tag{1.2}$$

which differs from the Newtonian prediction by a factor of two. This distinction was crucial because it provided one clear observational tests of General Relativity (Misner et al., 1973; Straumann, 2013).

The deflection of light by gravity is one of the most direct observational consequences of General Relativity. It was first confirmed during the 1919 solar eclipse through the apparent displacement of stars observed close to the solar limb (for a historical review, see Will, 2015). The possibility that a foreground object could act as a gravitational lens was later discussed by Einstein (1936), who showed that a foreground star could lens a more distant background star, although he noted that the probability of observing such an event would be extremely small. Shortly afterward, Zwicky (1937) recognized that galaxies, because of their much larger masses, would be far more efficient gravitational lenses. He also pointed out that gravitational lenses could act as natural telescopes, magnifying distant sources, and as probes of the mass distribution in galaxies and clusters.

The discovery of the first multiply imaged quasar, QSO 0957+561, shown in Figure 1.1, by Walsh et al. (1979), established gravitational lensing as a powerful observational tool. Since then, lensing has become an essential technique in astrophysics and cosmology. Depending on the mass of the lens, the alignment between the observer, lens, and source, and the angular resolution of the observations, gravitational lensing can produce several observable effects, including multiple images, arcs, Einstein rings, magnification, and distortions of the background source. These effects encode information about both the background source and the intervening mass distribution (Schneider et al., 1992; Meylan et al., 2006; Shajib et al., 2024; Vegetti et al., 2024; Vernardos et al., 2024).

In the strong-lensing regime, the projected mass density of the lens is sufficiently high to generate multiple images of the same background source. This regime is particularly important because the image positions, flux ratios, shapes, and extended arcs provide direct constraints on the projected gravitational potential of the lens. Strong lensing therefore provides a way to study both distant sources, which can be magnified by the lens, and the mass distribution of the foreground deflector, including its dark matter component.

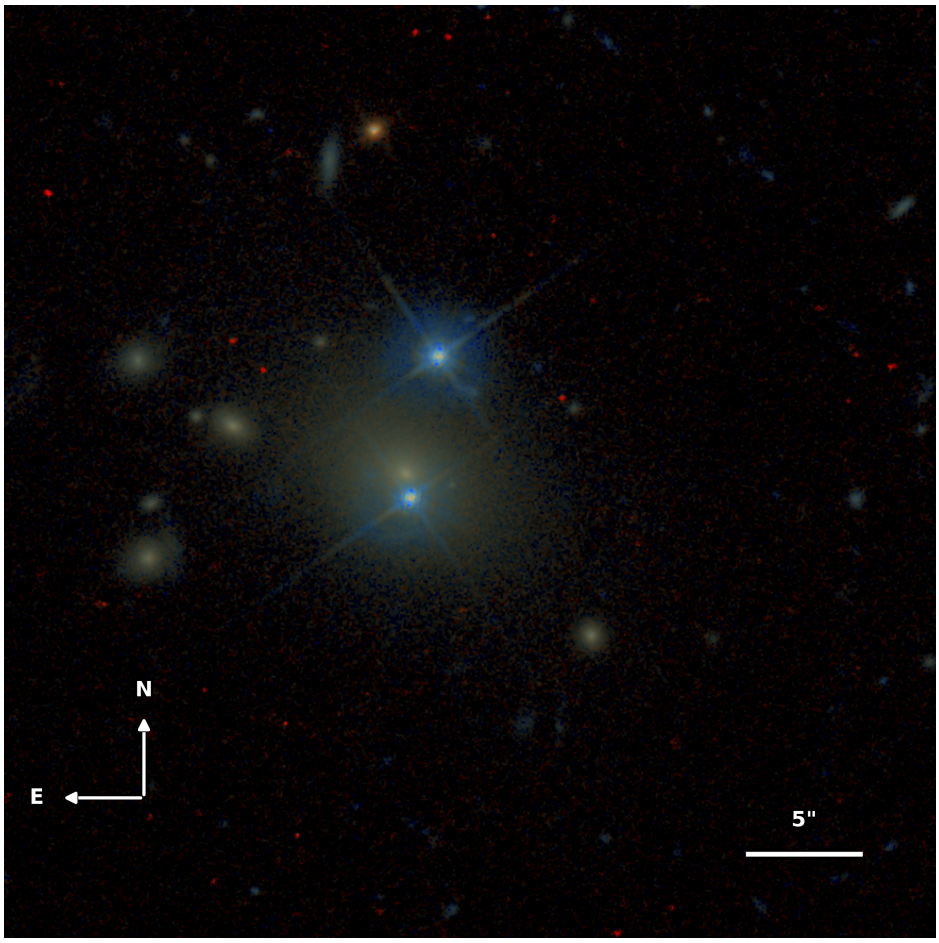


Figure 1.1: Composite image of QSO 0957+561 constructed from observations in the V and I filters. The image shows the field of view of a single HST WFPC2 wide-field CCD centered on the lensed quasar images.

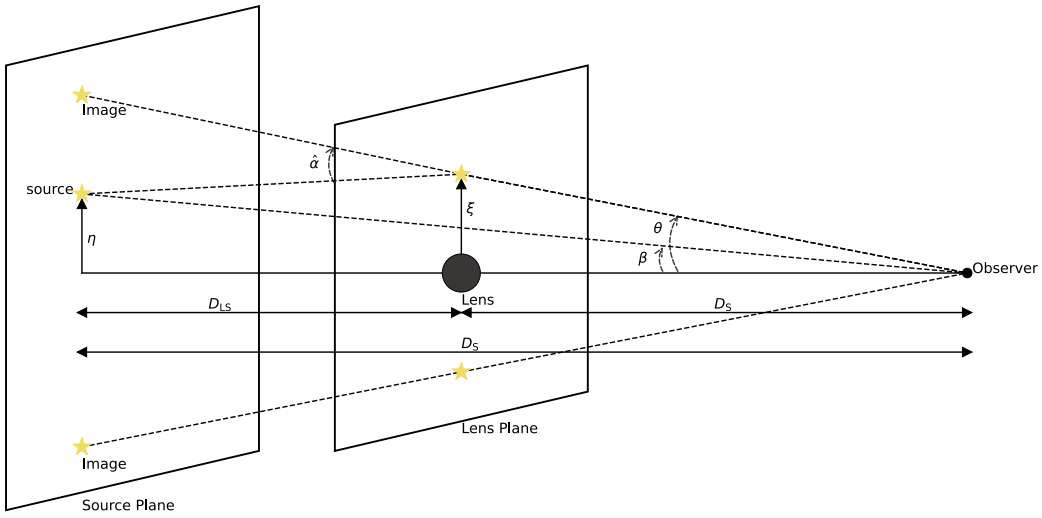


Figure 1.2: Sketch of a typical gravitational lens system. A background source is deflected by the gravitational potential of a foreground lens, producing one or more images on the sky. The angles $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ denote the angular positions of the unlensed source and the observed image, respectively, as seen by the observer. The angle $\boldsymbol{\alpha}$ is the physical deflection angle produced at the lens plane, while the reduced deflection angle is $\hat{\boldsymbol{\alpha}} = (D_S/D_{LS})\boldsymbol{\alpha}$. The corresponding physical positions in the source and lens planes are $\boldsymbol{\eta}$ and $\boldsymbol{\xi}$, and D_L, D_S, and D_{LS} are the angular-diameter distances between observer and lens, observer and source, and lens and source, respectively.

Although the full description of light propagation in a gravitational field requires General Relativity, most astrophysical lensing systems can be described using a simpler approximate framework. This approximation is valid in the weak-field regime, where $|\Phi|/c^2 \ll 1$, and in the small-angle limit. Galaxy-scale lenses satisfy these conditions: their gravitational potentials are weak, and their typical image separations and deflection angles are of the order of arcseconds. In addition, the physical extent of the deflecting mass along the line of sight is much smaller than the angular-diameter distances between the observer, lens, and source. Under these conditions, the mass distribution of the lens can be projected onto a two-dimensional lens plane, and the true curved trajectory of the light ray can be approximated by two straight segments joined at the lens position. This is known as the thin-lens approximation (Refsdal, 1964b; Schneider et al., 1992; Meylan et al., 2006; Congdon & Keeton, 2018).

For an intuitive interpretation of the phenomenon, we consider the geometry shown in Figure 1.2. A background source is located on the source plane, while a foreground galaxy acts as the lens. As light from the source travels toward the ob-

server, it is deflected by the gravitational potential of the lens galaxy. Therefore, the observed angular position of an image does not necessarily coincide with the intrinsic angular position of the source. If the projected surface mass density of the lens is sufficiently high and the source, lens, and observer are suitably aligned (see Section 1.1.7 for details), a single source can be observed as multiple images or extended arcs. In the case of nearly perfect axial alignment with a symmetrical lens, the images merge into a ring-like structure around the lens, known as an Einstein ring (like the one presented in Figure 1.3). In the following, we introduce the basic formalism used to describe this effect.

1.1.1 The deflection angle

We already introduced the deflection angle but this time we will make a proper definition of it, considering the deflection of a light ray by the exterior gravitational field of a spherically symmetric mass M. If the impact parameter ξ of the ray is much larger than the Schwarzschild radius of the mass,

$$\xi \gg R_S \equiv \frac{2GM}{c^2}, \tag{1.3}$$

then General Relativity predicts that the deflection angle is

$$\hat{\alpha} = \frac{4GM}{c^2 \xi}. \tag{1.4}$$

The condition $\xi \gg R_S$ implies that the deflection angle is small, $\hat{\alpha} \ll 1$, and that the gravitational field is weak, $\phi/c^2 \ll 1$. In this regime, the field equations of General Relativity can be linearized, and the deflection produced by a collection of mass elements can be obtained by summing the contributions from each component.

For an extended three-dimensional mass distribution with volume density $\rho(\mathbf{r})$, the mass can be decomposed into elements $\mathrm{d}m = \rho(\mathbf{r})\,\mathrm{d}V$. We choose coordinates such that, far from the lens, the incoming light ray propagates along the r_3 direction, while the two-dimensional vector $\boldsymbol{\xi} = (\xi_1, \xi_2)$ describes the position of the ray on the lens plane. In the thin-lens approximation, the deflected ray is treated as approximately straight while it crosses the lens. The impact vector of the ray relative to a mass element located at $\mathbf{r}' = (\boldsymbol{\xi}', r_3')$ is then $\boldsymbol{\xi} - \boldsymbol{\xi}'$, independent of r_3'. The total deflection angle is therefore

$$\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi}) = \frac{4G}{c^2} \int\int \mathrm{d}^2\boldsymbol{\xi}' \mathrm{d}r_3' \, \rho(\boldsymbol{\xi}', r_3') \frac{\boldsymbol{\xi} - \boldsymbol{\xi}'}{|\boldsymbol{\xi} - \boldsymbol{\xi}'|^2}. \tag{1.5}$$

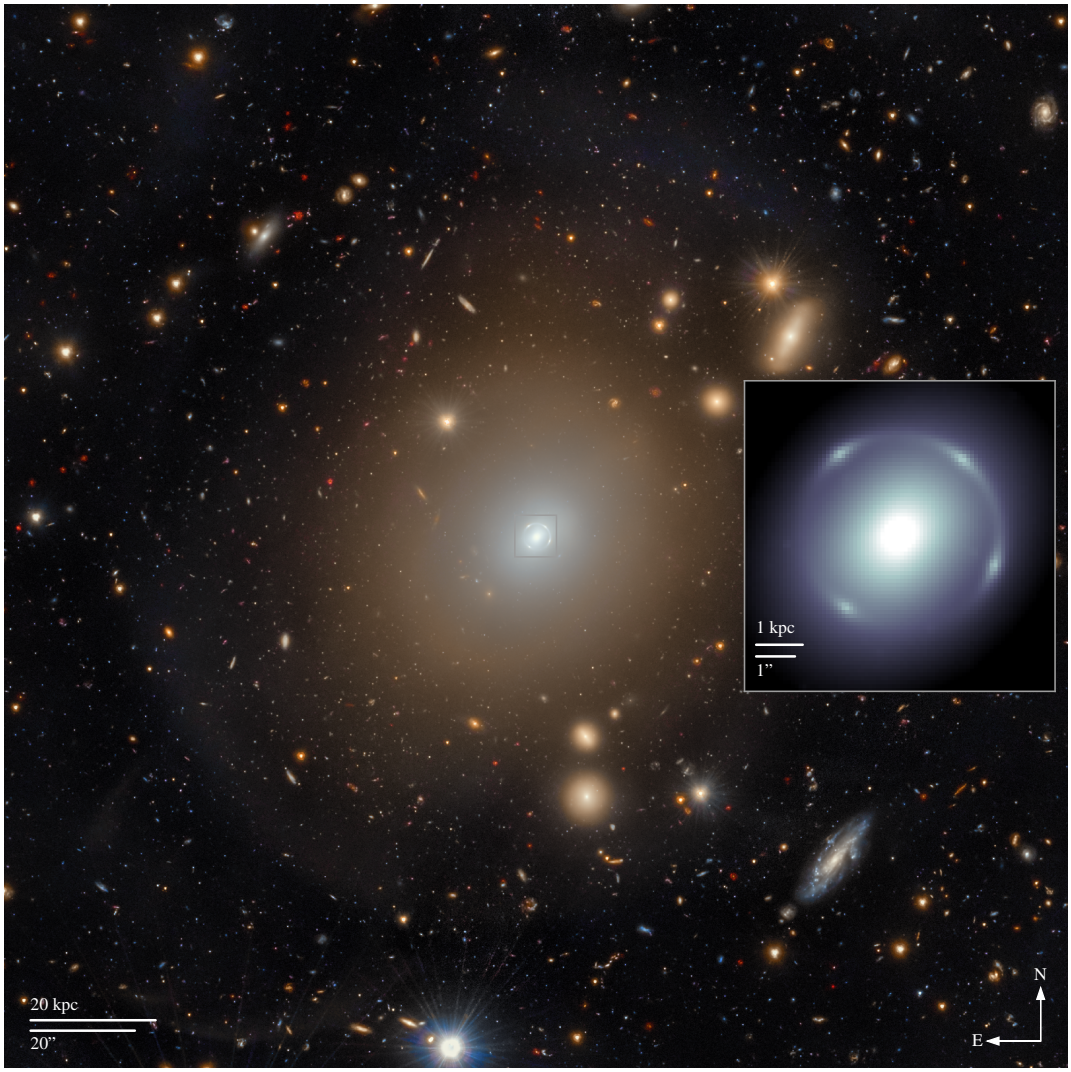


Figure 1.3: Euclid imaging of Altieri's lens, an example of an Einstein ring (O'Riordan et al., 2025). The main panel shows a composite false-colour image constructed from Euclid VIS and NISP observations. The VIS instrument provides high-resolution optical imaging, while NISP provides lower-resolution near-infrared photometry in the Y_E, J_E, and H_E bands. In this composite, the broadband VIS I_E image defines the brightness structure, and the NISP bands provide the colour information. The central light of the lens galaxy has been suppressed to enhance the visibility of the lensed arc. The inset shows the central 8 arcsec using only the higher-resolution VIS data, corresponding to the square marked in the main panel. The angular scale and the physical scale at the lens redshift are shown in each frame.

We define the projected surface mass density as

$$\Sigma(\boldsymbol{\xi}) \equiv \int \rho(\boldsymbol{\xi}, r_3)\, \mathrm{d}r_3. \tag{1.6}$$

The deflection angle produced by an arbitrary projected mass distribution is then

$$\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi}) = \frac{4G}{c^2} \int \mathrm{d}^2\boldsymbol{\xi}', \Sigma(\boldsymbol{\xi}') \frac{\boldsymbol{\xi} - \boldsymbol{\xi}'}{|\boldsymbol{\xi} - \boldsymbol{\xi}'|^2}. \tag{1.7}$$

Thus, within the thin-lens approximation, the relevant information about the deflector is encoded in the projected mass distribution $\Sigma(\boldsymbol{\xi})$ (Refsdal, 1964b; Liebes, 1964; Schneider et al., 1992; Meylan et al., 2006).

1.1.2 The lens equation

Let us now return to the geometry shown in Figure 1.2. A mass concentration at redshift z_L deflects light from a source at redshift z_S. We denote by D_L, D_S, and D_LS the angular-diameter distances from the observer to the lens, from the observer to the source, and from the lens to the source, respectively. If there are no additional deflectors along the line of sight, and if the extent of the lens along the line of sight is small compared with these distances, the light path can be approximated by two straight segments joined by a kink at the lens plane. The magnitude and direction of this kink are described by the deflection angle $\hat{\boldsymbol{\alpha}}$.

The lens equation relates the true position of the source to the observed angular positions of its images. Let $\boldsymbol{\eta}$ denote the two-dimensional physical position of the source on the source plane, and let $\boldsymbol{\xi}$ denote the physical position of the light ray on the lens plane. Using the small-angle approximation, $\sin\hat{\alpha} \simeq \hat{\alpha} \simeq \tan\hat{\alpha}$, the lens geometry gives

$$\boldsymbol{\eta} = \frac{D_\mathrm{S}}{D_\mathrm{L}}\boldsymbol{\xi} - D_\mathrm{LS}\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi}). \tag{1.8}$$

Introducing angular coordinates

$$\boldsymbol{\beta} = \frac{\boldsymbol{\eta}}{D_\mathrm{S}}, \qquad \boldsymbol{\theta} = \frac{\boldsymbol{\xi}}{D_\mathrm{L}}, \tag{1.9}$$

Equation 1.8 becomes the lens equation

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \frac{D_\mathrm{LS}}{D_\mathrm{S}}\hat{\boldsymbol{\alpha}}(D_\mathrm{L}\boldsymbol{\theta}) \equiv \boldsymbol{\theta} - \boldsymbol{\alpha}(\boldsymbol{\theta}), \tag{1.10}$$

where

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \frac{D_\mathrm{LS}}{D_\mathrm{S}}\hat{\boldsymbol{\alpha}}(D_\mathrm{L}\boldsymbol{\theta}) \tag{1.11}$$

is the scaled deflection angle. Equation 1.10 defines a mapping from the image plane to the source plane. A source located at $\boldsymbol{\beta}$ is observed at angular positions $\boldsymbol{\theta}$ that satisfy the lens equation (Equation 1.10). If the equation admits more than one solution for the same source position, the source is multiply imaged.

It is useful to introduce the dimensionless surface mass density, or convergence,

$$\kappa(\boldsymbol{\theta}) = \frac{\Sigma(D_{\rm L}\boldsymbol{\theta})}{\Sigma_{\rm cr}}, \tag{1.12}$$

where the critical surface mass density, value that sets the density scale at which the lens mapping can become singular (i.e, separates strong from weak lensing, Schneider et al., 1992)

$$\Sigma_{\rm cr} = \frac{c^2}{4\pi G}\frac{D_{\rm S}}{D_{\rm L}D_{\rm LS}}. \tag{1.13}$$

The convergence measures the projected surface mass density of the lens in units of the critical surface mass density. It quantifies the isotropic focusing produced by the lens: convergence causes a bundle of neighbouring light rays to focus, resulting in an isotropic change in the apparent size and magnification of the image.

Using Equations 1.7–1.13, the scaled deflection angle can be written in terms of κ as

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \frac{1}{\pi}\int_{\mathbb{R}^2} {\rm d}^2\boldsymbol{\theta}', \kappa(\boldsymbol{\theta}')\frac{\boldsymbol{\theta}-\boldsymbol{\theta}'}{|\boldsymbol{\theta}-\boldsymbol{\theta}'|^2}. \tag{1.14}$$

Since

$$\nabla \ln|\boldsymbol{\theta}| = \frac{\boldsymbol{\theta}}{|\boldsymbol{\theta}|^2}, \tag{1.15}$$

the scaled deflection angle can be expressed as the gradient of a scalar lensing potential,

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \nabla\psi(\boldsymbol{\theta}), \tag{1.16}$$

where

$$\psi(\boldsymbol{\theta}) = \frac{1}{\pi}\int_{\mathbb{R}^2} {\rm d}^2\boldsymbol{\theta}', \kappa(\boldsymbol{\theta}')\ln|\boldsymbol{\theta}-\boldsymbol{\theta}'|. \tag{1.17}$$

Using $\nabla^2 \ln|\boldsymbol{\theta}| = 2\pi\delta_{\rm D}(\boldsymbol{\theta})$, where $\delta_{\rm D}$ is the two-dimensional Dirac delta function, the lensing potential satisfies the Poisson equation

$$\nabla^2\psi = 2\kappa. \tag{1.18}$$

Therefore, given the lens potential ψ, one can easily compute the scaled deflection angle and the dimensionless surface mass density through differentiations (Schneider et al., 1992; Meylan et al., 2006; Mediavilla et al., 2016).

1.1.3 Cosmological distances

Until this point, we have introduced the lens equation without explicitly discussing the nature of the distances that appear in it. The distances in Equation 1.10 are angular-diameter distances. The angular-diameter distance, D_{A}, is defined as the ratio between the physical transverse size of an object and its angular size on the sky. It is therefore the appropriate distance for converting angular separations measured in telescope images into proper transverse separations at the lens or source redshift. The angular-diameter distance is related to the transverse comoving distance, D_{M}, by

$$D_{\mathrm{A}}(z) = \frac{D_{\mathrm{M}}(z)}{1+z}. \tag{1.19}$$

In gravitational lensing, the lens and the source are located at different redshifts. For a flat Universe, the angular-diameter distance between two redshifts z_1 and z_2, with $z_2 > z_1$, is given by

$$D_{\mathrm{A}}(z_1, z_2) = \frac{1}{1+z_2}\left[D_{\mathrm{M}}(z_2) - D_{\mathrm{M}}(z_1)\right]. \tag{1.20}$$

Thus, the relevant distances in the lensing geometry are D_{L}, D_{S}, and D_{LS}, corresponding respectively to the angular-diameter distances from the observer to the lens, from the observer to the source, and from the lens to the source.

Gravitational lensing is therefore directly connected to cosmology because the strength of the lens depends not only on the mass distribution of the deflector, but also on the geometry of the Universe. This geometrical dependence enters through angular-diameter distance ratios, such as the factor $D_{\mathrm{LS}}/D_{\mathrm{S}}$ that appears in the scaled deflection angle in Equation 1.11. These distances depend on the assumed cosmological model and on parameters such as the Hubble constant H_0, the matter density parameter, and the dark-energy density parameter. Consequently, lensing observables such as image positions, angular separations, Einstein radii, and inferred enclosed masses must be interpreted within a cosmological framework (Refsdal, 1964b; Hogg, 1999; Suyu et al., 2010, 2013; Birrer et al., 2024).

1.1.4 Magnification and distortion

The solutions $\boldsymbol{\theta}$ of the lens equation (Equation 1.10) give the angular positions of the images of a source located at $\boldsymbol{\beta}$. However, the lens mapping also changes the apparent shape and size of the source like is shown in the Figure 1.4 [1]. Since gravitational

[1] Original Figure 1.4 https://en.wikipedia.org/wiki/Gravitational_lensing_formalism

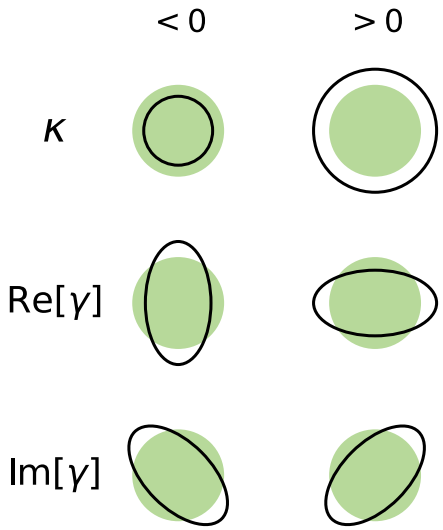


Figure 1.4: Distortions produced by the convergence κ and the two components of the shear, γ_1 and γ_2. The original circular source is shown in green, while the corresponding lensed image is shown in black. The convergence term produces an isotropic change in size, with $\kappa > 0$ magnifying the image and $\kappa < 0$ demagnifying it. The shear components produce anisotropic distortions: γ_1 stretches the image along the horizontal or vertical axis, while γ_2 stretches it along axes rotated by 45°. Adapted from Wikipedia; see also Schneider et al. (1992); Meylan et al. (2006).

lensing conserves surface brightness in the absence of absorption or emission, the observed change in flux is caused by a change in the apparent solid angle subtended by the image. If $I^s(\boldsymbol{\beta})$ is the surface brightness distribution of the source, then the observed surface brightness is

$$I(\boldsymbol{\theta}) = I^s\left[\boldsymbol{\beta}(\boldsymbol{\theta})\right]. \tag{1.21}$$

If the source is much smaller than the angular scale over which the lens properties vary, the lens mapping can be linearized locally. The local distortion is described by the Jacobian matrix

$$\mathcal{A}(\boldsymbol{\theta}) = \frac{\partial \boldsymbol{\beta}}{\partial \boldsymbol{\theta}} = \left(\delta_{ij} - \frac{\partial^2 \psi}{\partial\theta_i \partial\theta_j}\right) = \begin{pmatrix} 1-\kappa-\gamma_1 & -\gamma_2 \\ -\gamma_2 & 1-\kappa+\gamma_1 \end{pmatrix}. \tag{1.22}$$

Here, the shear components are

$$\gamma_1 = \frac{1}{2}\left(\frac{\partial^2\psi}{\partial\theta_1^2} - \frac{\partial^2\psi}{\partial\theta_2^2}\right), \qquad \gamma_2 = \frac{\partial^2\psi}{\partial\theta_1\partial\theta_2}, \tag{1.23}$$

and the complex shear is

$$\gamma \equiv \gamma_1 + i\gamma_2 = |\gamma| e^{2i\varphi}. \tag{1.24}$$

The convergence κ describes the isotropic focusing produced by the projected surface mass density, while the shear γ describes the anisotropic stretching produced by the tidal gravitational field, as is clearly seen in Figure 1.4.

For a point $\boldsymbol{\theta}_0$ within an image, corresponding to a source-plane point $\boldsymbol{\beta}_0 = \boldsymbol{\beta}(\boldsymbol{\theta}_0)$, the locally linearized lens equation gives

$$I(\boldsymbol{\theta}) = I^s\left[\boldsymbol{\beta}_0 + \mathcal{A}(\boldsymbol{\theta}_0)\cdot(\boldsymbol{\theta}-\boldsymbol{\theta}_0)\right]. \tag{1.25}$$

Thus, a source with circular isophotes is locally mapped into an ellipse. The inverse of the Jacobian,

$$\mathcal{M}(\boldsymbol{\theta}) = \mathcal{A}^{-1}(\boldsymbol{\theta}), \tag{1.26}$$

is the magnification tensor. The scalar magnification is given by the inverse determinant of the Jacobian,

$$\mu(\boldsymbol{\theta}) = \frac{1}{\det \mathcal{A}} = \frac{1}{(1-\kappa)^2 - |\gamma|^2}. \tag{1.27}$$

Images are therefore highly magnified where $\det \mathcal{A}$ approaches zero.

The images are therefore distorted in both shape and size. The shape distortion is produced by the shear, while the change in apparent size and flux is produced by the combined effect of convergence and shear (Schneider et al., 1992; Meylan et al., 2006).

1.1.5 Critical curves and caustics

The magnification introduced above becomes very large when the determinant of the Jacobian matrix approaches zero. The set of points in the image plane where this condition is exactly satisfied defines the critical curves of the lens. Therefore, these curves mark the locations where the lens mapping becomes locally singular and where the magnification formally diverges for an ideal point source. In real observations, this divergence is limited by the finite size of the source, the instrumental resolution, and other physical effects (Narayan & Bartelmann, 1996; Liaudat et al., 2023).

Critical curves are defined in the image plane, but their corresponding mapping through the lens equation defines curves in the source plane known as caustics. Caustics are especially useful because they determine how the number and location of lensed images change as the source moves across the source plane. When a point source moves towards (outwards) the lens by crossing a caustic, the number of images changes, usually through the creation (inward) or destruction (outward) of a pair of images (see Figure 1.5; Liebes, 1964; Schneider et al., 1992; Meylan et al., 2006).

1.1.6 Simple mass distributions

The lensing quantities introduced above, such as the deflection angle, convergence, shear, magnification, and critical curves, are determined by the projected mass distribution of the deflector. In practice, however, the true mass distribution of a galaxy is not known a priori. For this reason, gravitational-lens modelling usually relies on simplified parametric descriptions that capture the main properties of the lens while keeping the number of free parameters manageable (Refsdal, 1964b; Schneider et al., 1992; Keeton, 2001a; Meylan et al., 2006).

These models provide a controlled way to connect the observed image configuration with the underlying lens mass distribution. As illustrated in Figure 1.5, one can move from very simple circular models to more realistic angular structures. The point-mass lens represents the limiting case in which all the mass is concentrated at a single position (Einstein, 1936). The singular isothermal sphere (SIS) introduces an extended but circularly symmetric galaxy-scale mass distribution (Turner et al., 1984). Additional angular structure can be included either through an external shear term, which approximates tidal perturbations from the environment, or also through the singular isothermal ellipsoid (SIE), which adds ellipticity to the projected mass distribution (Kassiola & Kovner, 1993; Kormann et al., 1994).

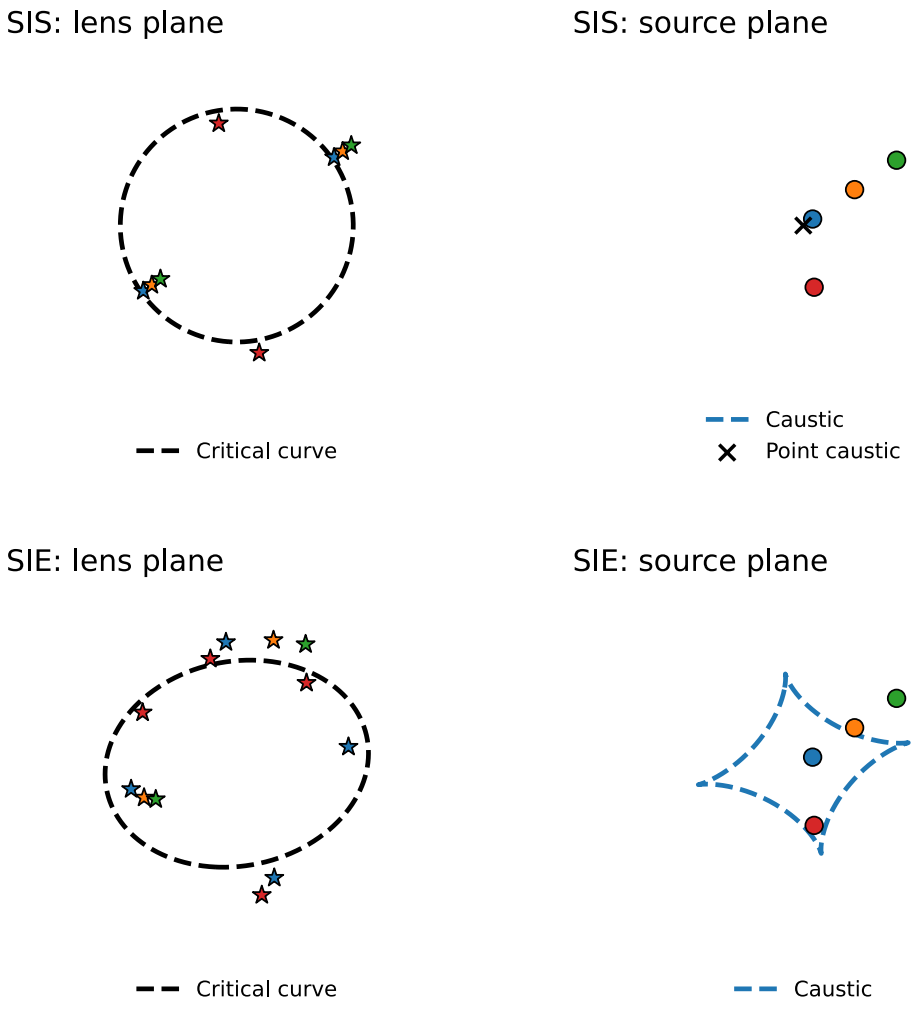


Figure 1.5: Examples of image configurations produced by simple lens mass models. The upper panels show a singular isothermal sphere (SIS), while the bottom panels show a singular isothermal ellipsoid (SIE). In each case, different source positions are shown relative to the caustics in the source plane, together with the corresponding image positions in the image plane. The comparison shows how changing the symmetry of the mass distribution changes the caustic structure, and the resulting number and arrangement of images.

Point-mass lens

The simplest lens model is the point-mass lens. In this case, the deflector is assumed to be compact compared with the angular scale over which the light is deflected. Although this model is too idealized to describe the full mass distribution of a galaxy, it provides a useful starting point for introducing the characteristic lensing scale and for describing microlensing by stars or other compact objects (Einstein, 1936; Liebes, 1964). For a point lens of mass M, the reduced deflection angle is

$$\alpha(\theta) = \frac{D_{\rm LS}}{D_{\rm L} D_{\rm S}} \frac{4GM}{c^2 \theta}. \tag{1.28}$$

The angular scale at which the source, lens, and observer are perfectly aligned is the Einstein radius,

$$\theta_{\rm E} = \left(\frac{4GM}{c^2} \frac{D_{\rm LS}}{D_{\rm L} D_{\rm S}} \right)^{1/2}. \tag{1.29}$$

This quantity defines the natural angular scale of the lens and depends on both the lens mass and the angular-diameter distance geometry. For a point-mass lens, a source that is not exactly aligned with the lens produces two image solutions; in the perfectly aligned case, these images merge into an Einstein ring (see, e.g., Treu, 2010; Congdon & Keeton, 2018, and references therein).

Critical curves and caustics, as illustrated in Figure 1.5, provide the connection between the mass distribution of a lens and the image configurations that it can produce. Critical curves are loci in the image plane where the lens mapping becomes singular and the magnification formally diverges, while caustics are the corresponding curves mapped into the source plane. The number and arrangement of the lensed images are determined by the position of the source relative to these caustics. For circularly symmetric models, such as the SIS, the caustic structure is simple. When angular structure is introduced, either through external shear or intrinsic ellipticity, the caustics become more extended and can develop the diamond-like shape characteristic of four-image configurations.

Singular isothermal sphere

A useful first approximation for galaxy-scale lenses is the singular isothermal sphere. In this model, the mass is not concentrated at a single point but is distributed over an extended region with spherical symmetry. The three-dimensional mass density profile is

$$\rho(r) = \frac{\sigma^2}{2\pi G r^2}, \tag{1.30}$$

where σ is the one-dimensional velocity dispersion. This profile corresponds to an isothermal system in dynamical equilibrium and produces a flat rotation curve, a property that approximately describes the total mass distribution of many massive early-type galaxies over the radial range probed by strong lensing (Blandford & Kovner, 1988; Schneider et al., 1992). For a SIS, the corresponding Einstein radius is

$$\theta_{\rm E} = 4\pi \left(\frac{\sigma}{c}\right)^2 \frac{D_{\rm LS}}{D_{\rm S}}, \tag{1.31}$$

This expression shows that, for a fixed lens-source geometry, the main lensing scale of the SIS is set by the velocity dispersion $\sigma_{\rm SIS}$.

Although real lens galaxies are generally elliptical rather than spherical, the SIS remains widely used because many lensing observables, particularly in two-image systems, are primarily sensitive to the circularly averaged mass distribution near the Einstein radius. In such systems, the dominant constraint is the enclosed projected mass, which is well characterized by the velocity dispersion. As a result, the SIS often provides a good description of the image separation and overall lensing geometry despite its simplified symmetry. Furthermore, its analytical simplicity makes it a useful starting point for lens modeling and for developing physical intuition about strong-lensing phenomena (Grillo et al., 2008).

Because of its spherical symmetry, the SIS is useful for understanding the main lensing scale and for performing simple calculations. However, this same symmetry also limits the model: a pure SIS cannot reproduce the full angular structure observed in many strong-lens systems. A common extension is to add external shear, which represents the tidal effect of nearby galaxies, groups, clusters, or line-of-sight structure. The SIS plus external shear model therefore provides an intermediate step between circularly symmetric lenses and fully non-spherical mass distributions (Keeton, 2001a; Lee & Kim, 2014).

Singular isothermal ellipsoid

Real lens galaxies are not perfectly spherical (Masters, 2025). Their projected mass distributions are usually flattened, in the sense that the contours of constant projected surface mass density are elliptical rather than circular. The singular isothermal

ellipsoid extends the SIS by introducing ellipticity into the projected mass distribution while preserving the approximately isothermal radial dependence (Kassiola & Kovner, 1993; Kormann et al., 1994).

In lensing applications, the SIE is commonly written in terms of the convergence $\kappa = \Sigma/\Sigma_{\rm crit}$ (Section 1.1.2),

$$\kappa(x', y') = \frac{\theta_{\rm E}}{2R_{\rm ell}}, \tag{1.32}$$

where

$$R_{\rm ell} = \sqrt{qx'^2 + \frac{y'^2}{q}}. \tag{1.33}$$

Here x' and y' are coordinates aligned with the major and minor axes of the projected mass distribution, q is the projected axis ratio, and $\theta_{\rm E}$ is the circularized Einstein radius. For $q = 1$, the elliptical radius reduces to the circular radius, $R_{\rm ell} = R$, and the SIE becomes equivalent to the SIS, with

$$\kappa(R) = \frac{\theta_{\rm E}}{2R}. \tag{1.34}$$

The Einstein radius is therefore the characteristic angular scale of the model. In the spherical limit, it is related to the velocity dispersion by

$$\theta_{\rm E} = 4\pi \left(\frac{\sigma_{\rm SIE}}{c}\right)^2 \frac{D_{\rm LS}}{D_{\rm S}}, \tag{1.35}$$

where $\sigma_{\rm SIE}$ is the velocity-dispersion normalization. For elliptical models, different conventions may absorb factors depending on q into the definition of the lens strength, but the parameter still plays the role of the Einstein-radius scale (Keeton, 2001a).

This angular structure makes the SIE more suitable for modelling realistic galaxy-scale lenses than the circular SIS. In particular, the ellipticity changes the critical curves and caustics, allowing the model to generate four-image configurations when the source lies inside the tangential caustic. For this reason, the SIE is commonly used as a starting point for modelling quadruply imaged quasars.

In many applications, the SIE is combined with an external shear term to account for tidal perturbations from nearby galaxies, galaxy groups, clusters, or line-of-sight structure that are not explicitly included in the main lens model. This combination provides a flexible but still relatively simple description of the lens potential and is widely used in quasar strong-lensing analyses (Keeton, 2001b; Birrer et al., 2015; Nightingale et al., 2021).

Together, these models form a sequence of increasing complexity. The point-mass lens establishes the basic scale associated with compact deflectors and microlensing,

while the SIS introduces an extended, circularly symmetric galaxy-scale mass distribution. External shear accounts for tidal perturbations from the lens environment, whereas the SIE incorporates intrinsic ellipticity in the projected mass distribution. More flexible models can be introduced when the data require a more detailed description of the lens potential. These include elliptical power-law profiles, in which the radial density slope is allowed to vary, and composite models that treat the baryonic and dark-matter components separately. Free-form or pixelated reconstructions provide even greater flexibility by representing the lens mass distribution on a spatial grid rather than imposing a predefined analytic profile. They can therefore recover localized departures from smooth models, such as mass substructure, asymmetries, radial variations, or complex environmental contributions, although this additional freedom generally requires regularization to prevent overfitting and to obtain physically plausible solutions (Barkana, 1998; Tessore et al., 2015; Suyu et al., 2010; Vegetti & Koopmans, 2009; Saha & Williams, 2004; Galan et al., 2022).

1.1.7 Necessary and sufficient conditions for multiple imaging

A projected mass distribution described by the convergence κ may or may not be strong enough to produce multiple images of a background source. Two useful criteria can be established (Schneider et al., 1992; Meylan et al., 2006):

1. An isolated transparent lens can produce multiple images if, and only if, there exists at least one point $\boldsymbol{\theta}$ for which

$$\det \mathcal{A}(\boldsymbol{\theta}) < 0. \tag{1.36}$$

   If $\det \mathcal{A} > 0$ everywhere, the lens mapping is globally invertible and multiple imaging cannot occur. If $\det \mathcal{A} < 0$ somewhere, the lens mapping changes parity, indicating that the source-plane mapping is no longer one-to-one. In this case, multiple images of the same source must be produced.

2. A sufficient, although not necessary, condition for the possible occurrence of multiple imaging is that there exists a point $\boldsymbol{\theta}$ such that

$$\kappa(\boldsymbol{\theta}) > 1. \tag{1.37}$$

   This condition means that the local surface mass density exceeds the critical surface mass density. Hence, $\Sigma_{\rm cr}$ provides the natural scale for the appearance of strong-lensing features such as arcs and multiple images.

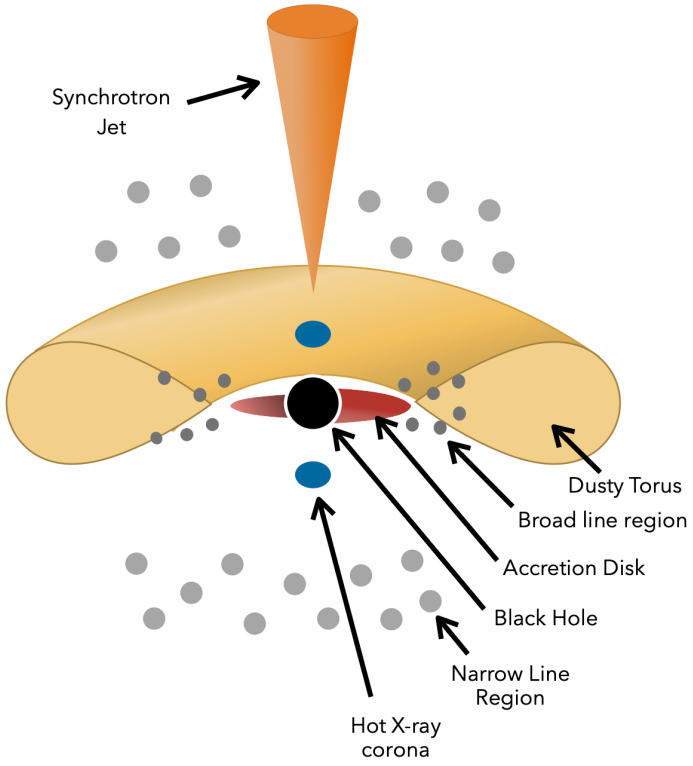


Figure 1.6: Classical schematic representation of the main structural components of an AGN in the center of the host galaxy (figure not at scale). The central SMBH and accretion disk are surrounded by the broad-line region, dusty circumlunar torus, the narrow-line region, and, in some systems, a relativistic jet. Obtain from (Thorne et al., 2022)

These criteria are closely connected to the critical curves and caustics introduced in Section 1.1.5. As illustrated in Figure 1.5, the resulting image multiplicity is determined by the source position relative to the caustic structure: sources outside the relevant caustics produce fewer images, while sources enclosed by them can give rise to multiple-image configurations.

Because Σ_{cr} depends on the angular-diameter distances to the lens and source, the strength of a lens depends on the source redshift (1.13). For a fixed physical surface mass density Σ, changing the source redshift changes Σ_{cr} and therefore modifies the positions of the critical curves and caustics. This is why, lensing with sources at different redshifts can be produce strong-lensing features at different angular separations from the galaxy center (Gavazzi et al., 2008; Shajib et al., 2024). Gravitational lensing is therefore a complex effect that can nevertheless be described by a compact formalism.

Its dependence on the projected mass distribution, the relative alignment between observer, lens, and source, and the angular-diameter distances makes it a powerful tracer of astrophysical mass distributions and cosmological applications (Refsdal, 1964a; Tdcosmo Collaboration et al., 2025; Pierel et al., 2026).

Up to this point, we focused on the theoretical foundations of the lensing effect. In the next section, we introduce the main source objects of study in this thesis: active galactic nuclei

1.2 Active Galactic Nuclei

Active galactic nuclei (AGN) are among the most luminous and energetic persistent sources in the Universe and they are schematically represented as is shown in Figure 1.6 with their different components. Since their discovery, they have played a central role in extragalactic astronomy, providing a direct observational link between galaxy evolution, supermassive black-hole growth, and energetic feedback processes (for historical and physical reviews, see Kellermann, 2015; Netzer, 2015; Padovani, 2017). An AGN is powered by the accretion of gas onto a central supermassive black hole (SMBH), with typical black-hole masses above $M_{\rm BH} \gtrsim 10^5, M_\odot$ and accretion luminosities that can reach a significant fraction of the Eddington limit (Padovani, 2017). The Eddington luminosity, obtained by balancing radiation pressure and gravity for ionized hydrogen, is given by

$$L_{\rm Edd} \simeq 1.26 \times 10^{38} \left(\frac{M_{\rm BH}}{M_\odot} \right) \text{erg, s}^{-1}. \tag{1.38}$$

The ratio between the bolometric AGN luminosity and this limit,

$$\lambda_{\rm Edd} = \frac{L_{\rm AGN}}{L_{\rm Edd}}, \tag{1.39}$$

is commonly used as a dimensionless measure of the accretion state of the system. AGN span a wide range of Eddington ratios, from very low values in weakly accreting systems to values approaching unity in rapidly accreting quasars (QSO) the most powerful kind of AGN. The ratio for AGN is commonly found at $\lambda_{\rm Edd} \sim 0.01$–$1$, with typical values of order $\lambda_{\rm Edd} \sim 0.1$, although the exact range depends on luminosity, redshift, and selection effects (Shen et al., 2011, 2013). In this sense, luminous quasars are generally associated with relatively high Eddington ratios, while lower-luminosity

AGN and LINER-like systems may accrete at much smaller fractions of the Eddington limit (Kollmeier et al., 2006; Ho, 2009).

The central engine of an AGN is not directly resolved in most observations. Nevertheless, a broad physical picture has emerged from spectroscopy, variability, reverberation mapping, polarimetry, X-ray observations, infrared emission, and radio studies (Antonucci, 1993; Urry & Padovani, 1995; Netzer, 2015; Padovani, 2017; Cackett et al., 2021). In this picture, the SMBH is surrounded by an accretion flow that converts gravitational potential energy into radiation. The radiative efficiency of this process is usually described by

$$\eta = \frac{L_{\rm AGN}}{\dot{M}c^2}, \tag{1.40}$$

where $\dot{M}$ is the mass accretion rate. A radiatively efficient flow converts a significant fraction of the accreted rest-mass energy into radiation, with a commonly adopted value $\eta \sim 0.1$ for a standard thin disk (Shakura & Sunyaev, 1973). Such flows are expected when the accretion rate is high enough for the gas to cool efficiently, typically above $\dot{M}/\dot{M}_{\rm Edd} \sim 10^{-2}$, although the exact transition depends on the accretion state and black-hole properties (Fabian, 2012; Yuan & Narayan, 2014; Reynolds, 2021). In radiatively efficient AGN, this accretion flow is commonly described as an optically thick, geometrically thin disk that emits a substantial fraction of the AGN's bolometric luminosity, primarily at optical and ultraviolet wavelengths.

1.2.1 The AGN structure

The different emitting regions of quasars span a wide range of physical scales, from the immediate surroundings of the supermassive black hole (SMBH) to the host-galaxy environment. The innermost regions, such as the X-ray corona and the optical/UV accretion disk, are extremely compact. As illustrated in Figure 1.7, the X-ray corona is thought to be located at radii of only a few to a few tens of gravitational radii, where $r_{\rm g} = GM_{\rm BH}/c^2$, and to produce much of the X-ray emission through inverse-Compton scattering of disk photons (Netzer, 2015; Cackett et al., 2021). Observational studies typically infer X-ray-emitting regions within $\sim$ 10–20, $r_{\rm g}$, while broad-band X-ray spectroscopy suggests hot electron plasmas with characteristic temperatures of order $kT_{\rm e} \sim$ 50–200, keV, with typical values close to $kT_{\rm e} \sim 100$, keV (Reis & Miller, 2013; Ricci et al., 2018). The optical/UV continuum-emitting accretion disk is larger, but still compact on galactic scales, with characteristic sizes of order light-days to tens

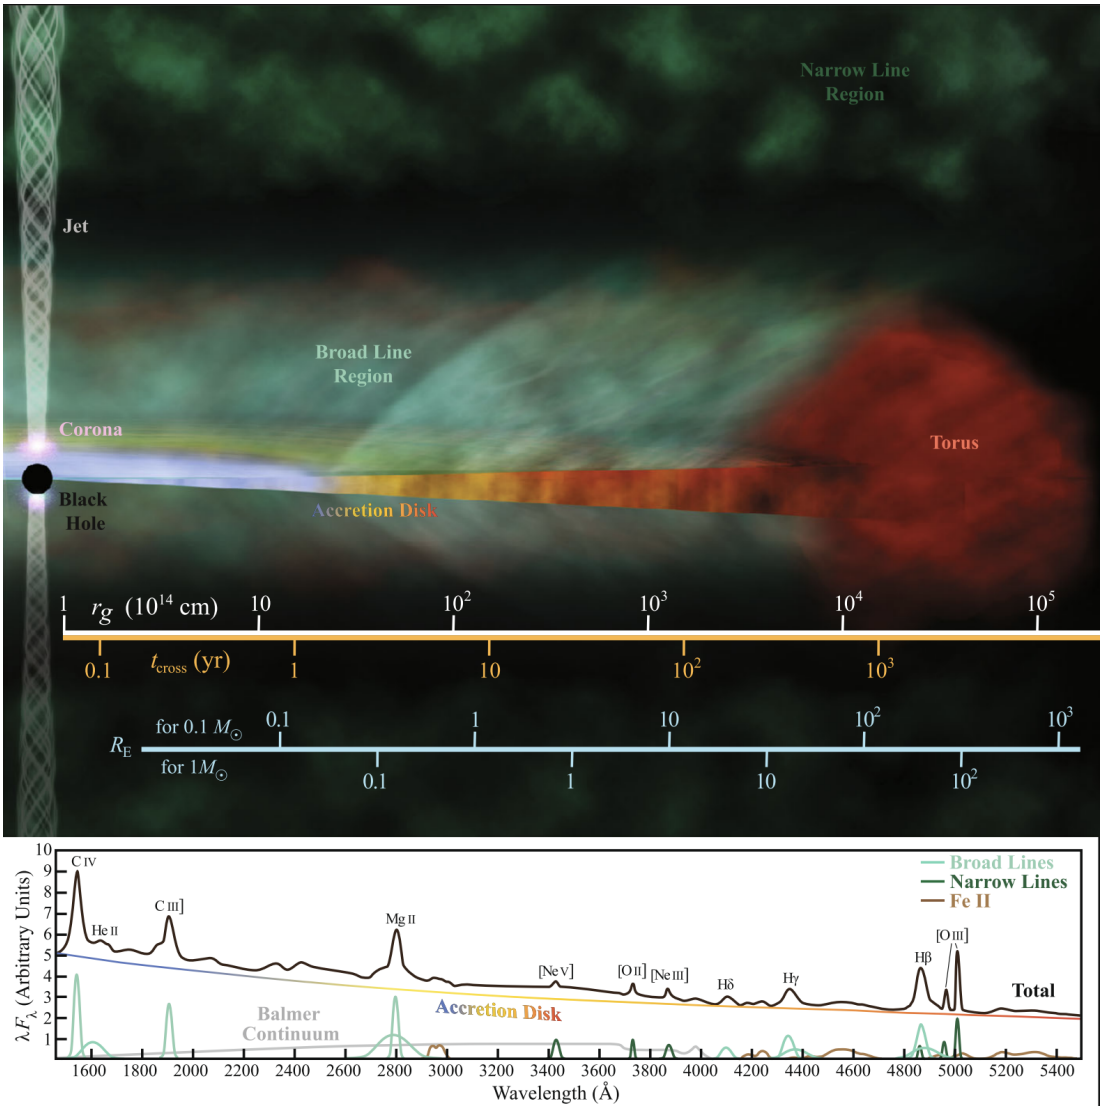


Figure 1.7: Top: Schematic view of the main emitting regions of a quasar, with approximate characteristic sizes in cm. Bottom: Optical AGN spectrum showing how the continuum, broad emission lines, and narrow emission lines trace the accretion disk, BLR, and NLR, respectively. Compact regions, such as the X-ray corona and the optical/UV accretion disk, are the most sensitive to microlensing, whereas extended components, such as the dusty torus and the narrow-line region, are much less affected and can be used as reference emission. From Vernardos et al. (2024).

of light-days in luminous quasars[2] (Fabian et al., 2015; Morgan et al., 2010; Jiménez-Vicente et al., 2012a; Fian et al., 2018). Surrounding the accretion disk is the broad-line region (BLR), which is composed of dense, dust-free gas moving at high velocities. This gas produces permitted emission lines with widths of the order of thousands of $\mathrm{km\,s^{-1}}$ (Shen et al., 2011; Rakshit et al., 2017; Cackett et al., 2021). Reverberation mapping shows that the BLR size increases with luminosity, with characteristic radii ranging from light-days to light-months and, in luminous quasars, reaching scales of hundreds of light-days. These distances are still close enough to the central engine for the gas dynamics to trace the gravitational potential of the SMBH (Kaspi et al., 2004; Bentz et al., 2013; Gravity Collaboration et al., 2018; Cackett et al., 2021). Farther from the SMBH, dust can survive and reprocess the ultraviolet and optical radiation from the central source into infrared emission. This dusty circumnuclear material, often referred to as the torus, begins near the dust sublimation radius, typically at tens to hundreds of light-days in Seyfert-like AGN and at larger, sub-parsec to parsec scales in luminous quasars. Its mid-infrared emission can extend over parsec scales (García-Burillo et al., 2021; GRAVITY Collaboration et al., 2024). The torus is not expected to be a smooth and static structure, but rather a clumpy and possibly wind-like medium whose geometry, covering factor, and physical properties may vary from object to object (Antonucci, 1993; Urry & Padovani, 1995; Ramos Almeida et al., 2011; Netzer, 2015; Padovani, 2017).

On still larger scales, the narrow-line region (NLR) traces lower-density ionized gas extending from the inner few parsecs to hundreds or thousands of parsecs, and in luminous systems it can reach kiloparsec scales (Mullaney et al., 2013; Kakkad et al., 2018). This gas often lies along ionization cones shaped by the obscuring material and by outflows (Mullaney et al., 2013; Harrison et al., 2014; Calderone et al., 2017). Narrow emission lines usually have widths of a few hundred $\mathrm{km, s^{-1}}$, typically FWHM $\sim$ 200–500, $\mathrm{km, s^{-1}}$, although disturbed or outflowing gas can produce broader components with velocities of several hundred up to $\sim$ 1000–1500, $\mathrm{km, s^{-1}}$ in luminous AGN (Mullaney et al., 2013; Harrison et al., 2014). In radio-loud AGN, part of the accretion power can also be carried away by relativistic jets, which may dominate the emission from radio to γ-ray wavelengths when viewed close to the jet axis (Urry & Padovani, 1995; Padovani, 2017).

[2] 1 light-day $= 2.590 \times 10^{15}$ cm $= 8.39 \times 10^{-4}$ pc.

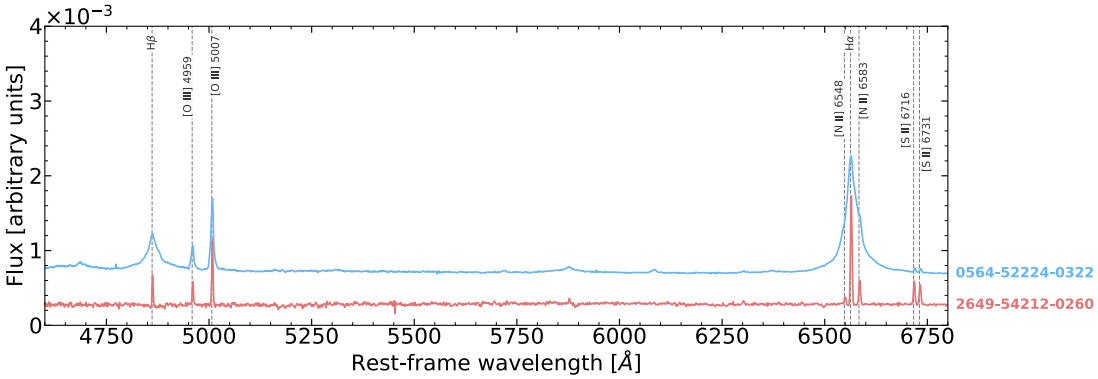


Figure 1.8: Example optical SDSS spectra of type-I and type-II AGN. The type-I spectrum, shown in blue, displays broad permitted emission lines produced in the broad-line region, while the type-II spectrum, shown in red, is dominated by narrow emission lines because the direct view of the central continuum source and broad-line region is obscured.

1.2.2 AGN classification

This structured view of AGN naturally explains part of their observational diversity. In the classical unification scenario, many differences between AGN classes are interpreted as orientation effects (Antonucci, 1993; Urry & Padovani, 1995; Ricci, 2026). If the line of sight to the central engine and BLR is relatively unobscured, the object is observed as a type-I AGN, showing a bright non-stellar continuum and broad permitted emission lines. If the central regions are blocked by dusty material, the same intrinsic source may be observed as a type-II AGN, where the optical spectrum lacks broad lines and is dominated by narrow emission lines from the NLR (see the example in Figure 1.8). This orientation-based interpretation has been highly successful, but it is incomplete. Parameters such as luminosity, accretion rate, torus covering factor, dust content, host-galaxy inclination, outflows, and evolutionary stage can also influence the observed AGN type (Ramos Almeida et al., 2011; Bianchi et al., 2012; Netzer, 2015; Padovani, 2017).

AGN classification is therefore both physically useful and observationally limited. Type-I AGN include Seyfert 1 galaxies and quasars, with the distinction depending mainly on luminosity. In current usage, Seyfert galaxies are often associated with low-to-moderate AGN luminosities, approximately $L_{\rm bol} \sim 10^{42}$–$10^{45}\,{\rm erg\,s^{-1}}$, while quasars are usually identified with higher-luminosity systems, roughly $L_{\rm bol} \gtrsim 10^{45}\,{\rm erg\,s^{-1}}$ (Hickox & Alexander, 2018). Their spectra show broad permitted and semi-forbidden emission lines, such as Hα, Hβ, Mg II, C III], and C IV, together with narrow forbid-

den lines when the narrow-line region is sufficiently strong. Type-II AGN, in contrast, show only narrow emission lines in direct light, although hidden broad-line regions can sometimes be detected through polarized emission (Antonucci, 1993). Intermediate classes, such as type 1.5, type 1.8, and type 1.9 AGN, are used when the relative strength of the broad and narrow Balmer components changes gradually (Mickaelian et al., 2021). For example, type 1.9 AGN show a broad component in Hα but little or no detectable broad component in Hβ (Winkler, 1992). This diversity shows that AGN types should not be interpreted as sharply separated physical classes, but rather as observational manifestations of a complex central structure. A visual example of the spectral differences between type-I and type-II AGN is shown in Figure 1.8.

As was mention above quasars represent the high-luminosity regime of the AGN population their optical and ultraviolet spectra can be dominated by a bright accretion-disk continuum and strong broad emission lines, making them visible over cosmological distances (see Figure 1.7 bottom). Because of this high luminosity, quasars are key probes of SMBH growth across cosmic time and are widely used in large spectroscopic surveys (Schneider et al., 2010; Shen et al., 2011; Lyke et al., 2020; DESI Collaboration et al., 2024). However, even in luminous quasars the observed spectral properties are not uniform. Some objects show strong Fe II emission, broad absorption features, large C IV blueshifts, weak narrow lines, or unusually weak broad emission lines (Boroson & Green, 1992; Crenshaw et al., 2003; Shen et al., 2011; Vietri et al., 2018). These trends are often connected to differences in Eddington ratio, BLR structure, orientation, and disk-wind strength, and are commonly discussed in the context of quasar main-sequence studies and the Population A/B phenomenology. Weak-line quasars, for example, have unusually faint or absent ultraviolet broad emission lines and may reflect unusual BLR covering factors, ionizing continua, or accretion states (Diamond-Stanic et al., 2009; Shemmer et al., 2010; Plotkin et al., 2015). These sources illustrate that the connection between accretion physics, BLR structure, and observed line emission is still not fully understood. The identification and characterization of AGN are also affected by observational selection effects. Optical and ultraviolet surveys are particularly efficient at finding unobscured, radiatively efficient type-I AGN, but they can miss obscured systems, reddened quasars, low-luminosity AGN, and host-dominated nuclei (Richards et al., 2002, 2006; Ross et al., 2012; Padovani, 2017). In low-luminosity or partially obscured sources, the stellar light of the host galaxy can dominate the observed optical spectrum, diluting the AGN continuum and reducing the equivalent widths of emission lines (see Figure 1.9). Stellar

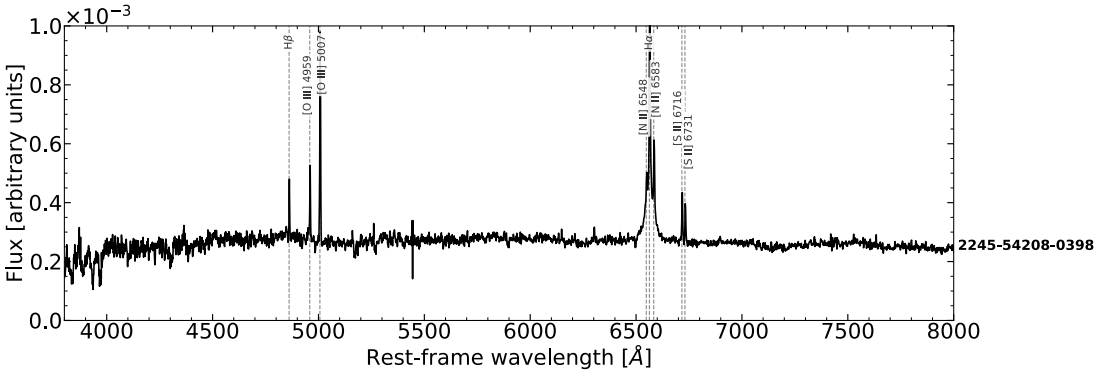


Figure 1.9: Example of an optical SDSS spectrum host-dominated AGN in the restframe. In this regime, stellar continuum and absorption features from the host galaxy can dominate part of the spectrum, making the nuclear AGN contribution more difficult to isolate.

absorption features can also affect the measurement of Balmer lines, while star formation can contribute narrow emission lines that complicate classical diagnostic diagrams (Baldwin et al., 1981; Kauffmann et al., 2003; Kewley et al., 2006). These effects introduce biases in luminosity estimates, line-width measurements, black-hole mass estimates, and Eddington ratios when the host contribution is not properly modelled (Vestergaard & Peterson, 2006; Shen et al., 2011; Popović, 2020).

A further complication is that AGN spectra are produced by several spatially distinct components that are blended in unresolved observations. The accretion disk, BLR, NLR, dusty material, outflows, and host galaxy can all contribute to the same observed wavelength range. Broad Fe II multiplets can form a pseudo-continuum that overlaps with Balmer lines and Mg II, making continuum placement and line decomposition non-trivial (Boroson & Green, 1992; Véron-Cetty et al., 2004; Kovačević et al., 2010). Ionized outflows can produce asymmetric wings and velocity shifts in lines such as [O III] λ5007, while disk winds can strongly affect high-ionization ultraviolet lines such as C IV (Mullaney et al., 2013; Harrison et al., 2014; Vietri et al., 2018). Therefore, although spectroscopy is one of the most powerful tools for studying AGN, the interpretation of AGN spectra requires careful separation of the different physical components contributing to the observed emission (this will be discussed in detail in Chapter 3).

Despite decades of progress, important aspects of AGN structure remain uncertain. Although the different emitting components can be identified through their spectral signatures, as illustrated in Figure 1.7 bottom panel, they remain spatially unresolved in most systems, even with current high-resolution facilities. Event Hori-

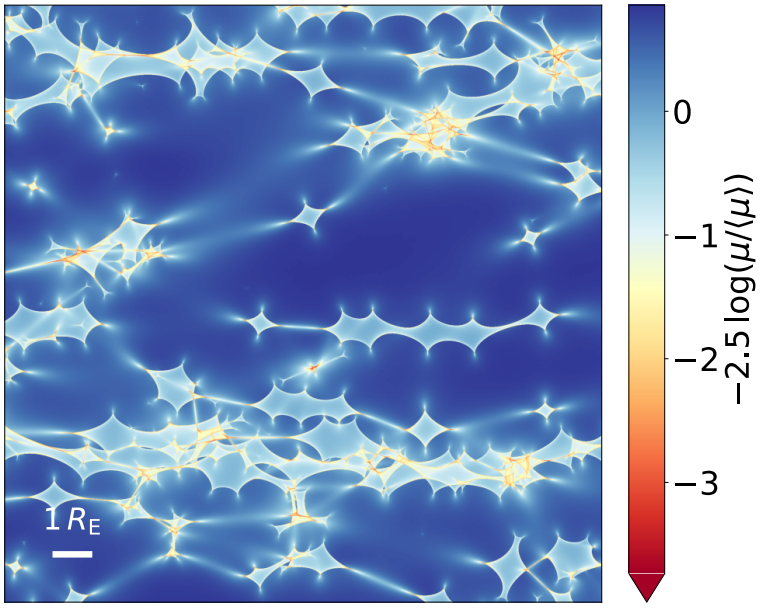


Figure 1.10: Example of a microlensing magnification map computed for $\kappa = \gamma = 0.3$, assuming a mean microlens mass of $\langle M_\star \rangle = 0.3\,M_\odot$. The map covers a region of $15\,R_E$ on each side and has a size of 2000×2000 pixels with smooth matter fraction = 0 (100% stars). It shows the caustic network produced by compact objects in the lens galaxy and represents the magnification pattern experienced by a background source in the source plane. Compact sources are sensitive to individual caustics and can therefore show large magnification fluctuations, whereas extended sources average over the caustic pattern and produce smoother microlensing signals. Colors indicate the microlensing magnification, with red regions corresponding to the highest magnifications.

zon Telescope observations have resolved horizon-scale structure only in the nearest SMBHs, such as M87* and Sgr A*, highlighting the exceptional angular resolution required to image structures on scales of a few gravitational radii (Event Horizon Telescope Collaboration et al., 2019, 2022). As a result, key questions remain regarding the size and temperature profile of the accretion disk, the geometry and kinematics of the BLR, and the connection between the continuum- and line-emitting regions (e.g. Jiménez-Vicente et al., 2014; Cornachione & Morgan, 2020; Vernardos et al., 2024). Quasar microlensing provides one of the few observational methods capable of probing these unresolved scales, and therefore motivates the discussion developed in the following section.

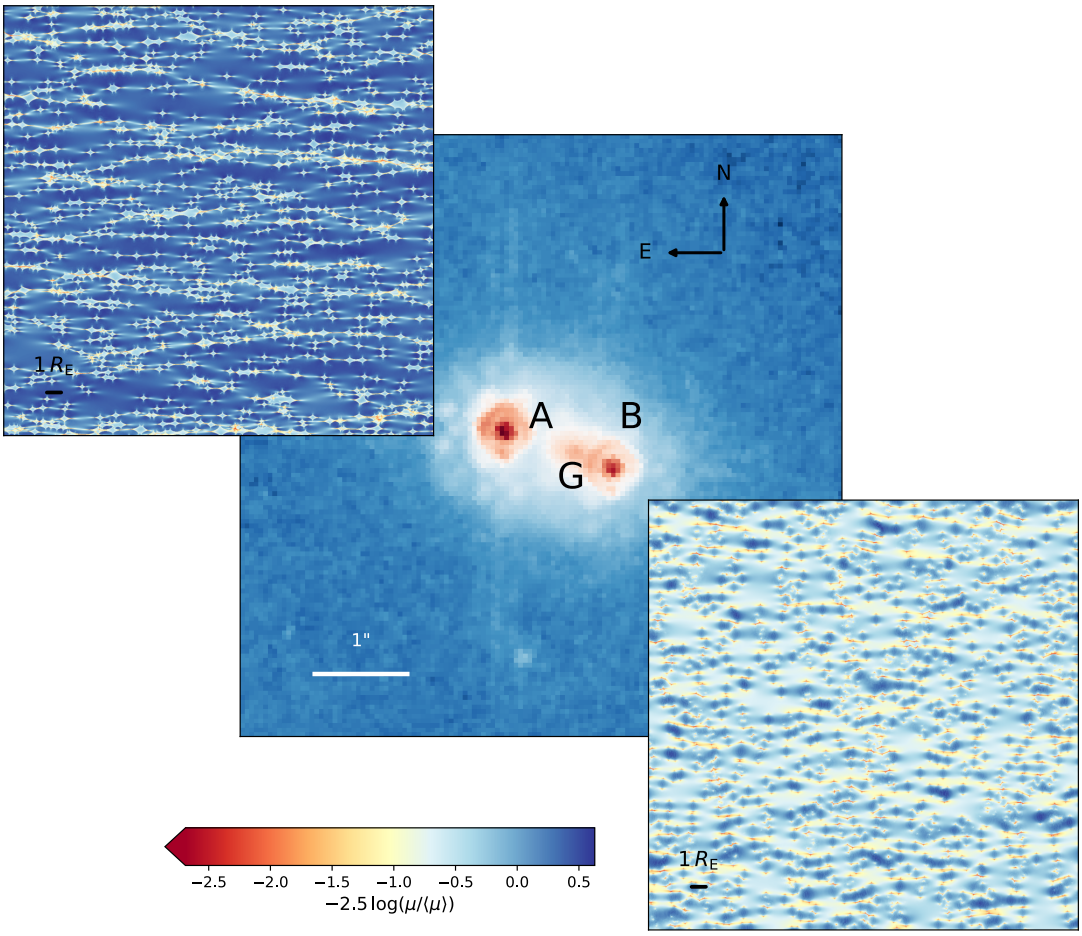


Figure 1.11: HST/NICMOS F160W image of the doubly imaged lensed quasar QJ0158−4325 (Faure et al., 2009), together with the corresponding microlensing magnification maps for images A and B. Each magnification map has a size of 1300×1300 pixels and spans $30\,R_{\rm E}$ in the source plane. The maps were generated assuming a smooth-matter fraction of $s = 0.2$. For image A, the local convergence and shear are $\kappa = 0.38$ and $\gamma = 0.38$, while for image B they are $\kappa = 0.76$ and $\gamma = 0.76$, respectively.

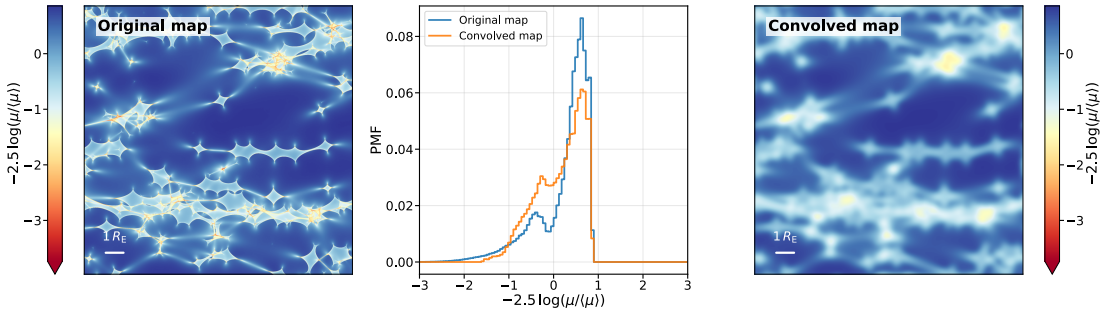


Figure 1.12: Comparison between the original and convolved microlensing magnification maps. Left: original microlensing magnification map, computed for $\kappa = \gamma = 0.3$ and a mean microlens mass of $\langle M_\star \rangle = 0.3\,M_\odot$. The map is the same as in Figure 1.10, spanning 15 R_E with 2000×2000 pixels. Center: probability distribution of the magnification values before and after convolution, showing how finite-source smoothing reduces the contrast of the magnification map. Right: map convolved with a finite source profile of characteristic size $R_S = 0.2\,R_E$, corresponding to approximately 0.8 light-days for $R_E \simeq 10^{16}$ cm. This size is comparable to the expected scale of the optical accretion disk. The convolution smooths the small-scale caustic structure because the finite source averages over the magnification pattern.

1.3 Quasar microlensing

Gravitationally lensed quasars provide a unique way to probe the unresolved structure of AGN. In these systems, the macroscopic gravitational potential of a foreground galaxy produces multiple images of a background quasar, as described in Section 1.1. In addition to this smooth lensing potential, individual stars and compact objects in the lens galaxy introduce small-scale magnification perturbations. This effect, known as gravitational microlensing, generates a network of caustics in the source plane that can differentially magnify the emission from different regions of the quasar (Paczynski, 1986; Schneider et al., 1992; Wambsganss, 2006).

The key property that makes microlensing useful for quasar studies is its dependence on source size. As discussed in Section 1.2.1, quasar emission arises from regions with very different characteristic scales. Compact regions, such as the X-ray corona and the optical/UV accretion disk, are strongly affected by the microlensing magnification pattern. In contrast, more extended regions, such as the BLR, dusty torus, and NLR, average over many caustics and therefore show weaker microlensing signatures. Since the optical/UV accretion disk is much smaller than most line- and dust-emitting regions, microlensing can be used to constrain its characteristic size and wavelength dependence on microarcsecond scales, far below the direct angular reso-

lution of current telescopes (Guerras et al., 2013a; Jiménez-Vicente et al., 2014; Cornachione & Morgan, 2020; Vernardos et al., 2024).

As the source, lens, and observer move relative to one another, different parts of the quasar sample different regions of the magnification pattern (Kayser et al., 1986; Mosquera & Kochanek, 2011). This relative motion can produce time-variable flux ratios between the multiple lensed images, as well as wavelength-dependent or velocity-dependent changes in the observed spectra. These signatures provide information about the spatial extent of the emitting regions and make microlensing a powerful tool for studying the internal structure of quasars (Fian et al., 2018; Jiménez-Vicente et al., 2014; Guerras et al., 2013a; Fian et al., 2024).

1.3.1 Magnification maps

To model quasar microlensing quantitatively, the combined effect of many stellar microlenses is commonly represented using microlensing magnification maps. These maps describe the spatial distribution of magnification in the source plane for a given macro-image and encode the network of micro-caustics produced by stars and other compact objects in the lens galaxy (Paczynski, 1986; Schneider et al., 1992; Wambsganss, 2006; Mediavilla et al., 2006; Vernardos & Fluke, 2013; Vernardos et al., 2024). An example is shown in Figure 1.10. As the source, lens, and observer move relative to one another, different parts of the quasar sample different regions of this pattern and experience different levels of amplification.

Each magnification map is computed for the local lensing conditions of a given macro-image. These conditions are set by the macro-model parameters introduced in Sections 1.1.4–1.1.6, together with the fraction of the surface mass density in compact objects. They determine the density, orientation, and morphology of the micro-caustic network. As a result, different macro-images of the same lensed quasar can exhibit different microlensing properties, depending on their location within the lens potential. This is illustrated in Figure 1.11 for the doubly imaged quasar QJ0158−4325 (Faure et al., 2009), where images A and B have different values of (κ) and (γ), leading to distinct micro-caustic networks in their corresponding magnification maps, since the local convergence and shear determine the density, orientation, and overall structure of the micro-caustics. .

The physical scaling of a magnification map depends on the characteristic mass of the microlenses. If the microlens mass changes from M to M', the physical length

scales in the map scale approximately as $\sqrt{M'/M}$ (Wambsganss, 2006). This mass–length scaling is important when converting map units into physical source sizes. However, the resulting microlensing signal is not determined by this scaling alone. The detailed shape of the caustic network, and therefore the probability and amplitude of magnification fluctuations, also depends on the local lensing conditions, the smooth matter fraction, the stellar mass function, and the effective transverse velocity of the system (Gil-Merino & Lewis, 2005; Mediavilla et al., 2015, 2016; Vernardos et al., 2024).

Magnification maps can also be used to model finite-size sources by convolving the microlensing magnification pattern with a source surface-brightness profile (Kochanek, 2004; Mortonson et al., 2005; Congdon et al., 2007). A compact source samples the magnification pattern locally and can therefore produce a broad magnification probability distribution, including a relatively high probability of strong magnification events. In contrast, an extended source averages over many caustics, smoothing the magnification pattern and suppressing the most extreme magnification values. This behaviour is illustrated in Figure 1.12, where the caustic structure becomes less prominent after convolution with a finite-size source.

The strong dependence of the microlensing signal on source size makes magnification-map analysis particularly useful for studying quasars. Compact regions, such as the X-ray corona and the optical/UV continuum-emitting accretion disk, can show strong microlensing variability, while more extended regions, such as the BLR, are expected to show weaker and smoother microlensing signatures (Morgan et al., 2010; Mosquera & Kochanek, 2011; Jiménez-Vicente et al., 2014). Thus, by comparing the microlensing behaviour of different spectral components, magnification maps provide a framework for connecting observed flux anomalies and variability to the physical size of quasar emission regions.

1.3.2 Microlensing as a probe of AGN structure

The dependence of microlensing on source size makes it a powerful probe of AGN structure. For a point-mass microlens, the relevant physical scale is the Einstein radius projected onto the source plane (Section 1.1.6),

$$R_{\rm E} = D_{\rm S}\theta_{\rm E} = \left(\frac{4G\langle M\rangle}{c^2}\frac{D_{\rm S}D_{\rm LS}}{D_{\rm L}}\right)^{1/2}, \tag{1.41}$$

where $D_{\rm L}$, $D_{\rm S}$, and $D_{\rm LS}$ are angular-diameter distances, and $\langle M \rangle$ is the mean mass of the microlenses. For example, for $z_{\rm L} = 0.5$, $z_{\rm S} = 2.0$, and $\langle M \rangle = 0.3\,M_{\odot}$, this scale is $R_{\rm E} \simeq 2.9 \times 10^{16}\,{\rm cm}$, or approximately 11 light-days. This is comparable to, or larger than, the expected optical/UV accretion-disk size in quasars, as shown schematically in Figure 1.7. A source is therefore compact, in the microlensing sense, when its projected size is smaller than or comparable to $R_{\rm E}$.

Photometric microlensing uses brightness differences, flux-ratio anomalies, or time-variable flux ratios between the multiple lensed images to infer the size of the continuum-emitting region. In this approach, the observables are obtained from imaging or light-curve photometry rather than from a spectral decomposition. In single-epoch studies, the observed image flux ratios are compared with those predicted by a macro-lens model, and the residuals are interpreted as possible microlensing after accounting for other effects that can modify the flux ratios, including differential extinction, intrinsic variability combined with time delays (differences in arriving times between images Refsdal, 1964a), uncertainties in the macro-lens model, such as those associated with the assumed mass profile, its parameters, and the contribution of external shear, and substructure in the lens galaxy (Pooley et al., 2007; Blackburne et al., 2011).

In multi-epoch photometric monitoring, the light curves of the different lensed images are shifted by their time delays and corrected for the macro-magnification. After this correction, the intrinsic quasar variability should be common to all images, while the remaining uncorrelated variability is attributed to microlensing. The standard quantitative approach is to compare the observed microlensing light curves with simulated light curves generated from magnification maps, allowing the source size, effective transverse velocity, and microlens mass scale to be constrained statistically (Kochanek, 2004). This method has been applied to long-term optical monitoring of lensed quasars to measure accretion-disk sizes and their relation with black-hole mass (Morgan et al., 2010, 2018), as well as to individual systems with multi-band or long-term monitoring (Hainline et al., 2013; Cornachione & Morgan, 2020; Goicoechea & Shalyapin, 2025). The temporal signal arises as the source moves across the magnification pattern, as illustrated in Figure 1.14, where a source trajectory crosses different regions of the caustic network. Repeating the analysis at different wavelengths makes it possible to constrain the size–wavelength relation of the source and, consequently, the temperature profile of the accretion disk (Blackburne et al., 2011; Morgan et al., 2010; Cornachione & Morgan, 2020).

A related but distinct approach is single-epoch spectroscopic continuum microlensing. In this case, spectra of the different lensed images are used to measure continuum flux ratios and compare them with a low-microlensing reference, commonly taken from the cores of broad emission lines or from narrow emission lines. Since these line-emitting regions are more extended than the optical/UV continuum (Figure 1.7), their flux ratios provide a baseline against which the continuum microlensing signal can be measured (Motta et al., 2012; Guerras et al., 2013a; Rojas et al., 2020). This method is especially useful for separating continuum microlensing from differential extinction (Motta et al., 2002; Muñoz et al., 2011) and for studying chromatic microlensing (e.g., magnificaiton change trough wavelenght Blackburne et al., 2011), because the continuum flux ratio can be measured as a function of wavelength. By comparing the observed wavelength-dependent microlensing amplitudes with simulated magnification distributions from microlensing maps, one can infer the characteristic size of the accretion disk and its size–wavelength relation (Jiménez-Vicente et al., 2014; Cornachione & Morgan, 2020).

Spectroscopic microlensing generalizes this idea by comparing the magnification of different spectral components. The continuum, Fe II pseudo-continuum, broad emission lines, and narrow emission lines originate on different spatial scales and therefore respond differently to microlensing (Abajas et al., 2005; Mediavilla et al., 2011; Sluse et al., 2012; Guerras et al., 2013a,b). The optical/UV continuum is usually the most compact component in the observed optical spectrum and is expected to be the most strongly affected. The Fe II pseudo-continuum and the broad-line region occupy intermediate scales, so their microlensing signals are generally weaker than that of the continuum but stronger than those of more extended components. Applying microlensing analyses to these spectral features can therefore constrain the relative sizes of the continuum, Fe II, and BLR-emitting regions (Guerras et al., 2013a; Jiménez-Vicente et al., 2014; Jiménez-Vicente & Mediavilla, 2025).

The broad-line region is particularly useful because different velocity components of a broad emission line may arise from spatially distinct regions. Although the BLR as a whole may be only weakly microlensed, the high-velocity line wings can be more strongly affected than the line core if the wing-emitting gas is located closer to the central continuum source (e.g. Abajas et al., 2002; Sluse et al., 2012; Guerras et al., 2013a; Motta et al., 2017; Fian et al., 2021; Hutsemékers et al., 2024). This differential magnification can distort the line profile and thereby constrain not only the size of the BLR, but also its geometry and kinematics. The narrow-line region, in contrast, is usually

sufficiently extended to be only weakly affected by microlensing and is therefore often used as a reference component.

Microlensing therefore acts as a natural spatial filter: each component is magnified according to its characteristic size, with smaller regions generally producing stronger microlensing signals. By comparing the microlensing signal in the X-ray continuum, optical/UV continuum, Fe II pseudo-continuum, broad emission lines, and narrow emission lines, one can infer the relative spatial scales of the emitting regions. This makes quasar microlensing one of the few observational methods capable of probing microarcsecond angular scales, corresponding to light-day to light-month physical scales in distant AGN, far below the angular resolution of current telescopes (Guerras et al., 2013a; Jiménez-Vicente et al., 2014; Cornachione & Morgan, 2020; Vernardos et al., 2024).

The connection with spectral modeling is therefore essential. To interpret spectroscopic microlensing, the observed spectrum must be separated into physically meaningful components: the AGN continuum, Fe II emission, broad and narrow emission lines, possible outflow components, and, when relevant, the host-galaxy stellar continuum. This decomposition cannot assume that all components have the same relative strength in every lensed image, because microlensing can change the shape and relative contribution of compact spectral components from one image to another. This is particularly important for host-dominated or low-contrast AGN, where the stellar contribution can dilute the nuclear signal and bias the measurement of the components used in microlensing analyses. Robust spectral decomposition is therefore the basis for identifying which parts of the spectrum are compact enough to be microlensed and which parts can be treated as extended reference emission.

In this context, AGN spectral modeling and quasar microlensing are complementary tools. Spectral decomposition identifies the physical components that contribute to the observed emission, while microlensing provides information about their relative spatial scales. This complementarity is especially important for lensed quasars, where microlensing transforms unresolved spectral differences into measurable constraints on the accretion disk, the BLR, and other compact AGN regions.

1.3.3 High-magnification events

One of the most informative manifestations of quasar microlensing is the occurrence of high-magnification events. These events are commonly associated with the passage

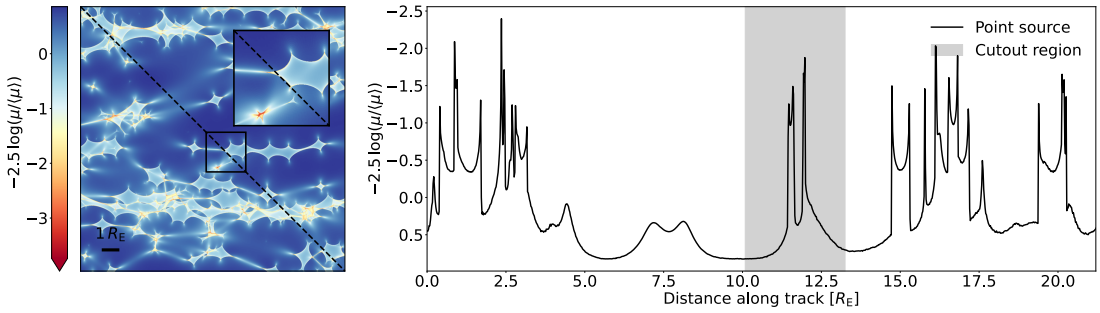


Figure 1.13: Same microlensing magnification map as in Figure 1.10, computed for $\kappa = \gamma = 0.3$ and a mean microlens mass of $\langle M_\star \rangle = 0.3, M_\odot$, with an example source trajectory overlaid. The right panel shows a zoomed-in view of the caustic-crossing event highlighted in grey, where the source passes through a high-magnification region and produces a clear increase in the microlensing signal.

of a compact source close to, or across, a micro-caustic in the magnification pattern, corresponding to the high-magnification regions shown in grey in Figure 1.13. During such an event, the observed brightness of one lensed image can increase by several tenths of a magnitude, and in some cases by more than one magnitude, over a finite period of time, while the other macro-images may remain comparatively less affected. Because the caustic pattern effectively scans the source, high-magnification events can provide particularly strong constraints on the source size, surface-brightness profile, and, in favorable cases, the structure of the innermost accretion disk (Anguita et al., 2008; Mediavilla et al., 2015; Vernardos et al., 2024).

The probability, amplitude, and duration of high-magnification events depend on the structure of the magnification map, the size of the source, and the effective transverse velocity between the observer, lens, and source. Compact sources can produce sharper and higher-amplitude events because they respond strongly to individual caustics as can be seen in Figure 1.14. Larger sources smooth the caustic structure and therefore produce weaker and broader variations. The event timescale is therefore not only a property of the lens system, but also depends on the physical size of the emitting region being microlensed and on the effective velocity with which the source crosses the caustic network (Mediavilla et al., 2015, 2016).

High-magnification events are especially useful because they can contain more spatial information than low-amplitude microlensing variability. In the idealized case of a fold-caustic crossing, the observed light curve can be interpreted as the convolution of the caustic magnification profile with the surface-brightness distribution of

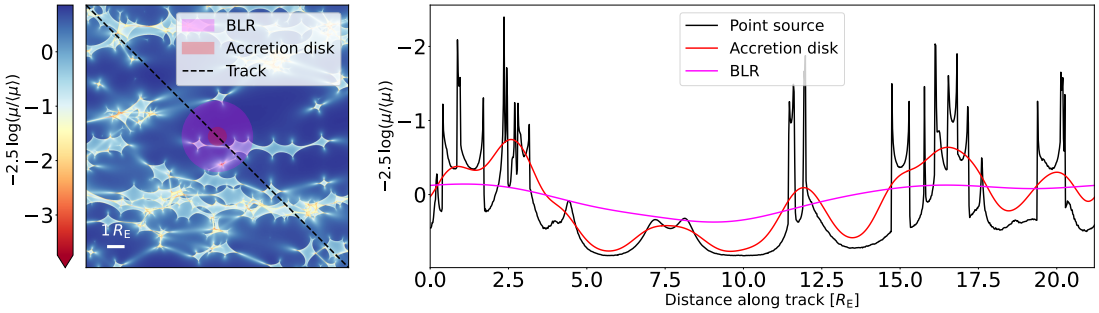


Figure 1.14: Same microlensing magnification map as in Figure 1.10, computed for $\kappa = \gamma = 0.3$ and a mean microlens mass of $\langle M_\star \rangle = 0.3\, M_\odot$, with an example source trajectory overlaid. The track illustrates how a source moving across the magnification pattern samples different regions of the caustic network, producing a microlensing light curve whose variability depends on the source size and on the caustics crossed along the path. Compact emitting regions, such as the accretion disk, are more strongly affected by individual caustics and can experience larger magnification fluctuations, whereas more extended regions average over the caustic pattern and produce smoother, lower-amplitude microlensing signals.

the source. Such events are rare and difficult to identify unambiguously, but some of the best-studied examples occur in the Einstein Cross Q2237+0305 (see Figure 1.15). In this system, high-magnification microlensing events have been modeled as probable caustic crossings and used to constrain the half-light radius of the accretion disk, as well as to search for possible signatures of structure in the innermost disk region (Anguita et al., 2008; Mediavilla et al., 2015).

However, real microlensing events can be affected by complex caustic networks, intrinsic quasar variability, observational cadence, photometric uncertainties, and the uncertain effective transverse velocity. These effects must be considered when interpreting high-magnification events (HME).

The importance of high-magnification events is expected to increase with large time-domain surveys, The work of Neira et al. (2020) in where using simulations they estimate hundred events per year to be observed by surveys like LSST (Ivezić et al., 2019). The discovery and monitoring of larger samples of lensed quasars will make it possible to identify ongoing microlensing events and trigger spectroscopic or multi-wavelength follow-up observations during the phases of highest information content. Such observations are particularly valuable because they can probe the most compact regions of quasars, including the accretion disk and possibly the X-ray corona, on spatial scales that cannot be resolved directly (Neira et al., 2020; Fagin et al., 2024; Neira et al., 2025). Thus, high-magnification events provide a dynamic and complementary

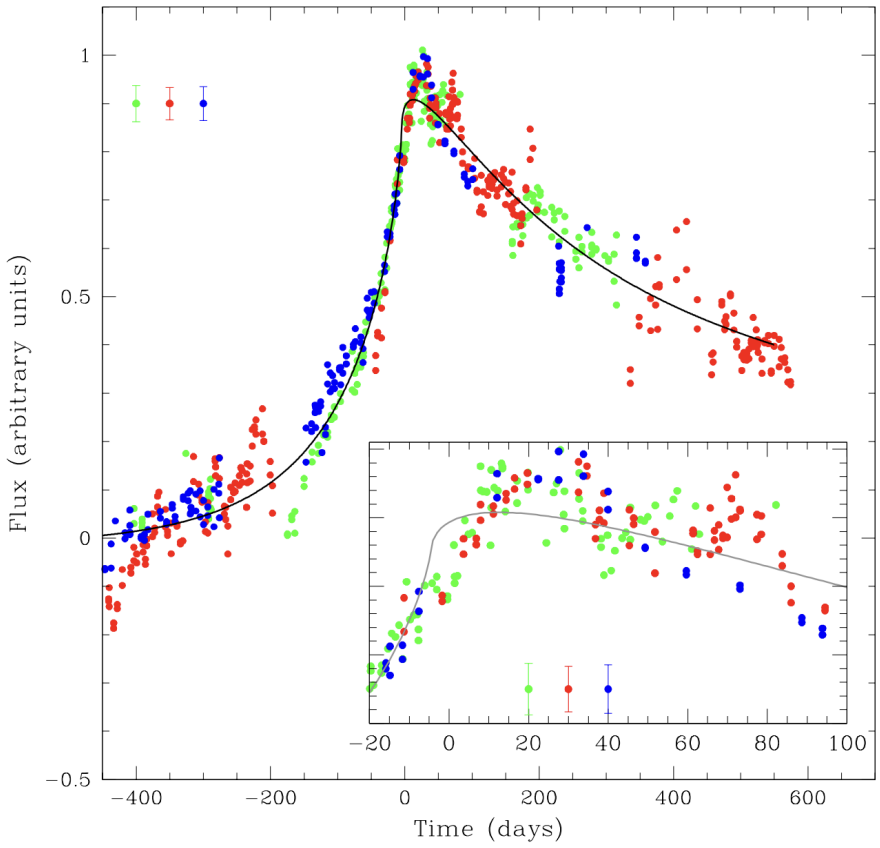


Figure 1.15: Global comparison of the three HMEs and analytical fit to a single caustic crossing in Q2237+0305. The inset is a zoom in the region of the maximum. Dots correspond to the experimental light curves of each of the three events. From Mediavilla et al. (2015)

route to connect microlensing variability with the physical structure of quasars.

Throughout this introduction, we have presented the main topics that motivate this thesis, linking the study of AGN with gravitational microlensing as a powerful tool to investigate their unresolved inner regions.

1.4 Motivation

The number of known gravitationally lensed quasars has increased substantially during the last decade. Around fifteen years ago, microlensing population studies were based on samples of only of order one hundred systems; for example, Mosquera & Kochanek (2011) estimated microlensing properties for a sample of 87 lensed quasars. Today, the number of known lensed quasars is of the order of several hundred (e.g., 347 confirmed systems; Ducourant et al., 2026), with many additional candidates reported by recent searches (e.g. Lemon et al., 2023; Euclid Collaboration et al., 2025; Karp et al., 2026). This population is expected to grow further with current and forthcoming wide-field surveys and missions, such as LSST (Ivezić et al., 2019), Euclid (Euclid Collaboration et al., 2025), and the large spectroscopic datasets produced by facilities such as the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al., 2016a, 2024) and the 4-metre Multi-Object Spectroscopic Telescope (4MOST; de Jong et al., 2019) provide a major opportunity to move beyond individual-object studies toward statistical analyses of quasar microlensing.

Microlensing is one of the few techniques capable of probing the innermost regions of quasars on spatial scales far below the angular resolution of current telescopes (Vernardos et al., 2024). In particular, it can constrain the size of the accretion disk, test the wavelength dependence of the continuum-emitting region, and probe the relative spatial scales of the BLR, Fe II emission, and other compact components (Guerras et al., 2013a; Jiménez-Vicente et al., 2014; Cornachione & Morgan, 2020; Fian et al., 2021). It can also identify high-magnification events, in which the passage of the source close to a microlensing caustic produces a strong and time-dependent amplification. Such events are especially informative because they can provide tight constraints on the source size and surface-brightness profile (Anguita et al., 2008; Mediavilla et al., 2015; Neira et al., 2025). However, despite its potential, quasar microlensing is still limited by the relatively small number of well-monitored lensed quasars and the complexity of predicting, for an individual system, when a strong microlensing event will occur (Schechter et al., 2014; Neira et al., 2020). This limitation is particularly relevant

in the Rubin-LSST era: although LSST will monitor large samples of lensed quasars, its nominal cadence may not be sufficient to fully resolve the light curves of short high-magnification events. Therefore, identifying systems with a high probability of undergoing observable microlensing variability will be essential for triggering complementary high-cadence photometric or spectroscopic follow-up (Neira et al., 2020; Fagin et al., 2025).

However, the probability of observing high-magnification microlensing events has often been estimated either from synthetic populations of lensed quasars (e.g. Neira et al., 2025) or from simplified estimates based on characteristic microlensing scales and timescales (Mosquera & Kochanek, 2011). Synthetic approaches are powerful because they can forecast the number of events expected in future surveys, but they necessarily depend on assumptions about the intrinsic population of lensed quasars, the lens-galaxy population, and the survey selection function (e.g. Oguri & Marshall, 2010; Yue et al., 2022). These assumptions are important because current samples of known lensed quasars are not purely representative: lens discovery is affected by image separation, angular resolution, depth, contrast between the quasar images and the lens galaxy, and the ability to obtain reliable redshifts for the deflector (Lemon et al., 2024). These effects can bias the set of systems for which detailed lens models and microlensing predictions are available.

Simplified point-source or scale-based approaches are useful for estimating characteristic microlensing timescales and for identifying promising systems for monitoring. However, they do not fully describe the finite size and surface-brightness profile of the quasar emission regions. They also cannot capture the full effect of the granular stellar mass distribution in the lens galaxy, which produces the complex magnification patterns sampled by an extended moving source (Wambsganss, 2006; Millon et al., 2020; Vernardos et al., 2024). Therefore, in the LSST era, it will be important not only to estimate how many high-magnification events may occur, but also to determine which individual lensed quasars are most likely to show observable microlensing variability and over what timescales. This information is essential for optimizing complementary high-cadence photometric or spectroscopic follow-up, since the nominal LSST cadence alone may not be sufficient to fully characterize the light curves of the most informative events (Fagin et al., 2025; Neira et al., 2025).

Spectroscopy provides a complementary and more physically informative framework for detecting and interpreting microlensing. The observed microlensing signal is obtained by comparing the flux ratios of different lensed images after accounting for

the macrolens magnification, extinction, time delays, and intrinsic variability. In spectroscopic microlensing, this comparison is performed separately for different spectral components. Compact components, such as the accretion-disk continuum, are expected to be more strongly microlensed, whereas extended components, such as the NLR, are usually weakly affected and can be used as reference emission (Abajas et al., 2005; Mediavilla et al., 2011; Guerras et al., 2013a,b). Therefore, the microlensing amplitude inferred from a spectrum depends directly on how well the AGN components are separated.

Although several works have used spectroscopic samples to study microlensing statistically, these samples have generally been compiled from heterogeneous observations rather than from a homogeneous and systematic spectroscopic survey of all known lensed quasars. For example, Mediavilla et al. (2009) used literature spectra to measure continuum-to-emission-line flux-ratio anomalies in multiple image pairs, while Jiménez-Vicente et al. (2012b) used these measurements to infer the average accretion-disk size from a sample of lensed quasars. More recent studies have extended this approach to broad emission lines and BLR structure, but still rely on spectra obtained from different epochs and observational programs (e.g. Fian et al., 2018, 2021). This motivates the need for more homogeneous spectroscopic analyses, where the continuum, Fe II, broad-line, and narrow-line components are modeled consistently across the sample. This connection makes AGN spectral modeling a central requirement for microlensing studies. To measure a reliable microlensing signal, the observed spectrum must be decomposed into physically meaningful contributions, including the continuum, broad emission lines, narrow emission lines, Fe II pseudo-continuum, possible outflow components, and, when present, the host-galaxy and lens-galaxy light. If these components are blended or incorrectly modeled, the inferred flux ratios can mix compact and extended emission, biasing the microlensing measurement. This is particularly important for large future samples, where manual modeling is not feasible and the analysis must be automated, flexible, and robust.

Current methodologies have made significant progress in AGN spectral decomposition and microlensing analysis, but several limitations remain. Some approaches are optimized for individual systems rather than large samples (Sexton et al., 2021; Cappellari, 2023), while others do not flexibly account for contamination by the lens galaxy, host galaxy, or complex line components (Ilić et al., 2023; Ren et al., 2024). At the same time, upcoming spectroscopic surveys will produce data volumes that require efficient tools capable of treating heterogeneous spectra in a consistent way

(Merloni et al., 2019; Bauer et al., 2023; DESI Collaboration et al., 2024). This motivates the development of analysis methods that combine physical flexibility with computational scalability.

The present work is motivated by two main objectives. First, we aim to quantify how often high-magnification microlensing events can occur in known lensed quasars, using microlensing simulations that account for the finite size of the emitting regions and the magnification patterns produced by compact objects in the lens galaxy. This objective connects the observed population of lensed quasars with the expected frequency and timescale of events that can be targeted by future monitoring and follow-up campaigns.

Second, we aim to develop a robust and efficient tool to characterize quasar structure from spectroscopic data. This tool is designed to separate the main AGN components that enter microlensing analyses, including the continuum, BLR, NLR, Fe II emission, and possible host- or lens-galaxy contamination. By improving the measurement of the spectral components used to compute microlensing signals, this approach provides a direct connection between AGN modeling and the physical interpretation of quasar microlensing.

Together, these two objectives address both the statistical and methodological challenges of the coming era of lensed-quasar studies. The first objective evaluates the occurrence of microlensing events in the known population, while the second provides the spectral-modeling framework required to interpret such events in large and heterogeneous samples.

CHAPTER 2

Microlensing time-scales and flux magnification probabilities of a sample of 204 lensed quasars

Gravitationally lensed quasars constitute one of the most powerful laboratories in modern astronomy, the theory associated to them is presented in Section 1.1. In particular, when high-magnification events (HMEs) occur, these systems become especially valuable for probing both the structure of quasars and the mass distribution of the lensing galaxies. However, the frequency of such events, and the probability of observing them in known lensed quasars, has remained an open question and has been the subject of considerable discussion over the past decades.

In this chapter, we investigate how frequently microlensing events may occur and quantify the probability associated with HMEs. To address this problem, we complement the standard estimates based on Einstein-ring crossing times with a more realistic treatment that accounts for the optical depth of lensed quasar images. In particular, we compute magnification maps for every image in our sample, which allow us to estimate not only the relative probability of different levels of microlensing activity, but also their characteristic timescales. To obtain statistically significant magnification probabilities for a large number of systems, we use the fast multipole method–inverse

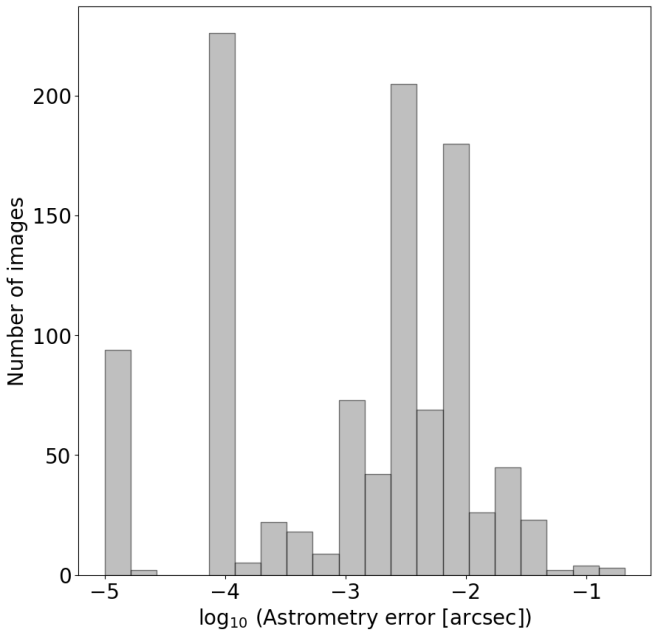


Figure 2.1: Astrometric error distribution for the sample of images modeled.

polygon mapping technique (FMM–IPM; Jiménez-Vicente & Mediavilla, 2022), which makes feasible the efficient computation of sufficiently large microlensing magnification maps for a broad sample of lensed images.

Beyond the event frequency itself, the amplitude of the microlensing-induced flux variations is another key property that can be explored through these magnification maps. As a first step toward this analysis, we perform a homogeneous and simple parametric automatic modeling of all gravitationally lensed quasars in the sample, with the aim of minimizing systematics associated with the lens-modeling procedure. This approach also provides a complete and uniform set of lens-model parameters, which is valuable for studying the global properties of the population of lensed quasar systems.

This work was already accepted for publication in Astronomy & Astrophysics and the chapter is organized as follows. Section 2.1 presents the dataset. Section 2.2 describes the macro-modeling procedure, the computation of magnification maps, and the derivation of the relevant timescales and probabilities. The results are presented in Section 2.3 and discussed in Section 2.4. Finally, the main conclusions are summarized in Section 2.5. Throughout this chapter, we adopt a cosmology with $\Omega_m = 0.3$ and $\Omega_\Lambda = 0.7$.

2.1 Data

The first uniform and high-resolution image database of gravitationally lensed quasars was provided in the 1990s by the CfA–Arizona Space Telescope LEns Survey[1] (CASTLES; Muñoz et al., 1998). Using the Hubble Space Telescope (HST), broad-band images were obtained for $\sim$ 100 systems, primarily in the filters F160W, F814W, and F555W, including 15 galaxy-scale quadruply lensed quasars.

Decades later, in 2015, the Strong Lensing Insights into the Dark Energy Survey (STRIDES; Treu et al., 2018) initiated a new search for lensed quasars in the Dark Energy Survey The Dark Energy Survey Collaboration (2005), aiming to discover new quadruply lensed quasars and to select the best candidates for a time-delay follow-up and H_0 estimation. The initial search was subsequently expanded to incorporate Gaia data (Krone-Martins et al., 2018), which enabled the exclusion of stars as contaminants (Lemon et al., 2017; Agnello et al., 2018; Lemon et al., 2023).

A few years ago, the Gravitationally Lensed Quasar Database (GLQD) listed the names and references of most known lensed quasars. However, the information provided for the systems was not homogeneous, and the page has since been taken down. New projects, such as the Strong Lensing Database (SLED[2]), have emerged to fill this gap, aiming to provide a unified and comprehensive database of strong lensing systems. However, it still does not provide complete information (e.g., photometry in multiple bands for the images) for all the known systems.

Considering this inhomogeneity, we construct a dataset to identify the largest collection of systems containing photometric information for the images, relative astrometry, and the redshifts of both the lens and source. To achieve this, we began our search with the systems published in the CASTLES database, where objects are reported with their photometric data in the HST bands, relative astrometry, and redshift information for the source and lens. Next, we turned to the GLQD and search the literature for photometric and relative astrometric data on the latest discovered systems. We also included systems not listed in either of the databases above. Since the redshift estimates for the source and lens are not always reported in the same publication, we searched for the most recent publicly available values. Our search produced an output sample of 314 known systems, each with different levels of information. This census can be accessed online[3].

[1] https://lweb.cfa.harvard.edu/castles/
[2] https://sled.amnh.org/
[3] https://github.com/felavila/qumas/blob/main/examples/how_to_use_census.ipynb

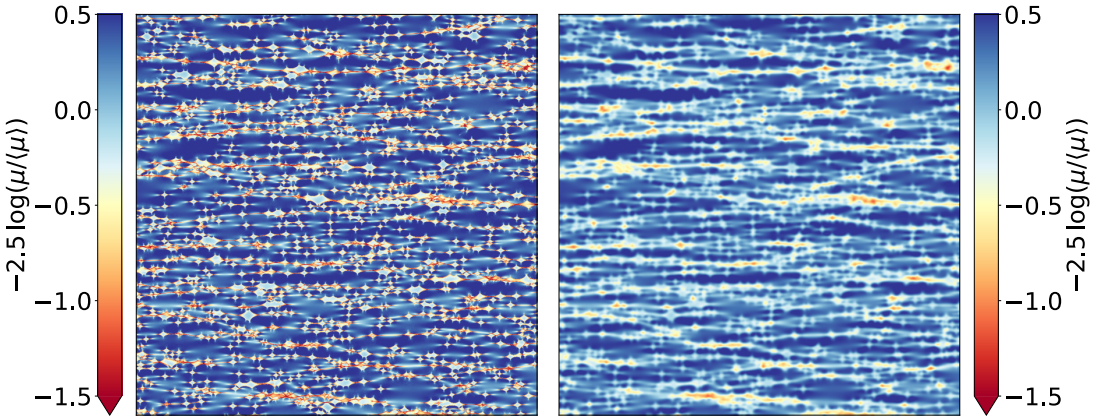


Figure 2.2: Microlensing magnification patterns produced by stars in the lensing galaxy corresponding to image A of the system DES2158−5812. *Left:* Original magnification map. The color scale encodes different magnification levels, while the diamond-shaped curves trace the caustics. *Right:* Same map after convolution with a Gaussian profile with a sigma size of 3.72 [pix], showing the effective magnification pattern. Both panels have a size of 1200 × 1200 pixels with a pixel size of 1.1 [pix/light-days] and 60 R_E. The maps were produced using the code developed in Jiménez-Vicente & Mediavilla (2022).

In this chapter, we also compare our estimates of microlensing magnification probabilities with the observed microlensing histogram presented by Mediavilla et al. (2009, hereafter MED09; see also Mediavilla et al. 2024). This empirical histogram is constructed from single-epoch microlensing measurements of image pairs and, therefore, it does not preserve the sign of the magnification. For this reason, we also computed the probability distributions for the unsigned value of the microlensing magnification.

2.2 Methods

2.2.1 Macro lens modeling

Over the last decades, different lens modeling software (Lefor et al., 2013) with a variety of approaches have been applied to study individual systems (Shajib et al., 2022), meanwhile, the use of machine learning techniques is starting to appear to tackle the upcoming thousands of objects that will be discovered (Andika et al., 2025; Erickson et al., 2025). For example, to model complex systems using high-resolution images (i.e., quadruply lensed quasars with extended arcs), several codes exist, in particular

the two well-tested modeling codes `Lenstronomy` (Birrer et al., 2015; Birrer & Amara, 2018) and `Glee` (Gravitational Lens Efficient Explorer; Suyu et al., 2010). Both codes were used in auto-modeling task (see Schmidt et al., 2023; Ertl et al., 2023), concluding that for some of the systems, human supervision is still required. In fact, a few of the modeling codes offer a native automatic process. One example is `AutoLens` (Nightingale et al., 2021). This code uses a non-linear search to determine the set of light and mass profile parameters that best fit the data using images of the objects (see more details in `pyautolens` docs [4]).

Another difficulty is that the kind of observational data required by the most complex codes (e.g., line-of-sight velocity dispersion, high-resolution images Knabel et al. 2025) are not available for a significant part of the known quasar lenses.

As our aim here has been to study as many systems as possible using a homogeneous procedure, we adopted a straightforward approach based solely on photometric data from the lensed quasar images and on the relative positions between the images and the lens galaxy, which are typically available in the literature. We selected `Lensmodel` (Keeton, 2001b, 2011), a parametric gravitational lens modeling code that the community has extensively tested for over two decades. This code supports a wide range of analytic mass distributions and can fit lens systems using constraints from image positions and flux ratios. The modeling can be performed either in the image plane or in the source plane. In our case, we adopted the image-plane approach, in which the code minimizes a standard chi-square statistic of the form,

$$\chi^2 = \sum_i \frac{|\theta_{i,\mathrm{mod}} - \theta_{i,\mathrm{obs}}|^2}{\sigma_{i,\theta}^2} + \sum_i \frac{(f_{i,\mathrm{mod}} - f_{i,\mathrm{obs}})^2}{\sigma_{i,f}^2}, \tag{2.1}$$

where $\theta_{i,\mathrm{mod}}$ and $\theta_{i,\mathrm{obs}}$ are the modeled and observed image positions (in arcseconds) respectively, with positional uncertainties, $\sigma_{i,\theta}$. In addition, $f_{i,\mathrm{mod}}$ and $f_{i,\mathrm{obs}}$ are the modeled and observed fluxes with flux uncertainties, $\sigma_{i,f}$. The goal is to determine the set of lens model parameters that minimizes χ^2, thereby yielding the best match between the model predictions and the observational constraints.

Building on this well-established code, we developed quasar microlensing analysis (QuMAS)[5]. It is a Python wrapper that implements an automatic modeling pipeline and the methodologies presented in this work. This framework enables fully automated lens modeling, allowing us to systematically analyze a large number of systems with a consistent methodology. Below, we detail the steps of this automatic

[4] `https://pyautolens.readthedocs.io/`

[5] `https://github.com/felavila/qumas`

pipeline. In Sect. 2.3.1, we compare our results with those from more complex or sophisticated models to assess how much a single model that captures the essence of the phenomenon, can differ from a more elaborated one. Next, we explain our automatic modeling procedure and its application to our current data sample. The steps of the procedure as as follows.

We start by modeling the 314 systems using the most recent and accurate photometry and astrometry available in the I band ($\lambda = 8140$ Å) or similar. The first filter checks whether the system contains known relative astrometry and photometry for all images, which results in a sample of 204 systems that can be modeled. The initial model applied is a singular isothermal sphere (SIS, Keeton, 2001a), the simplest mass distribution and a common first approximation in many cases. The initial conditions for image uncertainties are derived from the data census. The galaxy position uncertainty is fixed at 10^{-3} arcseconds, providing a strong constraint, and the fluxes are also taken from the census with an associated error of 20%. This uncertainty in the fluxes accounts for the effects of microlensing in one or more images, which can introduce additional magnification. The value is chosen as a conservative estimate based on the literature (Morgan et al., 2010; Vives-Arias et al., 2016; Morgan et al., 2018), where the reported uncertainties range between 10% and 50%.

To evaluate the convergence of the model, in addition to a χ^2 value, we also assess the agreement between the model and the data by computing the maximum separation between the calculated and the observed image positions($\Delta\theta_{\rm max}$), defined as

$$\Delta\boldsymbol{\theta}_i = \boldsymbol{\theta}_{i,\rm mod} - \boldsymbol{\theta}_{i,\rm obs}, \tag{2.2}$$

$$\Delta\theta_i = |\Delta\boldsymbol{\theta}_i|\,, \tag{2.3}$$

$$\Delta\theta_{\rm max} = \max \Delta\theta_i. \tag{2.4}$$

This choice was motivated by the potential uncertainties associated with flux measurements and our threshold for determining the quality of the modeling. We consider a model successful if $\Delta\theta_{\rm max} \leq 10^{-2}$ arcseconds. This threshold is based on the error distribution of the census images (Fig. 2.1). Although the specific median positional error was 10^{-3} arcseconds, we adopted a threshold five times larger to allow for a simpler model.

If the model is successful, we take the values of the convergence (κ), which represents the dimensionless surface mass density of the lens; the shear (γ), which is associated with the distortion induced by the tidal gravitational field of the lens and its environment, along with the Einstein radius, given its proportionality to the lens

galaxy mass, as provided by `Lensmodel`. We then estimate the uncertainties using the Markov chain Monte Carlo (MCMC) method implemented in `Lensmodel`. The MCMC is initialized using the modeling parameters as Gaussian priors, and is run with ten chains and a maximum of 10,000 steps.

If the condition ($\Delta\theta_{\rm max} \leq 10^{-2}$ arcseconds) is not satisfied, the system will be re-modeled based on the number of images. For doubles, we consider models such as singular isothermal ellipsoid (SIE, Kormann et al., 1994), or SIS with γ, where shear represents an external gravitational perturbation introducing ellipticity. For quads, the modeling options expand to include SIE, SIS+γ, power-law mass distributions (POW, Keeton, 2001a), SIE+γ, or POW+γ. The larger number of images in quads allows for testing a greater variety of mass distributions due to the increased number of free parameters.

If the system cannot be successfully modeled, we return to step 3 and modify the initial conditions. We first increase the images flux error to 50%, keeping the lens galaxy position error unchanged. Next, we set the flux error to 20% and increase the lens galaxy position error to 0.1 arcseconds. These changes help address scenarios where external shear effects are insufficient to characterize distortions and the centroid of the potential does not coincide with the center of the light distribution of the lens galaxy. Finally, we run a model with a flux error of 50% and a lens galaxy position error of 0.1 arcseconds. If no modeling option succeeds after all these attempts, to preserve the maximum number of systems in the sample, we end up selecting the model with the smallest $\Delta\theta_{\rm max}$ and `Lensmodel` χ^2.

2.2.2 Microlensing timescales

The relative motion between the observer, lens, and source results in the time-dependent microlensing effect. Two temporal scales can be associated with microlensing (for a recent review see Sect. 3.5.6 of Vernardos et al., 2024,). One is the basic temporal scale standard, t_E, which is defined as the time it takes for the source to cross the Einstein radius of a single, isolated microlens (R_E) and can be written as

$$t_E = \frac{R_E}{v}, \tag{2.5}$$

where v corresponds to the effective velocity of the source, and R_E is, defined as

$$R_E = D_{OS}\left[\frac{4GM_\star}{c^2}\frac{D_{LS}}{D_{OL}D_{OS}}\right]^{1/2}, \tag{2.6}$$

where D_{OS}, D_{LS}, and D_{OL} correspond to the angular diameter distances between obs-source, lens-source, obs-lens, G is the gravitational constant, c is the speed of light, and $M_\star$ corresponds to the mass of the star acting as a microlens. To calculate R_E, we assume a value of $M_\star = 0.3 M_\odot$. Two reasons drive this assumption: 1) this is the value estimated by Holtzman et al. (1998) observing the Baade's window in the Galactic bulge using Hubble space telescope; and 2) this value has been used largely in previous microlensing studies (see Mao, 2012; Jiménez-Vicente et al., 2014; Fian et al., 2018; Paic et al., 2022).

According to M&K, the t_E timescale is expected to be on the order of decades (see Smith et al., 2009; Bonaca et al., 2017). However, if the source crosses a compact region of very high magnification such as one of the caustic curves (Blandford & Narayan, 1992) shown in Fig. 2.2, it is possible to observe microlensing variations in shorter timescales; this timescale is known as the crossing timescale, t_S, and defined as

$$t_S = \frac{R_S}{v}. \tag{2.7}$$

Here, R_S represents the radius of the light source, and it can be estimated assuming a simple thin-disk model (Shakura & Sunyaev, 1973), with a temperature profile of $T \propto R^{-3/4}$ thus (M&K),

$$R_S = \frac{3.4 \times 10^{15}}{\sqrt{\cos i}} \frac{D_{OS}}{r_H} \left(\frac{\lambda}{\mu\mathrm{m}}\right)^{3/2} \left(\frac{zpt}{3631\,\mathrm{Jy}}\right)^{1/2} 10^{-0.2\,(m-19)}\, h^{-1}\,\mathrm{cm}, \tag{2.8}$$

where i is the average inclination of the quasar disk, zpt is the zero point, which depends on the photometric system and broad band used, λ is the reference wavelength corresponding to the mean of the band in which the observations were made, and m is the source aparent magnitude after removing the gravitational lens magnification (i.e., demagnified magnitude). The values for zpt and λ are not always available in the literature, so we obtained them from the Filter Profile Service (Rodrigo & Solano, 2020)[6] using the specified instrument and band. For objects where this information is not explicitly provided, we assumed a Vega system zpt.

Following other authors, we adopt $i = \pi/3$, which corresponds to the average inclination of known systems (see Wildy & Czerny, 2017). Since not all systems have observations at the same wavelength (λ_{obs}), we rescaled all the obtained R_S using the thin-disk model approximation, $R_\lambda \propto \lambda^{4/3}$, to a reference rest-frame wavelength of $\lambda = 0.250\ \mu$m. Following Morgan et al. (2010), we use the equation,

[6] https://svo2.cab.inta-csic.es/svo/theory/fps3/

$$R_\lambda = R_S \cdot \left(\frac{(1+z_s)\cdot\lambda}{\lambda_{obs}} \right)^{(4/3)}, \tag{2.9}$$

where R_λ corresponds to the size of the radius at the rest-frame, λ, and R_S is the radius obtained from Eq. 2.8, with λ_{obs} as the wavelength of the observed band.

Next, we introduce the velocity model used to estimate the scales (see Eqs. 2.5 and 2.7), also known as the effective velocity of the source (v). This is calculated considering the relative motions of observer, lens, and source (Kayser et al., 1986), using Eq. 8 from Mediavilla et al. (2016):

$$v = \sqrt{V_1 + V_2 + V_3 + V_4}. \tag{2.10}$$

Hence, we have

$$V_1 = \left(\frac{v_{\rm CMB}}{1+z_L} \frac{D_{LS}}{D_{OL}} \right)^2, \tag{2.11a}$$

$$V_2 = \left(\frac{\sqrt{2}\,\sigma_{\rm pec}(z_L)}{1+z_L} \frac{D_{OS}}{D_{OL}} \right)^2, \tag{2.11b}$$

$$V_3 = \left(\frac{\sqrt{2}\,\sigma_{\rm pec}(z_S)}{1+z_S} \right)^2, \tag{2.11c}$$

$$V_4 = 2\left(\frac{\sqrt{2}\,\sigma_*}{1+z_S} \frac{D_{OS}}{D_{OL}} \right). \tag{2.11d}$$

In Eq. 2.11a, $v_{\rm CMB}$ represents the projection of the CMB dipole velocity. This velocity depends on both the velocity and coordinates of the CMB, as well as the coordinates of the object. It is calculated following Poindexter & Kochanek (2010), assuming a velocity of CMB of 369.82 ± 0.11 km/s at right ascension and declination coordinates, respectively, of RA $167^\circ.942 \pm 0^\circ.007$ and DEC $-6^\circ.944 \pm 0^\circ.007$ (Planck Collaboration et al., 2020).

The peculiar velocity ($\sigma_{\rm pec}$) in Eqs. 2.11b and 2.11c refers to the velocity at which a galaxy deviates from the Hubble flow at a given redshift. We can estimate this by following the steps of Mediavilla et al. (2016), assuming the linear regime in a ΛCDM cosmology. In this framework, $\sigma_{\rm pec}$ can be expressed in terms of the cosmological growth rate factor $f(z)$ as

$$\sigma_{\rm pec}(z) = \frac{\sigma_{\rm pec}(0)}{(1+z)^{1/2}} \frac{f(z)}{f(0)}, \tag{2.12}$$

adopting the peculiar velocity at redshift zero from Heß & Kitaura (2016) $\sigma_{\rm pec}(0) = 439 \pm 69$ kms^{-1}. Then, $f(z)$ can be calculated using

$$f(z) = \Omega_m^{\gamma}(z), \tag{2.13}$$

where $\Omega_m(z)$ is the matter density parameter and γ is the growth index, with $\gamma = 0.55$ for the lambda cold dark matter cosmology (ΛCDM; (Avila et al., 2022)).. Finally, σ_* in Eq. 2.11c represents the velocity dispersion of the stars in the lens galaxy. It is estimated using the SIE mass distribution, following previous works (Grillo et al., 2008; Treu et al., 2009) and expressed as

$$\sigma_* = \sqrt{\frac{c^2 D_{OS} \theta_E}{4\pi D_{LS}}}, \tag{2.14}$$

In those cases where the lens or source redshift is unknown, we follow Coe & Moustakas (2009); Redlich et al. (2012); Robertson et al. (2020) and similar works, where they are typically assumed to be $z_L \sim 0.5$ and $z_S \sim 2.0$, respectively. This assumption allows us to be consistent with previous works and the distribution of redshift in our sample (e.g., Treu et al., 2018; Schmidt et al., 2023)

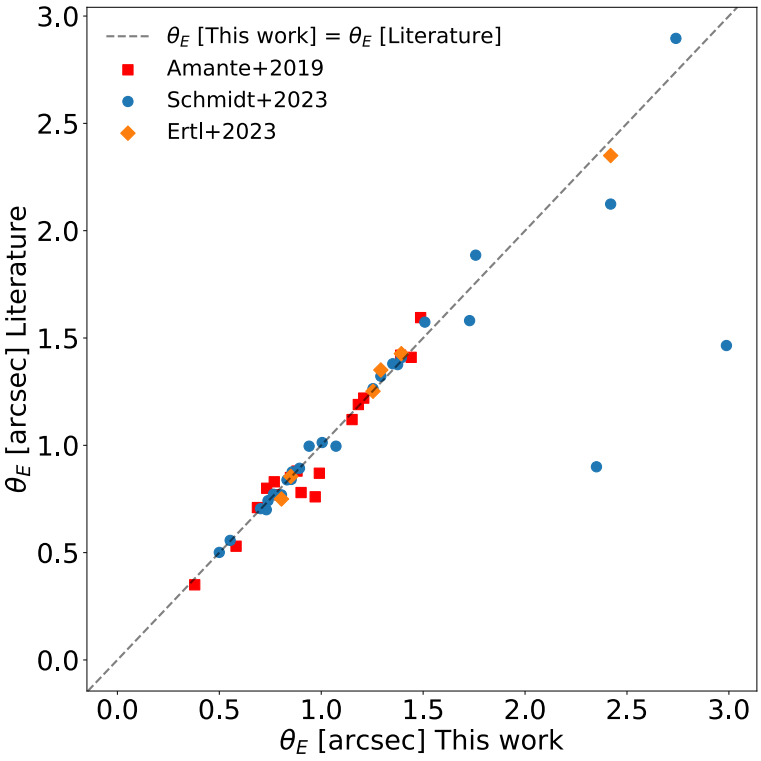


Figure 2.3: Comparison of θ_E in this work vs. values in the literature for modeled lens quasars. The systems that are clearly far from a 1:1 relation are SDSS 1155+6346 and HS 0818+1227.

2.2.3 Microlensing magnification maps

The Einstein radius defined in the previous section applies to a single, isolated microlens. However, the microlensing of lensed quasar images is associated with a relatively high optical depth of the microlenses, and cooperative effects between them are expected. To estimate the probability and timescales of microlensing in this scenario, we chose to use microlensing magnification maps. The probability density functions (PDFs) of microlensing magnification obtained from these maps provide information about the probability of brightness anomalies with respect to the intrinsic flux-ratios between images.

The analysis of magnification maps is a vital component of microlensing studies. In recent years, developments in this field have advanced significantly, particularly through the incorporation of GPU-based methods (e.g., Weisenbach, 2025) and large public databases such as GERLUMPH (Vernardos et al., 2014, 2015), which contain thousands of high-resolution magnification maps, as well as through updates to classical techniques, such as the FMM–IPM Jiménez-Vicente & Mediavilla (2022), an improved version of the classical inverse polygon mapping method (Mediavilla et al., 2006, 2011). In this study, we employ the FMM–IPM, which has been used extensively in microlensing analyses (e.g., Jiménez-Vicente & Mediavilla, 2025; Fian et al., 2024; Esteban-Gutiérrez et al., 2023), and whose magnification maps can be queried online [7].

The parameters used to generate the magnification maps include the convergence, κ, and shear, γ, obtained from the lens modeling (their distributions are shown in Fig. 2.4); the fraction of the total convergence in microlenses, defined as $\alpha = \kappa_* / \kappa_{\rm tot}$; and the microlens Einstein radius (see Eq. 2.6). The magnification maps have a physical size of 60 Einstein radii, with pixel dimensions ranging from 400 (Q2237+030) to 3800 (WISE2329-1258), corresponding to an average scale of approximately 1.1 light-days per pixel. This resolution is sufficient to avoid introducing additional biases in our results, as demonstrated by Vernardos & Fluke (2013).

Magnification maps are generated for two values of α. We select $\alpha = 0.1$ and $\alpha = 0.2$, taking into consideration a low stellar mass fraction that has been reported in some microlensing studies, with values of around a few percent to 10% (Mediavilla et al., 2009), consistent with the notion that lens galaxies are largely dominated by dark matter at the radius of the quasar images. Then, in more recent analyses that include

[7] see https://gloton.ugr.es/microlensing/

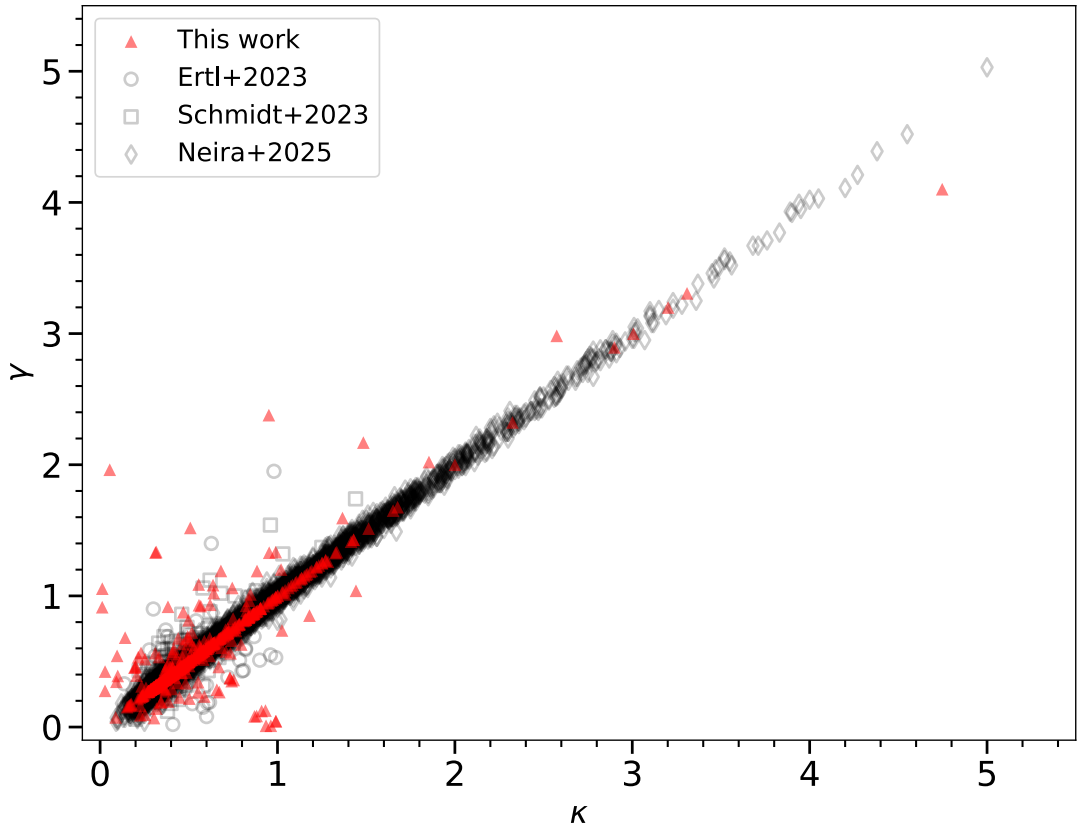


Figure 2.4: Values of κ and γ for our sample (red filled diamonds), compared with literature samples: unfilled black diamonds Neira et al. (2025), unfilled circles Ertl et al. (2023), and unfilled black squares Schmidt et al. (2023).

finite-source effects and larger samples, higher stellar mass fractions are indicated, converging on $\alpha \sim 0.2$ as a representative value near the Einstein radius (Jiménez-Vicente et al., 2015). This result is also in line with strong-lensing and dynamical studies of early-type lens galaxies, which suggest that stars contribute roughly 10–30% of the surface mass density at a few effective radii (e.g., Vernardos et al., 2024). Thus, by adopting $\alpha = 0.1$ and 0.2, we were able to bracket the most plausible range of stellar mass fractions inferred for lensing galaxies, while also ensuring comparability with previous microlensing studies that commonly use these values as benchmarks. One example of such a magnification map is shown in Fig. 2.2.

To account for the magnification of a finite-size source, maps were convolved with a Gaussian source profile of standard deviation r_s, $I(r) \propto \exp\left(-r^2/2r_s^2\right)$. As shown by Mortonson et al. (2005), the specific shape of the brightness profile does not strongly affect the statistical magnification properties, but just its half-light radius, which for a

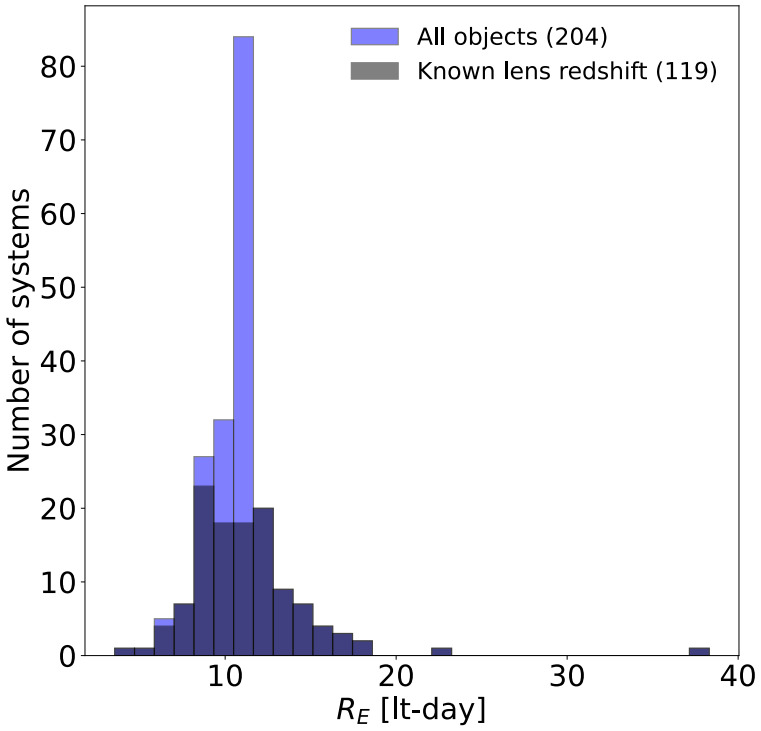


Figure 2.5: Histogram of the Einstein radii at the source plane for the systems. Blue shows the distribution of all the objects. Grey shows the distribution for the systems with the redshift of the lens known

Gaussian profile is $R_{1/2} = 1.18r_s$ and $R_{1/2} = 2.44R_S$ (i.e., $r_s = 2.07R_S$). The r_s values for each system are inferred from thin-disk estimates of R_S, rescaled to the median size derived from microlensing observations (see Sect. 2.3.2).

For each system image, magnification maps are generated using the corresponding κ and γ values. The physical size of the maps varies among systems to preserve a fixed pixel resolution in light-days. We consider two values of the microlensing mass fraction, $\alpha = 0.1$ and 0.2, and three source-size scalings, $0.3 \times r_s$, $1 \times r_s$, and $2 \times r_s$. For each set of (α, r_s), this yields 520 magnification maps for the full sample consisting of 149 doubles, 53 quads, and 2 five-image systems. This gives a total of 3120 convolved maps for all the pairs (α, r_s). These maps can be used to study the effects of α and r_s on the magnification probability distributions and combined with observational microlensing constraints, they allow us to evaluate the likelihood of different parameter combinations (Sect. 2.3.3).

2.3 Results

2.3.1 Homogeneous estimate of the Einstein radius

The homogeneous modeling of gravitational lensing systems provides an estimate of the parameters that define our model. The Einstein radius is the most relevant for our study, given its proportionality with the lens galaxy mass. To ascertain the robustness of our estimations, we compare them with those obtained in the literature and with other codes.

Comparison with previous studies not based on automatic modeling

To make a comparison with the literature, we use the sample of lensed quasars selected by Amante et al. 2020, which compiles the θ_E values individually obtained for each system (see Fig. 2.3). The agreement is, in general, excellent, except for two systems. One of them, SDSS1155+6346 (Pindor et al., 2004), is a double lens quasar, and the model selected in the literature was chosen from Rojas et al. (2014). In their work, they selected a SIS+γ model while our model is an SIE for this system. The discrepancy might be related to the difference in mass distribution selection. The other case is HS0818+1227 (Hagen & Reimers, 2000). This system is a double and the model selected, as cited in Leier et al. (2011), is affected by a close external galaxy, which impacts the estimation of the Einstein radius.

Comparison with previous automatic modeling studies

The newest discovered quadruple systems have been automatically modeled using HST high-resolution images. Schmidt et al. (2023) shows the results for thirty quadruply lensed quasars obtained with `LENSTRONOMY`. We modeled 29 of the 30. The missing model system corresponds to J1721+8842, the first confirmed ZigZag lens (Dux et al., 2025), consisting of two lens galaxies at different redshifts. For the remaining 29, we obtained a good agreement for 27, with a difference in the $R_E < 0.01$ arcsec (see Fig. 2.3). The other two systems are J0343-2828 (Lemon in prep, which has a fifth image) and 2M1310-1714 (Lucey et al. (2018), which also possesses a fifth image). We also compare the κ and γ obtained in our models with those obtained in Schmidt et al. (2023) we find $\Delta\kappa_{\rm med} \simeq 0.01$ and $\Delta\gamma_{\rm med} \simeq -0.01$. Given the corresponding scatters, $\sigma_{\kappa,\rm mad} = 0.09$ and $\sigma_{\gamma,\rm mad} = 0.09$, these offsets correspond to $0.16\,\sigma$ in κ and $0.06\,\sigma$ in γ, indicating a good agreement with previous results.

On the other hand, Ertl et al. (2023) modeled 9 of the 30 systems using GLEE. Figure 2.3 demonstrates that eight of the nine systems (we also remove J1721+8842) analyzed in that study can be reproduced with excellent consistency (within 2 σ for seven cases).

2.3.2 Microlensing timescales

Our sample can be divided into two main populations: one for systems with known redshifts (119) and the other with unknown redshifts for the lens (85). As previously noted, the microlensing timescale depends on the Einstein radius of the microlens, the source size, and the effective velocity, all measured at the source plane (see Sect. 2.2.2 for more detail). In Fig. 2.5, we present the histogram of the Einstein radii distribution for the two samples, which have medians of $10.900 \pm 0.138\,\mathrm{lt}-\mathrm{day}$ for the known redshift sample and $10.790 \pm 0.070\,\mathrm{lt}-\mathrm{day}$ for the entire sample assuming $z_L = 0.5$ for the systems with unknown redshift. The 1 σ uncertainty was inferred from the Monte Carlo (MC) sampling of the results.

In the case of the radius of the accretion disk (R_S), it is known that the size obtained by S&S is not in agreement with observational estimates (obtained from reverberation mapping and/or microlensing) and usually underestimates the size of the accretion disk (Morgan et al., 2010; Jiménez-Vicente et al., 2012a; Fausnaugh et al., 2018; Jha et al., 2022). To account for this discrepancy, we first calculated the median and standard deviation of the microlensing-based values obtained by Cornachione & Morgan (2020) and Jiménez-Vicente et al. (2012a). Then we rescaled all the values to a rest frame of $\lambda_{rest} = 2500\mathring{A}$. Using an MC sampling to estimate the 1 σ uncertainties, we obtained a representative experimental median value of $R_S^{exp} = 1.77 \pm 0.41\ \mathrm{lt}-\mathrm{day}$. Then we used the ratio between this quantity and the median of our theoretical estimates of R_S, based on Eq. 2.8, $R_S^{theor} = 0.23 \pm 0.003\ \mathrm{lt}-\mathrm{day}$, to obtain microlensing corrected R_S. The resulting R_S values are shown in Fig. 2.6, with a median of $2.23 \pm 0.06\,\mathrm{lt}-\mathrm{day}$ at 2500 Å after rescaling.

The last ingredient to calculate the timescales is the effective velocity of the source (see Sect. 2.2.2); we present our results in Fig. 2.7. The median value of the histogram is 807 ± 17 km/s for the systems with known redshifts and 778 ± 10 km/s for the whole sample.

Finally, to make a comparison with M&K results, we first need to rescale our median-corrected source sizes to the wavelength $8140\mathrm{Å}/(1+z_S)$, which is the original wavelength that M&K use to present their results. Our values are given in Fig. 2.8.

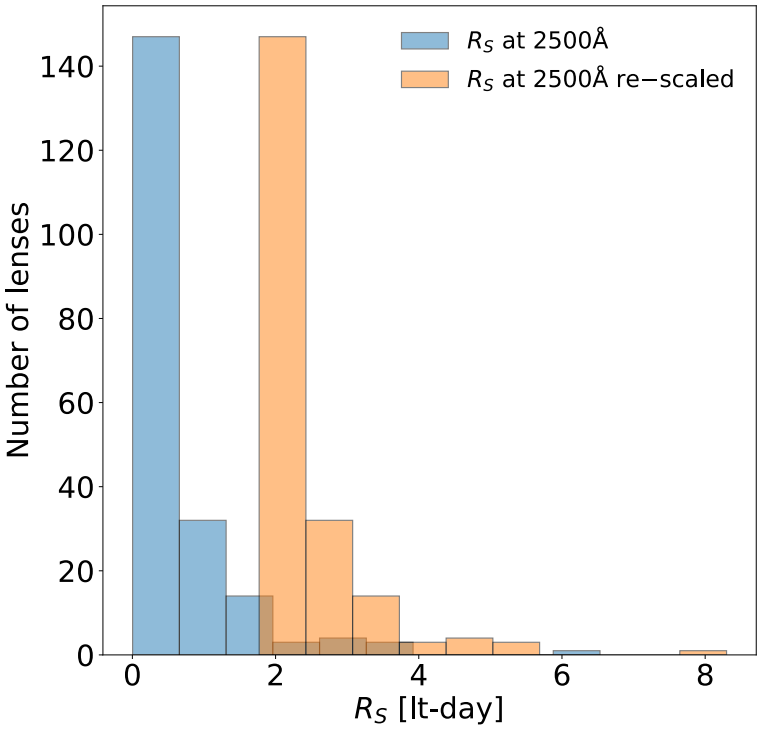


Figure 2.6: Histogram of the S&S based sizes of the source, R_S, before and after applying a microlensing size correction (204 systems).

From these histograms, we can derive, with 1 σ uncertainties, that the mean Einstein ring-crossing time is 11.29 ± 0.05 years for the entire sample and 10.82 ± 0.20 years for the systems with known lens redshifts; whereas M&K obtained a value of 20.6 years. The difference between the two estimates is mainly related to using the updated value for the peculiar velocity at zero redshift by Heß & Kitaura (2016). We notice that Mediavilla et al. (2024), who used a similar approach to estimate the peculiar velocities, obtained an average Einstein radius of 9.4 years, in good agreement with our mean value. On the other hand, for the source crossing timescale, t_S, we obtained a mean value of 2.59 ± 0.07 and 2.59 ± 0.09 years for known redshift lens systems and all the sample, respectively. This stands in contrast to the 0.61 years calculated by M&K and this discrepancy arises from both the rescaling of the theoretical sizes and the updated effective velocity.

2.3.3 Microlensing magnification statistics

To explore the possible impact of the source size, r_s ($r_s = 2.07 R_S$), and of the fraction of mass in microlenses, α, on the microlensing magnification statistics, we calculated (for

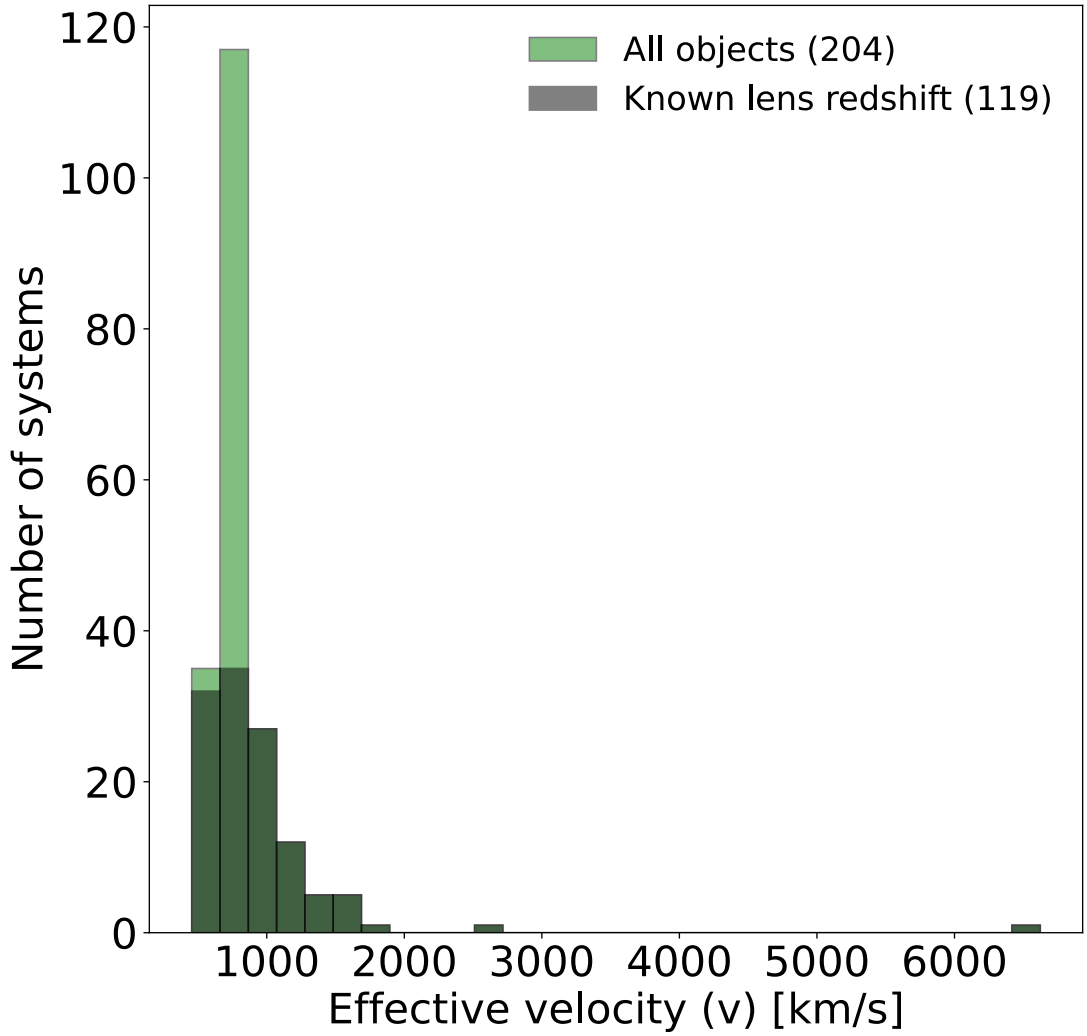


Figure 2.7: Histogram of the distribution of the effective velocities (Eq. 2.10).

each image of the systems in the sample) magnification maps for two different values of the fraction of mass in microlenses (see Sect. 2.2.3). These maps were convolved with Gaussian sources of three different sizes, obtaining 2×3 different magnification map models (α, r_s) and their corresponding magnification histograms. Then, for each of the six individual PDFs obtained for each image, we calculated the mean of the unsigned microlensing magnification values [8]. In Fig. 2.9, we present the histograms of these means (520 maps each (α, r_s) pair), giving the probability of obtaining an

[8] As discussed in Sect. 2, the empirical histogram of microlensing magnifications from Mediavilla et al. (2009, 2024), which we use to compare predictions with observations, is based on single-epoch measurements that do not determine the sign and are derived from image pairs, whereas our values are derived from individual quasar images.

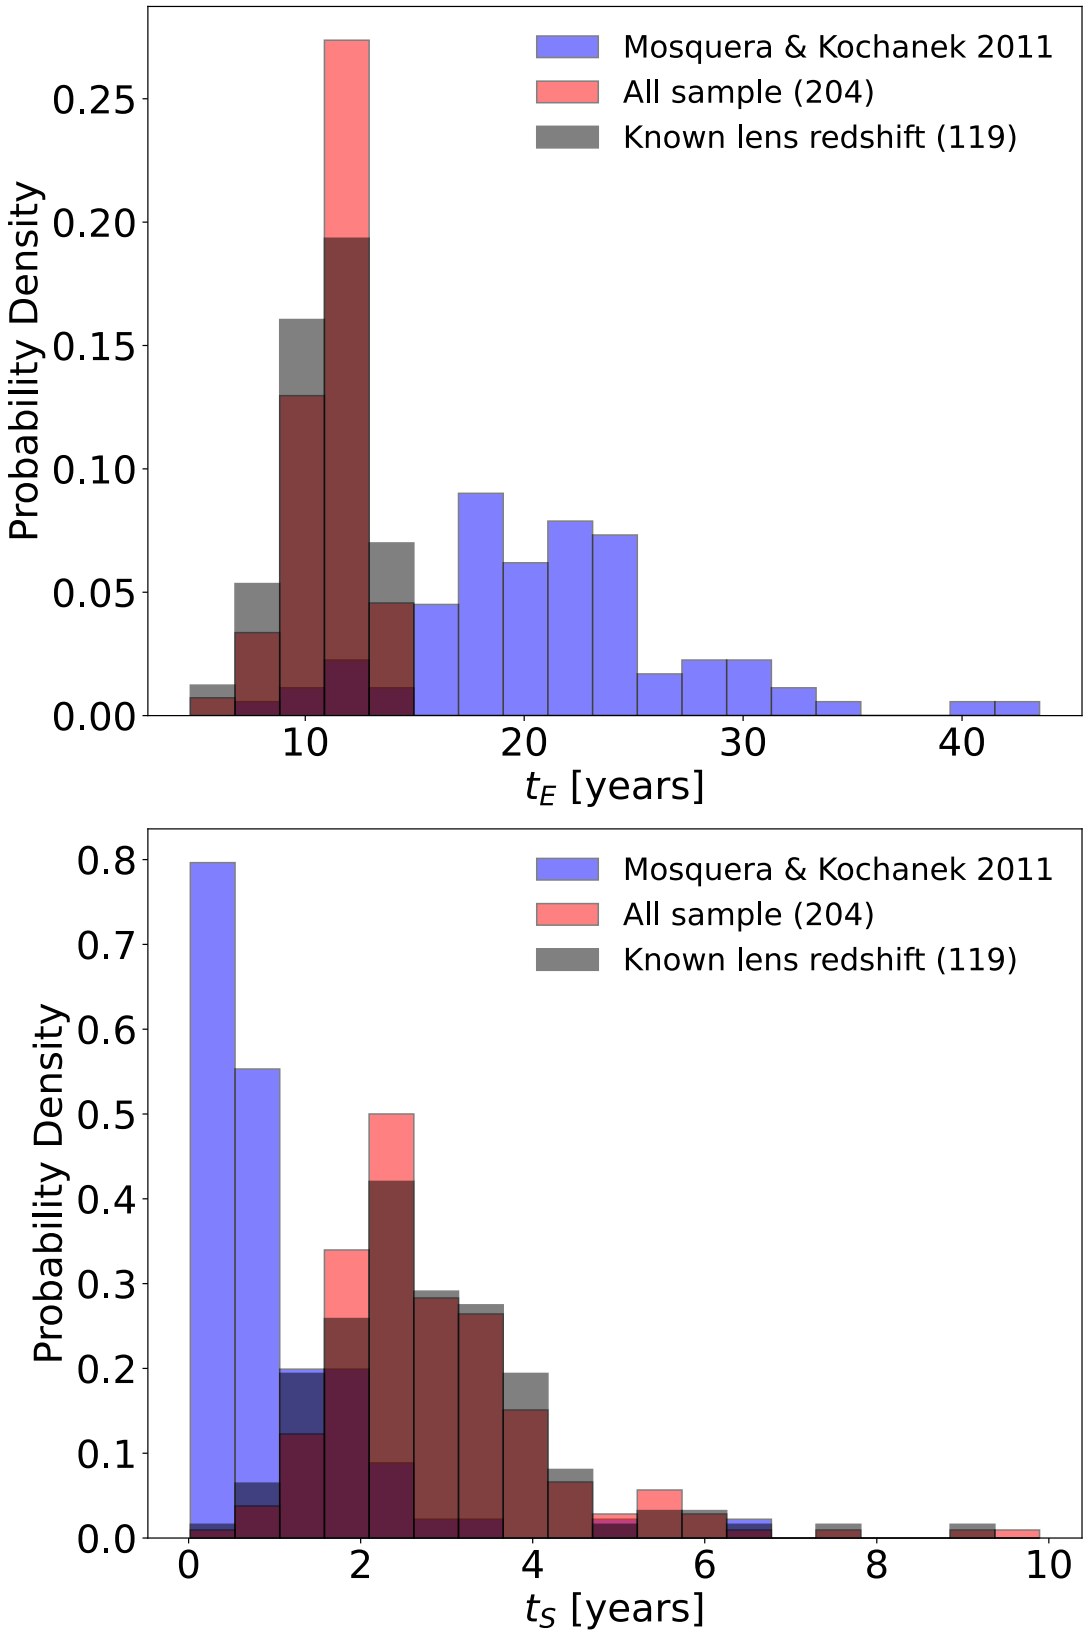


Figure 2.8: *Top*: Distribution of Einstein timescales. *Bottom*: Distribution of source timescales.

Table 2.1: Mean microlensing statistics for different stellar mass fractions and source-size scaling factors.

α	$\times \{r_s\}$	Mean $\lvert\Delta\text{mag}\rvert$	$\hat{\chi}^2$	Mean prob. $\Delta\text{mag} < -0.32$	Mean prob. $\Delta\text{mag} > 0.32$
0.1	0.3	0.227	9.362	12.193	20.043
0.2	0.3	0.320	8.497	15.162	29.586
0.1	1.0	0.153	9.203	6.309	17.815
0.2	1.0	0.217	6.430	8.646	26.310
0.1	2.0	0.096	16.316	1.923	15.984
0.2	2.0	0.158	8.650	3.188	23.834

The table lists the mean values of $\lvert\Delta\text{mag}\rvert$, the probability of $\Delta\text{mag} < -0.32$, and the probability of $\Delta\text{mag} > 0.32$ for different values of α and source-size scaling factors $\times \{r_s\}$. The samples were normalized so that bin counts sum to 44. The quantity $\hat{\chi}^2$ was calculated using Eq. (2.15).

(unsigned) value of microlensing magnification if we were to randomly select one image from the sample. The mean values of these histograms are listed in Table 2.1. As it can be inferred from Fig. 2.9, an important difference between the six histograms is the significance of the tail, which is basically ruled by the size of the source (the smaller the size the heavier the tail). This difference may be relevant because there is no evidence of a fat tail in the histograms of observed microlensing magnification (see histogram MED09 in Fig. 1 of Mediavilla et al., 2024). To explore this question in detail, we plot the six histograms in Fig. 2.10 with a bin size matching that of Mediavilla et al. (2024). This is done for all the images in our sets, for each (α,r_s) pair. According to this figure, larger values of the fraction of mass in microlenses and small sizes tend to predict larger magnifications, which were not observed in the MED09 sample. In contrast, small fractions of mass in microlenses and large sizes tend to overpredict low magnifications.

To quantitatively compare the different histograms with MED09, we used the $\hat{\chi}^2$ statistics that takes into account the contribution to the sample variance of both, the modeled and the observed histograms,

$$\hat{\chi}^2(\alpha, r_s) = \sum_i \frac{(O_i - H_i)^2}{O_i + H_i}. \tag{2.15}$$

Here, O_i and H_i correspond to each bin from the observed histogram, and the modeled one, respectively, from Mediavilla et al. (2024)

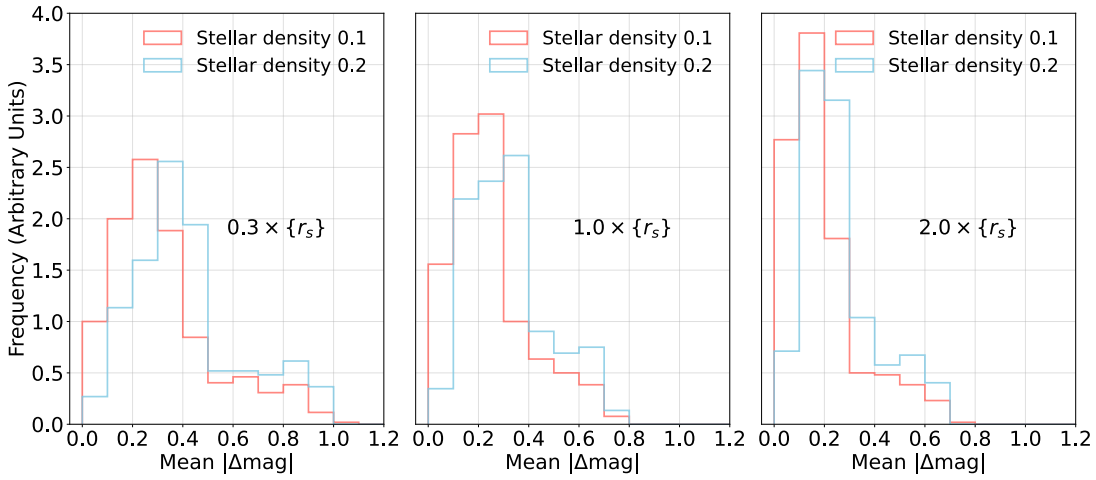


Figure 2.9: Probability distribution of unsigned microlensing magnification for three different scaling factor values for $0.3 \times \{r_s\}$, $1.0 \times \{r_s\}$, and $2.0 \times \{r_s\}$, with two different stellar densities (α).

The $\hat{\chi}^2(\alpha, r_s)$ values obtained are listed in Table 2.1 and the relative likelihoods of the models, calculated from $L(\alpha, r_s) \propto \exp\left(-\hat{\chi}^2(\alpha, r_s)/2\right)$, are represented in Fig. 2.11. There is a clear minimum in $\hat{\chi}^2(\alpha, r_s)$, with the maximum probability given in Fig. 2.11. This corresponds to the model $\alpha = 0.2$ and set of sizes $1 \times \{r_s\}$. In Fig. 2.11 we also plot the contour enclosing a 68% of the probability, which leaves outside the models corresponding to sizes $2 \times \{r_s\}$. We notice that there is a slight covariance between size and fraction of mass in microlenses in the sense that large sizes can be compensated by high mass fractions and vice versa. Marginalizing on the mass fraction, α, we can obtain an estimate of the expected values for the size, $(1.0 \pm 0.5) \times \{r_s\}$, which, taking into account our median value $R_s = 2.23$ light-days, results in $r_s = 2.07 R_S = 4.6 \pm 2.3$ light-days and in an half-light radius of $R_{1/2} = 2.44 R_S = 5.4 \pm 2.7$ light-days. Marginalizing on size, we can infer a lower limit for the fraction of mass in microlenses of $\alpha \geq 0.15$ (0.12) in the 68% (90%) confidence interval.

To go further with the statistical analysis of the histograms, we can focus on high magnifications, which are associated to interesting events such as caustic crossings. In Fig. 2.12 we show the cumulative distribution functions (CDFs) separating the magnifications ($\Delta m < 0$) and demagnifications ($\Delta m > 0$). The most distinctive feature of these cumulative distributions is the size dependence of the probability of high magnifications. The integrated probability of events with $\Delta m < -0.40$ for sizes $0.3 \times \{r_s\}$ is a non-negligible 20% (10%) for $\alpha = 0.2$ (0.1) while these high magnifications are very unlikely for larger sizes.

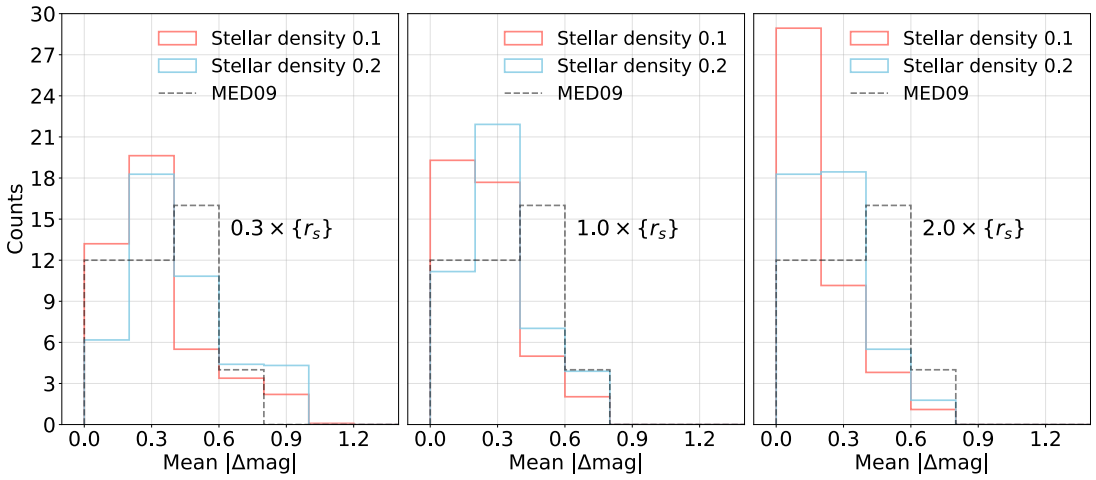


Figure 2.10: Histogram of unsigned microlensing magnification for three different scaling factors values for r_s, with two different stellar densities (α), comparing the results from Mediavilla et al. (2024) and this work, with all histograms normalized to the same total count.

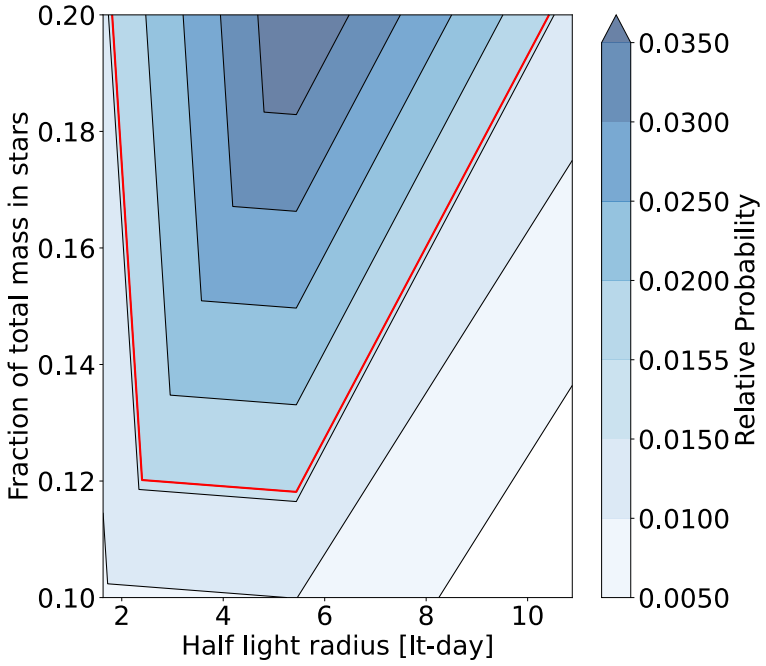


Figure 2.11: Relative probability distribution ($\propto e^{-\chi^2/2}$) in the (half light radius [light days], fraction of total mass in stars) parameter space. The red contour marks the 68% confidence region, enclosing the highest-probability values.

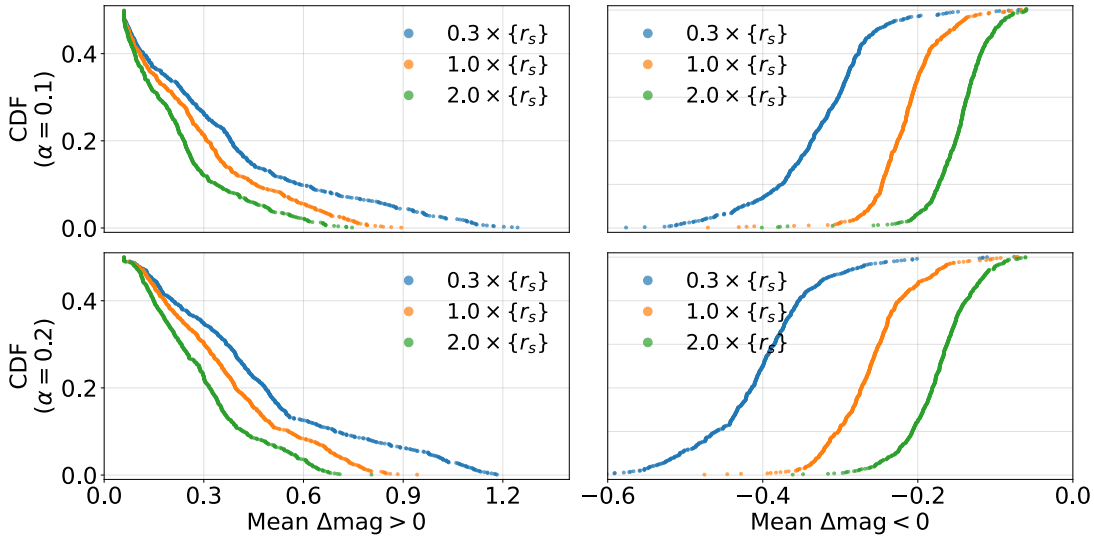


Figure 2.12: Cumulative distribution functions of magnitude changes during microlensing events. *Left column*: Probability of observing a demagnification event with $\Delta\text{mag} > 0$. *Right column*: Magnification event with $\Delta\text{mag} < 0$. *Top row*: Stellar fraction of $\alpha = 0.1$. *Bottom row*: $\alpha = 0.2$.

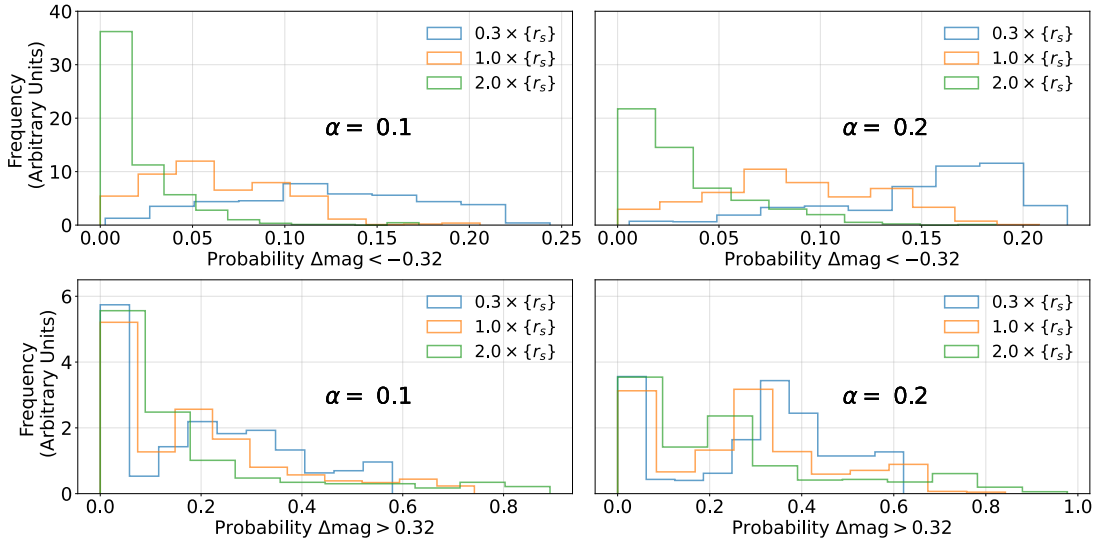


Figure 2.13: *Top*: Histogram of the probabilities of obtaining a magnification measurement of less than $\Delta\text{mag} = -0.32$ for different stellar densities and values of r_s. *Bottom*: Density histogram of the probabilities of obtaining a demagnification greater than $\Delta\text{mag} = 0.32$ for different stellar densities and values of r_s.

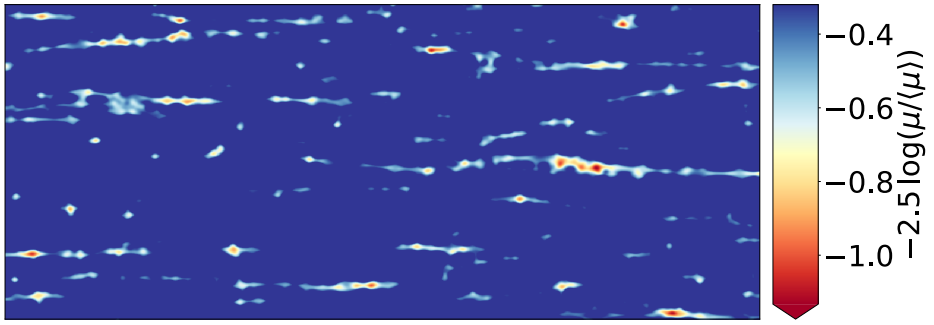


Figure 2.14: Same convolved microlensing magnification map as in Fig. 2.2 (*right*), With a zoom-in between pixel 0 and 400 in the Y axis and full length (1200) in the X axis. Here, the magnification map is masked for values over -0.32, highlighting the regions of strongest magnification within the caustic network.

To estimate the probability of an HME, in Fig. 2.13 we plot the probability of having a magnification[9] $\Delta m \leq -0.32$ and a demagnification $\Delta m \geq 0.32$, respectively. As can be seen in this figure, the probabilities for magnification are dependent on the source size and the fraction of mass in microlenses. On the contrary, there is no significant dependence in the case of demagnification. The mean probabilities for the case ($\alpha = 0.2, 1 \times \{r_s\}$) are 9% ($\Delta m \leq -0.32$) and 26% ($\Delta m \geq 0.32$). When microlensing is detected from a flux ratio, it is not possible, a priori, to separate magnification from demagnification. Consequently, the probability of observing microlensing magnification of an amplitude, $|\Delta m| \geq 0.32$, can be significantly larger (even 35%); however, in at most one-third of the cases, the large degree of magnification might ultimately be shown to be related to a singular magnification event such as a caustic crossing.

2.4 Discussion

2.4.1 Average size of the quasar source and properties of the microlenses population.

We found a good matching between the experimental microlensing data from MED09 and the histogram of microlensing magnifications corresponding to the case $\alpha = 0.2$ and average half-light radius $R_{1/2} = 5.4$ light-days modeled by us for the available sample of lenses. These values are in good agreement with previous results from

[9] We select the umbral magnification $\Delta m = -0.32$, as it is the value at the Einstein ring of an isolated particle.

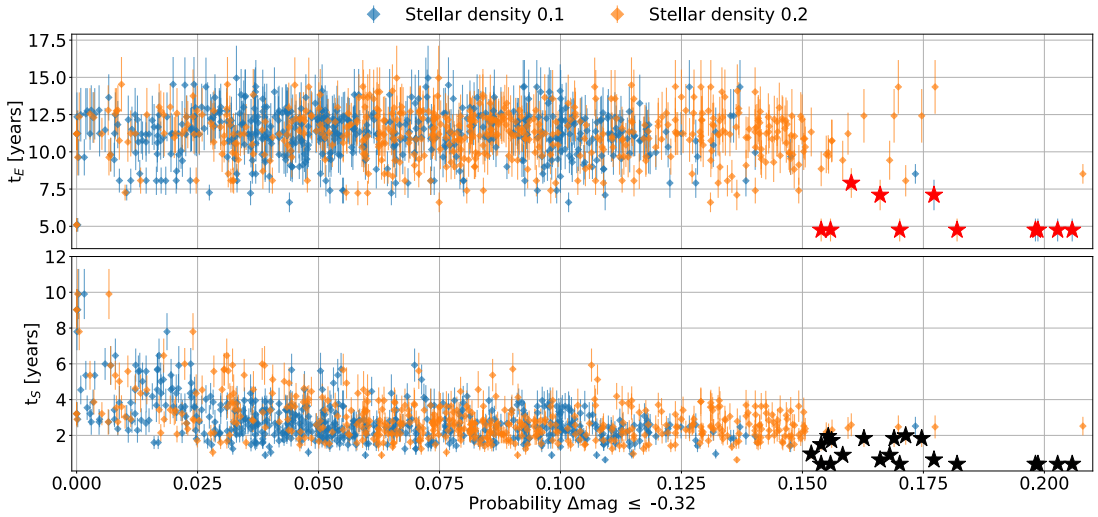


Figure 2.15: *Top*: Probability of $\Delta\text{mag} \leq -0.32$ versus Einstein radius crossing time (t_E [years]), with blue (orange) representing values obtained for a stellar density of 0.1 (0.2) and a 1.0 $\times \{r_s\}$, the red stars represent the object that shows the smaller Einstein radius crossing time and high magnification probability that correspond to DESJ2038-4008 and Q2237+030. *Bottom*: Probability or $\Delta\text{mag} \leq -0.32$ versus source crossing time (t_S [years]), with blue (orange) representing values obtained for a stellar density of 0.1 (0.2) and a $R_S \times 1$, in black stars the objects that show the smallest source crossing time and hight probability that are CY2201-3201, WG2100-4452, SDSSJ1138+0314, SDSSJ0821+4542, J0343-2828, WG0214-2105, and Q2237+030.

microlensing and reverberation mapping studies. Specifically, from the comparison between histograms, we obtained an estimate of the source size, $R_{1/2} = 5.4 \pm 2.7$ light-days, which is only slightly affected by the degeneracy with the fraction of mass in microlenses that has been pointed out in other studies. This value for the size is in good agreement with the average size estimate related to reverberation mapping observations (see the discussion in Mediavilla et al., 2017). This result is quite robust, as a significant reduction in size would result in a fat tail of the histogram above $\Delta m \geq -0.5$ not observed in single epoch microlensing measurements (MED09) or in other observations based on light curves (see the histogram PEAKS in Mediavilla et al., 2024). With respect to the mass fraction in microlenses, we estimated a lower limit, $\alpha \geq 0.15$, also in good agreement with previous studies (see, e.g., Esteban-Gutiérrez et al., 2022, 2023). To establish more stringent limits, larger values for α should be considered, which implies a computational challenge. We defer this study for future work, although values of α above 0.3 are unexpected.

Regarding the possible existence of a population of intermediate mass black holes (IMBHs) contributing to micro- or millilensing, we note that by virtue of the mass-length degeneracy (see Sect. 1.1), for relatively high microlensing masses ($\gtrsim 10 M_{\odot}$) the quasars will behave as point sources (see, e.g., Esteban-Gutiérrez et al., 2023; Heydenreich et al., 2024). Consequently, a fat tail of high magnifications should be present in the histogram of observations. The absence of this tail excludes the existence of a significant population of BHs above this mass.

Finally, the good matching found between the histogram of magnifications modeled from the full sample of lens systems available and the histogram of observed microlensing magnifications (inferred from just a subset of systems) suggests two key points: 1) the histogram MED09 is statistically representative of microlensing magnification properties of lensed quasar systems and 2) the procedure we followed to infer the magnification statistics (i.e., the lens modeling used to obtain κ and γ for each lensed image) is reliable.

A caveat can be made regarding the comparison with MED09. These data correspond to differential microlensing between pairs of images, which may introduce some broadening of the histogram relative to single-image microlensing. However, the good agreement between MED09 and the distribution of microlensing peaks from light curves (based on single images; see Mediavilla et al., 2024) indicates that this broadening is small compared to other factors, such as intrinsic variability or observational uncertainties.

2.4.2 Magnification statistics and flux-ratio anomalies.

A crucial step in many studies related to multiple imaged quasars is the modeling of the lens. This procedure needs to be very precise, particularly with respect to the prediction of gravitational time delays between the images, which may be used to solve the current tension in the determination of the Hubble constant. However, calculated models very often present strong differences between their predicted flux ratios and the observed ones. These anomalies are often attributed to microlensing. However, the reported model anomalies (see Fig. 1 in Mediavilla et al., 2024) are very high with mean values above 0.5 magnitudes and fat tails reaching even 2.5 magnitudes. According to our statistical study (see Fig. 2.9), such high levels of magnification have not been predicted even for unexpectedly small source sizes. Thus, we confirm for the present sample the results of previous simulations (Esteban-Gutiérrez et al., 2022; Awad et al., 2023). We provide microlensing magnification histograms for each individual image, which can be used to estimate the potential impact of microlensing on flux ratio anomalies as a criterion for rejecting or accepting a lens model.

2.4.3 Microlensing timescales and frequency of HME.

One main motivation of the statistical study of lens systems is to know the probability of one lens system undergoing high microlensing magnification. In M&K, the answer to this question was based on the comparison between the Einstein and source crossing times. These authors concluded that half of the lensed images will be active with the source lying in the caustic ridges.

Using our microlensing magnification statistics, we can now quantitatively set both the threshold to define an HME and the probability of exceeding it. In Fig. 2.13, we show the probabilities histogram of having microlensing magnification larger than -0.32, which is the threshold to be within the Einstein ring in the case of an isolated microlens. As shown in Fig. 2.13, the average probability is around 9% (the average probability of demagnification, $\Delta m \gtrsim 0.32$, associated with smooth variability, is about 26%). We can compare our high magnification probability with the work of Neira et al. (2025), where they estimated the number of HME events with magnification larger than 0.3 Δmag by year for simulated light curves, obtaining ~ 60 for 1250 lens quasar images, which is 4.9% per yr, this difference can be associated to the differences in sample and methodology.

It is also interesting to explore the possible correlation between the HME proba-

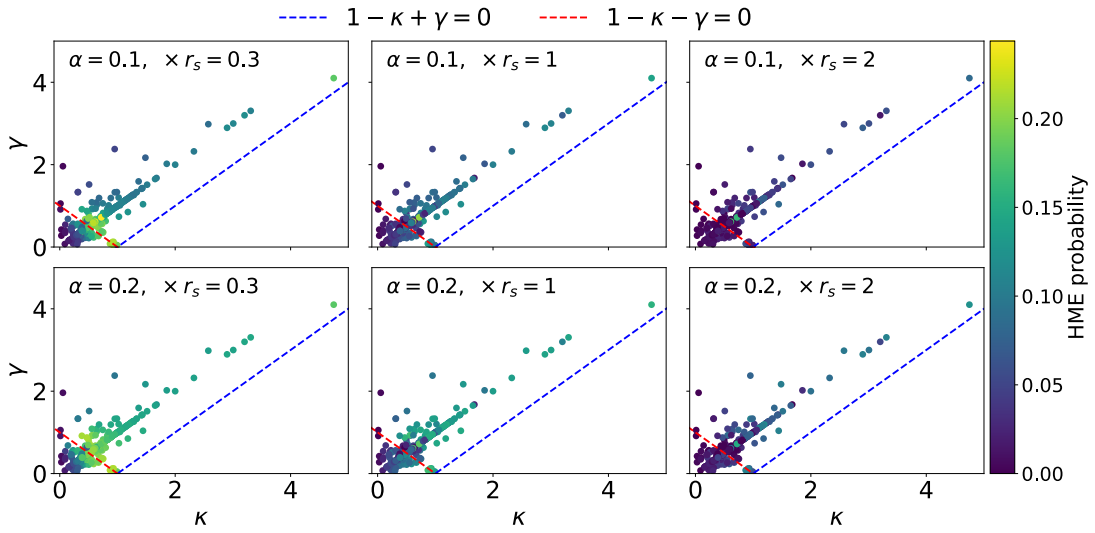


Figure 2.16: Correlation between the values of κ and γ in our sample, with the color bar representing the HME probability. The panels are separated by α and by multiples of r_s.

bility and the macro magnification of the image. With this aim, in Fig. 2.16 we plot the HME probability (i.e., the probability that $\Delta\text{mag} < -0.32$) versus κ and γ for each image in the sample. From this figure, we can observe (for the smallest size, $\times r_s = 0.3$, considered in the calculations) a clear correlation of the HME probability with the high magnification region defined by the $1-\kappa-\gamma=0$ curve. This correlation disappears for larger values of the source size. The explanation of the correlation observed for $\times r_s = 0.3$ is that a large macro-magnification of the image induce more structure in the microlensing magnification map, increasing the contrast between high magnification ridges and demagnified valleys (see, e.g., Weisenbach et al., 2021, 2024; Neira et al., 2025) and, hence, the probability of HMEs. However, this structure is washed out when convolved with a source of a relatively large size.

On the other hand the origin of the discrepancy with the ∼50% obtained by M&K when comparing t_E and t_S is that for optical depth close to 1, the real magnification map cannot be understood in terms of a single microdeflector. This is because cooperative effects among microlenses give rise to the highly magnified caustic networks at expenses of other regions, which become demagnified; namely, the scale of regions in the magnification map fulfilling $\Delta m = -0.32$ is smaller than one Einstein radius (see Sect. 1.1).

In Fig. 2.14, we show a magnification map where the low magnification pixels with $\Delta m > -0.32$ have been masked. We notice that the remaining regions appear

to be the caustics and inner to caustic regions, and that they occupy just a fraction of the whole map. As discussed above, the apparent "caustics on top of caustics" pattern reflects the densification and overlap of microcaustics for high-macro-magnified images (see Fig. 2.16), where the caustic network becomes connected and corrugated (see, e.g., Weisenbach et al., 2021, 2024; Neira et al., 2025). We confirmed these results by inspecting magnification maps of 10 images uniformly selected according to the distribution of lensed images in Fig. 2.15 (top).

In Fig. 2.15 we show, for each one of the images of lens systems in our sample, the probability of microlensing magnification above -0.32 mag vs. crossing times. If we are interested in caustic scanning of the accretion disk or the innermost region of the BLR (Fian et al., 2024), we can select the objects with the highest probability and smaller source crossing time as candidates; for instance, CY2201-3201, WG2100-4452, SDSSJ1138+0314, SDSSJ0821+4542, J0343-2828, WG0214-2105, and Q2237+030. If we are interested in estimating any pseudo-periodical property of microlensing (e.g., the caustic crossing frequency related to the peculiar velocity of the lens galaxy, see Mediavilla et al., 2016), we should instead select systems with smaller Einstein radius crossing times; for instance: DESJ2038-4008 and Q2237+030.

2.5 Conclusions

With the aim of conducting a statistical study of some relevant properties of quasar microlensing, we collected data from all known lensed quasar systems and performed a comprehensive and homogeneous modeling of gravitational lensing systems to estimate key parameters. In particular, we used the Einstein radius of the lensed systems, along with the convergence (κ) and shear (γ) of the lensed quasar images. From κ and γ, we obtained microlensing magnification maps for all the images using two values for stellar density (0.1 and 0.2). We combined the available data and calculations to estimate the microlensing magnification probabilities and crossing times. Our main results are as follows.

1. Our lens fitting procedure displayed a high level of agreement with previous studies. In particular, when compared with the sample of lensed quasars compiled by Amante et al. (2020), we found that 88% of the systems are consistent within our estimated uncertainties. Further comparisons with automatic modeling studies, such as those performed by Schmidt et al. (2023), who used `LENSTRONOMY` on 30 quadruply lensed quasars, showed a good concordance for 21 out of 29 common systems. Simi-

larly, when compared with models from GLEE, our results were in agreement to within $2\,\sigma$ for seven out of eight systems, reinforcing the reliability of our modeling approach.

2. We compared the observed microlensing magnification histogram with the magnification histogram of our sample of 204 lensed quasar systems modeled for different values of the fraction of mass in microlenses, α, and the size of the quasar source. We obtained a Bayesian estimate of the average half-light radius of $R_{1/2} = 5.4 \pm 2.7$ light days and a lower limit of $\alpha \geq 0.15$, in excellent agreement with previous studies based on microlensing and reverberation mapping.

3. According to the histogram of mean magnifications, it is unlikely that microlensing would lead to overall flux-ratio anomalies above $\sim$0.3 mag. Thus, lens modeling leading to larger flux-ratio anomalies should be regarded with caution; particularly in the context of Hubble constant estimate from lensed quasar monitoring). The individual microlensing magnification PDFs we provide here are available upon request for each image. These data can be used to cross-check lens models by comparing the resulting flux-ratio anomalies with the predicted microlensing.

4. We found a median timescale of $t_E = 11.29 \pm 0.05$ years for the Einstein radius crossing time and $t_S = 2.59 \pm 0.07$ years for the source size crossing time after accounting for uncertainties through MC simulations. These values differ from those reported by M&K, who found median values of 20 years and 0.61 years, respectively. The discrepancies can be mainly attributed to differences in the methodologies for calculating effective velocities and to our rescaling of the theoretical accretion disk sizes to account for recent measurements based on reverberation mapping and microlensing.

5. Beyond the single particle-based analysis of the Einstein radius crossing time, the study of microlensing magnification maps and histograms enables us to obtain more detailed and quantitative information about the probability of observing some interesting phenomenology related to microlensing, such as caustic crossing. Specifically, due to the formation of caustics networks, the size of the regions with magnification $\Delta m \leq -0.32$ (corresponding to the region inner to the Einstein radius of one isolated particle) is significantly smaller than the Einstein radius. We find that these regions correspond to the caustic and (partly) to the inner to caustic locations. We also find that these regions occupy, on average, 9% of the magnification maps (i.e., if you randomly pick an image from our sample, this is the rough probability of this image being engaged in a past or future caustic crossing event). In agreement with previous studies (see, e.g., Weisenbach et al., 2021, 2024; Neira et al., 2025), we find a correlation between the macro-magnification of the image and the probability of HMEs for

sources of relatively small sizes.

6. Combining the crossing times with the probabilities of a high magnification event, we were able to determine the best candidates for dissecting the structure of accretion disks through caustic scanning and studying the statistics of pseudo-periodical magnification events (e.g., caustic or zero crossings). This work provides a homogeneous statistical framework for linking lens-model parameters with microlensing observables in the current population of known lensed quasars. The resulting magnification maps, probability distributions, and timescales provide a reference for interpreting flux-ratio anomalies and for selecting the most promising systems for future microlensing monitoring. This approach will be particularly valuable in the context of the larger lensed quasar samples expected from ongoing and upcoming wide-field surveys.

2.6 Data availability

In the body of the chapter, we presented figures and results for the scales, effective velocity, source sizes, and magnification map statistics. We provide an online appendix at `https://github.com/felavila/microlensing-Maps-Post-Analysis-results`.

Chapter 3

Spectral Handling and Estimation of AGN Parameters: The first AGN fitting GPU-based code

Gravitationally lensed quasars are a powerful astronomical tool. Nevertheless, fully exploiting their potential requires large amounts of data and careful analysis. This is especially relevant for lensed-quasar spectroscopy, where the observed spectra can be affected by several sources of contamination, ranging from host-galaxy light from the source quasar to variations produced by stars in the lens galaxy acting as microlenses on the individual quasar images. Microlensing can differentially magnify emission arising from regions of different physical sizes, strongly affecting the compact continuum source while producing weaker or more complex effects on the broad- and narrow-line regions. As a result, it can modify continuum slopes, equivalent widths, emission-line profiles, and flux ratios between the lensed images. Accurate spectral decomposition and careful modelling are therefore required to distinguish intrinsic quasar properties from lensing-induced distortions. Before addressing the additional complexity introduced by lensing, however, it is first necessary to characterize the quasar spectra themselves accurately. Motivated by this need, we developed a tool designed to analyse these objects in a robust and efficient manner, with the ultimate

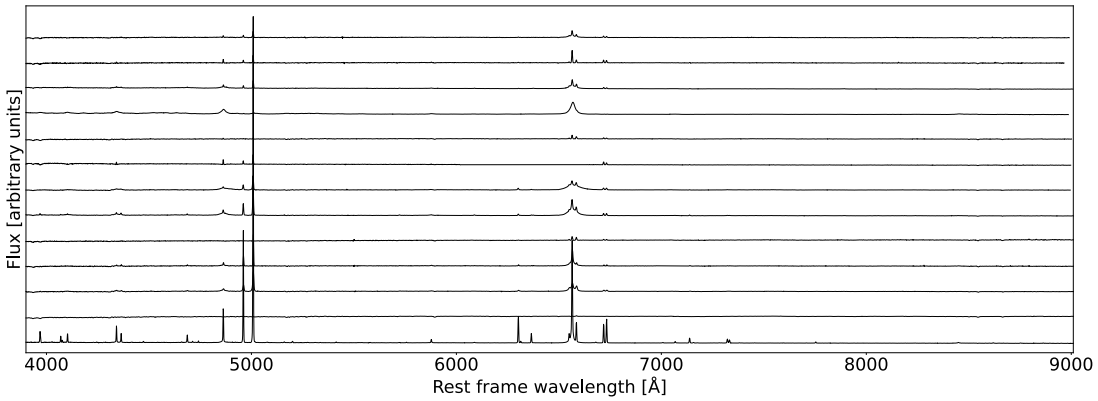


Figure 3.1: Multiple spectra observed with SDSS in rest frame

goal of applying it to gravitationally lensed quasars.

In this context, we present `Spectral Handling and Estimation of AGN Parameters` (`SHEAP`)[1] , a GPU-enabled spectral decomposition tool designed to analyze large samples of optical AGN spectra and deliver reliable parameter estimates with substantially reduced runtimes. To validate `SHEAP`, we apply it to $\sim$ 2000 AGN spectra, performing a consistent multi-component decomposition over a broad wavelength range (1100–7000 Å) and across different AGN populations and spectral resolutions.

This chapter is structured as follows. In Section 3.1, we provide a detailed overview of the optical and ultraviolet spectral characteristics of AGN. In Section 3.2, we review the main methodologies used for optical spectral modeling. Section 3.3 describes the methodology implemented in this work, while Section 3.4 presents the validation tests used to assess the reliability of the code. Finally, in Section 3.6, we discuss the results and present the main conclusions. Throughout this work, we adopt a flat ΛCDM cosmology with $\Omega_\Lambda = 0.7$, $\Omega_M = 0.3$, and $H_0 = 70\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$.

3.1 The optical ultraviolet spectra of AGN

In Figure 3.1, we present a set of rest-frame AGN spectra that illustrates the variety of spectral structures observed in these sources. This figure provides a small example of the diversity that can be found in optical/UV AGN spectra.

The optical/UV range covers rest-frame wavelengths from the near-ultraviolet to the optical, approximately $\lambda \sim$ 1000–9000 Å. This interval includes several of the

[1] `https://sheap.readthedocs.io/`

most important AGN spectral features, such as Lyα, C IV, Mg II, the Balmer series, Fe II blends, [O III], Hα, and other narrow and broad emission lines. These features trace different physical regions of the AGN, including the accretion disk, the BLR, and the NLR, and therefore provide some of the most widely used diagnostics of AGN physics (e.g. Boroson & Green, 1992; Shen et al., 2011; Padovani, 2017).

The optical/UV continuum emission in unobscured AGN is mainly associated with the accretion disk, while the broad emission lines are produced in the BLR. Because dust in the circumnuclear environment, especially in the torus, can absorb and scatter optical/UV radiation, this emission is directly visible only when the line of sight toward the nucleus is relatively unobscured. As a result, optical/UV-selected AGN samples are naturally biased toward unobscured sources, which are commonly classified as type I AGN. The physical interpretation of this classification, and its connection with orientation-based unification schemes, is discussed in Section 1.2.2.

Optical/UV selection also tends to favor radiatively efficient, non-jetted AGN. Only a fraction of jetted sources, such as radio-loud quasars and blazars, show a standard accretion disk and a prominent BLR in the optical/UV. Conversely, this selection can miss a large population of obscured AGN, usually classified as type II sources, although some of them can still be identified through their narrow emission lines. It can also miss moderately obscured AGN, low-luminosity AGN in which the host-galaxy light dominates over the nuclear emission, and AGN whose colors overlap with the stellar locus. The latter is particularly relevant at redshifts around $z \sim 2.5$ and $z \sim 3.5$, where quasar colors become similar to those of stars in broad-band optical surveys (Richards et al., 2002, 2006; Ross et al., 2012).

Despite these selection effects, the optical/UV band remains fundamental for AGN studies. First, optical/UV spectra provide direct measurements of the continuum and emission-line properties used to study the accretion disk, the BLR, and the NLR. These measurements are essential for estimating SMBH masses through reverberation mapping and single-epoch virial methods, as well as for characterizing the diversity of AGN spectra (Vestergaard & Peterson, 2006; Shen et al., 2011; Popović, 2020). Second, large optical spectroscopic surveys have produced extensive AGN catalogs, enabling statistical studies of AGN demographics, luminosity functions, black-hole growth, and spectral evolution across cosmic time (e.g. Schneider et al., 2010; Lyke et al., 2020; DESI Collaboration et al., 2024).

Large optical and near-infrared spectroscopic surveys have therefore been central to modern AGN research. The Sloan Digital Sky Survey (SDSS) has provided one

of the largest homogeneous samples of optically selected quasars, enabling statistical studies of quasar demographics, emission-line properties, black-hole masses, and accretion rates (e.g. Schneider et al., 2010; Shen et al., 2011; Lyke et al., 2020). More recently, the Dark Energy Spectroscopic Instrument (DESI) has extended this effort by obtaining spectra for millions of galaxies and quasars over a wide redshift range, providing an unprecedented dataset for studying AGN and galaxy evolution (DESI Collaboration et al., 2016b, 2024).

In the coming years, 4-metre Multi-Object Spectroscopic Telescope (4MOST; de Jong et al., 2019) will further expand the spectroscopic view of AGN populations through several complementary extragalactic surveys. In particular, the 4MOST X-ray AGN survey will follow up eROSITA-selected X-ray AGN, providing a large and highly complete spectroscopic sample of accreting SMBHs (Merloni et al., 2019). This effort is complemented by the 4MOST IR AGN survey, which targets mid-infrared-selected obscured AGN candidates, and by the Chilean AGN/Galaxy Extragalactic Survey (ChANGES), which will select AGN based on optical variability and spectral energy distribution properties (Bauer et al., 2023). Together, these 4MOST surveys will provide large and diverse AGN samples, spanning unobscured, obscured, X-ray-selected, infrared-selected, and variability-selected populations.

3.2 Modeling the optical spectra

The modeling of AGN and the development of analysis tools to infer their physical properties have evolved substantially, from the earliest spectroscopic classifications (Khachikian & Weedman, 1974) to modern large-scale characterization pipelines (Wu & Shen, 2022; Moustakas et al., 2023). Nevertheless, many analysis pipelines have historically been developed for individual projects and have remained private (Calderone et al., 2017). This began to change over the past $\sim$15 years, as open-source practices became increasingly common in astronomy (Weiner et al., 2009), leading to the release of multiple community tools that adopt different methodological approaches to modeling AGN spectra.

A milestone was the release of `QSFit` (Calderone et al., 2017) and its Python implementation, `PyQSOFit` (Guo et al., 2018), which have been used to characterize large samples, including 750414 AGN spectra (Wu & Shen, 2022) observed with the SDSS (Lyke et al., 2020). Since then, additional tools have appeared to address specific challenges. For instance, because Fe II emission can strongly affect continuum place-

ment and measurements of blended lines, dedicated methodologies and automated pipelines such as FANTASY (Ilić et al., 2020; Rakić, 2022; Ilić et al., 2023) have been developed to improve Fe II modeling. More broadly, the need for robust host–AGN separation and more realistic uncertainty estimates has motivated approaches that incorporate full stellar-continuum and kinematic modeling, such as Bayesian AGN Decomposition Analysis (BADASS; Sexton et al., 2019, 2021). BADASS integrates Penalized PiXel-Fitting (pPXF) (Cappellari & Emsellem, 2004; Cappellari, 2012, 2017, 2023), a widely used tool for galaxy spectra, to model the host-galaxy component in AGN observations.

The release of DESI is already delivering spectra at unprecedented volume, with DR1 providing spectra for $\gtrsim$ 18 million unique targets, forcing analysis frameworks to prioritize automation, robustness, and well-characterized systematics (e.g., pipeline and value-added products, Siudek et al., 2024). This has driven survey-optimized approaches such as FastSpecFit (Moustakas et al., 2023), as well as AGN-focused decomposition efforts such as DAS (Du, 2024), which integrate instrument- and survey-level information directly into the fitting process. More broadly, the community expectation is that next-generation multi-object facilities will push these requirements further (e.g., the discussion of survey-scale spectroscopy and end-to-end data products in Mainieri et al., 2024).

The need for scalability, which means the ability to process large and heterogeneous spectroscopic datasets in an automated and computationally efficient way, is becoming even more acute in the era of survey spectroscopy. In the coming years, the start of operations of 4MOST will further extend this landscape by delivering an exceptionally large AGN census: the 4MOST AGN survey is designed to provide physical properties for up to $\sim 10^6$ AGN, enabling population-scale tests of SMBH growth and AGN–galaxy co-evolution (de Jong et al., 2019; Merloni et al., 2019; Bauer et al., 2023).

3.3 Methods

In this section, we present the main steps and strategies of our approach to modeling AGN spectra with SHEAP, a Python 3 package built on JAX (Bradbury et al., 2021). JAX is a numerical computing library that provides automatic differentiation, just-in-time compilation, vectorization, and parallel execution on central processing units (CPUs) and graphics processing units (GPUs), thereby reducing computational time and im-

proving efficiency.

3.3.1 Data pre-processing

Before delving into the modeling procedure of our method, it is important to explore the data structure and the corrections required to obtain physically representative results.

Data handling & structure

Our method operates on data stored as a three-dimensional array $A_{i,j,k}$, where:

- $i \in \{0, \ldots, N-1\}$ indexes the spectra in the batch,
- $j \in \{0, 1, 2\}$ indexes the data component, corresponding to wavelength, flux density, and per-pixel observational uncertainty (σ),
- $k \in \{0, \ldots, L-1\}$ indexes pixels within a spectrum.

Given the complexity of spectral modeling (e.g. Cappellari, 2017; Guo et al., 2018; Ilić et al., 2023), an array-based, batched data layout enables efficient batch processing and GPU-based vectorization during model evaluation and optimization. By expressing the computation in terms of array-shaped operations, the data-parallel structure of the problem is made explicit to the underlying compiler and accelerator backends. This facilitates optimized memory access patterns (for more details, see Bradbury et al., 2021) and significantly accelerates both forward-model evaluation and automatic differentiation, particularly in large parameter spaces and in workflows that require repeated evaluations, such as gradient-based optimization and posterior sampling (as in our case; see Section 3.3.3 Dalalyan, 2017; Ma et al., 2019).

Spectral corrections

Our method assumes that the input data are provided in the rest frame, although it can also apply standard corrections to observed-frame spectra. We first transform the spectra to the rest frame[2] (Hogg et al., 2002; Blanton & Roweis, 2007). Since the observed spectra are also affected by Galactic dust extinction (Fitzpatrick, 1999), we estimate the line-of-sight color excess, E(B-V), from the source coordinates using the

[2] $(\lambda_{\rm rest} = \lambda_{\rm obs}/(1+z))$ and $(F_\lambda^{\rm rest}(\lambda_{\rm rest}) = (1+z)F_\lambda^{\rm obs}(\lambda_{\rm obs}))$

`sfdmap`[3] Python package, which is based on the full-sky dust maps of Schlegel et al. (1998). We then correct the spectra using the extinction law of Cardelli et al. (1989), adopting ($R_V \equiv A_V/E(B-V) = 3.1$), where ($A_V$) is the extinction in the (V) band. The Cardelli law provides a convenient and widely adopted parameterization of the average Milky Way extinction curve over a broad wavelength range, while ($R_V = 3.1$) is characteristic of the diffuse Galactic interstellar medium. Comparable preprocessing procedures, including Galactic-extinction correction and transformation to the rest frame, are also implemented in other AGN spectral-analysis tools (e.g., Ilić et al., 2023; Guo et al., 2018). Finally, the E(B-V) values derived from the Schlegel et al. maps are recalibrated following Schlafly & Finkbeiner (2011).

3.3.2 Modeling with `SHEAP`

The `SHEAP` modeling pipeline constructs a spectral model (SM) that depends on a defined wavelength range, this interval determines the set of physical components that will appear in an AGN spectrum model as host-galaxy starlight, Fe II emission, non-thermal continuum, and emission lines (EL). Then, for a general spectral model:

$$F_{\rm SM}(\lambda_{\rm rest}) = F_{\rm host}(\lambda_{\rm rest}) + F_{\rm Fe}(\lambda_{\rm rest}) + F_{\rm cont}(\lambda_{\rm rest}) + F_{\rm Balmer}(\lambda_{\rm rest}) + F_{\rm EL}(\lambda_{\rm rest}), \tag{3.1}$$

where $F_{\rm SM}(\lambda_{\rm rest})$ is the total modeled flux density of the spectrum in the rest frame. The individual terms represent the different spectral components contributing to the observed emission: $F_{\rm host}(\lambda_{\rm rest})$ is the stellar continuum from the host galaxy, $F_{\rm Fe}(\lambda_{\rm rest})$ accounts for the blended Fe II emission complexes, $F_{\rm cont}(\lambda_{\rm rest})$ represents the featureless AGN continuum emission, $F_{\rm Balmer}(\lambda_{\rm rest})$ describes the Balmer continuum emission, and $F_{\rm EL}(\lambda_{\rm rest})$ represents the sum of all emission-line component (e.g narrow, broad, outflow, winds).

The host galaxy model

It is well known that, at $z \lesssim 1$, the stellar continuum and absorption features of the host galaxy can contribute significantly to the observed AGN spectrum (e.g., Greene & Ho, 2005; Jalan et al., 2023; Ren et al., 2024). To account for this effect, we implemented a lightweight host-galaxy module inspired by the penalized Pixel-Fitting

[3] https://github.com/kbarbary/sfdmap

(pPXF) methodology (Cappellari & Emsellem, 2004; Cappellari, 2017), designed specifically for AGN+host spectra.

In our approach, the host-galaxy stellar emission is modeled as a linear combination of single stellar population (SSP) templates spanning a grid of stellar ages and metallicities (Vazdekis et al., 2010). The composite template is convolved with a Gaussian line-of-sight velocity distribution (LOSVD) to reproduce the broadening induced by the stellar motions within the galaxy (van der Marel & Franx, 1993). The convolution is performed efficiently in Fourier space (Brigham, 1974), ensuring consistent spectral broadening while keeping the kinematic description minimal.

In contrast to full spectral-fitting frameworks such as pPXF, which can simultaneously recover detailed stellar kinematics and population properties using penalized likelihoods and Gauss–Hermite expansions (e.g., Cappellari & Emsellem, 2004; Cid Fernandes et al., 2005), our current implementation restricts the LOSVD to a purely Gaussian kernel. In this approximation, the LOSVD is parameterized only by the mean line-of-sight velocity V and the velocity dispersion σ, without including higher-order Gauss–Hermite moments. Similarly, the method does not attempt to infer non-parametric star-formation histories. This design choice prioritizes numerical stability and computational speed in regimes where the host contribution is typically subdominant and strongly blended with the AGN continuum and emission-line components. We constructed a custom cube of simple stellar population (SSP) templates from the E-MILES library (Vazdekis et al., 2016) using the "Tune SSP Models" web tool[4]. The original library spans 1680–49,999.4 Å; from which we selected the 1680–8950 Å interval, representative of the typical wavelength coverage of UV/optical instruments. For consistency, we adopted the parameter choices used by Cappellari (2023) for the "Tune SSP Models". Accordingly, the templates were generated with linear sampling, $\Delta\lambda = 0.9\,\text{Å}\,\text{pix}^{-1}$, and have a native spectral resolution of FWHM $\simeq$ 2.51–5 Å across the selected wavelength range. We adopted the Padova+00 isochrones (Girardi et al., 2000), a unimodal initial mass function with a slope of 2.30, the base Fe abundance pattern, and the full E-MILES grids in metallicity, [M/H]=-2.32 to 0.22, and age, 0.1–10 Gyr.

[4] https://research.iac.es/proyecto/miles/pages/webtools/tune-ssp-models.php

The AGN Continuum model

Our method provides flexibility in the continuum component, allowing the background model to be chosen according to the physical nature of the source and the wavelength range under analysis. Among the available options are a linear baseline for small spectral windows with nearly flat local continua (Vanden Berk et al., 2001), a power-law continuum for the featureless accretion-disk emission commonly observed in AGN (e.g., Richards et al., 2006), and a broken power law for wider wavelength ranges where a single slope does not adequately describe the continuum shape (Dong et al., 2008).

The Balmer emission model

We model the Balmer emission as a set components (Sulentic et al., 2000) that shape the pseudo-continuum in the near-ultraviolet (UV) region that will be have velocity shift independent. The Balmer continuum, produced by bound–free recombination of hydrogen in partially optically thick gas, becomes significant shortward of the Balmer edge at 3646 Å (Sulentic et al., 2000; Dong et al., 2011).

Meanwhile, at longer wavelengths (3700–4000 Å) emission from high-order Balmer lines (typically $n \gtrsim 7$) blends together due to Doppler broadening and instrumental resolution, forming a quasi-continuum that further elevates the background level. This blended Balmer emission is particularly prominent in luminous and high-redshift AGNs, where it can substantially affect the apparent continuum shape and the measurement of nearby spectral line features (Dietrich et al., 2002). Formally, we split the total Balmer contribution as:

$$F_{\rm Balmer}(\lambda_{\rm rest}) \;=\; F_{\rm PB}\,(\lambda_{\rm rest}) \;+\; F_{\rm HighBalmer}(\lambda_{\rm rest}), \tag{3.2}$$

where $F_{\rm PB}(\lambda_{\rm rest})$ represents a parametric model of the Balmer continuum, defined following Grandi (1982); Dietrich et al. (2002), and $F_{\rm HighBalmer}(\lambda_{\rm rest})$ accounts for the high-order Balmer emission. The latter is modeled using an empirical template ($T_{\rm HBalmer}(\lambda)$), constructed from the blend of high-(n) Balmer emission lines that merge blueward of (4000 Å) (Storey & Hummer, 1995). In this work, we adopt the updated version of this template presented by Bernal et al. (2025) (hereafter, Bernal+25), which extends the modeled series to (n>7). The template is first convolved with a Gaussian kernel to reproduce the observed line broadening and can subsequently be shifted in velocity and scaled by a multiplicative factor to match the observed spectrum.

The Fe II treatment

The Fe II emission is a prominent source of contamination and degeneracy in the study of AGN spectra, arising from a multitude of blended transitions of singly ionized iron. It spans a broad wavelength range, from the ultraviolet (e.g., around Mg II; Vestergaard & Wilkes, 2001; Popović et al., 2019), through the optical (e.g., near Hβ and Hα; Véron-Cetty et al., 2004; Park et al., 2022), to the near-infrared (e.g., around Paγ; Rudy et al., 2000; Landt et al., 2008; Garcia-Rissmann et al., 2012). This pseudo-continuum frequently overlaps with key diagnostic lines and can significantly bias measurements of broad-line widths and fluxes if not properly accounted for. Accurate modeling and subtraction of both Fe II and Fe III emission are therefore essential for recovering the intrinsic properties of the BLR and for robustly measuring the Balmer and Mg II line parameters.

To model this component, we adopt the classical template-fitting approach, in which an empirical Fe template is scaled, broadened through convolution with a Gaussian kernel, and allowed to shift in velocity to match the data (this component being independent of the others), following the same procedure described in Section 3.3.2. In particular, we use the Fe II templates from Boroson & Green (1992) and Vestergaard & Wilkes (2001), following implementations such as Calderone et al. (2017), which provide continuous coverage from 1074 to 7000 Å by combining UV and optical Fe emission complexes. In this scheme, no explicit separation into "broad" and "narrow" Fe components is introduced; instead, the template is treated as a single blended complex whose effective kinematics are captured by the broadening and velocity shift. Nevertheless, we note that the use of templates may be insufficiently flexible in particularly complex regions, such as the vicinity of Hα. For this reason, we also implemented the empirical Fe modeling strategy available in `FANTASY`, which includes a full model of Fe II emission based solely on atomic data. This model includes a total of 283 transitions, divided into 17 atomic groups, over the wavelength range 3700–11,000 Å (see Ilić et al., 2023, for details). In this approach, pseudo-templates are generated for the different Fe atomic groups and tied together in velocity shift and width.

Emission-lines treatment

Emission-line fitting is a central component of AGN spectral modeling, since line profiles trace the kinematic signatures of different emitting regions, including broad

and narrow components, as well as winds and outflows. Our method is designed to efficiently fit complex models comprising multiple profile families (e.g., Gaussian, Lorentzian, and Voigt), multiple kinematic components within a given line, and multiple spectral regions simultaneously. It also accounts for the kinematics of Fe II emission, the stellar component, and Balmer emission. As part of the modeling strategy, we adopt a predefined list of emission lines that distinguishes between transitions known to exhibit broad emission, transitions typically associated only with narrow emission, and transitions that can display both narrow and broad components (Ilić et al., 2023). This scheme helps to define a consistent and physically motivated line configuration for each fitted region. The parameter bounds for these regions are set according to the values presented in Table A.1, ensuring that the fitted parameters remain physically meaningful. In addition, we incorporate a set of known atomic constraints into the main fitting structure, as in the cases of [O III]$\lambda\lambda$4959, 5007 and [N II] $\lambda\lambda$6548, 6583(Vestergaard & Peterson, 2006; Mejía-Restrepo et al., 2016)

3.3.3 Optimization Strategy

Over the past decades, optimization strategies for non-linear spectral fitting have evolved substantially. Early and widely adopted approaches rely on classical non-linear least-squares solvers such as the Levenberg–Marquardt (LM) algorithm (Moré, 1978), which remains a standard choice in many fitting packages (e.g., `LMFIT`; Newville et al. 2014) and `MPFIT` (Markwardt, 2009). In parallel, posterior exploration via Monte Carlo Markov Chain (MCMC) samplers, such as `emcee` (Foreman-Mackey et al., 2013a,b), has become common when robust uncertainty quantification is required. More recently, constrained non-linear least-squares minimizers have also been introduced in widely used spectral-fitting frameworks (e.g., `capfit` within `pPXF`; Cappellari 2023 and references therein), reflecting the need to handle increasingly flexible and high-dimensional models.

In this work, we adopt a first-order, gradient-based optimization strategy within the `JAX` framework. With the model defined for the different spectral regions (Eq. 3.1), we impose physically motivated bounds that depend on the kinematic component, summarized in Table A.1. The optimization in `SHEAP` is performed using the Adam optimizer (Kingma & Ba, 2014), implemented through the `Optax` library (DeepMind et al., 2020), that depends on two hyper parameters: the learning rate and the number of iterations. This choice enables efficient and adaptive gradient updates, while lever-

aging automatic differentiation and accelerated compilation for fast model evaluation and differentiation.

Our minimization strategy accounts for fixed and bounded parameters (as well as tied relations where applicable). Parameter bounds are enforced via smooth re-parameterizations to an unconstrained optimization space, avoiding hard projections that can hinder gradient-based methods. This approach is conceptually similar to the bound-handling strategy adopted in `LMFIT` (Newville et al., 2014), where bounds are implemented through internal parameter transformations. This design keeps the optimizer agnostic to the underlying spectral model and ensures that constraints are applied consistently across components and regions.

Residuals are evaluated in a weighted manner using the spectral uncertainties, so that high-S/N pixels carry greater influence in the fit. The normalized residual vector is defined as

$$r_i = \frac{f_{\mathrm{model},i} - f_{\mathrm{obs},i}}{\sigma_i}, \tag{3.3}$$

where $f_{\mathrm{model},i}$ corresponds to the model evaluated at pixel i, $f_{\mathrm{obs},i}$ is the observed flux at pixel i, and σ_i is the uncertainty by pixel. We adopt a robust residual loss based on the log-cosh function,

$$\mathcal{L}_{\mathrm{residual}} = \langle \log\cosh\left(r_i\right) \rangle \, . \tag{3.4}$$

The log-cosh function provides a smooth and outlier-tolerant alternative to a purely quadratic (L_2) loss: it behaves approximately quadratically for small residuals, while transitioning to an approximately linear (L_1-like) penalty at large $|r_i|$. As a result, isolated deviant pixels (e.g., residual sky features, imperfect masking, or artifacts) have reduced leverage on the solution, while the loss remains differentiable and therefore well-suited for gradient-based optimization (see Saleh & Ehsanes Saleh, 2022, for a detailed discussion).

It is important to highlight that this procedure is applied consistently to the full batch of fitted spectra, using the same model configuration for all objects. Because `SHEAP` is implemented to run natively in parallel (i.e., the same computation is evaluated over a batch of spectra in a single compiled call), the optimization is carried out simultaneously across the batch rather than spectrum-by-spectrum.

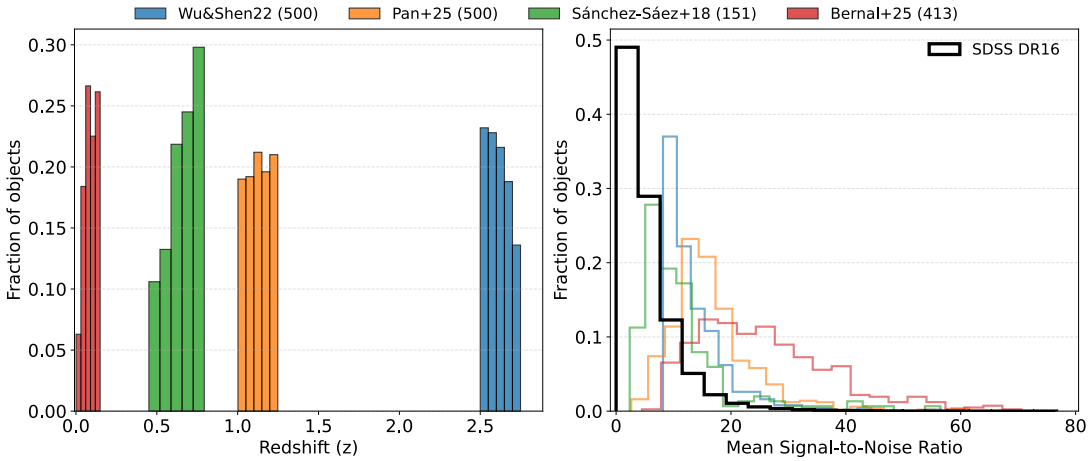


Figure 3.2: Left: Redshift distribution by sample. Right: Mean signal-to-noise ratio (S/N) distribution by sample and the distribution of S/N from DR16 in black. Sample names are shown at the top of the figure, with the number of objects in parentheses,

3.3.4 Uncertainty Estimation

Because AGN spectral decomposition is intrinsically nonlinear, the best-fit parameters alone may not capture the full correlation structure among parameters, nor the impact of local degeneracies between continuum components and emission-line profiles. To estimate uncertainties beyond local curvature approximations, `SHEAP` provides a Monte Carlo (MC) resampling mode based on perturbation sampling (e.g., Press et al., 1992). In this approach, the observed spectrum is repeatedly perturbed according to its per-pixel uncertainties, and the full model is re-fit for each perturbed realization. The ensemble of re-fits yields an empirical distribution of model parameters that naturally captures parameter covariances and certain nonlinear effects (Andrae, 2010; Shen et al., 2011).

Each MC realization is initialized at the same initial conditions as the main fitting. This choice accelerates convergence while still allowing the optimizer to explore nearby basins of the loss landscape when the likelihood surface is multimodal or exhibits strong degeneracies. Similar "flux-randomization + refitting" strategies are commonly used in quasar spectral fitting pipelines to obtain robust measurement uncertainties (e.g., Guo et al., 2018; Wu & Shen, 2022).

3.3.5 Post-fit derived quantities

Beyond delivering a single best-fitting model, SHEAP is designed to support a consistent post-processing stage in which both best-fit parameters and sampled parameters can be propagated into physically relevant ("added-value") quantities. In practice, the code can (i) evaluate derived quantities directly from the best-fit solution and/or (ii) evaluate them for each sample obtained (see Section 3.3.4), thereby producing posterior distributions for quantities such as line fluxes and luminosities, monochromatic continuum luminosities, EW, and measurements of post-processed (combined) broad-line profiles.

This design ensures that uncertainties in derived quantities are derived from the same modeling assumptions used during the fit (e.g., the adopted continuum and pseudo-continuum definitions, kinematic ties, and line-profile parameterization), while enabling downstream analyses to use either point estimates or full posterior samples. In the next sections, we explain how these are defined.

Additional products derived by SHEAP

A central goal of SHEAP is to translate spectral fits into physically meaningful and reproducible measurements. For each fitted spectrum and, when available, for each MC realization, the code provides a set of value-added products, including:

- Emission-line products (per line and, when applicable, for post-processed combined profiles): integrated line fluxes and line luminosities; EW; and kinematic/profile measurements such as FWHM and velocity shift, together with the other line parameters returned by the fit (e.g., amplitudes and profile-shape parameters, depending on the adopted parameterization).

- Monochromatic continuum luminosities: λL_λ evaluated at the standard wavelengths used in AGN applications (e.g., 5100, 3000, and 1350 Å)

For emission lines, EW is computed using the standard definition relative to a local level. In a decomposition context, this pseudo-continuum must be defined consistently (e.g., an AGN power-law plus additional continuum-like components such as Fe host-galaxy starlight, and/or the Balmer continuum, depending on the adopted model; Calderone et al., 2017; Bernal et al., 2025). For cosmological distances, we use the cosmological module of astropy (Astropy Collaboration et al., 2013, 2018, 2022).

Our method also allows the main broad-line component to be estimated using different approaches, either by combining all broad components or by combining only those that are kinematically consistent (see Appendix B).

In this way, our uncertainty-estimation framework (see Section 3.3.4) provides a natural mechanism for propagating uncertainties to all post-processing quantities, including luminosities, EW, combined-profile FWHM, and single-epoch SMBH mass estimates.

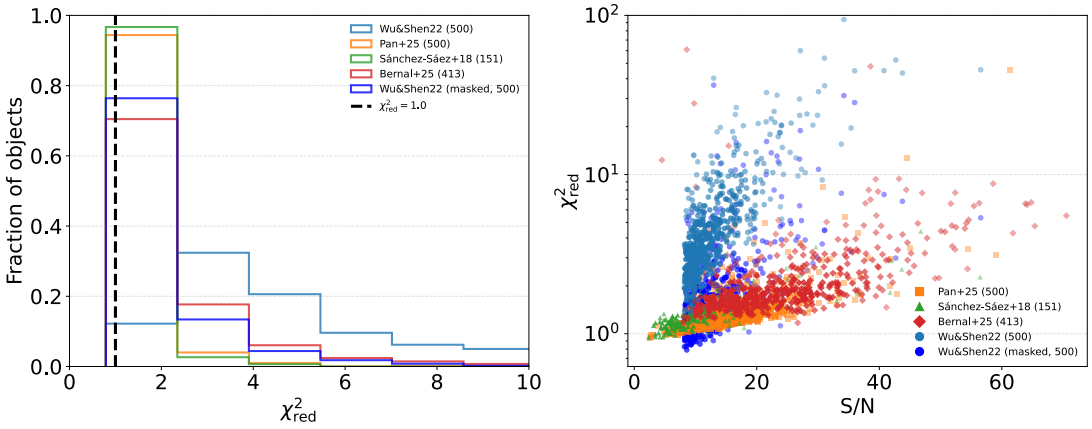


Figure 3.3: Right: Distribution of reduced χ^2 by sample, normalized to the fraction of objects in each sample. The vertical step line marks $\chi^2_{\rm red} = 1$. Left: Relation between signal-to-noise ratio (S/N) and the reduced chi-square, $\chi^2_{\rm red}$. Each color represent a different sample.

3.4 Performance validation of SHEAP

To evaluate the performance of our method across different spectral regions, resolutions, and AGN populations, we perform a series of tests by comparing our results with those from previously published studies. To avoid bias toward any single methodology, we include comparisons spanning a variety of independent approaches.

In these tests, all emission-line profiles are modeled with Gaussian components, and the full list of emission lines used in each model is given in Table C.1. The underlying AGN continuum is represented by a power law and, as in Ilić et al. (2023), we adopt a single baseline model for each sample. This ensures that differences within each sample primarily reflect data quality and population diversity, rather than changes in the model configuration.

For the results presented in this paper, we optimize each model for 2000 iterations

using a learning rate of 0.01 (see Section 3.3.3). Uncertainties are estimated using our framework for 50 realizations (see Section 3.3.4).

For each compared parameter (X), we define the logarithmic difference as

$$\Delta_X = \log_{10}(X_{\rm SHEAP}) - \log_{10}(X_{\rm lit}), \tag{3.5}$$

and we quantify the agreement using the fraction of objects satisfying ($|\Delta_X| \leq 0.3$) dex, where the adopted threshold corresponds to agreement within approximately a factor of two in linear space and provides a practical tolerance for the intrinsic scatter and systematic uncertainties associated with the different measurement methods. We complement this metric with the median offset and the normalized median absolute deviation (NMAD), which characterize the systematic and random differences between our measurements and those from the literature.

All computations were performed on a workstation equipped with two Intel Xeon Gold 5220R (2.20,GHz) CPUs and an NVIDIA RTX A5000 GPU.

In Fig. 3.2, we present the signal-to-noise ratio distributions for the different samples, together with their corresponding redshift distributions.

3.4.1 The Wu&Shen22 sample

We first test the performance of `SHEAP` against `PyQSOFit` (Guo et al., 2018), as applied by Wu&Shen22 to the DR16Q catalog of 750414 broad-line quasars over $0.1 \leq z \leq 6$. We select a stratified subsample with $2.0 \leq z \leq 2.25$ to test the C IV region, corresponding to a rest-frame wavelength coverage of 950–2968 Å. To reduce redshift-related systematics, we require consistency among the catalog redshift estimates, enforcing $|\Delta z| \leq 0.01$.

As described in Section 3.3.1, `SHEAP` operates on array-based inputs. Since SDSS spectra can have different wavelength coverages and array lengths (Bolton et al., 2012), we resample the selected spectra onto a common wavelength grid using `SpectRes` (Carnall, 2017), which conserves flux density and propagates uncertainties during rebinning. For this test, we fit the 1,100–2,750 Å range using a power-law continuum and Gaussian emission-line components, adopting three broad components for the C IV region, resulting in a model with 41 parameters.

Examples of the best-fit models are shown in Figs. D.1 and D.2. The main source of increased $\chi^2_{\rm red}$ is absorption from the Lyα forest at $\lambda_{\rm rest} \lesssim 1300$ Å. After masking this region, the $\chi^2_{\rm red}$ distribution becomes significantly narrower, with the median and

scatter decreasing from 4.3 ± 3.7 to 1.5 ± 0.95 (Fig. 3.3). We find no strong dependence of fit quality on S/N, although the highest-S/N objects show slightly larger $\chi^2_{\rm red}$ values.

We compare our measurements with the C IV FWHM and L_{1350} values reported by Wu&Shen22, combining broad components following the same strategy described in Appendix B. The results are shown in Fig. 3.4. For $\mathrm{FWHM_{C\,IV}}$, 96.40% of the objects lie within the ± 0.3 dex band, with a median bias of $+0.04$ dex and NMAD $= 0.07$ dex. For L_{1350}, the agreement is tighter, with 99.8% of objects within ± 0.3 dex, a median bias of $+0.01$ dex, and NMAD $= 0.02$ dex. The broader scatter in $\mathrm{FWHM_{C\,IV}}$ is expected because the broad profile is modeled with multiple components that are later combined.

The full distributions of the logarithmic differences are shown in Fig. 3.5A, and their dependence on S/N is shown in Fig. 3.6A. Overall, we find good agreement with Wu&Shen22, with no strong increase in scatter toward lower S/N, although the largest outliers occur preferentially at the lowest S/N.

3.4.2 The Pan+25 sample

The second test uses DESI DR1 spectra, which provide higher spectral resolution than SDSS and a uniform wavelength sampling across all spectra. We draw a subsample from the *Lya Working Group Quasar Redshift Catalog v1.7* (hereafter the *BHMass Catalog*; Pan et al., 2025), which contains spectral properties for 490,468 objects over $0.6 \leq z \leq 1.6$ and covers the Mg II region. The catalog reports measurements from two independent pipelines, `FastSpecFit` (`FSF`; Moustakas et al., 2023) and `DASpec` (`DAS`; Du, 2024).

We select 500 sources over $0.60 \leq z \leq 0.622$ and download their spectra following the official DESI data access procedure.[5] Each spectrum contains 7781 pixels, and the redshift and S/N distributions are shown in Fig. 3.2. We compare our measurements with the catalog values of Mg II FWHM, L_{3000}, and $R_{\rm Fe\,II}$, where

$$R_{\rm Fe\,II} = \frac{F_{\rm Fe}}{F_{\rm Mg\,II}}, \tag{3.6}$$

with $F_{\rm Fe}$ measured over 2250–2650 Å.

Since this test focuses on the Mg II region, we fit the spectra over 1750–4400 Å. The model includes a power-law continuum, Fe II emission, the Balmer continuum,

[5] `https://data.desi.lbl.gov/doc/access/`

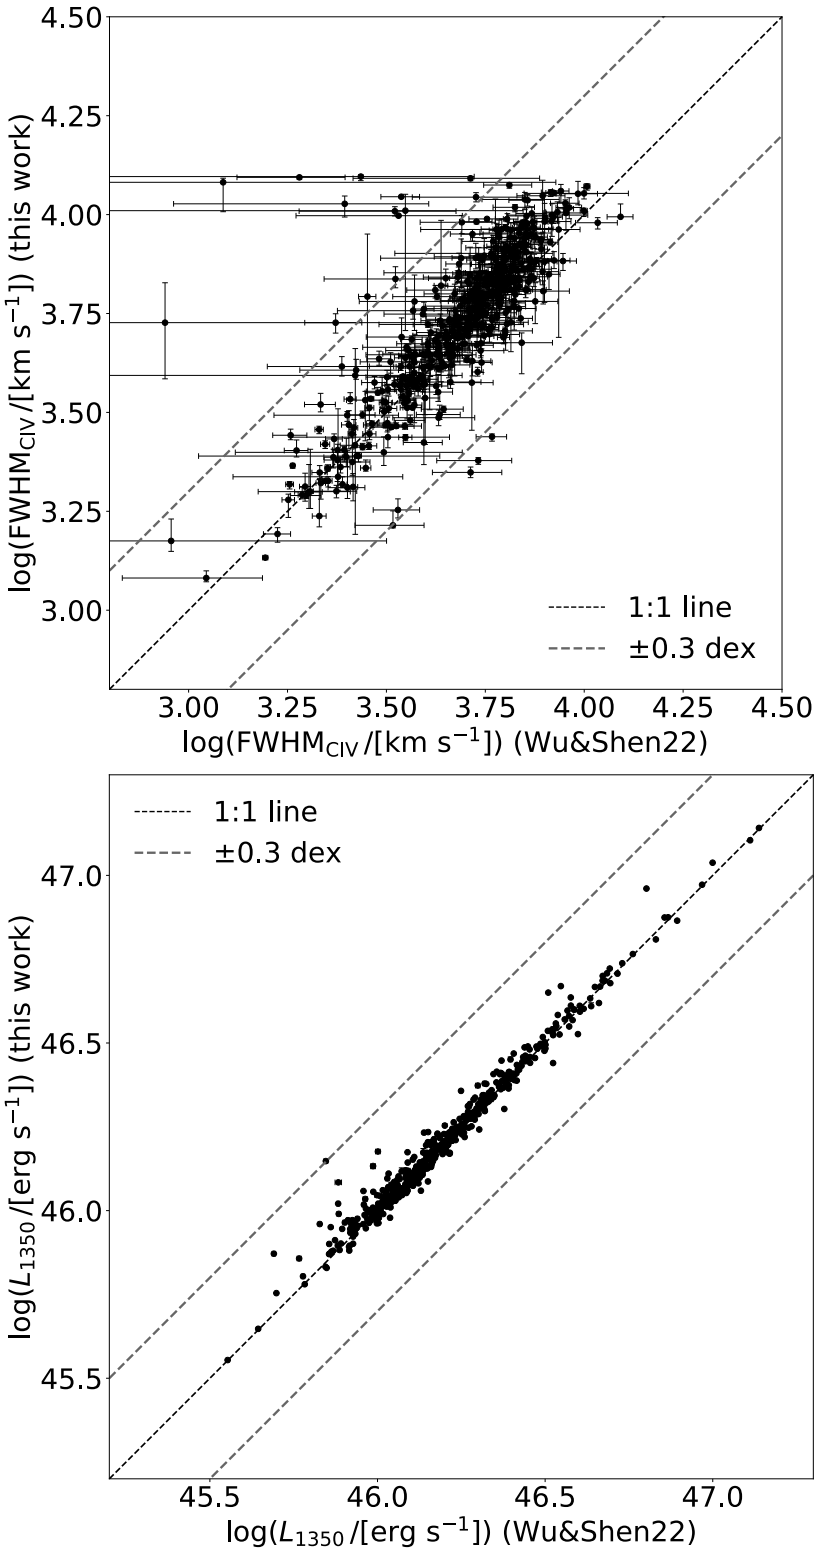


Figure 3.4: Comparison between our measurements and those of Wu&Shen22. Top: C IV emission-line FWHM in logarithmic scale. Bottom: Monochromatic continuum luminosity at 1350 Å in logarithmic scale. In both panels, the black dashed line indicates the 1:1 relation, while the gray dashed lines mark the ±0.3 dex region.

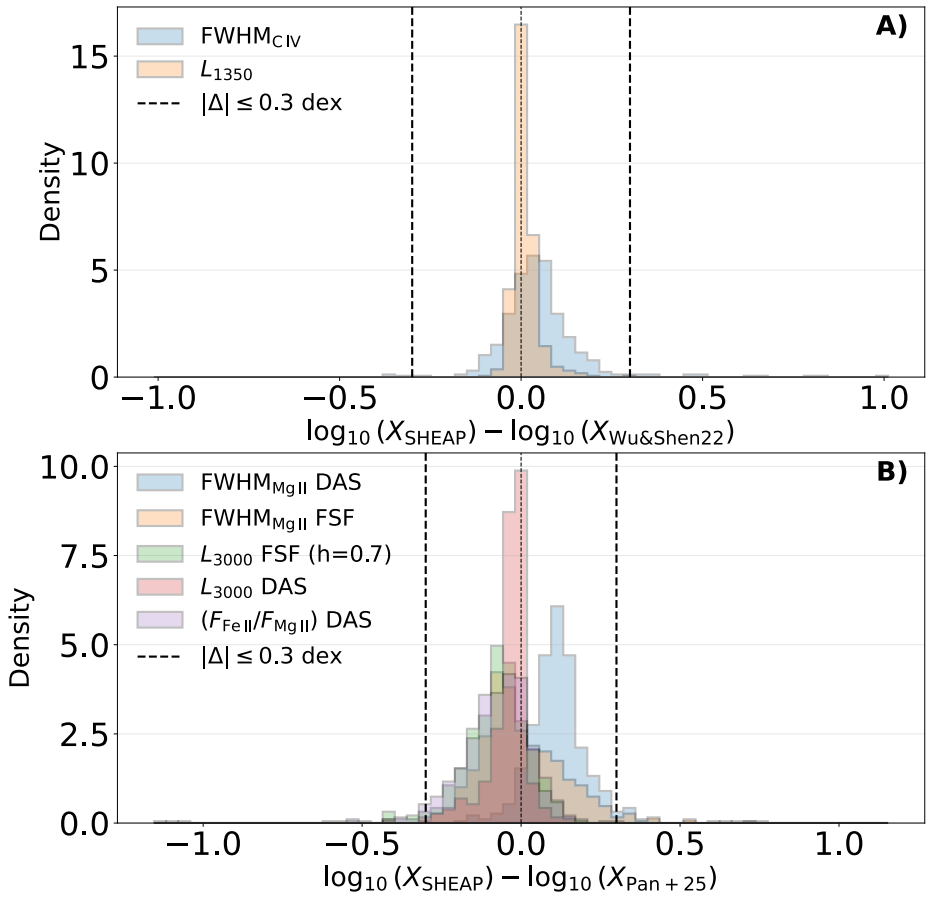


Figure 3.5: Distributions of the logarithmic differences, $\log_{10}(X_{\rm SHEAP}) - \log_{10}(X_{\rm lit})$, for the set of spectral parameters compared in each sample. Panel A shows against Wu&Shen22, while Panel B shows against Pan et al. (2025). In each panel, the legend reports the compared parameters. The grey dashed line indicates the 1:1 relation, while the black dashed-lines demark the $\pm$0.3 dex.

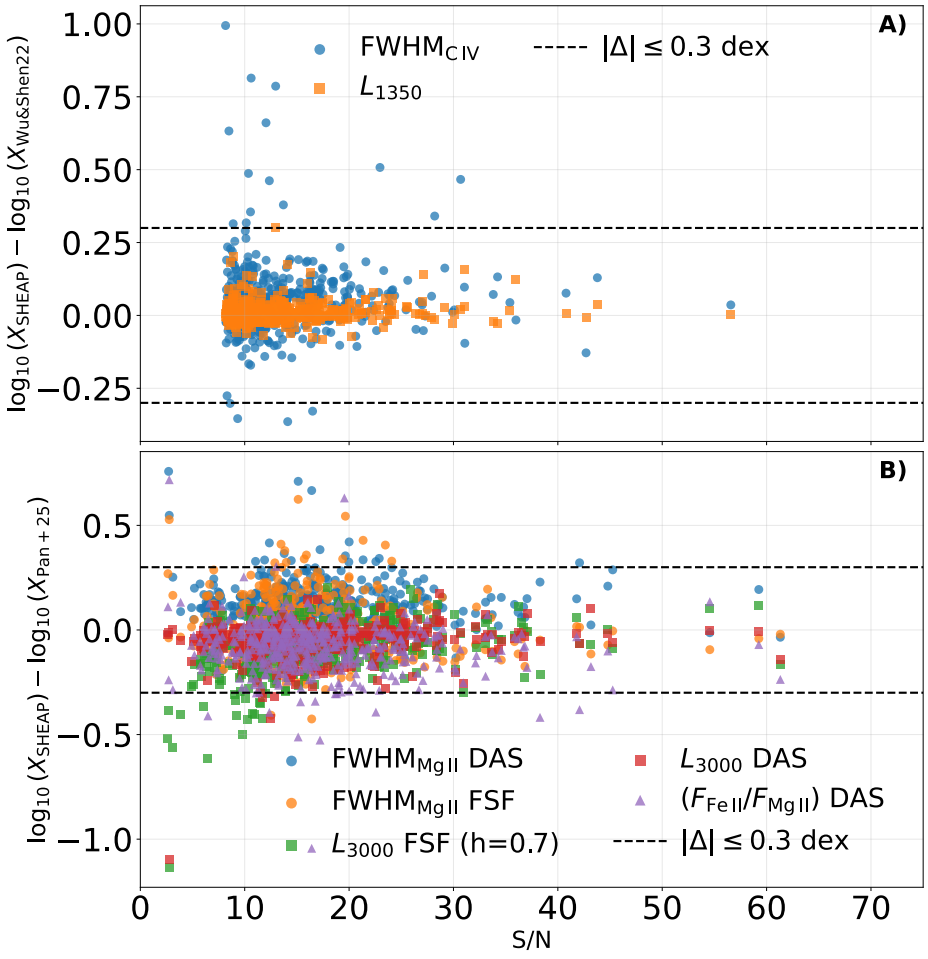


Figure 3.6: Differences between our measurements and those from the literature as a function of signal-to-noise ratio (S/N; x-axis). The y-axis shows $\log_{10}(X_{\rm SHEAP}) - \log_{10}(X_{\rm lit})$. Panel A against Wu&Shen22, and Panel B against Pan+25. The legend identifies each parameter, and the horizontal dashed lines mark the ± 0.3 dex band.

higher-order Balmer lines, and one broad Gaussian component for the emission lines. Although Pan+25 model Mg II with multiple broad and narrow components, we find that adding a second broad Gaussian does not significantly improve the fit quality. We therefore retain the simplest 26-parameter model for this validation test.

The fit-quality distribution gives a median $\chi^2_{\rm red} = 1.3 \pm 0.25$, with more than 90% of the objects having $\chi^2_{\rm red} \leq 2$ (Fig. 3.3). We observe a mild increase in $\chi^2_{\rm red}$ toward higher S/N, likely because higher-quality spectra reveal small model imperfections more clearly. Representative spectra close to the median S/N and median $\chi^2_{\rm red}$ are shown in Figs. D.3 and D.4.

The comparison with the catalog measurements is shown in Fig. 3.7. For $\mathrm{FWHM_{Mg\,II}}$, 96.6% of the DAS values and 97.0% of the FSF values lie within the ± 0.3 dex band, with median biases of $+0.12$ and -0.02 dex, respectively. For L_{3000}, the agreement is tighter: 99.20% of the DAS values and 96.00% of the corrected FSF values lie within the band, with median biases of -0.02 and -0.07 dex, respectively. These results indicate that the continuum luminosity is robustly recovered, while the FWHM comparison is more sensitive to differences in the adopted Mg II decomposition and Fe II treatment.

Finally, we compare $R_{\rm Fe\,II}$ against the DAS measurements only. We find that 96.80% of the objects lie within the ± 0.3 dex band, with a median bias of -0.07 dex and NMAD $= 0.01$ dex. The systematic underestimation of $R_{\rm Fe\,II}$ is consistent with the broader trends shown in Fig. 3.5B, where our method generally returns slightly smaller values than the catalog, except for quantities directly tied to the Mg II profile. As shown in Fig. 3.6B, most outliers occur at low S/N, with no strong increase in scatter toward higher S/N.

3.4.3 The Sánchez-Sáez+18 sample

In Sánchez-Sáez et al. (2018) (hereafter, Sánchez-Sáez+18), the authors analyzed 2,345 SDSS AGN spectra using the methodology of Mejía-Restrepo et al. (2016). For 224 sources, they performed detailed fits to the Hβ and Mg II regions, including host-galaxy subtraction through the *k-star* method (Greene & Ho, 2005; Kim et al., 2006). From this set, we retain 151 objects with $0.45 \leq z \leq 0.8$, which provides rest-frame coverage from 3000 to 5000 Å and enables a simultaneous analysis of the Hβ and Mg II regions. The S/N and redshift distributions are shown in Fig. 3.2. Since this sample is based on SDSS spectra, we resample it as described in Section 3.4.1.

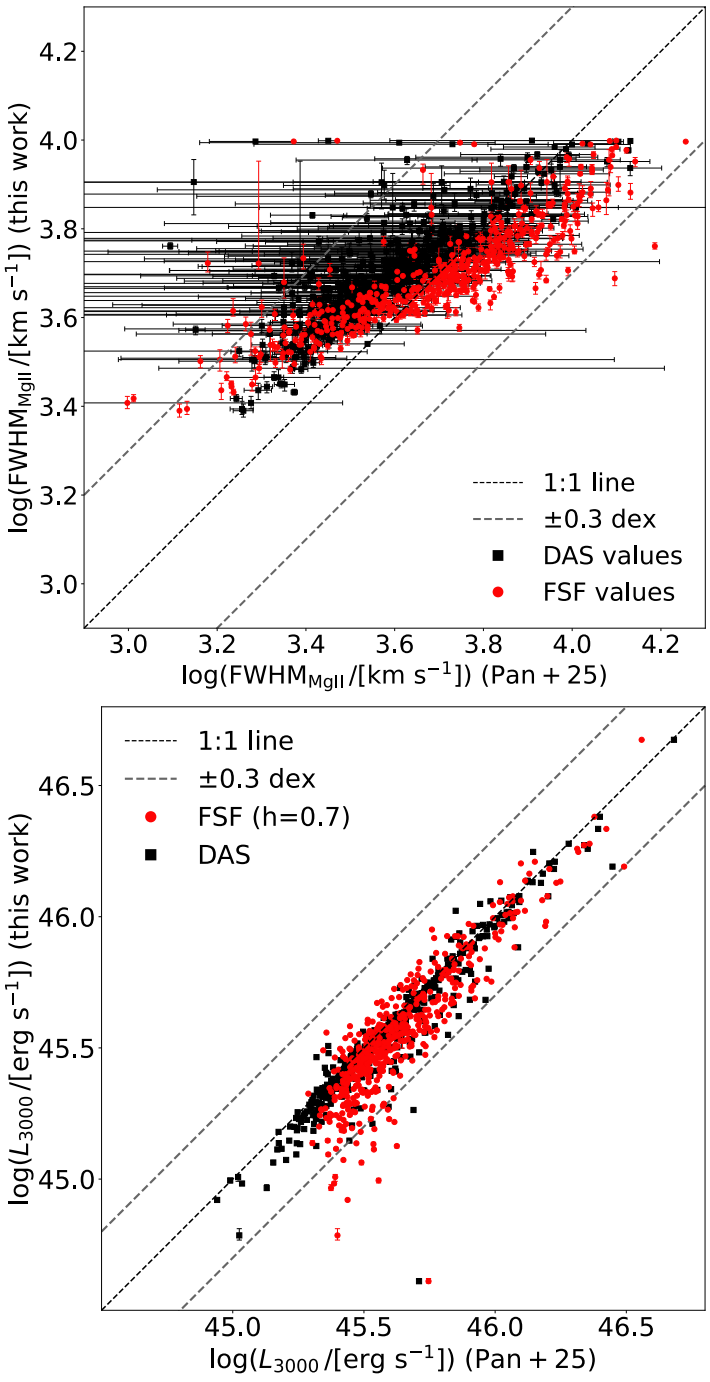


Figure 3.7: Comparison between our measurements and those of Pan+25. Top: Mg II emission-line FWHM in logarithmic scale. Bottom: Monochromatic continuum luminosity at 3000 Å in logarithmic scale. In both panels, the black dashed line indicates the 1:1 relation, while the gray lines mark the ±0.3 dex region. Red circles represent values computed with FSF, while black squares correspond to DAS.

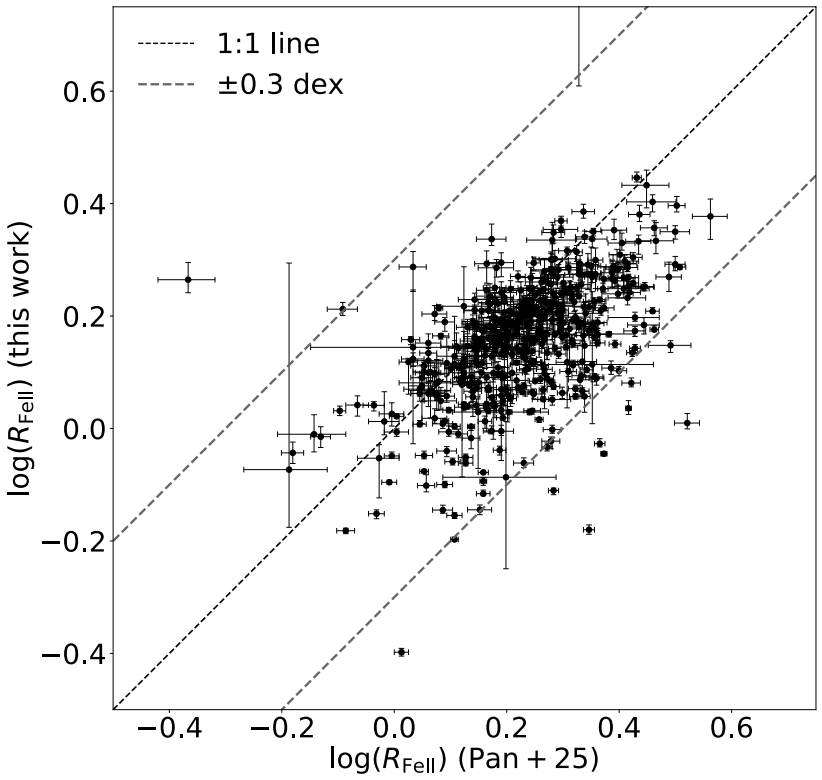


Figure 3.8: Comparison between our results and Pan et al. (2025) for the $R_{\rm FeII}$ in logarithmic scale. The black dashed line indicates the 1:1 relation, while the grey lines demark the ± 0.3 dex.

We compare our results with the reported FWHM values for Hβ and Mg II, together with the monochromatic continuum luminosities at 5100 and 3000 Å. The reference measurements were obtained by modeling the spectra with a double-Gaussian description of the broad emission lines, a narrow component, and Fe II emission, after subtracting the host-galaxy contribution based on the Ca II K absorption feature. In our analysis, we fit the common 2450–5500 Å window using two broad components plus a narrow component, Fe II emission, Balmer continuum and higher-order Balmer lines, possible outflow components, a power-law continuum, and our host-galaxy modeling method (Section 3.3.2). The resulting model contains 421 parameters.

This sample has the lowest S/N distribution among our tests. Nevertheless, the fit quality remains good, with a median $\chi^2_{\rm red} = 1.2 \pm 0.2$ and no objects exceeding $\chi^2_{\rm red} \sim 3$ (Fig. 3.3). As in the previous samples, $\chi^2_{\rm red}$ shows a mild increase with S/N, while low-S/N spectra remain close to unity. Representative fits near the median S/N and median $\chi^2_{\rm red}$ are shown in Figs. D.5 and D.6. Even in the lowest-S/N cases, the main components are recovered, although Fe II is less tightly constrained.

For this sample line widths are measured from the combined broad components using the "combining all broad components" approach (Section B) in both methods. The comparison is shown in Fig. 3.9Top. For $\mathrm{FWHM}_{\rm Hfi}$, 95.36% of the objects fall within the $\pm$0.3 dex band, with a median bias of $+0.01$ dex and NMAD $= 0.08$ dex. For $\mathrm{FWHM}_{\rm Mg\,II}$, considering the 117 objects with available catalog values, 92.30% lie within the band, with a median bias of $+0.08$ dex and NMAD $= 0.08$ dex. The somewhat larger Mg II values may reflect the degeneracy between Fe II emission and the Mg II wings.

The continuum-luminosity comparison is shown in Fig. 3.9Bottom. For L_{5100}, 77.48% of the objects lie within the $\pm$0.3 dex band, with a median bias of -0.13 dex and NMAD $= 0.17$ dex. This offset is likely driven by differences in host-galaxy subtraction, since the *k-star* method relies on the Ca II K absorption feature, whereas our approach uses a broader set of spectral features. For L_{3000}, the agreement is stronger: 99.12% of the 117 objects with available catalog values fall within the band, with a median bias of -0.10 dex and NMAD $= 0.07$ dex.

Overall, the parameter-difference distributions are shown in Fig. 3.10A, and their dependence on S/N is shown in Fig. 3.11A. The largest discrepancies occur preferentially at low S/N, but we do not find evidence for a strong systematic trend with S/N.

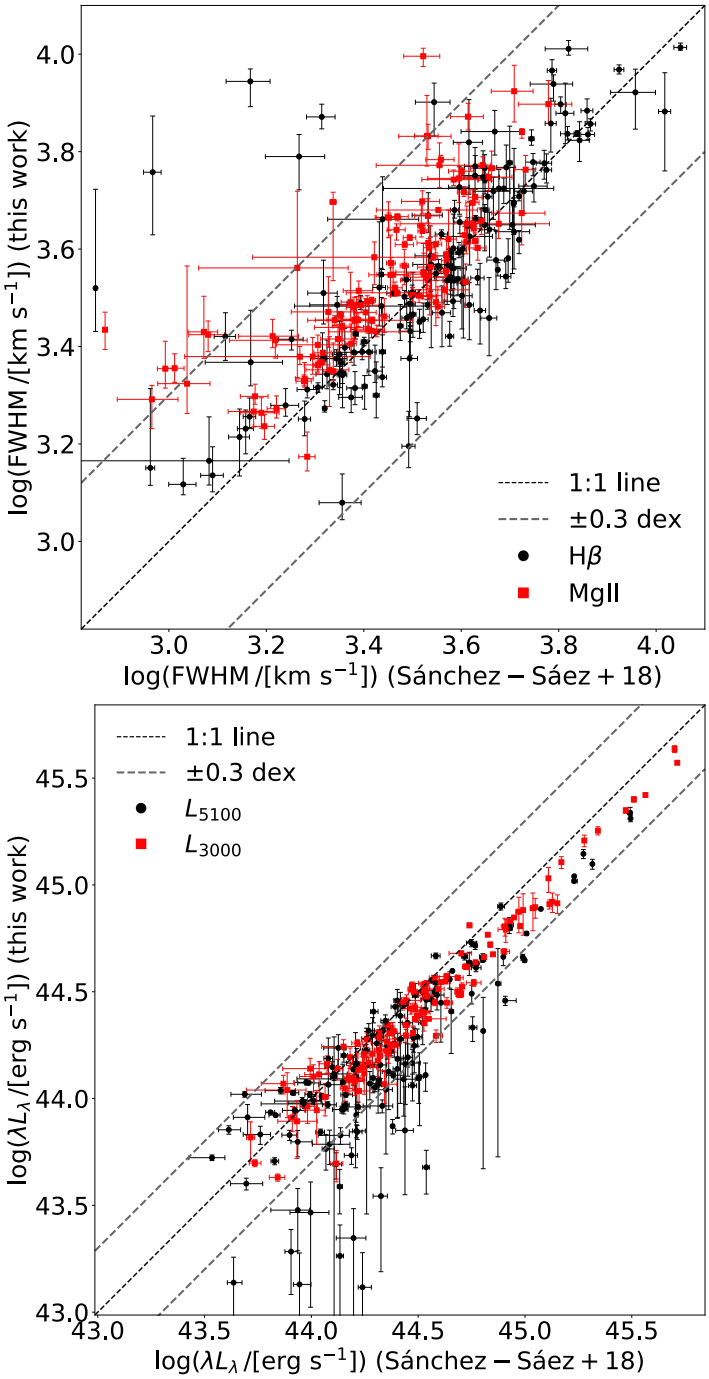


Figure 3.9: Comparison between our measurements and those of Sánchez-Sáez+18. Top: comparison of the emission-line FWHM in logarithmic scale. Bottom: comparison of the monochromatic continuum luminosity in logarithmic scale. In both panels, the black dashed line indicates the 1:1 relation, while the gray lines mark the ±0.3 dex region. In the top panel, black circles correspond to Hβ measurements, while red squares represent Mg II. In the bottom panel, black circles correspond to L_{5100} measurements, while red squares represent L_{3000}.

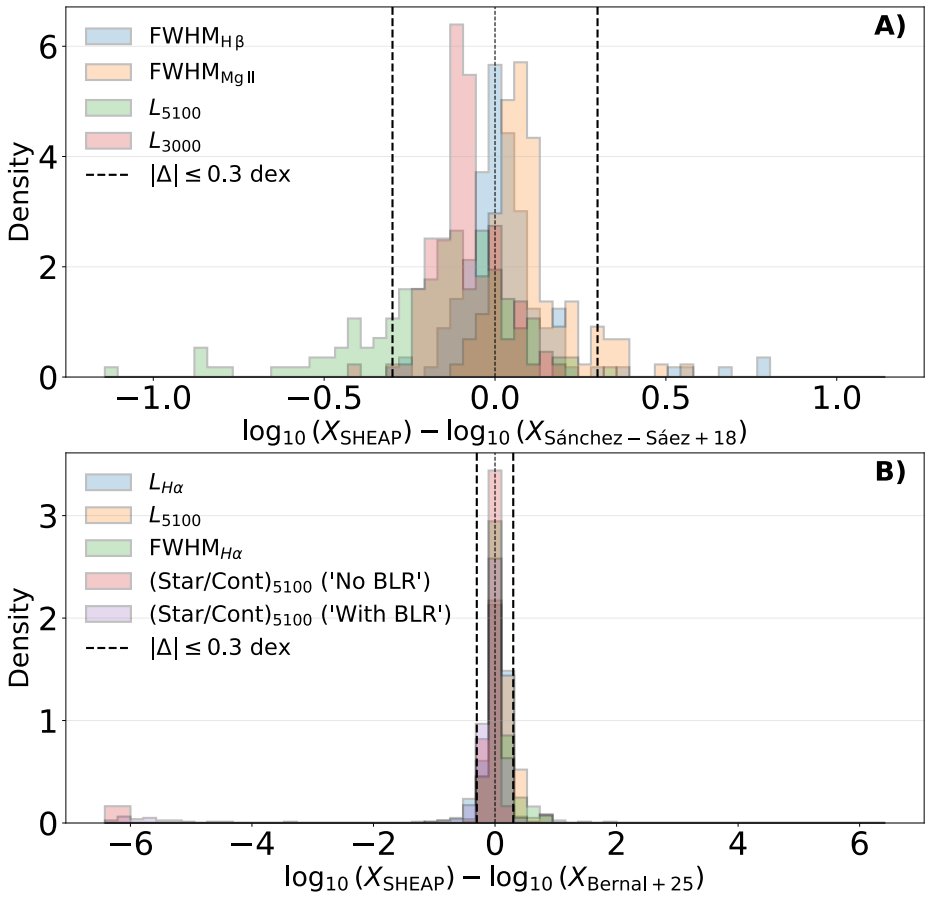


Figure 3.10: Same as Fig. 3.5 but this time Panel A shows the comparison against Sánchez-Sáez+18, meanwhile Panel B shows the comparison against Bernal+25.

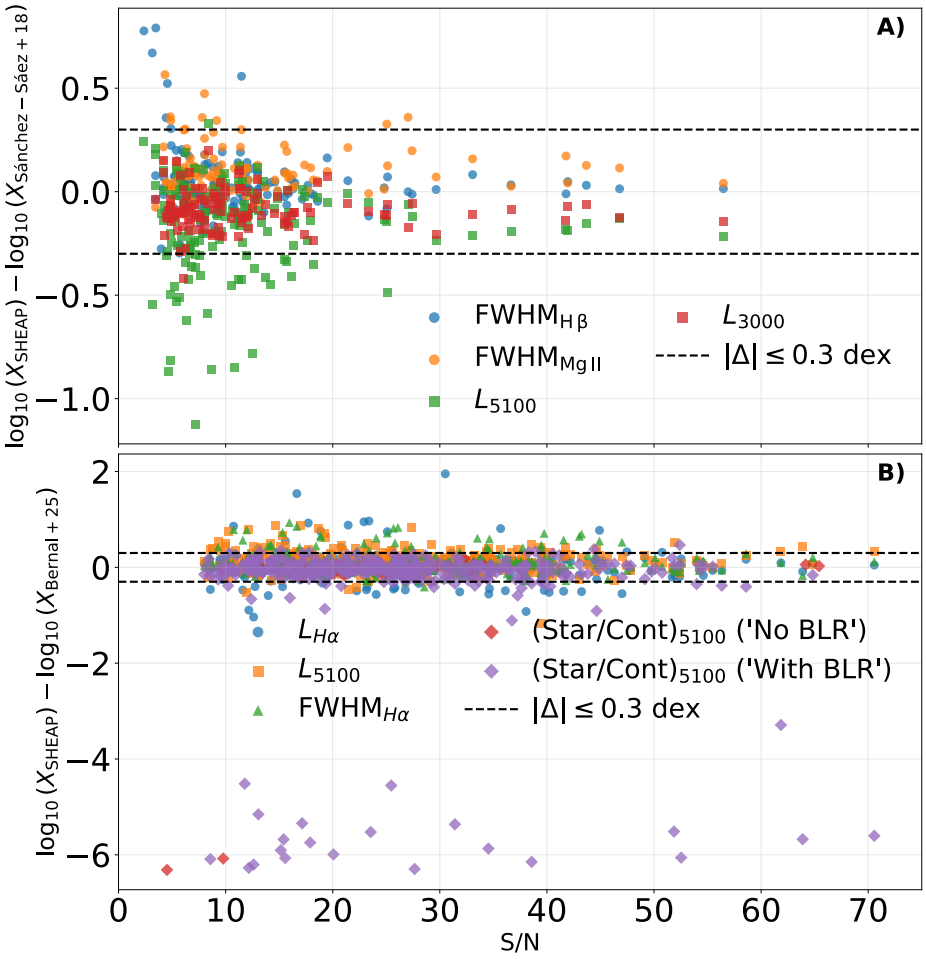


Figure 3.11: Same as Fig. 3.6 but this time Panel A compares against Sánchez-Sáez+18, and Panel B compares against Bernal+25.

3.4.4 The Bernal+25 sample

To evaluate the performance of our method in a low-redshift, host-dominated sample, we use the dataset presented by Bernal+25. This sample spans $0.01 \leq z \leq 0.14$ and was analyzed with a customized pPXF wrapper to decompose the AGN and host-galaxy components. From the publicly available catalog[6], we construct a sample of 413 objects, whose redshift and S/N distributions are shown in Fig. 3.2. The spectra cover 3559–10,394 Å in the observed frame and, since they are based on SDSS data products, are resampled as described in Section 3.4.1.

We compare the host-to-continuum ratio at 5100 Å $(\mathrm{Star/Cont})_{5100}$, the monochromatic luminosity L_{5100}, the FWHM of the broad Hα component, and the broad Hα luminosity. To obtain comparable measurements, we adopt a model similar to that of Bernal+25, including two broad components, a narrow component, Fe II emission, Balmer emission, and a host-galaxy component. We fit the rest-frame wavelength range 3800–8000 Å, resulting in a model with 427 parameters.

Despite the complexity of these spectra, the reduced chi-square distribution has a median of $\chi^2_{\mathrm{red}} = 1.8 \pm 0.9$ (Fig. 3.3). As in the other samples, χ^2_{red} increases mildly with S/N, likely because higher-S/N spectra reveal small model imperfections more clearly. A representative fit near the median S/N is shown in Fig. D.7, where the host-galaxy component is well reproduced.

For $(\mathrm{Star/Cont})_{5100}$, 86.90% of the objects lie within the ± 0.3 dex band, with a median bias of -0.04 dex and NMAD $= 0.129$ dex (Fig. 3.12). When separating the sample according to the BLR classification of Bernal+25, the agreement is 86.42% for the BLR subsample and 93.10% for the no-BLR subsample. A small number of extreme outliers are present, but visual inspection shows that most BLR outliers do not exhibit a clear host-galaxy contribution, while the no-BLR outliers have unreliable fits with $\chi^2_{\mathrm{red}} > 10$.

For the remaining comparisons, we restrict the analysis to the 384 objects with a detected broad component based on Bernal+25 criteria. The Hα FWHM and luminosity comparisons are shown in Fig. 3.13. For $\mathrm{FWHM}_{\mathrm{H}\alpha}$, using the kinematic-combination method (Section B), 88.80% of the objects fall within the ± 0.3 dex band, with a median bias of $+0.04$ dex and NMAD $= 0.11$ dex. For the broad Hα luminosity, 87.76% fall within the band, with a median bias of $+0.04$ dex and NMAD $= 0.15$ dex. In both cases, the largest discrepancies occur preferentially in systems with low equivalent

[6] https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/694/A127

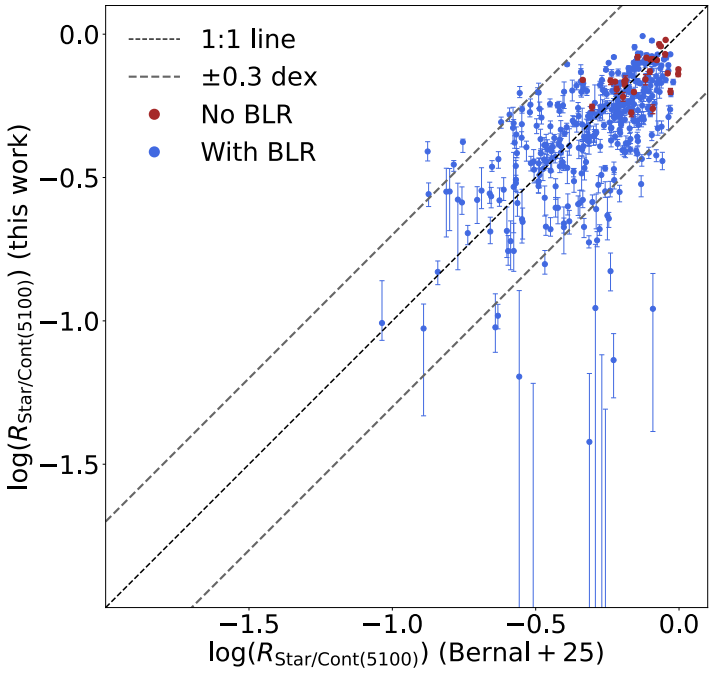


Figure 3.12: Comparison between the stellar-to-continuum flux ratio at 5100 Å derived by Bernal+25 and the values obtained in this work. Red points show all objects in the sample, while blue points highlight sources with a confirmed broad emission-line component. The dashed line marks the 1:1 relation, and the shaded region indicates the ±0.3 dex agreement band.

widths (e.g., EW ¡ 10 Å). while objects with EW $\gtrsim$ 100 Å show the tightest agreement.

Finally, for L_{5100}, 84.33% of the objects lie within the ±0.3 dex band, with a median bias of +0.08 dex and NMAD = 0.15 dex (Fig. 3.14). The parameter-difference distribution in Fig. 3.10B shows that our values are systematically larger, consistent with differences in the AGN–host decomposition: relative to Bernal+25, our fits assign a smaller contribution to the stellar component, increasing the inferred AGN continuum luminosity at 5100 Å.

3.5 Computational time

Another important test is the computational time required to fit the different samples. The measured execution times are reported in Table 3.1. For this analysis, we separately measured the time required for the two main computational tasks. First, we recorded the time needed by our method to fit each complete sample. Second, we measured the time required to perform the Monte Carlo iterations and derive poste-

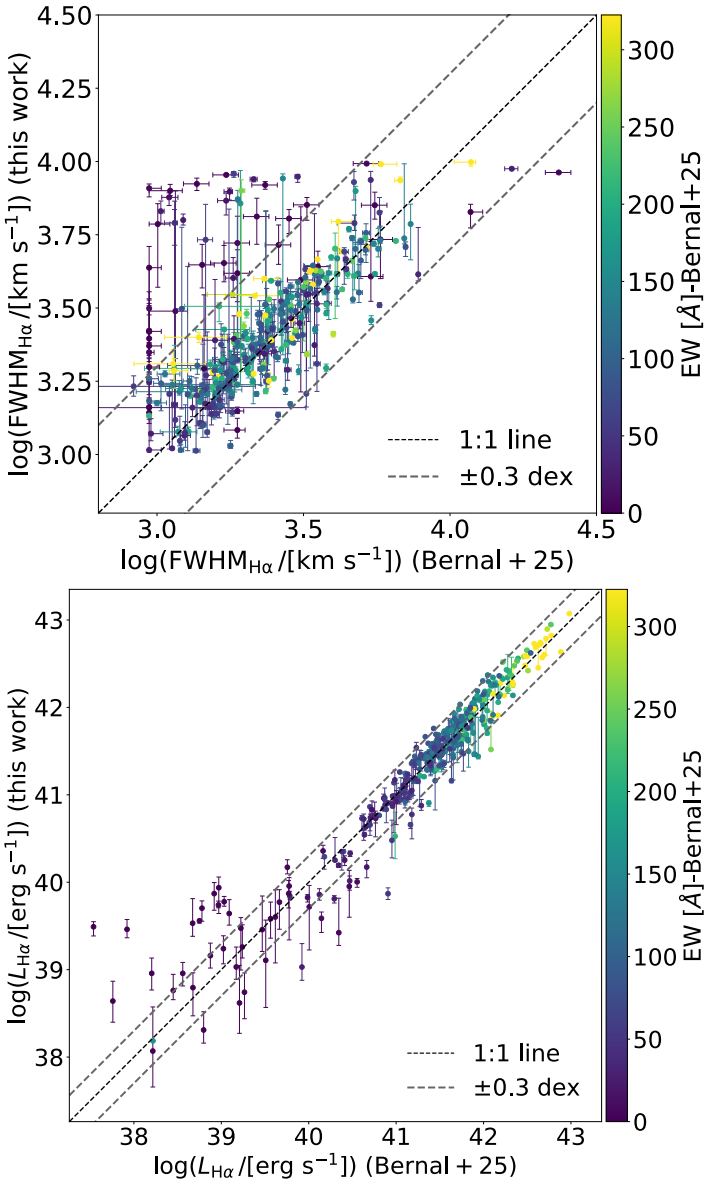


Figure 3.13: Comparison between our measurements and those of Bernal+25. Top: Hα emission-line FWHM in logarithmic scale. Bottom: Hα emission-line luminosity in logarithmic scale. In both panels, the black dashed line indicates the 1:1 relation, while the gray lines mark the $\pm$0.3 dex region. Point colors represent the EW reported by Bernal+25.

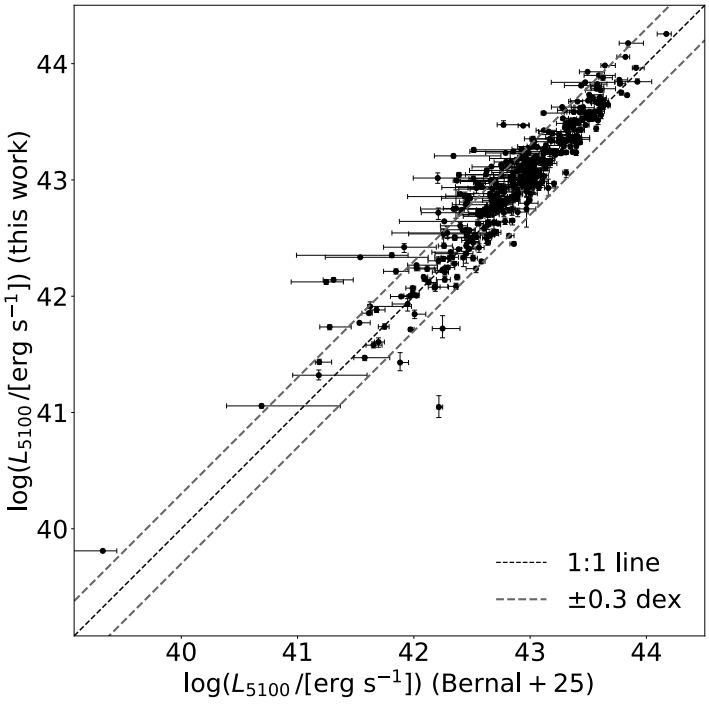


Figure 3.14: Comparison between our results and Bernal+25 for the monochromatic continuum luminosity at 5100 Å in log scale.The black dashed line indicates the 1:1 relation, while the grey lines demark the ± 0.3 dex.

Table 3.1: Summary of fit quality, model complexity, and computational performance for the literature comparison samples.

Sample	$N_{\rm spec}$	$\widetilde{\rm S/N}$	N_λ	$N_{\rm par}$	$\widetilde{\chi}^2_{\rm red}$	$t_{\rm fit}$ [min]	$t_{\rm MC}$ [min]
Wu&Shen+22 (C IV)	500	11.93	3809	41	$4.3^{+5.5}_{-1.8}(1.5^{+1.4}_{-0.5})$	0.2	5.3
Pan+25 (Mg II)	500	15.25	7781	26	$1.3^{+0.4}_{-0.1}$	0.9	42.1
Sánchez-Sáez+18 (Hβ/Mg II)	151	9.33	3808	421	$1.2^{+0.3}_{-0.1}$	0.9	16.9
Bernal+25 (Hα)	413	24.13	2878	427	$1.8^{+1.4}_{-0.4}$	1.3	35.5
Bernal et al. (2026) (pPXF iteration)	413	–	–	–	–	$130.8^{+123.9}_{-41.3}$	–

Table 3.2: $\widetilde{\chi}^2_{\rm red}$ denotes the median reduced chi-square, reported together with the 16th and 84th percentiles as asymmetric uncertainties. For Wu&Shen+22, we report the unmasked value, with the value obtained after masking $\lambda < 1300$ Å given in parentheses. $\widetilde{\rm S/N}$ is the median signal-to-noise ratio of the sample. $N_{\rm par}$ is the number of free parameters used in the fit, corresponding to the length of the optimized parameter vector. N_λ denotes the number of pixels per spectrum. $t_{\rm fit}$ is the runtime for a single best-fit optimization per spectrum, and $t_{\rm MC}$ is the runtime for the Monte Carlo posterior stage per spectrum.

rior distributions for the value-added products described in Section 3.3.5.

For all samples, the fitting time ranges from 0.2 to 1.3 minutes, corresponding to less than one second per object, independently of the sample considered. The Monte Carlo analysis requires between 5.3 and 42.1 minutes, with a non-linear dependence on the characteristics of the sample. Achieving this level of computational performance is primarily enabled by GPU-accelerated programming (Dupourqué et al., 2024). For comparison, we also include the execution times reported by Bernal et al. (2026), who obtained a total fitting time of $130.8^{+123.9}_{-41.3}$ minutes for the same sample. Compared with this result, our method achieves a computational speed-up of approximately 101 times.

Another relevant result of this analysis is that the main source of computational cost appears to be the number of spectral pixels rather than the number of free parameters included in the minimization. This can be illustrated by comparing the samples of (Pan+25) and (Bernal+25). Their models contain 26 and 427 parameters, respectively, corresponding to a difference of 401 parameters. Nevertheless, their fitting times differ by only 0.3 minutes, while their spectra differ by approximately 3900 pixels. This comparison indicates that the number of spectral pixels has a stronger impact on the computational cost than the number of model parameters.

3.6 Discussion & Conclusions

In this work, we presented SHEAP, a GPU-enabled AGN spectral-fitting code designed for efficient, flexible, and scalable analysis of large spectroscopic datasets. We validated its performance using four literature-based comparison samples spanning different redshift ranges, spectral regions, resolutions, and AGN populations, including SDSS and DESI spectra covering the C IV, Mg II, Hβ, and Hα regions. This strategy demonstrates that SHEAP performs robustly under heterogeneous observational and modelling conditions, rather than being tuned to a single homogeneous dataset.

Overall, SHEAP provides statistically acceptable fits across all samples. The reduced chi-square distributions are typically centered near unity, with median values of $\tilde{\chi}^2_{\rm red} \sim 1.2$–$1.8$ for the Pan+25, Sánchez-Sáez+18, and Bernal+25 samples. The Wu&Shen+22 sample initially shows a broader distribution, but after masking the region at $\lambda_{\rm rest} < 1300$ Å, affected by the Lyα forest, the median $\tilde{\chi}^2_{\rm red}$ improves from $4.3^{+5.5}_{-1.8}$ to $1.5^{+1.4}_{-0.5}$. Across all samples, we do not find evidence for a strong degradation of fit quality toward lower S/N; instead, the mild increase of $\chi^2_{\rm red}$ at high S/N likely

reflects the increased visibility of subtle spectral features and modelling imperfections.

The comparison with literature measurements shows that SHEAP recovers the main AGN spectral parameters with generally good agreement. The fraction of objects within the ± 0.3 dex band is typically above 90% for line-width and continuum-luminosity measurements in the Wu&Shen+22, Pan+25, and Sánchez-Sáez+18 samples, including $\mathrm{FWHM_{C\,IV}}$, $\mathrm{FWHM_{Mg\,II}}$, $\mathrm{FWHM_{H\beta}}$, L_{1350}, L_{3000}, and $R_{\mathrm{Fe\,II}}$. Lower agreement is found for host-sensitive quantities, such as L_{5100} and $(\mathrm{Star/Cont})_{5100}$, particularly in the Sánchez-Sáez+18 and Bernal+25 samples, where differences in host-galaxy subtraction and AGN–host decomposition become more important.

The main residual discrepancies are therefore linked to methodological differences rather than to a single systematic bias in SHEAP. In the UV and Mg II regions, differences are mainly associated with the parameterization and combination of broad-line components, as well as the treatment of Fe II emission. In the optical region, the largest offsets arise in continuum-related quantities, especially L_{5100}, where the balance between stellar and AGN continuum components becomes critical.

A major result of this work is that this level of agreement is achieved at low computational cost. The runtime depends more strongly on the number of pixels per spectrum and the number of spectra being fitted than on the formal number of free parameters, allowing models with hundreds of parameters to remain tractable. In absolute terms, the best-fit optimization requires between 0.2 and 1.3 minutes for the tested subsamples, whereas the Monte Carlo stage requires between 5.3 and 42.1 minutes per tested subsample, this including 50 Monte Carlo realizations plus the initial best-fit run.

As a direct benchmark, our implementation requires 1.3 minutes for the fitting stage, compared to the $130.8^{+123.9}_{-41.3}$ minutes reported by Bernal et al. (2026) for their pPXF-based iterative pipeline. This corresponds to an improvement factor of approximately 101 times, or $\sim 1.0\%$ of the reported fitting time. This comparison should be interpreted with caution, since differences in hardware and computational environment may also contribute to the runtime difference.

Taken together, these results show that SHEAP can recover key AGN spectral properties across heterogeneous datasets while substantially reducing computational cost. Its GPU-accelerated, gradient-based framework makes it well suited for upcoming large spectroscopic surveys, where scalability, reproducibility, and flexible model configurations are essential. This flexibility also makes SHEAP particularly suitable for gravitationally lensed quasars, where microlensing can differentially magnify emis-

sion arising from regions of different physical sizes and thereby alter the continuum, broad emission lines, and their relative spectral components. By allowing the model configuration to be adapted to individual systems and spectral regions, SHEAP can help separate intrinsic quasar emission from microlensing-induced distortions and support systematic studies of the structure of quasar accretion discs and broad-line regions. Future developments will extend this framework through more flexible line-profile parameterizations, applications to integral-field spectroscopy data cubes, and adaptations to related problems such as galaxy stellar-population analysis and dynamical spectral modelling.

3.7 Data availability

In the body of the thesis, we presented Figures and results for a series of comparisons. We include an online appendix that reproduces all the results presented in this work. We provide an online appendix at https://github.com/felavila/Results_paper_SHEAP.

CHAPTER 4

Conclusions

Throughout this thesis, we have investigated the role of microlensing in strongly lensed quasars and developed new methodologies to analyze these systems in a reliable, flexible, and computationally efficient way. The work presented here is motivated by two closely connected problems. First, gravitational microlensing is one of the few observational techniques capable of probing the unresolved inner regions of quasars, including the optical/UV accretion disk and the inner broad-line region (BLR), on spatial scales far below the angular resolution of current telescopes. Second, the rapidly increasing volume of imaging and spectroscopic data from current and forthcoming surveys requires automated tools that can model lensed quasars and active galactic nuclei (AGN) spectra on a large scale.

The first part of this thesis focused on the statistical characterization of microlensing variability in known lensed quasars. This work has already been accepted for publication in *Astronomy & Astrophysics* and represents one of the largest samples of lensed quasars analyzed within this framework. To carry out this study, we developed an automated lens-modeling pipeline that provides the macromodel quantities required to characterize the microlensing perturbations affecting each lensed image. These quantities were then used to generate microlensing magnification maps that account for the granular stellar mass distribution in the lens galaxy, rather than treating microlensing only as a point-source or isolated-lens phenomenon.

This approach allowed us to estimate the probability of microlensing variability,

characterize the typical timescales of the events, and place constraints on physical quantities such as the half-light radius of the quasar emission region and the stellar mass fraction in the lens galaxy. In particular, the use of magnification maps is essential because the observed microlensing signal does not depend only on the Einstein radius of an isolated microlens, but also on the convergence, shear, stellar mass fraction, source size, and caustic structure experienced by each lensed image.

For the source-size scales considered in this work, we find characteristic microlensing timescales of $t_{\rm E} = 11.29 \pm 0.05$ yr and $t_{\rm S} = 2.59 \pm 0.07$ yr. Here, $t_{\rm E}$ is the Einstein-radius crossing time, defined as the time required for the effective transverse motion of the system to cross one projected Einstein radius of a typical microlens. This timescale provides a useful global reference scale for microlensing variability, because it is set by the characteristic length scale of the microlensing magnification pattern. However, $t_{\rm E}$ does not necessarily correspond to the duration of an observed microlensing event affecting a finite quasar emission region.

The second timescale, $t_{\rm S}$, is the source crossing time. It is defined as the time required for the effective transverse motion to cross the characteristic size of the source, such as the accretion-disk half-light radius. Since the continuum-emitting region is usually smaller than the projected Einstein radius, $t_{\rm S}$ is generally shorter than $t_{\rm E}$. It is therefore more closely related to the timescale over which a compact quasar component can show significant microlensing variability, especially during caustic-crossing or high-magnification events (HME).

Another finding of our study concerns the occurrence of HMEs. Despite their scientific value, well-sampled HMEs remain rare and have been observed almost exclusively in the Einstein Cross, Q2237+0305 (Mediavilla et al., 2015). In our analysis, we estimate an average HME probability of approximately 9% across the analyzed sample. This value appears modest compared with some simulation-based predictions for future LSST-like monitoring. For example, Neira et al. (2025) predict approximately 60 HMEs per year among ground-detectable lensed quasars with image separations larger than $1''$ and sufficiently bright images, whereas Fagin et al. (2025) estimate that LSST-like light curves could enable the successful prediction of approximately 55% of HME peaks, corresponding to tens to hundreds of predicted events per year. However, these estimates are not directly comparable to our per-image probability because they differ in sample size, selection criteria, monitoring duration, cadence, and event definition. Therefore, rather than demonstrating a lower intrinsic HME rate, our result highlights the need to account carefully for survey selection and temporal coverage

when comparing observed probabilities with predictions for future surveys.

These events are particularly important because they occur when a compact emitting region approaches or crosses a caustic in the microlensing magnification pattern. In this situation, the caustic acts as a natural spatial scanner of the source. Therefore, HME can provide especially strong constraints on the size and surface-brightness profile of the accretion disk and, potentially, on the spatial structure of the inner BLR. The relevant timescales also make target selection important: in our sample, we find a mean Einstein-radius crossing time of 11.29 ± 0.05 yr and a mean source-crossing time of 2.59 ± 0.07 yr, while earlier scale-based estimates for 87 lensed quasars found a median Einstein-radius crossing time of 20.6 yr and a median source-crossing time of (7.3) months (Mosquera & Kochanek, 2011).

In addition, we obtain a characteristic half-light radius of $R_{1/2} = 5.4 \pm 2.7$ light-days, and infer a lower limit on the stellar mass fraction of $\alpha \geq 0.15$. Together with the HME probabilities and timescales, these results make our sample valuable for selecting promising targets for future high-magnification-event monitoring. This is especially relevant in the LSST era, where survey photometry will identify variability in large samples, but complementary high-cadence photometric or spectroscopic follow-up will be needed to fully exploit the most informative events.

These constraints highlight the power of microlensing as a quantitative probe of quasar structure and demonstrate the importance of including realistic magnification patterns when interpreting the variability of lensed quasars.

The results of the first part show that microlensing cannot be fully characterized using overly simplified point-source approximations. Point-source estimates can provide useful characteristic scales, but they neglect the finite size of the emitting regions and the smoothing effect produced when an extended source crosses a complex caustic network. In contrast, realistic simulations based on microlensing magnification maps allow the source size, caustic morphology, and macromodel properties of each lensed image to be included self-consistently. This is crucial for identifying promising systems for future monitoring campaigns and for interpreting HME in the context of time-domain surveys such as LSST.

The second part of this thesis focused on studying the structure of AGNs through optical spectroscopy. In particular, the spectroscopic analysis of gravitationally lensed quasars is essential for disentangling the intrinsic properties of the background quasar from the effects introduced by the lens galaxy, including stellar microlensing. Gravitational magnification also provides access to intrinsically fainter or more distant

quasars that may be underrepresented in conventional spectroscopic surveys. However, microlensing by stars in the lens galaxy can differentially modify the continuum and emission-line profiles of the observed spectra. Before these effects can be interpreted reliably, a robust method for decomposing and characterizing AGN spectra is therefore required. This need, together with the rapidly increasing number of spectra provided by surveys such as SDSS, DESI, and 4MOST, motivates the development of tools capable of performing both classical AGN spectral fitting and more complex tasks, such as separating microlensing signatures from the intrinsic spectra of lensed quasars.

This work is directly connected to microlensing because its signal is produced by the differential magnification of spatially distinct emitting components rather than by the uniform magnification of the quasar as a whole. The compact continuum-emitting accretion disk is expected to be strongly affected by microlensing, whereas more extended components, such as the narrow-line region, should be only weakly affected and can therefore provide a reference emission. The broad-line region occupies an intermediate spatial regime and may exhibit velocity-dependent microlensing signatures. Consequently, a reliable decomposition of the observed spectrum is necessary to isolate the continuum, broad and narrow emission lines, Fe,II complexes, host-galaxy contribution, and possible outflow components.

To address this challenge, we developed `SHEAP`, a new GPU-accelerated tool designed to analyze AGN spectra rapidly, reliably, and flexibly. The tool enables the different components of an AGN spectrum to be modeled simultaneously while retaining the computational efficiency required for large samples. We validated the methodology using 2,000 spectra and compared the recovered measurements with those obtained using established spectral-analysis codes. The results were broadly consistent across the different methods, while `SHEAP` required substantially less computational time. These findings demonstrate that GPU-based implementations provide a practical route for scaling detailed AGN spectral modeling to the large datasets produced by current and forthcoming spectroscopic surveys, including 4MOST (Bauer et al., 2023; Merloni et al., 2019; Collett et al., 2023).

Taken together, the two parts of this thesis that provide complementary approaches to the study of quasar structure. The microlensing analysis constrains the spatial scales and variability expected in lensed quasars, while the spectroscopic methodology provides a way to separate and characterize the physical components that may be differentially affected by microlensing. In this sense, the work presented

here connects gravitational-lensing variability with the broader problem of AGN spectral decomposition. Spectral modeling identifies the physical components that contribute to the observed emission, and microlensing provides a way to transform their unresolved spatial structure into measurable magnification differences.

More generally, this thesis shows that the next generation of lensed-quasar and AGN studies will require methods that are both physically motivated and computationally scalable. The increasing number of known lensed quasars, together with the arrival of large time-domain and spectroscopic surveys such as LSST, DESI, and 4MOST, will make it possible to study quasar structure statistically rather than only through individual systems. However, this opportunity also requires analysis tools capable of dealing with heterogeneous data, complex spectral components, lens-galaxy contamination, intrinsic variability, time delays, and large sample sizes.

The methodologies developed in this work represent a step in this direction. By combining realistic microlensing simulations with efficient AGN spectral modeling, this thesis provides a framework for studying the inner structure of quasars in the era of large astronomical surveys. In addition to developing new methodologies and testing new techniques, this work provides the scientific community with open-source software that can be used, extended, and adapted for future applications.

Future extensions of this work should include applying these methods to newly discovered lensed quasars, incorporating multi-epoch spectroscopic monitoring, and jointly modeling microlensing variability and spectral decomposition. Multi-epoch spectroscopy is particularly valuable for multiply imaged quasars because persistent spectral features can provide stable reference signals, while temporal variations can help distinguish intrinsic quasar variability from microlensing effects that evolve independently in each image. These capabilities create a clear synergy between the characterization of quasar structure and the study of microlensing variability.

CHAPTER 5

Future work

The microlensing signatures observed in lensed AGNs and the analysis of their optical spectra are closely connected. Microlensing can differentially magnify the continuum-emitting accretion disk and different spatial and kinematic components of the broad-line region (BLR), including the high-velocity wings of broad emission lines. By contrast, the more extended narrow-line region (NLR) is generally expected to be only weakly affected and can therefore provide a reference emission. In this context, the two main projects developed in this thesis are linked by a common goal: understanding the unresolved inner structure of quasars.

Nevertheless, both lines of research can also be pursued independently, giving this thesis relevance to both microlensing studies and AGN spectral analysis. From the microlensing perspective, the magnification-map analysis could be extended to stellar mass fractions up to $\alpha \sim 0.4$. Because the current result, $\alpha \geq 0.15$, represents a lower limit, expanding the grid would determine whether the likelihood closes at higher values, avoid artificially truncated constraints, and better quantify degeneracies with source size and effective transverse velocity. Variations in α also modify the caustic network and the resulting magnification distributions, affecting the predicted occurrence, amplitude, and duration of HMEs (Schechter & Wambsganss, 2002).

A broader range of $\alpha = \kappa_*/\kappa$ would also help constrain the IMF normalization of lens galaxies. Different IMFs imply different stellar mass-to-light ratios and therefore different stellar surface mass densities at the image positions. Microlensing measure-

ments could thus distinguish between relatively light, Milky-Way-like IMFs and heavier, Salpeter-like IMFs while helping to break the degeneracy between IMF normalization and dark-matter fraction (Oguri et al., 2014; Schechter et al., 2014; Jiménez-Vicente & Mediavilla, 2019). Applied to larger samples, this approach could test whether the IMF varies with galaxy mass, velocity dispersion, radius, or redshift. However, constraining the detailed IMF shape would require additional information about the microlens mass distribution from light-curve timescales and complementary photometric or dynamical observations.

The magnification maps could additionally be used to generate and analyze synthetic light curves for individual systems, enabling a more detailed characterization of HME amplitudes, durations, and detectability in current and future monitoring surveys. This work could therefore serve as the foundation for a larger population-level project using the substantial new samples of lensed systems expected from Euclid, the Nancy Grace Roman Space Telescope, and the Vera C. Rubin Observatory Legacy Survey of Space and Time. Euclid is expected to provide large samples suitable for quasar microlensing analyses (Vernardos & Tsagkatakis, 2019); Roman should deliver a large and complementary sample of strong lenses, including lensed quasars (Weiner et al., 2020; Hamolli et al., 2023); and Rubin LSST is expected to monitor hundreds to thousands of lensed quasars over its ten-year survey (Oguri & Marshall, 2010; Fagin et al., 2025). These samples will enable substantially stronger population-level constraints on α, HME rates, and their dependence on lens and source properties.

The `SHEAP` project can also be developed as an independent line of research. In particular, it is well suited for application to much larger datasets, such as those expected from the Chilean AGN/Galaxy Extragalactic Survey (Bauer et al., 2023) and future facilities such as the Wide-field Spectroscopic Telescope (Bacon et al., 2024). Its GPU-based implementation makes it especially valuable for large-scale spectral analysis, for which both computational efficiency and model flexibility are required. Moreover, the methodology could be extended to integral-field spectroscopy (IFS), enabling spatially resolved studies of AGNs, host-galaxy contamination, outflows, and complex emission-line structures. An additional advantage of GPU acceleration is its ability to make uncertainty and systematic-error analyses computationally feasible. Large numbers of Monte Carlo realizations and alternative modeling configurations could be fitted efficiently, allowing statistical uncertainties to be propagated and the effects of modeling assumptions to be quantified for both individual spectra and complete IFS data cubes.

Finally, combining the two projects opens an especially promising avenue for future research. With a large spectroscopic sample of lensed quasars selected to have high probabilities of undergoing HMEs, `SHEAP` could be used to characterize the associated spectral variations in detail. Such an analysis would reveal how microlensing affects different spectral components, including the accretion-disk continuum, Fe II emission, and the velocity-resolved BLR, while using the NLR as a comparatively stable reference. This would directly connect microlensing variability with the physical structure of quasars, allowing future time-domain and spectroscopic surveys not only to detect microlensing events but also to recover information about the spatial and kinematic structure of otherwise unresolved AGN regions.

Appendix A

Kinematic region limits

Table A.1: SHEAP kinematic region limits.

Region	FWHM limits [km s^{-1}]	$\pm v_{\rm shift}$ [km s^{-1}]	References
Broad	1500–10000	3000	Shen et al. 2011, Rakshit et al. 2017
Narrow	100–500	500	Mullaney et al. 2013, Calderone et al. 2017
Outflow	500–1500	500	Mullaney et al. 2013, Circosta et al. 2018, Harrison et al. 2018
Winds	3000–15000	8000	Richards et al. 2011, Coatman et al. 2016, 2017
Host	100–1000	2896	Bruzual & Charlot 2003, Sánchez-Blázquez et al. 2006, Cappellari & Emsellem 2004
Fe	100–7000	3000	Vestergaard & Wilkes 2001, Kovačević et al. 2010
Balmer	100–11000	500	Jin et al. 2012, Kovačević et al. 2014

Table A.2: Columns list the region name, the allowed FWHM range, the maximum absolute velocity offset, and the references used to define or motivate these limits.

Appendix B

Broad Component Combination Methods

Broad emission lines in Type 1 AGNs are often modeled with multiple components (e.g., multiple Gaussians or others profiles, see Section 3.3.2) to reproduce asymmetric wings, intermediate-width structure, or blended sub-features. While this improves fit flexibility, it introduces the practical question of how to report a single set of broad-line measurements (e.g., velocity shift, line luminosity, and FWHM) for subsequent analysis.

In `SHEAP`, broad-line observables can be measured on reconstructed broad-line profiles obtained under different component-combination schemes. In particular, the FWHM is measured as the full width of a reconstructed broad-line profile at half the maximum of that profile peak. Because multi-component broad lines admit more than one physically motivated definition of the "main" broad profile, `SHEAP` computes broad-line measurements for two complementary strategies: (i) combining all fitted broad components, and (ii) combining only those broad components that are kinematically consistent with the systemic reference frame (as defined by the narrow lines). For the analyses presented in this work, we calculate both sets of measurements and compare them as appropriate to the sample under study; ultimately, the choice depends on the scientific application.

Combining all broad components. This method assumes that the set of broad components included in the model is an adequate description of the physical broad-line region signal. All broad components associated with a given transition are added, and a single set of broad-line observables is measured from the resulting total profile. This approach mirrors common practice in automated quasar-fitting tools, in which multiple broad subcomponents are treated collectively when reporting global broad-line properties (e.g., Guo et al., 2018).

Combining kinematically reasonable components This alternative method does not assume that all fitted broad components are physically associated with the main virialized broad-line emission. Instead, it uses the systemic reference frame defined by narrow emission lines to determine which broad components are plausibly associated with the AGN systemic velocity. Narrow lines are widely used as kinematic anchors in AGN spectral decomposition and quasar catalogs (e.g., Shen et al., 2011). In `SHEAP`, a broad component is included in the "main" broad profile only if its velocity shift relative to the narrow-line systemic velocity falls within a predefined threshold ($|v_{\rm shift}| \leq v_{\rm th}$).

In our implementation we adopt, for example, $v_{\rm th} = 150$ km s^{-1} following the criterion used in (Bernal et al., 2026, 2025). The combined profile is then formed by summing only the selected components, and all reported broad-line measurements (including FWHM, centroid shift, and integrated flux) are derived from that kinematically filtered total profile.

APPENDIX C

Emission lines

Table C.1: Emission lines used by the sample. Counts: Wu&Shen22 (broad=3); Pan+25 (broad=1); Sánchez-Sáez+18 (broad=2, narrow=1, outflow=1); Bernal+25 (broad=2, narrow=1).

Region	Wu&Shen22	Pan+25	Sánchez-Sáez+18	Bernal+25
broad	Lyα λ 1215.670	AlIII λ 1858.753	HeIIb λ 2733.290	Hϵ λ 3970.072
	NV λ 1240.812	SiIII] λ 1892.000	MgII λ 2799.000	Hδ λ 4101.742
	OI+SiII λ 1304.350	CIII] λ 1908.730	HeIa λ 2945.106	HeIc λ 4143.761
	SiIV λ 1393.755	HeIIb λ 2733.290	HeIb λ 3187.745	Hg λ 4340.471
	OIV] λ 1402.770	MgII λ 2799.000	HeIIc λ 3203.100	HeId λ 4471.479
	CIV λ 1549.479	HeIa λ 2945.106	Hϵ λ 3970.072	HeIId λ 4685.710
	HeIIa λ 1640.420	HeIb λ 3187.745	Hδ λ 4101.742	Hβ λ 4861.333
	AlIII λ 1858.753	HeIIc λ 3203.100	HeIc λ 4143.761	HeIe λ 5877.254
	SiIII] λ 1892.000	Hϵ λ 3970.072	Hg λ 4340.471	NaIa λ 5890.000
	CIII] λ 1908.730	Hδ λ 4101.742	HeId λ 4471.479	NaI1a λ 5896.000
	HeIIb λ 2733.290	HeIc λ 4143.761	HeIId λ 4685.710	OIa λ 6046.440
		Hg λ 4340.471	Hβ λ 4861.333	Hα λ 6562.820
				OIb λ 7002.230
				HeIf λ 7065.196
				OIc λ 7254.448
				HeIg λ 7281.349
				OId λ 7774.170
				HeIh λ 7816.136
narrow	–	–	HeIIb λ 2733.290	[NeIII] λ 3868.760
			HeIa λ 2945.106	Hϵ λ 3970.072
			HeIb λ 3187.745	Hδ λ 4101.742
			HeIIc λ 3203.100	HeIc λ 4143.761
			[NeV]a λ 3346.790	Hg λ 4340.471
			[NeV]b λ 3426.850	OIIIa λ 4363.214
			[OII] λ 3728.000	HeId λ 4471.479
			[NeIII] λ 3868.760	HeIId λ 4685.710
			Hϵ λ 3970.072	Hβ λ 4861.333
			Hδ λ 4101.742	OIIIb λ 4958.896
			HeIc λ 4143.761	OIIIc λ 5006.803
			Hg λ 4340.471	HeIe λ 5877.254
			OIIIa λ 4363.214	OIa λ 6300.304
			HeId λ 4471.479	OIb λ 6363.776
			HeIId λ 4685.710	NIIa λ 6548.050
			Hβ λ 4861.333	Hα λ 6562.820
			OIIIb λ 4958.896	NIIb λ 6583.460
			OIIIc λ 5006.803	SIIa λ 6716.440
				SIIb λ 6730.810
				HeIf λ 7065.196
				HeIg λ 7281.349
				HeIh λ 7816.136
outflow	–	–	[NeV]a λ 3346.790	–
			[NeV]b λ 3426.850	
			[OII] λ 3728.000	
			[NeIII] λ 3868.760	
			OIIIb λ 4958.896	
			OIIIc λ 5006.803	

Appendix D

Emission lines fitting examples

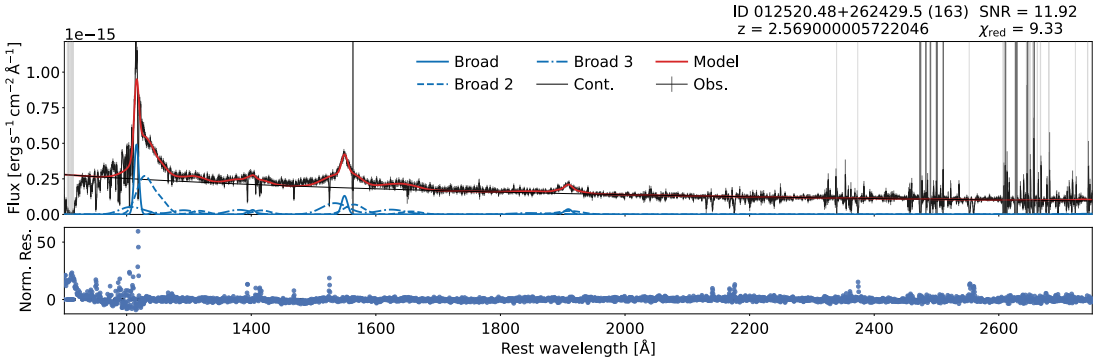


Figure D.1: Example of fited spectra for the Wu&Shen sample for the object with median SN

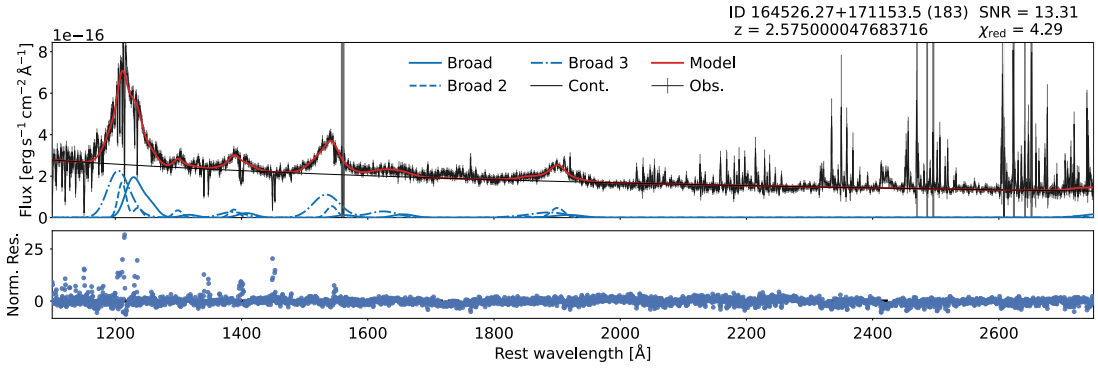


Figure D.2: Example of fited spectra for the Wu&Shen sample for the object with median χ^2_{reduced}

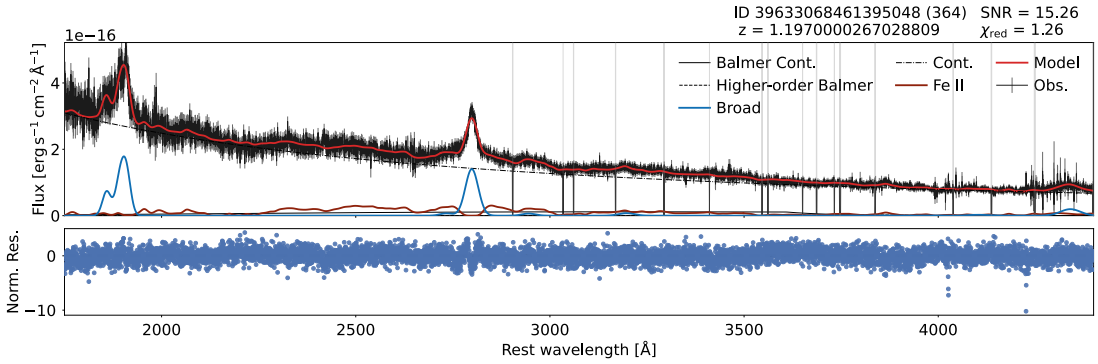


Figure D.3: Example of fited spectra for the Pan+25 sample for the object with median SN

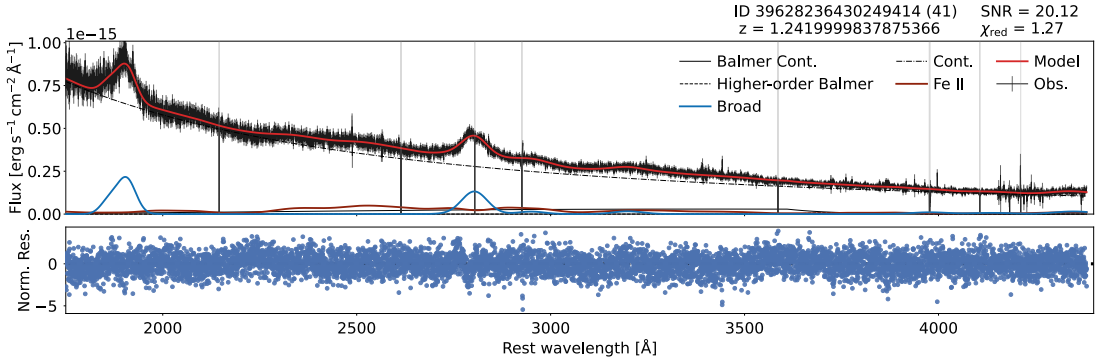


Figure D.4: Example of fited spectra for the Pan+25 sample for the object with median χ^2_{reduced}

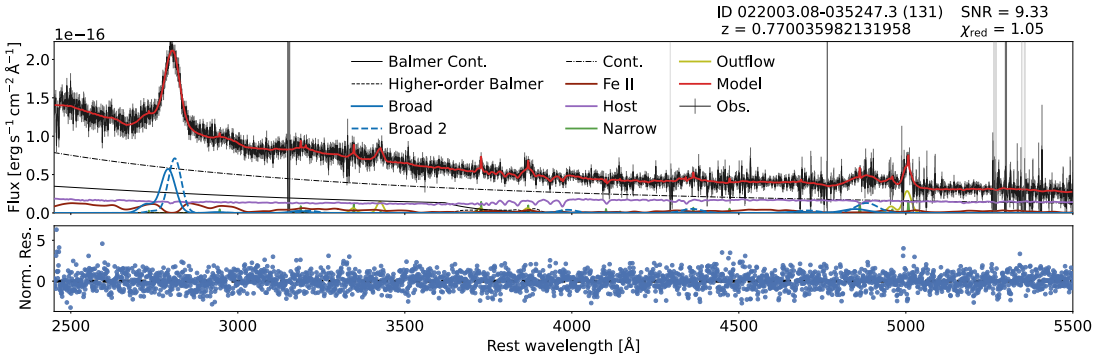


Figure D.5: Example of fited spectra for the Sánchez-Sáez+18 sample for the object with median SN

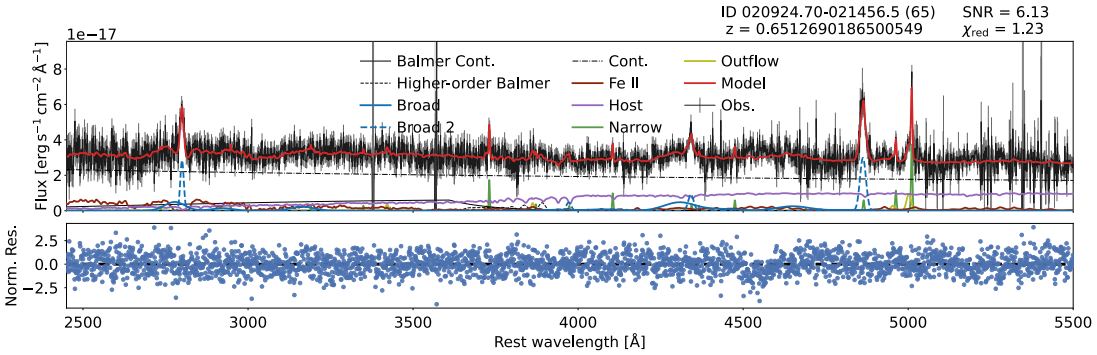


Figure D.6: Example of fited spectra for the Sánchez-Sáez+18 sample for the object with median χ^2_{reduced}

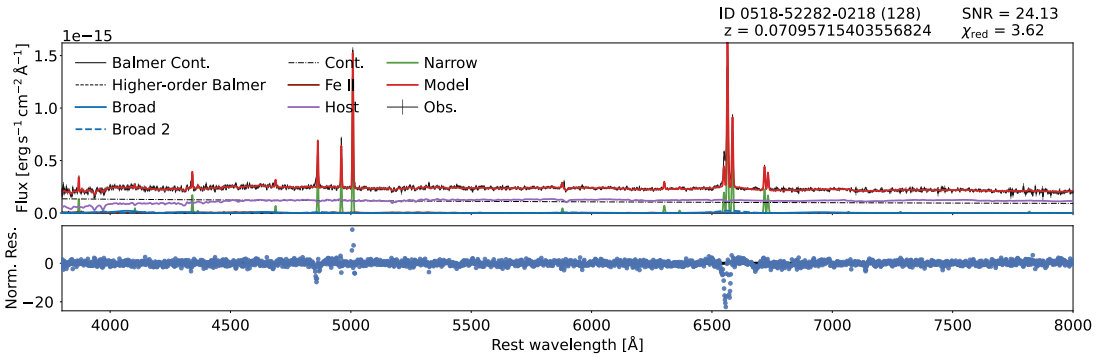


Figure D.7: Example of fited spectra for the Bernal+25 sample for the object with median SN

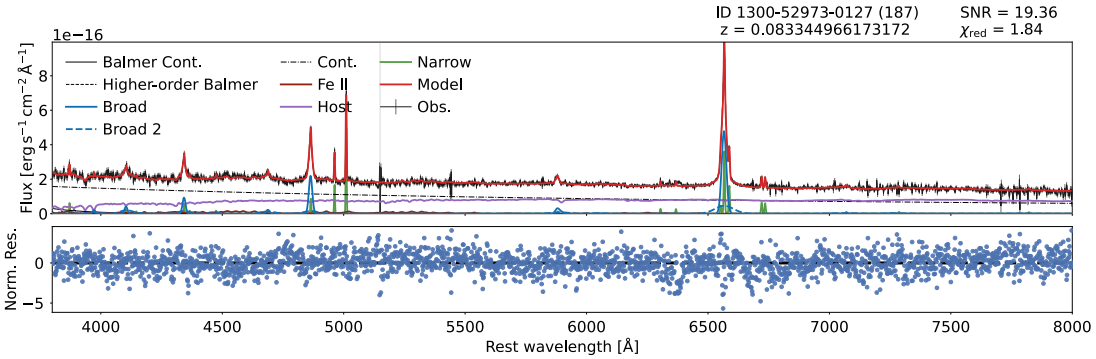


Figure D.8: Example of fited spectra for the Bernal+25 sample for the object with median χ^2_{reduced}

Bibliography

Abajas C., Mediavilla E., Muñoz J. A., Popović L. Č., 2005, Memorie della Societa Astronomica Italiana Supplementi, 7, 48

Abajas C., Mediavilla E., Muñoz J. A., Popović L. Č., Oscoz A., 2002, ApJ, 576, 640

Agnello A., Grillo C., Jones T., Treu T., Bonamigo M., Suyu S. H., 2018, MNRAS, 474, 3391

Amante M. H., Magaña J., Motta V., García-Aspeitia M. A., Verdugo T., 2020, MNRAS, 498, 6013

Andika I. T., Schuldt S., Suyu S. H., Bag S., Cañameras R., Melo A., Grillo C., Chan J. H. H., 2025, A&A, 694, A227

Andrae R., 2010, arXiv e-prints, arXiv:1009.2755

Anguita T., Schmidt R. W., Turner E. L., Wambsganss J., Webster R. L., Loomis K. A., Long D., McMillan R., 2008, A&A, 480, 327

Antonucci R., 1993, ARA&A, 31, 473

Astropy Collaboration et al., 2022, ApJ, 935, 167

Astropy Collaboration et al., 2018, AJ, 156, 123

Astropy Collaboration et al., 2013, A&A, 558, A33

Avila F., Bernui A., Nunes R. C., de Carvalho E., Novaes C. P., 2022, MNRAS, 509, 2994

Awad P., Chan J. H. H., Millon M., Courbin F., Paic E., 2023, A&A, 673, A88

Bacon R. et al., 2024, arXiv e-prints, arXiv:2405.12518

Baldwin J. A., Wampler E. J., Burbidge E. M., 1981, ApJ, 243, 76

Barkana R., 1998, ApJ, 502, 531

Bauer F. E. et al., 2023, The Messenger, 190, 34

Bentz M. C. et al., 2013, ApJ, 767, 149

Bergamini P. et al., 2023, ApJ, 952, 84

Bernal S. et al., 2026, A&A, 707, A206

Bernal S., Sánchez-Sáez P., Arévalo P., Bauer F. E., Lira P., Sotomayor B., 2025, A&A, 694, A127

Bianchi S., Maiolino R., Risaliti G., 2012, Advances in Astronomy, 2012, 782030

Birrer S., Amara A., 2018, Physics of the Dark Universe, 22, 189

Birrer S., Amara A., Refregier A., 2015, ApJ, 813, 102

Birrer S. et al., 2024, Space Sci. Rev., 220, 48

Birrer S. et al., 2020, A&A, 643, A165

Blackburne J. A., Pooley D., Rappaport S., Schechter P. L., 2011, The Astrophysical Journal, 729, 34

Blandford R. D., Kovner I., 1988, Phys. Rev. A, 38, 4028

Blandford R. D., Narayan R., 1992, ARA&A, 30, 311

Blanton M. R., Roweis S., 2007, AJ, 133, 734

Bolton A. S. et al., 2012, AJ, 144, 144

Bonaca A., Conroy C., Wetzel A., Hopkins P. F., Kereš D., 2017, ApJ, 845, 101

Bond I. A. et al., 2004, ApJ, 606, L155

Boroson T. A., Green R. F., 1992, ApJS, 80, 109

Bradbury J. et al., 2021, JAX: Autograd and XLA, Astrophysics Source Code Library, record ascl:2111.002

Brigham E. O., 1974, The fast Fourier Transform

Bruzual G., Charlot S., 2003, MNRAS, 344, 1000

Cackett E. M., Bentz M. C., Kara E., 2021, iScience, 24, 102557

Calderone G., Nicastro L., Ghisellini G., Dotti M., Sbarrato T., Shankar F., Colpi M., 2017, MNRAS, 472, 4051

Cappellari M., 2012, pPXF: Penalized Pixel-Fitting stellar kinematics extraction, Astrophysics Source Code Library, record ascl:1210.002

Cappellari M., 2017, MNRAS, 466, 798

Cappellari M., 2023, MNRAS, 526, 3273

Cappellari M., Emsellem E., 2004, PASP, 116, 138

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

Carnall A. C., 2017, arXiv e-prints, arXiv:1705.05165

Cid Fernandes R., Mateus A., Sodré L., Stasińska G., Gomes J. M., 2005, MNRAS, 358, 363

Circosta C. et al., 2018, A&A, 620, A82

Coatman L., Hewett P. C., Banerji M., Richards G. T., 2016, MNRAS, 461, 647

Coatman L., Hewett P. C., Banerji M., Richards G. T., Hennawi J. F., Prochaska J. X., 2017, MNRAS, 465, 2120

Coe D., Moustakas L. A., 2009, ApJ, 706, 45

Collett T. E. et al., 2023, The Messenger, 190, 49

Congdon A. B., Keeton C. R., 2018, Principles of Gravitational Lensing: Light Deflection as a Probe of Astrophysics and Cosmology

Congdon A. B., Keeton C. R., Osmer S. J., 2007, MNRAS, 376, 263

Cornachione M. A., Morgan C. W., 2020, ApJ, 895, 93

Crenshaw D. M., Kraemer S. B., George I. M., 2003, ARA&A, 41, 117

Dalalyan A. S., 2017, arXiv e-prints, arXiv:1704.04752

de Jong R. S. et al., 2019, The Messenger, 175, 3

DeepMind et al., 2020, The DeepMind JAX Ecosystem

DESI Collaboration et al., 2024, AJ, 168, 58

DESI Collaboration et al., 2016a, arXiv e-prints, arXiv:1611.00036

DESI Collaboration et al., 2016b, arXiv e-prints, arXiv:1611.00036

Di Valentino E. et al., 2025, Physics of the Dark Universe, 49, 101965

Diamond-Stanic A. M. et al., 2009, ApJ, 699, 782

Dietrich M., Appenzeller I., Vestergaard M., Wagner S. J., 2002, ApJ, 564, 581

Dong X., Wang T., Wang J., Yuan W., Zhou H., Dai H., Zhang K., 2008, MNRAS, 383, 581

Dong X.-B., Wang J.-G., Ho L. C., Wang T.-G., Fan X., Wang H., Zhou H., Yuan W., 2011, ApJ, 736, 86

Du P., 2024, PuDu-Astro/DASpec: DASpec: A code for spectral decomposition of active galactic nuclei

Ducourant C. et al., 2026, A&A, 707, A345

Dupourqué S., Barret D., Diez C. M., Guillot S., Quintin E., 2024, A&A, 690, A317

Dux F. et al., 2025, A&A, 694, A300

Einstein A., 1936, Science, 84, 506

Erickson S. et al., 2025, AJ, 170, 44

Ertl S., Schuldt S., Suyu S. H., Schmidt T., Treu T., Birrer S., Shajib A. J., Sluse D., 2023, A&A, 672, A2

Esteban-Gutiérrez A., Agües-Paszkowsky N., Mediavilla E., Jiménez-Vicente J., Muñoz J. A., Heydenreich S., 2022, ApJ, 929, 123

Esteban-Gutiérrez A., Mediavilla E., Jiménez-Vicente J., Muñoz J. A., 2023, ApJ, 954, 172

Euclid Collaboration et al., 2025, arXiv e-prints, arXiv:2503.15326

Event Horizon Telescope Collaboration et al., 2022, ApJ, 930, L12

Event Horizon Telescope Collaboration et al., 2019, ApJ, 875, L1

Fabian A. C., 2012, ARA&A, 50, 455

Fabian A. C., Lohfink A., Kara E., Parker M. L., Vasudevan R., Reynolds C. S., 2015, MNRAS, 451, 4375

Fagin J. et al., 2025, ApJ, 981, 61

Fagin J. et al., 2024, ApJ, 965, 104

Faure C., Anguita T., Eigenbrod A., Kneib J.-P., Chantry V., Alloin D., Morgan N., Covone G., 2009, A&A, 496, 361

Fausnaugh M. M. et al., 2018, ApJ, 854, 107

Ferreras I., Saha P., Leier D., Courbin F., Falco E. E., 2010, MNRAS, 409, L30

Fian C., Guerras E., Mediavilla E., Jiménez-Vicente J., Muñoz J. A., Falco E. E., Motta V., Hanslmeier A., 2018, ApJ, 859, 50

Fian C. et al., 2021, A&A, 654, A70

Fian C., Muñoz J. A., Forés-Toribio R., Mediavilla E., Jiménez-Vicente J., Chelouche D., Kaspi S., Richards G. T., 2024, A&A, 682, A57

Fitzpatrick E. L., 1999, PASP, 111, 63

Foreman-Mackey D. et al., 2013a, emcee: The MCMC Hammer, Astrophysics Source Code Library, record ascl:1303.002

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013b, PASP, 125, 306

Galan A., Vernardos G., Peel A., Courbin F., Starck J.-L., 2022, A&A, 668, A155

García-Burillo S. et al., 2021, A&A, 652, A98

Garcia-Rissmann A., Rodríguez-Ardila A., Sigut T. A. A., Pradhan A. K., 2012, ApJ, 751, 7

Gaudi B. S., 2010, in Seager S., ed, Exoplanets. p. 79

Gavazzi R., Treu T., Koopmans L. V. E., Bolton A. S., Moustakas L. A., Burles S., Marshall P. J., 2008, ApJ, 677, 1046

Gil-Merino R., Lewis G. F., 2005, A&A, 437, L15

Girardi L., Bressan A., Bertelli G., Chiosi C., 2000, A&AS, 141, 371

Goicoechea L. J., Shalyapin V. N., 2025, A&A, 701, L13

Grandi S. A., 1982, ApJ, 255, 25

GRAVITY Collaboration et al., 2024, A&A, 690, A76

Gravity Collaboration et al., 2018, Nature, 563, 657

Greene J. E., Ho L. C., 2005, ApJ, 630, 122

Grillo C., Lombardi M., Bertin G., 2008, A&A, 477, 397

Guerras E., Mediavilla E., Jimenez-Vicente J., Kochanek C. S., Muñoz J. A., Falco E., Motta V., 2013a, ApJ, 764, 160

Guerras E., Mediavilla E., Jimenez-Vicente J., Kochanek C. S., Muñoz J. A., Falco E., Motta V., Rojas K., 2013b, ApJ, 778, 123

Guo H., Shen Y., Wang S., 2018, PyQSOFit: Python code to fit the spectrum of quasars, Astrophysics Source Code Library, record ascl:1809.008

Hagen H. J., Reimers D., 2000, A&A, 357, L29

Hainline L. J. et al., 2013, ApJ, 774, 69

Hamolli L., Hafizi M., De Paolis F., Guliqani E., 2023, Galaxies, 11, 71

Harrison C. M., Alexander D. M., Mullaney J. R., Swinbank A. M., 2014, MNRAS, 441, 3306

Harrison C. M., Costa T., Tadhunter C. N., Flütsch A., Kakkad D., Perna M., Vietri G., 2018, Nature Astronomy, 2, 198

Heß S., Kitaura F.-S., 2016, MNRAS, 456, 4247

Heydenreich S., Mediavilla E., Jiménez-Vicente J., Vives-Arias H., Muñoz J. A., 2024, A&A, 690, A307

Hickox R. C., Alexander D. M., 2018, ARA&A, 56, 625

Ho L. C., 2009, ApJ, 699, 626

Hogg D. W., 1999, arXiv e-prints, astro

Hogg D. W., Baldry I. K., Blanton M. R., Eisenstein D. J., 2002, arXiv e-prints, astro

Holtzman J. A., Watson A. M., Baum W. A., Grillmair C. J., Groth E. J., Light R. M., Lynds R., O'Neil J., Earl J., 1998, AJ, 115, 1946

Hu J.-P., Wang F.-Y., 2023, Universe, 9, 94

Hutsemékers D., Sluse D., Savić ., 2024, A&A, 691, A292

Ilić D. et al., 2020, A&A, 638, A13

Ilić D., Rakić N., Popović L. Č., 2023, ApJS, 267, 19

Ivezić Ž. et al., 2019, ApJ, 873, 111

Jalan P., Rakshit S., Woo J.-J., Kotilainen J., Stalin C. S., 2023, MNRAS, 521, L11

Jha V. K., Joshi R., Chand H., Wu X.-B., Ho L. C., Rastogi S., Ma Q., 2022, MNRAS, 511, 3005

Jiménez-Vicente J., Mediavilla E., 2019, ApJ, 885, 75

Jiménez-Vicente J., Mediavilla E., 2022, ApJ, 941, 80

Jiménez-Vicente J., Mediavilla E., 2025, MNRAS, 541, 1264

Jiménez-Vicente J., Mediavilla E., Kochanek C. S., Muñoz J. A., 2015, ApJ, 799, 149

Jiménez-Vicente J., Mediavilla E., Kochanek C. S., Muñoz J. A., Motta V., Falco E., Mosquera A. M., 2014, ApJ, 783, 47

Jiménez-Vicente J., Mediavilla E., Muñoz J. A., Kochanek C. S., 2012a, ApJ, 751, 106

Jiménez-Vicente J., Mediavilla E., Muñoz J. A., Kochanek C. S., 2012b, ApJ, 751, 106

Jin C., Ward M., Done C., Gelbord J., 2012, MNRAS, 420, 1825

Kakkad D. et al., 2018, A&A, 618, A6

Karp J. S. M. et al., 2026, ApJ, 999, 198

Kaspi V. M., Ransom S. M., Backer D. C., Ramachandran R., Demorest P., Arons J., Spitkovsky A., 2004, ApJ, 613, L137

Kassiola A., Kovner I., 1993, ApJ, 417, 450

Kauffmann G. et al., 2003, MNRAS, 346, 1055

Kayser R., Refsdal S., Stabell R., 1986, A&A, 166, 36

Keeton C. R., 2001a, arXiv e-prints, astro

Keeton C. R., 2001b, arXiv e-prints, astro

Keeton C. R., 2011, GRAVLENS: Computational Methods for Gravitational Lensing, Astrophysics Source Code Library, record ascl:1102.003

Kellermann K. I., 2015, in IAU Symposium, Vol. 313, Massaro F., Cheung C. C., Lopez E., Siemiginowska A., eds, Extragalactic Jets from Every Angle, p. 190

Kewley L. J., Groves B., Kauffmann G., Heckman T., 2006, MNRAS, 372, 961

Khachikian E. Y., Weedman D. W., 1974, ApJ, 192, 581

Kim M., Ho L. C., Im M., 2006, ApJ, 642, 702

Kingma D. P., Ba J., 2014, arXiv e-prints, arXiv:1412.6980

Knabel S., Treu T., Cappellari M., Shajib A. J., Chen C.-F., Birrer S., Bennert V. N., 2025, ApJ, 990, 51

Kochanek C. S., 2004, ApJ, 605, 58

Kollmeier J. A. et al., 2006, ApJ, 648, 128

Koopmans L. V. E., Treu T., Bolton A. S., Burles S., Moustakas L. A., 2006, ApJ, 649, 599

Kormann R., Schneider P., Bartelmann M., 1994, A&A, 284, 285

Kovačević J., Popović L. Č., Dimitrijević M. S., 2010, ApJS, 189, 15

Kovačević J., Popović L. Č., Kollatschny W., 2014, Advances in Space Research, 54, 1347

Krone-Martins A. et al., 2018, A&A, 616, L11

Landt H., Bentz M. C., Ward M. J., Elvis M., Peterson B. M., Korista K. T., Karovska M., 2008, ApJS, 174, 282

Lee D.-W., Kim S.-J., 2014, MNRAS, 443, 328

Lefor A. T., Futamase T., Akhlaghi M., 2013, , 57, 1

Leier D., Ferreras I., Saha P., Charlot S., Bruzual G., La Barbera F., 2016, MNRAS, 459, 3677

Leier D., Ferreras I., Saha P., Falco E. E., 2011, ApJ, 740, 97

Lemon C. et al., 2023, MNRAS, 520, 3305

Lemon C. et al., 2024, Space Sci. Rev., 220, 23

Lemon C. A., Auger M. W., McMahon R. G., Koposov S. E., 2017, MNRAS, 472, 5023

Liaudat T. I., Starck J.-L., Kilbinger M., 2023, Frontiers in Astronomy and Space Sciences, 10, 1158213

Liebes S., 1964, Physical Review, 133, 835

Lucey J. R., Schechter P. L., Smith R. J., Anguita T., 2018, MNRAS, 476, 927

Lyke B. W. et al., 2020, ApJS, 250, 8

Ma L., Liu Y., Zhang X., Ye Y., Yin G., Johnson B. A., 2019, ISPRS Journal of Photogrammetry and Remote Sensing, 152, 166

Mahler G. et al., 2023, ApJ, 945, 49

Mainieri V. et al., 2024, arXiv e-prints, arXiv:2403.05398

Mao S., 2012, Research in Astronomy and Astrophysics, 12, 947

Markwardt C. B., 2009, in Astronomical Society of the Pacific Conference Series, Vol. 411, Bohlender D. A., Durand D., Dowler P., eds, Astronomical Data Analysis Software and Systems XVIII, p. 251

Massey R., Kitching T., Richard J., 2010, Reports on Progress in Physics, 73, 086901

Masters K., 2025, arXiv e-prints, arXiv:2502.09610

Mediavilla E. et al., 2026, ApJ, 997, 274

Mediavilla E., Jiménez-Vicente J., Motta V., 2024, AJ, 167, 171

Mediavilla E., Jiménez-vicente J., Muñoz J. A., Mediavilla T., 2015, ApJ, 814, L26

Mediavilla E., Jiménez-Vicente J., Muñoz J. A., Vives-Arias H., Calderón-Infante J., 2017, ApJ, 836, L18

Mediavilla E., Mediavilla T., Muñoz J. A., Ariza O., Lopez P., Gonzalez-Morcillo C., Jimenez-Vicente J., 2011, ApJ, 741, 42

Mediavilla E. et al., 2009, ApJ, 706, 1451

Mediavilla E., Muñoz J. A., Garzón F., Mahoney T. J., 2016, Astrophysical Applications of Gravitational Lensing

Mediavilla E., Muñoz J. A., Lopez P., Mediavilla T., Abajas C., Gonzalez-Morcillo C., Gil-Merino R., 2006, ApJ, 653, 942

Mejía-Restrepo J. E., Trakhtenbrot B., Lira P., Netzer H., Capellupo D. M., 2016, MNRAS, 460, 187

Merloni A. et al., 2019, The Messenger, 175, 42

Meylan G., Jetzer P., North P., Schneider P., Kochanek C. S., Wambsganss J., eds, 2006, Gravitational Lensing: Strong, Weak and Micro

Mickaelian A. M., Abrahamyan H. V., Paronyan G. M., Mikayelyan G. A., 2021, Frontiers in Astronomy and Space Sciences, 7, 82

Millon M. et al., 2020, A&A, 639, A101

Misner C. W., Thorne K. S., Wheeler J. A., 1973, Gravitation

Moré J. J., 1978, in Lecture Notes in Mathematics, Berlin Springer Verlag, Vol. 630. p. 105

Morgan C. W., Hyer G. E., Bonvin V., Mosquera A. M., Cornachione M., Courbin F., Kochanek C. S., Falco E. E., 2018, ApJ, 869, 106

Morgan C. W., Kochanek C. S., Morgan N. D., Falco E. E., 2010, ApJ, 712, 1129

Mortonson M. J., Schechter P. L., Wambsganss J., 2005, ApJ, 628, 594

Mosquera A. M., Kochanek C. S., 2011, ApJ, 738, 96

Motta V., Mediavilla E., Falco E., Muñoz J. A., 2012, ApJ, 755, 82

Motta V. et al., 2002, ApJ, 574, 719

Motta V., Mediavilla E., Rojas K., Falco E. E., Jiménez-Vicente J., Muñoz J. A., 2017, ApJ, 835, 132

Moustakas J., Buhler J., Scholte D., Dey B., Khederlarian A., 2023, FastSpecFit: Fast spectral synthesis and emission-line fitting of DESI spectra, Astrophysics Source Code Library, record ascl:2308.005

Muñoz J. A., Falco E. E., Kochanek C. S., Lehár J., McLeod B. A., Impey C. D., Rix H. W., Peng C. Y., 1998, Ap&SS, 263, 51

Muñoz J. A., Mediavilla E., Kochanek C. S., Falco E. E., Mosquera A. M., 2011, ApJ, 742, 67

Mullaney J. R., Alexander D. M., Fine S., Goulding A. D., Harrison C. M., Hickox R. C., 2013, MNRAS, 433, 622

Narayan R., Bartelmann M., 1996, arXiv e-prints, astro

Neira F., Anguita T., Vernardos G., 2020, MNRAS, 495, 544

Neira F., Anguita T., Vernardos G., 2025, A&A, 701, A35

Netzer H., 2015, ARA&A, 53, 365

Newville M., Stensitzki T., Allen D. B., Ingargiola A., 2014, LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python

Nightingale J. et al., 2021, The Journal of Open Source Software, 6, 2825

Oguri M., Marshall P. J., 2010, MNRAS, 405, 2579

Oguri M., Rusu C. E., Falco E. E., 2014, MNRAS, 439, 2494

O'Riordan C. M. et al., 2025, A&A, 694, A145

Paczynski B., 1986, ApJ, 304, 1

Padovani P., 2017, Frontiers in Astronomy and Space Sciences, 4, 35

Paic E., Vernardos G., Sluse D., Millon M., Courbin F., Chan J. H., Bonvin V., 2022, A&A, 659, A21

Pan Z. et al., 2025, ApJ, 987, 48

Park D., Barth A. J., Ho L. C., Laor A., 2022, ApJS, 258, 38

Pierel J. D. R. et al., 2026, ApJ, 998, 219

Pindor B. et al., 2004, AJ, 127, 1318

Planck Collaboration et al., 2020, A&A, 641, A1

Plotkin R. M. et al., 2015, ApJ, 805, 123

Poindexter S., Kochanek C. S., 2010, ApJ, 712, 658

Pooley D., Blackburne J. A., Rappaport S., Schechter P. L., 2007, ApJ, 661, 19

Popović L. Č., 2020, Open Astronomy, 29, 1

Popović L. Č., Kovačević-Dojčinović J., Marčeta-Mandić S., 2019, MNRAS, 484, 3180

Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, Numerical recipes in FORTRAN. The art of scientific computing

Rakić N., 2022, MNRAS, 516, 1624

Rakshit S., Stalin C. S., Chand H., Zhang X.-G., 2017, ApJS, 229, 39

Ramos Almeida C. et al., 2011, ApJ, 731, 92

Redlich M., Bartelmann M., Waizmann J. C., Fedeli C., 2012, A&A, 547, A66

Refsdal S., 1964a, MNRAS, 128, 307

Refsdal S., 1964b, MNRAS, 128, 295

Reis R. C., Miller J. M., 2013, ApJ, 769, L7

Rektsini N. E., Batista V., 2024, arXiv e-prints, arXiv:2407.06689

Ren W., Guo H., Shen Y., Silverman J. D., Burke C. J., Wang S., Wang J., 2024, ApJ, 974, 153

Reynolds C. S., 2021, ARA&A, 59, 117

Ricci C., 2026, arXiv e-prints, arXiv:2606.12511

Ricci C. et al., 2018, MNRAS, 480, 1819

Richards G. T. et al., 2002, AJ, 123, 2945

Richards G. T. et al., 2011, AJ, 141, 167

Richards G. T. et al., 2006, ApJS, 166, 470

Robertson A., Smith G. P., Massey R., Eke V., Jauzac M., Bianconi M., Ryczanowski D., 2020, MNRAS, 495, 3727

Rodrigo C., Solano E., 2020, in XIV.0 Scientific Meeting (virtual) of the Spanish Astronomical Society, p. 182

Rojas K., Motta V., Mediavilla E., Falco E., Jiménez-Vicente J., Muñoz J. A., 2014, ApJ, 797, 61

Rojas K., Motta V., Mediavilla E., Jiménez-Vicente J., Falco E., Fian C., 2020, ApJ, 890, 3

Ross N. et al., 2012, in American Astronomical Society Meeting Abstracts, Vol. 219, American Astronomical Society Meeting Abstracts #219, p. 209.06

Rudy R. J., Mazuk S., Puetter R. C., Hamann F., 2000, ApJ, 539, 166

Saha P., Williams L. L. R., 2004, AJ, 127, 2604

Saleh R. A., Ehsanes Saleh A. K. M., 2022, arXiv e-prints, arXiv:2208.04564

Sánchez-Blázquez P. et al., 2006, MNRAS, 371, 703

Sánchez-Sáez P., Lira P., Mejía-Restrepo J., Ho L. C., Arévalo P., Kim M., Cartier R., Coppi P., 2018, ApJ, 864, 87

Schechter P. L., Pooley D., Blackburne J. A., Wambsganss J., 2014, ApJ, 793, 96

Schechter P. L., Wambsganss J., 2002, ApJ, 580, 685

Schlafly E. F., Finkbeiner D. P., 2011, ApJ, 737, 103

Schlegel D. J., Finkbeiner D. P., Davis M., 1998, ApJ, 500, 525

Schmidt T. et al., 2023, MNRAS, 518, 1260

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

Schneider P., Ehlers J., Falco E. E., 1992, Gravitational Lenses

Sexton R. O., Canalizo G., Hiner K. D., Komossa S., Woo J.-H., Treister E., Hiner Dimassimo S. L., 2019, ApJ, 878, 101

Sexton R. O., Matzko W., Darden N., Canalizo G., Gorjian V., 2021, MNRAS, 500, 2871

Shajib A. J. et al., 2024, Space Sci. Rev., 220, 87

Shajib A. J. et al., 2022, A&A, 667, A123

Shakura N. I., Sunyaev R. A., 1973, A&A, 24, 337

Shemmer O. et al., 2010, ApJ, 722, L152

Shen Y. et al., 2013, ApJ, 778, 98

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

Siudek M. et al., 2024, A&A, 691, A308

Sluse D., Hutsemékers D., Courbin F., Meylan G., Wambsganss J., 2012, A&A, 544, A62

Smith M. C. et al., 2009, MNRAS, 399, 1223

Storey P. J., Hummer D. G., 1995, MNRAS, 272, 41

Straumann N., 2013, General Relativity

Sulentic J. W., Marziani P., Dultzin-Hacyan D., 2000, ARA&A, 38, 521

Suyu S. H. et al., 2013, ApJ, 766, 70

Suyu S. H., Marshall P. J., Auger M. W., Hilbert S., Blandford R. D., Koopmans L. V. E., Fassnacht C. D., Treu T., 2010, ApJ, 711, 201

Tdcosmo Collaboration et al., 2025, A&A, 704, A63

Tessore N., Bellagamba F., Metcalf R. B., 2015, Lensed: Forward parametric modelling of strong lenses, Astrophysics Source Code Library, record ascl:1505.026

The Dark Energy Survey Collaboration , 2005, arXiv e-prints, astro

Thorne J., Robotham A., Davies L., Bellstedt S., 2022, AGN Unification Diagram

Treu T., 2010, ARA&A, 48, 87

Treu T. et al., 2018, MNRAS, 481, 1041

Treu T., Gavazzi R., Gorecki A., Marshall P. J., Koopmans L. V. E., Bolton A. S., Moustakas L. A., Burles S., 2009, ApJ, 690, 670

Treu T., Koopmans L. V. E., 2004, ApJ, 611, 739

Turner E. L., Ostriker J. P., Gott J. R., III, 1984, ApJ, 284, 1

Urry C. M., Padovani P., 1995, PASP, 107, 803

van der Marel R. P., Franx M., 1993, ApJ, 407, 525

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

Vazdekis A., Koleva M., Ricciardelli E., Röck B., Falcón-Barroso J., 2016, MNRAS, 463, 3409

Vazdekis A., Sánchez-Blázquez P., Falcón-Barroso J., Cenarro A. J., Beasley M. A., Cardiel N., Gorgas J., Peletier R., 2010, in IAU Symposium, Vol. 262, Bruzual G. R., Charlot S., eds, Stellar Populations - Planning for the Next Decade, p. 65

Vegetti S. et al., 2024, Space Sci. Rev., 220, 58

Vegetti S., Koopmans L. V. E., 2009, MNRAS, 392, 945

Vegetti S., Lagattuta D. J., McKean J. P., Auger M. W., Fassnacht C. D., Koopmans L. V. E., 2012, Nature, 481, 341

Vernardos G., Fluke C. J., 2013, MNRAS, 434, 832

Vernardos G., Fluke C. J., Bate N. F., Croton D., 2014, ApJS, 211, 16

Vernardos G., Fluke C. J., Bate N. F., Croton D., Vohl D., 2015, ApJS, 217, 23

Vernardos G. et al., 2024, Space Sci. Rev., 220, 14

Vernardos G., Tsagkatakis G., 2019, MNRAS, 486, 1944

Véron-Cetty M.-P., Joly M., Véron P., 2004, A&A, 417, 515

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

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

Vietri G. et al., 2018, A&A, 617, A81

Vives-Arias H., Muñoz J. A., Kochanek C. S., Mediavilla E., Jiménez-Vicente J., 2016, ApJ, 831, 43

Walsh D., Carswell R. F., Weymann R. J., 1979, Nature, 279, 381

Wambsganss J., 2006, arXiv e-prints, astro

Weiner B. et al., 2009, in astro2010: The Astronomy and Astrophysics Decadal Survey, Vol. 2010, p. P61

Weiner C., Serjeant S., Sedgwick C., 2020, Research Notes of the American Astronomical Society, 4, 190

Weisenbach L., 2025, MNRAS, 541, 281

Weisenbach L., Anguita T., Miralda-Escudé J., Oguri M., Saha P., Schechter P. L., 2024, Space Sci. Rev., 220, 57

Weisenbach L., Schechter P. L., Pontula S., 2021, ApJ, 922, 70

Wildy C., Czerny B., 2017, Frontiers in Astronomy and Space Sciences, 4, 43

Will C. M., 2015, Classical and Quantum Gravity, 32, 124001

Winkler H., 1992, MNRAS, 257, 677

Wong K. C. et al., 2020, MNRAS, 498, 1420

Wu Q., Shen Y., 2022, ApJS, 263, 42

Yuan F., Narayan R., 2014, ARA&A, 52, 529

Yue M., Fan X., Yang J., Wang F., 2022, AJ, 163, 139

Zwicky F., 1937, Physical Review, 51, 679