Astrophysical probes of dark matter: opportunities and challenges from a strong gravitational lensing perspective

Wolfgang Enzi

München 2021

Astrophysical probes of dark matter: opportunities and challenges from a strong gravitational lensing perspective

Wolfgang Enzi

Dissertation
an der Fakultät für Physik
der Ludwig–Maximilians–Universität
München

vorgelegt von
Wolfgang Enzi
aus Viechtach, Deutschland

München, den 26. Februar 2021

Erstgutachter: PD Dr. Torsten Enßlin
Zweitgutachter: Prof. Dr. Jochen Weller
Tag der mündlichen Prüfung: 23. April 2021

Contents

List of Figures

List of Tables

List of Abbreviations

ABC	Approximate Bayesian Computation
BOSS	Baryon Oscillation Spectroscopic Survey
CDM	Cold Dark Matter
CMB	Cosmic Microwave Background
DES	Dark Energy Survey
ELT	Extremely Large Telescope
EELT	European Extremely Large Telescope
ETG	Early Type Galaxy
EPS	Extended Press-Schechter
FDM	Fuzzy Dark Matter
FWHM	Full Width at Half Maximum
GMT	Giant Magellan Telescope
HST	Hubble Space Telescope
NIR	Near-Infrared
IGM	Intergalactic Medium
JWST	James Webb Space Telescope
LMC	Large Magellanic Cloud
LSST	Large Synoptic Survey Telescope
MCMC	Markov Chain Monte Carlo
MW	Milky Way
NFW	Navarro-Frenk-White
PSF	Point Spread Function
SDSS	Sloan Digital Sky Survey
SNR	Signal-to-Noise Ratio
SPH	Smoothed Particle Hydrodynamics
SUSY	Supersymmetry
TMT	Thirty Meter Telescope
UV	Ultraviolet
VLBI	Very Long Baseline Interferometry
WDM	Warm Dark Matter
WIMP	Weakly Interacting Massive Particle

Conforming with the writing convention in our field, I chose the pronoun “we” throughout this thesis.

Zusammenfassung

Die Natur der Dunklen Materie ist eines der großen Mysterien der Astrophysik und Kosmologie. Obwohl es zahlreiche Belege für ihre Existenz gibt, die von den Rotationskurven der Galaxien bis zum Leistungsspektrum des kosmischen Mikrowellenhintergrunds reichen, existieren noch viele offene Fragen zu ihren Eigenschaften. Im Standardmodell der Kosmologie ist dunkle Materie kalt (CDM, für engl. cold dark matter), d.h. die Teilchen, aus denen sie besteht, zeigen im frühen Universum vernachlässigbare thermische Geschwindigkeiten. Dieses Modell sagt erfolgreich die Strukturbildung auf Skalen voraus, die größer als $\sim$ 1 Mpc sind. Es ist jedoch noch unklar, ob alternative Beschreibungen der Dunklen Materie erforderlich sind, um diesen Prozess auf kleineren, galaktischen und subgalaktischen Skalen zu beschreiben. Auf diesen Skalen treffen Modelle mit selbstwechselwirkender dunkler Materie (SIDM, für engl. self-interacting dark matter), unscharfer dunkler Materie (FDM, für engl. fuzzy dark matter) oder warmer dunkler Materie (WDM, für engl. warm dark matter) unterschiedliche Vorhersagen über die Menge und die Struktur von Halos aus dunkler Materie. Bis heute ist jedoch keines dieser Modelle ausgeschlossen. Für diese herausfordernde Aufgabe wurden mehrere astrophysikalische Beobachtungen vorgeschlagen.

Das erste Ziel dieser Arbeit ist die präzedenzlose Kombination der Analyse starker Gravitationslinsensysteme mit der des Lyman-α Waldes und der von Satelliten der Milchstraße. Wir bestimmen neue astrophysikalische Schranken für alternative Modelle der dunklen Materie, wie z.B. WDM und SIDM. Unsere Grenzwerte schließen das Dunkle Materie Modell der sterilen Neutrinos mit einer Masse von 7.1 keV für Lepton-Asymmetrien größer als $L_6 > 10$ sowohl auf dem 2-σ Level als auch mit einem Bayes-Faktor von 1/20 aus. Das ETHOS-4-Modell der selbstwechselwirkenden dunklen Materie ist bei einem Konfidenzniveau von 95 Prozent ebenfalls ausgeschlossen. Für jede astrophysikalische Untersuchungmethode diskutieren wir die systematischen Unsicherheiten, Einschränkungen und Verbesserungsmöglichkeiten.

Das zweite Ziel dieser Arbeit fokussiert sich auf Beobachtungen von starken Gravitationslinsen. Zuerst quantifizieren wir Quellen systematischer Fehler im Zusammenhang mit dem Linsenmodellierungsprozess. Traditionell gehen die meisten Modellierungstechniken für Gravitationslinsen davon aus, dass vereinfachte parametrische Formen die Massenverteilung der Linsen gut beschreiben. Hier untersuchen wir die Auswirkung komplexerer Massenverteilungen auf die Eigenschaften der rekonstruierten Hintergrundgalaxien, die Inferenz der Halo- und Subhalo-Massenfunktionen, und die Hubble-Konstante. Zu diesem Zweck analysieren wir eine Stichprobe von vierundzwanzig SLACS-ähnlichen Gravitationslinsen-

systemen mit Hintergrund- und Deflektorgalaxien, welche der Illustris-1 Simulation entnommen sind. Wir stellen fest, dass die Präsenz von Kernen in den Massendichteprofilen der simulierten Galaxien zu einer Unterschätzung der Quellengrößen um $\sim$ 50 Prozent und zu einer systematischen Überschatzung der abgeleiteten Hubble-Konstante von $\sim$ 25 Prozent führt. Unsere Arbeit unterstreicht die Notwendigkeit zusätzlicher Beobachtungen, um bestehende Entartungen im Rekonstruktionsprozess zu beseitigen. Wir finden keine Entartung zwischen der Komplexität des Linsenpotentials und der abgeleiteten Menge an Subhalos.

Zweitens untersuchen wir mehrere Beobachtungsstrategien, um zu bestimmen, welche davon in naher Zukunft die Natur der Dunklen Materie stärker einschränken können. Zu diesem Zweck analysieren wir reale und synthetische Beobachtungen des starken Gravitationslinseneffekts mit dem Atacama Large Millimeter/Submillimeter Array (ALMA) und dem James Webb Space Telescope (JWST). Wir bestimmen die Empfindlichkeitsfunktion, d.h. die niedrigsten Masse, welche Dunkle Materie Strukturen besitzen müssen, um nachgewiesen werden zu können. Aus dieser Berechnung ergibt sich, dass mehr Gravitationslinsensysteme erforderlich wären, um das CDM-Modell mit dem JWST zu testen, als derzeit bekannt sind. Stattdessen könnten ALMA-Beobachtungen des thermischen Kontinuums von Staub mit einer Winkelauflösung von $180 \times 220\text{mas}^2$ für 49 Gravitationslinsensysteme möglicherweise ausreichend sein, trotz nicht zu vernachlässigenden Beobachtungskosten von 69 Stunden. Langfristig erwarten wir, dass die große Anzahl starker Gravitationssysteme, welche beispielsweise von der Euclid-Mission in Verbindung mit hochauflösenden Follow-up-Beobachtungen mit dem EELT (engl. für European Extremely Large Telescope) entdeckt werden sollen, es ermöglichen wird, starke Schranken für die Natur der Dunklen Materie zu liefern.

Abstract

The nature of dark matter is one of the great mysteries of astrophysics and cosmology. While there exists a large amount of evidence for its existence, ranging from the rotation curves of galaxies to the power spectrum of the cosmic microwave background, there are still many open questions about its properties. In the standard ΛCDM model of cosmology, dark matter is cold (CDM), i.e. its particles show negligible thermal velocities in the early Universe. This model successfully predicts the formation of structure on scales larger than $\sim$ 1Mpc. However, it is still unclear whether alternative dark matter descriptions are necessary to describe this process on smaller, galactic and subgalactic scales. At these scales, self-interacting dark matter (SIDM), fuzzy dark matter (FDM), or warm dark matter (WDM) models make different predictions on the amount and the structure of dark matter haloes. To this date, however, none of these models has been ruled out. Several astrophysical observations have been proposed for this challenging task.

The first goal of this thesis is the unprecedented combination of the analysis of strong gravitational lens systems with those of the Lyman-α forest and the luminous satellites of the Milky Way, to derive novel astrophysical constraints on alternative dark matter models, such as thermal relic WDM and SIDM. Our derived limits rule out the 7.1 keV sterile neutrino dark matter model both at the 2-σ level and with a Bayes factor of 1/20 for Lepton asymmetries larger than $L_6 > 10$. The ETHOS-4 model of self-interacting dark matter is also ruled-out at the 95 per cent confidence level. For each probe, we also discuss its systematic biases, limitations and prospects for improvement.

The second goal of this thesis focuses further on strong gravitational lensing. First, we quantify sources of systematic errors related to the lens modelling process. Traditionally, most gravitational lens modelling techniques assume that simplified parametric forms well describe the lensing mass distribution. Here, we investigate the effect of more complex mass distributions on the properties of the reconstructed background source galaxies, the inference of the halo and subhalo mass functions, and the Hubble constant. To this end, we analyse a sample of twenty-four SLACS-like galaxy-galaxy strong gravitational lens systems with a background source and deflectors taken from the Illustris-1 simulation. We find that the presence of cored mass-density profiles in the simulated galaxies results in an underestimation of the source sizes by $\sim$ 50 per cent and a systematic bias on the inference of the Hubble constant of $\sim$ 25 per cent. Our work highlights the need for additional observational constraints to break existing degeneracies. We find no degeneracy between the complexity of the lensing potential and the inferred amount of subhaloes.

Second, we examine multiple observing strategies to investigate which one is more likely to provide improved constraints on the nature of dark matter in the short term future. To this end, we analyse real and synthetic strong gravitational lensing observations with the Atacama Large Millimeter/Submillimeter Array (ALMA) and James Webb Space Telescope (JWST). From the calculation of the sensitivity function, that is, the lowest detectable substructure and field-halo mass, we find that more gravitational lens systems than currently known would be required to test the CDM model with the JWST. Instead, ALMA observations of the thermal continuum with an angular resolution of 180×220 mas^2 could for 49 gravitational lens systems potentially be sufficient at the non-negligible observational cost of 69 hrs. In the long-term future, we expect that the large number of strong gravitational systems that are expected to be discovered by, for example, the Euclid mission coupled with high-angular-resolution follow-up observations with the European Extremely Large Telescope, will allow us to provide stringent constraints on the nature of dark matter.

Chapter 1

Introduction

In this thesis, we derive joint constraints on thermal relic dark matter models by combining the observations of strong gravitational lens systems, the Lyman-α forest, and the Milky Way satellites. We then discuss important systematic errors in these analyses and outline prospects for their future developments. In particular, we infer constraints on the half-mode length scale, which determines the scale at which structure is suppressed in thermal relic dark matter models. This scale can then be related to the mass of the thermal relic dark matter particle while being relatively independent of cosmological assumptions.

Focusing deeper on strong gravitational lensing, we then study potential systematic errors and future possibilities in more detail: first, we examine whether un-modelled complexity in the lensing potential beyond a regular elliptical mass profile can lead to systematic biases in the reconstruction of lenses. Using realistic lens galaxies from the state-of-the-art Illustris simulation, we find that the background source galaxy reconstruction is affected by the presence of a core in the mass density profile of the lens. In turn, this leads to a bias in the inference of the Hubble constant. Conversely, constraints on the amount of subhaloes are robust against the presence of complexity in the lens for the current data quality. Second, we quantify whether three dimensional (spatial plus spectral) observations with the Atacama Large Millimeter Array (ALMA) and the James Webb Space Telescope (JWST) can improve current constraints on the properties of dark matter.

This chapter provides the context of and the background information on the results and methods used in this thesis. In Section 1.1 we present a summary on dark matter, evidence for its existence, candidates of interest, astrophysical probes of it, and its connection to galaxy formation. We then provide a brief introduction to strong gravitational lensing in Section 1.2. As most of the analyses considered here are based and Bayesian statistics, we summarise its main concepts in Section 1.3. Finally, in Section 1.4, we provide a short outline of this thesis.

1.1 Dark matter

This section provides a summary on dark matter, whose nature remains one of the great mysteries of astrophysics and cosmology to this very day. We start by discussing observational

evidence for its existence, then introduce the linear evolution of dark matter density perturbations, gravitational collapse, the formation of haloes, and the halo mass function. Finally, we discuss dark matter candidates and end this section by presenting some of the most promising astrophysical probes of the properties of dark matter.

1.1.1 Observational evidence for the existence of dark matter

The existence of dark matter was first proposed by Zwicky (1933) to explain the orbits of galaxies within the Coma cluster. Later, Rubin & Ford (1970) and Rubin et al. (1978, 1985) similarly found that the rotation curves of stars and gas in the M31 galaxy and other spiral galaxies require more than the observed luminous matter to explain their orbital velocities.

It soon became established that the radial velocities of galaxy cluster members require an additional dark matter component (see e.g. Peebles 1993). As the baryon content of clusters is in agreement with the expectations derived from standard Big Bang Nucleosynthesis (BBN, see e.g. Dar 1995), this suggested the presence of dark matter that is non-baryonic. Further evidence then came from observations of weak and strong gravitational lensing, where the baryonic matter alone does not account for the observed lensing effect of elliptical galaxies and galaxy clusters (see e.g. Sanders et al. 1984, Fahlman et al. 1994, Kochanek 1995, Keeton 2001, Clowe et al. 2004, 2006). It furthermore became evident that dark matter leaves an imprint on the large scale structure of the Universe (see e.g. Frenk & White 2012). More recently, analyses of the temperature fluctuations in the cosmic microwave background (CMB) showed that dark matter constitutes around 83 per cent of the Universes matter content (see e.g. Hinshaw et al. 2013, Ade et al. 2016, Aghanim et al. 2020), rendering it a crucial ingredient to explain the evolution and formation of structure within the Universe.

1.1.2 Evolution of density perturbations and gravitational collapse

The temperature fluctuations observed in the CMB (their power spectrum is shown in Figure 1.1) are related to perturbations in the matter density distribution of the early Universe. This density distribution can be described as $\rho(\vec{x},t) = \bar{\rho}(t) \times (1 + \delta(\vec{x},t))$, with $\bar{\rho}(t)$ being the average mass density at the considered time t. The density contrast $\delta(\vec{x},t)$ is statistically well characterized by a power spectrum, $P(k)$ that is defined according to the following equation as a function of wave-number $\vec{k}$:

$$P(k)V^{-1}\delta^D(\vec{k}-\vec{k}') = \left\langle \delta(\vec{k},t)\delta^*(\vec{k}',t)\right\rangle_\delta , \tag{1.1}$$

where the angle brackets represent the calculation of an expectation value and V is the volume under consideration. The primordial density contrast closely follows Gaussian statistics and fluctuates around zero, which reflects the origin of these perturbations during inflation (see e.g. Peacock 1996, 2003). These small initial over- and under-densities form the seeds of the structure observed in the Universe today.

To describe their evolution, it is assumed that dark matter is a non-relativistic fluid with a velocity dispersion $\sigma_\mathrm{v}(t)$. In this context, $\sigma_\mathrm{v}(t) = 0$ corresponds to ideal cold dark matter

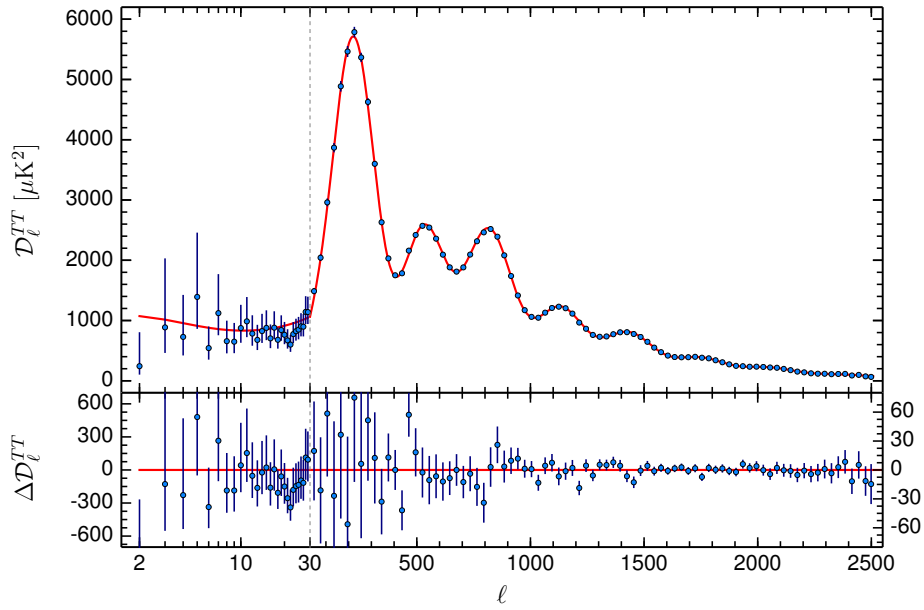

Figure 1.1: **The Planck 2015 temperature power spectrum.** The top panel shows the power spectrum of temperature fluctuations with its uncertainty, as determined from the analysis of the Planck data (blue points) and their best fit assuming the ΛCDM cosmology (red curve). The bottom panel shows residuals from the best fit. This figure is taken from Ade et al. (2016).

(CDM) and $\sigma_\mathrm{v}(t) > 0$ corresponds to warm dark matter (WDM) models. It is then possible to derive the linear evolution of the density contrast of dark matter under the influence of a gravitational potential. The resulting equations of motions within the framework of an expanding space are given as (see e.g. Schneider et al. 2013):

$$\left(\frac{\partial^2}{\partial t^2} + 2\frac{\dot{a}}{a}\frac{\partial}{\partial t}\right)\delta(\vec{k},t) = \left(4\pi G\bar{\rho}(t) - \frac{\sigma_\mathrm{v}^2(t)\vec{k}^2}{a^2}\right)\delta(\vec{k},t)\,, \tag{1.2}$$

with $a(t)$ being the scale factor relating comoving distances to proper ones and $\delta(\vec{k},t)$ being the Fourier transform of the density contrast. The second term on the right-hand side of equation (1.2) accounts for the free streaming of particles. Notice that the above relation only holds within the linear regime of structure formation and does no longer apply in the limit of high densities and streaming velocities.

Ignoring the $\frac{\dot{a}}{a}$ term of equation (1.2), which accounts for the influence of the expansion of the Universe, we find that the right-hand side must be positive for a density fluctuation to grow in amplitude. At a given time t, this defines the comoving Jeans scale $\lambda_\mathrm{J}(t)$ and the Jeans mass $M_\mathrm{J}(t)$, which are characteristic for gravitationally collapsing objects:

$$\lambda_\mathrm{J} = \frac{2\pi a}{k_\mathrm{J}(t)} = \sigma_\mathrm{v}(t)\sqrt{\frac{\pi}{G\bar{\rho}(t)}}\,, \tag{1.3}$$

$$M_{\rm J} = \frac{4\pi}{3}\bar{\rho}(t)\lambda_{\rm J}^3(t)\,. \tag{1.4}$$

Modes with $\lambda < \lambda_{\rm J}$ correspond to propagating acoustic waves, while modes with $\lambda > \lambda_{\rm J}$ correspond to growing stationary waves and represent gravitational instabilities (see Peacock 2003 or Mo et al. 2010 for a more detailed derivation).

1.1.3 Dark matter haloes

In this section, we give a short description of the formation and properties of dark matter haloes, followed by a discussion on their mass function.

The formation and structure of haloes

The gravitational collapse of collisionless dark matter over-densities leads to the formation of gravitationally bound dark matter haloes. These objects, which have been the focus of both numerical and observational studies, form in a bottom-up hierarchical manner, i.e. smaller structures form earlier than the larger ones. For a given cosmology, it is possible to determine the density of haloes at the time of spherical collapse t, i.e. $\rho_{\rm halo} = \Delta_{\rm halo}\bar{\rho}(t)$ (see e.g. Benson 2010, Mo et al. 2010), with $\Delta_{\rm halo} \approx 200$ (see.g. Davis et al. 1985, Eke et al. 1996). Assuming that this density is constant over the virialized region, one can then determine the virial radius of a halo of a given mass M as:

$$r_{\rm vir} = \left(\frac{3M}{4\pi\Delta_{\rm halo}\bar{\rho}(t)}\right)^{1/3}. \tag{1.5}$$

The resulting haloes show a remarkably self-similar structure independent of cosmology or mass (Mo et al. 2010). In particular, Navarro et al. (1996) determined from N-body simulations, that haloes show a universal mass density profile:

$$\rho_{\rm NFW}(r) = 4\frac{\rho_{\rm s}}{r/r_{\rm s}(1+r/r_{\rm s})^2}\,, \tag{1.6}$$

with $r_{\rm s}$ being the scale radius and $\rho_{\rm s}$ the density at this radius. The concentration of a halo is defined as $c = r_{\rm vir}/r_{\rm s}$ and depends on the mass, redshift and assembly history of a halo (Navarro et al. 1997, Bullock et al. 2001, Wechsler et al. 2002, Duffy et al. 2008). Simulations show that the NFW profile provides a good description of the mass distributions of haloes over a mass range of at least 20 orders of magnitude with a 10 per cent accuracy (Wang et al. 2020). The more the more general parameterization given by the Einasto profile (Einasto 1965) provides an improved accuracy of a few percent.

The halo mass functions

One of the key predictions of dark matter models is the halo mass function, which describes how the number density of haloes, n, changes with halo mass, m. The shape of this function

strongly depends on the properties of dark matter and therefore constitutes a promising tool to constrain dark matter models.

Considering the Gaussian nature of the primordial density perturbations, Press & Schechter (1974) derived an expression for the mass function $\frac{dn}{dm}$ by associating haloes with density peaks in the early Universe. To this end, they considered the variance of the density field smoothed by a top-hat filter with a characteristic radius R that, on average, contains a mass of m. For a given power spectrum one can derive that $\sigma^2(m) = \int dk P(k) W(k, R) k^2$, with $W(k, R)$ being the Fourier transformed top-hat filter. By calculating the probability that a density fluctuation of such size falls above the critical over-density $\delta_c(t)$, which is required for spherical collapse to happen at time t, they derived that the mass function is given by:

$$\frac{dn}{dm} = \left(\frac{2}{\pi}\right)^{1/2} \frac{\rho_0}{m^2} \frac{\delta_c(t)}{\sigma(m)} \times \left|\frac{d \ln \sigma(m)}{d \ln m}\right| \times \exp\left[-\frac{\delta_c^2(t)}{2\sigma^2(m)}\right], \tag{1.7}$$

where ρ_0 is the critical density of the Universe. Similar assumptions can be used to derive the merger rate and history as a function of redshift and mass (for more details, see Lacey & Cole 1993).

The mass function as derived by Press & Schechter (1974) required an unexplained factor of two in order for the total dark matter content of the Universe to be within collapsed haloes. This factor was later explained by an extension of the Press-Schechter formalism in the form of an excursion set theory (EPS, Bond et al. 1991). Moreover, by allowing for the elliptical collapse of haloes, Sheth et al. (2001) further improved upon previous descriptions.

At high mass densities and particle velocities the evolution of dark matter haloes becomes non-linear, preventing a completely analytical derivation of the mass function. N-body and hydrodynamical simulations become then necessary to determine the halo mass function and its evolution (see e.g. Jenkins et al. 2001, Lukić et al. 2007, Bocquet et al. 2016). This non-linearity is of particular importance for subhaloes as their number and properties can be significantly affected by the interactions with their host halo (see e.g. Springel et al. 2008).

Galaxy formation and dark matter

Both simulations (see e.g. Klypin & Shandarin 1983, Springel 2005, Heitmann et al. 2008, Marinacci et al. 2018) and observations (see e.g. Hamilton et al. 1991, Springel et al. 2005, Lavaux & Jasche 2015, Jasche et al. 2015) reveal that the formation of structures at large scales results in a network of haloes, voids, sheets and filaments, which are collectively referred to as the cosmic web. Within this web, galaxies have been shown to roughly follow the distribution of dark matter haloes (see e.g. Verde et al. 2002, Ross et al. 2015).

Galaxies form from the gas of the intergalactic medium (IGM) that falls and cools in the potential wells of the dark matter haloes (see e.g. Mo et al. 2010). The cooling of gas happens mainly via Bremsstrahlung, (de-)excitation of atoms and molecules, and inverse Compton scattering. The dominant cooling mechanism as a function of gas temperature is shown in Figure 1.2. Once the gas has cooled down to the typical temperature of around $10 - 20$K, the formation of stars begins (see e.g. Larson 2003).

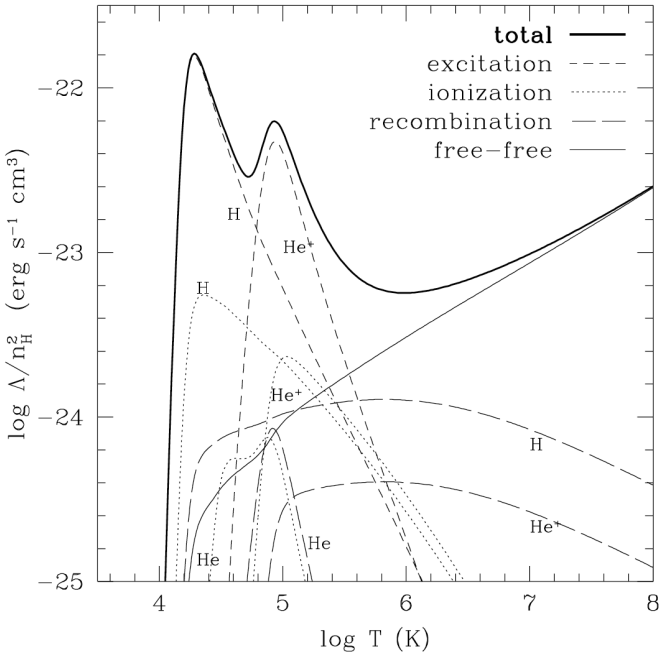

Figure 1.2: **Cooling rate of primordial gas as a function of temperature**. Shown are the total cooling rate and the contributions from individual cooling processes. This figure is taken from Thoul & Weinberg (1995).

The first galaxies are assumed to form at redshifts $z \approx 20 - 50$ (Tegmark et al. 1997, Gao et al. 2007). Their evolution across cosmic time depends then on a combination of several complicated phenomena: further accretion of gas, mergers with other galaxies, active-galactic-nuclei (AGN) and stellar feedback (see e.g. Nelson et al. 2015). Many of these processes operate on scales currently not resolved by simulations. It is therefore common practice to describe them via simplified sub-resolution models. The specific choice of these prescriptions significantly affects the resulting properties of simulated galaxies (e.g. Naab & Ostriker 2017).

Many astrophysical probes of dark matter target galaxies and their properties (see Section 1.1.5). However, the complicated interplay between dark and baryonic matter significantly affects the main properties of galaxies (e.g., densities, temperatures, kinematics, and morphologies) in a way that it is still not fully understood (e.g., Lagos et al. 2016, Despali et al. 2019b, Vogelsberger et al. 2014, Nelson et al. 2019). For this reason, astrophysical observations that are less affected by baryonic physics (such as strong gravitational lensing) are considered to be a more reliable probe of the nature of dark matter.

We refer to Benson (2010) and Mo et al. (2010) for a detailed description of galaxy formation and evolution processes.

1.1.4 Dark matter candidates

In this section, we give a non-exhaustive list of dark matter models. More detailed discussions of dark matter candidates are presented by Baltz (2004) or Bertone & Tait (2018). Dark matter models are usually motivated to either explain specific astrophysical observations or address open questions in particle physics. In the following, we will present examples of both cases. We highlight that dark matter does not need to consist of a single candidate exclusively, but could also be the combination of several different physical phenomena.

Thermal relic dark matter

One of the first theoretical discussions of the general class of thermally produced relics was presented by Scherrer & Turner (1986a,b). In particular, they derived an approximate description for the evolution of the relic abundance of Weakly Interacting Massive Particles (WIMPS) independent of their mass and interaction strength. Thermal relic dark matter is of particular interest due to the relation that connects the mass of the dark matter particle with the warmth of the dark matter model. Since this class of dark matter models is the focus of this thesis, we follow with a short description.

If dark matter is a thermal relic, it consists of particles that were produced in thermodynamic equilibrium with photons and other relativistic particles in the early Universe[1]. As the temperature of the Universe drops, dark matter decouples chemically and kinetically from the surrounding plasma (at the freeze-out time). Its density relative to the total entropy density of the Universe is then frozen in time (see e.g. Bertone et al. 2005) and it starts to free stream. This process results in the suppression of the dark matter power spectrum on small scales. In particular, those related to the free-streaming lengths and the size of the horizon at the time of decoupling. The warmer the dark matter particles, the larger the scale at which the suppression happens. In this context, cold dark matter and warm dark matter can be considered as belonging to a continuum spectrum of free-streaming length or particle mass. Figure 1.4 illustrates the difference in the number of subhaloes of a Milky Way-like galaxy, between a CDM simulation and a WDM simulation with a thermal relic particle mass of $m_{\rm th} = 2.0$ keV. The latter shows a strong suppression in the amount of low-mass haloes.

We notice that there exist observational limits on how strong such suppression can be. For example, hot dark matter models (HDM, Blumenthal et al. 1984, Primack & Blumenthal 1984), for which the particle mass is of the order of a few tens of eV, are ruled out by several large scale structure observations, e.g. the distribution of galaxies in the SDSS survey (e.g. Stoughton et al. 2002).

The cut-off in the WDM power-spectrum $P_{\rm WDM}(k)$ can be expressed as a modification to the CDM power spectrum, $P_{\rm CDM}(k)$, according to the following transfer function $T_{\rm matter}(k)$

[1]Physically motivated WDM-class particle models exhibit very different production mechanisms and are not in thermal equilibrium. However, in some cases, such as for the sterile neutrino, the shapes of their linear matter power spectrum can be well approximated by thermal relic models (see Lovell 2020, for an elaborate discussion on the opportunities and limitations of this approximation).

(see e.g. Bode et al. 2001):

$$T_{\rm matter}(k)^2 = \frac{P_{\rm WDM}(k)}{P_{\rm CDM}(k)} = \left((1+(\alpha k)^{2\mu_t})^{-5/\mu_t}\right)^2, \tag{1.8}$$

with the slope parameter $\mu_t = 1.12$ and the scale parameter α, which for a given thermal relic density $\Omega_{\rm th}$ and Hubble constant h is determined to be:

$$\alpha = 0.049\left(\frac{m_{\rm th}}{1\ {\rm keV}}\right)^{-1.11}\left(\frac{\Omega_{\rm th}}{0.25}\right)^{0.11}\left(\frac{\rm h}{0.7}\right)^{1.22} {\rm h}^{-1}\ {\rm Mpc}. \tag{1.9}$$

The half-mode scale $\lambda_{\rm hm}$, at which the transfer function drops to 1/2, is then given by,

$$\lambda_{\rm hm} = 2\pi\alpha\left[2^{\mu_t/5}-1\right]^{-\frac{1}{2\mu_t}}. \tag{1.10}$$

The mass corresponding to this length scale is referred to as the half-mode mass $M_{\rm hm}$, where $M_{\rm hm} = 0$ is equivalent to an idealized CDM model (showing no cut-off) and $M_{\rm hm} \sim 10^{-6}\ {\rm M}_\odot$ is predicted for CDM WIMPs (Green et al. 2004, Schneider et al. 2013).

The suppression of the power-spectrum manifests itself also in the mass function $\frac{d}{dm}n$ as follows (see e.g. Schneider et al. 2012, Lovell et al. 2014b):

$$\frac{d}{dm}n_{\rm WDM} = \frac{d}{dm}n_{\rm CDM}\left(1+\frac{M_{\rm hm}}{m}\right)^{\beta}, \tag{1.11}$$

with the logarithmic slope $\beta \approx -1.3$. We refer to Lovell (2020) for a more recent and general parametrisation. Figure 1.3 shows the subhalo mass function for several half-mode mass values. Following Nadler et al. (2019), a relation between the half-mode mass, $M_{\rm hm} = \bar{\rho}_{\rm m}\frac{4\pi}{3}(\lambda_{\rm hm}/2)^3$, and the thermal relic particle mass, $m_{\rm th}$, can then be derived using that $\bar{\rho}_{\rm m} = \Omega_{\rm th}\rho_{\rm crit} \approx \Omega_{\rm th} \times 2.77 \times 10^{11}\ M_\odot\ {\rm h}^2\ {\rm Mpc}^{-3}$ and the combination of equations (1.9) and (1.10):

$$m_{\rm th} = 2.32\left(\frac{M_{\rm hm}}{10^9 M_\odot}\right)^{-0.3}\left(\frac{\Omega_{\rm th}}{0.25}\right)^{0.4}\left(\frac{\rm h}{0.7}\right)^{0.8}\ {\rm keV}\,. \tag{1.12}$$

Self-interacting dark matter

As a natural consequence of extensions to the standard model of particle physics, self-interacting dark matter (SIDM) constitutes an important class of dark matter (see e.g Feng 2010). In SIDM models, dark matter particles exchange energy and momentum with each other, resulting in predictions that differ from those of the CDM model on the scales of galaxy clusters and below (Burkert 2000, Peter 2012).

In particular, dark matter-only N-body simulations have shown that SIDM haloes show shallower density profiles in their center and tend to be more spherical than their ΛCDM counterparts. (see e.g. Burkert 2000, Yoshida et al. 2000, Miralda-Escude 2002, Cyr-Racine et al. 2016, Vogelsberger et al. 2016). However, this simplistic picture changes when important processes of baryonic physics are taken into account (Davé et al. 2001). Smoothed

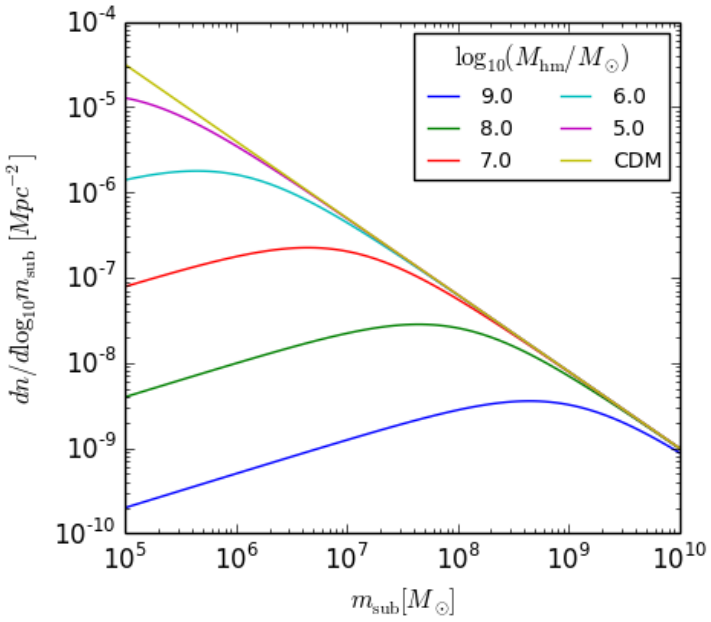

Figure 1.3: **The subhalo mass function for six thermal relic WDM models with different values of the half-mode mass.** The amplitude of the mass functions was set to 1 Mpc^{-2}. The main difference between CDM and WDM models is the number density of low-mass haloes.

particle hydrodynamical simulations reveal a diverse evolution of haloes that depends on their formation time, mass accretion history, baryonic content, and morphology (Vogelsberger et al. 2014, Despali et al. 2019b, Robertson et al. 2018, 2019). Furthermore, while the simplest SIDM scenarios assume that the self-interaction cross-section is constant, other more complex models allow for velocity-dependent interactions (see e.g Vogelsberger et al. 2016) or dark radiation (see e.g. Cyr-Racine et al. 2016), which change the predictions of an SIDM model significantly. Several astrophysical observations place constraints on the self-interaction cross-section of dark matter σ/m, among them the elliptical nature of clusters (see e.g. Miralda-Escude 2002, with $\sigma/m \lesssim 0.02$ cm^2/g) and the analysis of the cluster mergers (Randall et al. 2008, with $\sigma/m \lesssim 0.7 - 1.25$ cm^2/g).

It should also be noted that, like WDM models, some flavours of SIDM are also expected to suppress the number of low-mass haloes (Yoshida et al. 2000, Spergel & Steinhardt 2000, Vogelsberger et al. 2016).

Supersymmetric particles

One of the most studied extensions to the standard model of particle physics is Supersymmetry (SUSY). SUSY not only provides multiple possible candidates of dark matter particles, but also offers a solution to the hierarchy problem (see e.g. Shirman 2010, Csáki et al. 2018, Baer et al. 2020). This problem concerns the fine-tuning of the Higgs bosons potential so that the quadratic quantum corrections to its mass cancel each other out for high energy cutoffs.

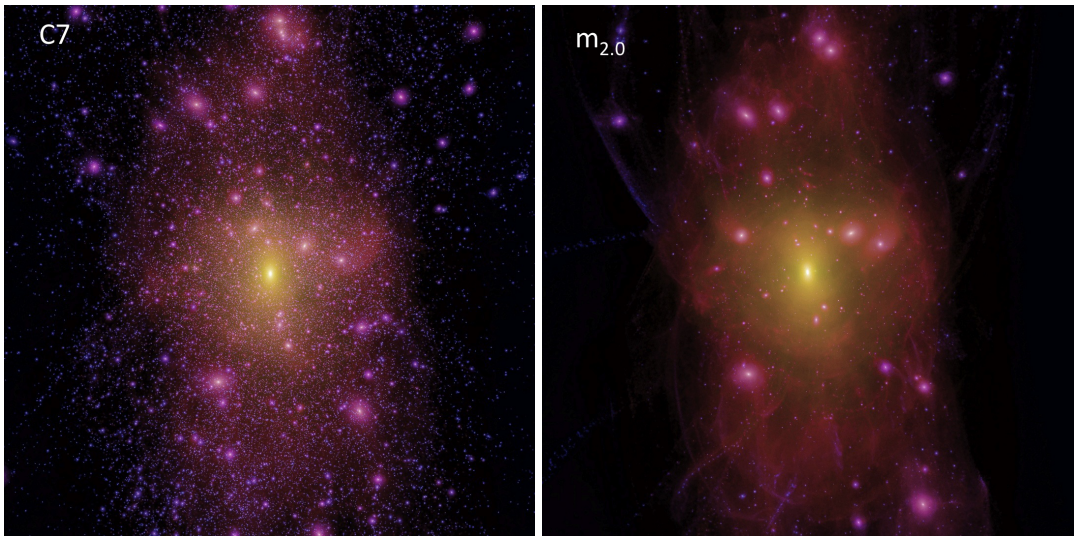

Figure 1.4: **N-body simulations of a Milky Way-like halo in CDM (left) and a thermal relic WDM model (right) in projection.** The image intensity scales with the projected squared dark matter density. The colour changes according to the density-weighted mean velocity dispersion. The side length of both images corresponds to 1.5 Mpc in both dimensions. This figure is taken from Lovell et al. (2014b), who considered simulations based upon the Aquarius project (Springel et al. 2008).

Under Supersymmetry, every bosonic particle has a fermionic particle counterpart and vice versa. The neutral gravitino and the neutralinos are of particular interest, as they constitute (depending on the SUSY model) the lightest and stable superpartners (Feng 2005). The spin 3/2 gravitino is the superpartner of the graviton and provides a candidate for both warm and cold dark matter. Depending on its mass and the specific SUSY model, gravitinos are mainly produced during inflation and the following period of reheating, via thermal freeze-out (see Section 1.1.4), or as a late decay product before and during BBN (Feng 2005, Ellis et al. 1982). The four spin 1/2 Fermion superpartners of the neutral gauge and Higgs bosons, the neutralinos, constitute further SUSY dark matter candidates (Feng 2005). Assuming these neutralinos are stable, thermal freeze-out is their most likely production mechanism. When the masses of the gauge field superpartners are of weak-scale order (< 300 GeV), the proper thermal relic density is automatically obtained (see e.g. Baltz 2004).

Thermally produced SUSY particles belong to the class of WIMPs, for which there exists extensive constraints from several experiments (see Figure 5 of Roszkowski et al. 2018): the direct detection of dark matter in nuclear scattering experiments (see e.g. Peter et al. 2013, Undagoitia & Rauch 2015), indirect detections from excess emission of Gamma and X-rays in galaxies (see e.g. Cirelli et al. 2011, Bulbul et al. 2014, Gaskins 2016) and particle accelerators such as the LHC (see e.g. Mitsou 2015).

An example of an astrophysical simulation based on WIMPs as a dark matter model is found in Wang et al. (2020).

Fuzzy Dark matter

A further dark matter model that has recently gained much attention is fuzzy dark matter (FDM, see Chan et al. 2020, Kulkarni & Ostriker 2020, May & Springel 2021, Nori & Baldi 2021). FDM consists of bosonic particles which can manifest their quantum nature on galactic scales. The reason for this are their low masses of $m_{\rm FDM} \approx 10^{-22}$ eV, which result in de Broglie wavelengths of $\sim$ 1kpc.

This property results in a suppression of the matter density on small scales, similar to WDM models, but does not require the free streaming of particles. It has further been shown that FDM models lead to the formation of cores in the centre of dark matter haloes (Hu et al. 2000, Macciò et al. 2012, Ferreira 2020). FDM simulations reveal that structure formation is furthermore affected by the interference patterns arising from its quantum nature (Schive et al. 2014, Nori & Baldi 2018, Zhang et al. 2018, Li et al. 2019, Mocz et al. 2019), thereby providing unique observational predictions that can distinguish it from WDM models.

Possible candidates for FDM are ultra-light Axion-like particles (Diez-Tejedor & Marsh 2017). Constraints on their mass can be derived from the rotation curve of the Milky Way (Maleki et al. 2020), Ultra-faint dwarf galaxies such as Eridanus-II (Marsh & Niemeyer 2019) and the Lyman-α forest (Iršič et al. 2017), which find lower limits around $m = 10^{-22} - 10^{20}$ eV.

Axions

The existence of Axions was originally motivated as an solution to the strong CP problem. This problem concerns the theoretically allowed, but observationally unobserved, violation of the combined symmetries of charge and parity in the strong nuclear force. Assuming that the parameter describing the strength of this violation is related to an additional dynamical field can provide a natural solution to the problem (Peccei & Quinn 1977, Weinberg 1978, Wilczek 1978) as its ground state can entail minimal violation.

The particle corresponding to this new field is the Axion, and since it satisfies many of the properties that are required for dark matter its cosmological implications have been studied and discussed extensively (see e.g. Sikivie 2008, Marsh 2016, 2017, Eggemeier et al. 2020). There has also been an intensive experimental effort to find Axion-like particles (see e.g. Graham et al. 2015). A rough estimate for the Axions mass range can be derived from cosmological and astrophysical observations. Stellar and supernovae cooling puts constraints on the axion mass. From the cooling of neutron stars, one can for example derive an upper limit of roughly $m_{\rm a} = 10^{-2} - 10^{-1}$ eV (see e.g. Sedrakian 2016). A lower bound of $m_{\rm a} = 10^{-5}$ eV can be found by requiring that the relic density of Axions must be compatible with the observed critical density of the Universe (Raffelt 1995).

Axion masses are expected to be significantly below the boundary that is expected to correspond to warm dark matter ($\approx$ keV). However, they are not produced thermally and would therefore constitute a cold dark matter candidate (Baltz 2004).

Sterile neutrinos

The origin of the low masses of neutrinos and the oscillation of flavours (see e.g. Fukuda et al. 1998) are unexplained phenomena of the standard model of particle physics (Boyarsky et al. 2019). Several different solutions have been proposed to explain them, many of which require additional particles such as the sterile neutrino.

Sterile neutrinos with a mass of $m_\nu \approx$ keV constitute an interesting (warm) dark matter candidate and have first been proposed as such by Dodelson & Widrow (1994). For a suitable choice of the active-sterile mixing parameter, these neutrinos can be stable for timescales that are longer than the Hubble time. The recent discovery of a 3.5 keV X-ray line in galaxy clusters and at the centre of the Milky Way and Andromeda brought particular attention to sterile neutrinos as a dark matter candidate (Boyarsky et al. 2014, Bulbul et al. 2014, Iakubovskyi 2014, Boyarsky et al. 2015).

Sterile neutrinos can be produced either resonantly, during reheating after inflation, or thermally in the early Universe. In this section, we focus on sterile neutrinos, which are resonantly produced at energies corresponding roughly to the energy scale where the Universe transitions from a quark-gluon plasma to hadrons ($\approx$10 MeV). The oscillation of active neutrinos attributes some probability to the transition to such sterile neutrinos. This probability significantly increases in the presence of a strong lepton asymmetry L_a, i.e. an excess in the density of leptons over anti-leptons (relative to the entropy density). The enhancement in the production rate in combination with the mass of the neutrino determines the sterile neutrino velocity and therefore, how warm the resulting dark matter is (Lovell et al. 2016).

Among the most important constraints on sterile neutrinos are strong gravitational lensing observations (Vegetti et al. 2018, Despali et al. 2019a), the Lyman-α forest (Yèche et al. 2017), the CMB and neutrino oscillation experiments (Bridle et al. 2017), the Milky Way satellites (Kennedy et al. 2014), the large scale structure, BBN (Hamann et al. 2012), and many more (see e.g. Drewes & Garbrecht 2017).

Recently, significant effort has been placed on running cosmological simulations employing sterile neutrino dark matter (e.g. Angus et al. 2014, Bozek et al. 2016, Lovell et al. 2017, Bozek et al. 2018).

1.1.5 Astrophysical constraints on thermal relic dark matter

In this thesis we focus on thermal relic warm dark matter and how it can be constrained with astrophysical observations. To achieve this purpose several methods have been proposed to measure the distribution of dark matter on galactic and subgalactic scales. In the following section, we give a non-exhaustive list of the most important methods.

The Lyman-alpha forest

The Lyman-α forest is one of the most studied features of optical spectra from Lyman-α emitting quasars at high redshifts. It is the main observational probe of the intergalactic medium (IGM; see Meiksin 2009, McQuinn 2016, for a review), and is used to constrain the

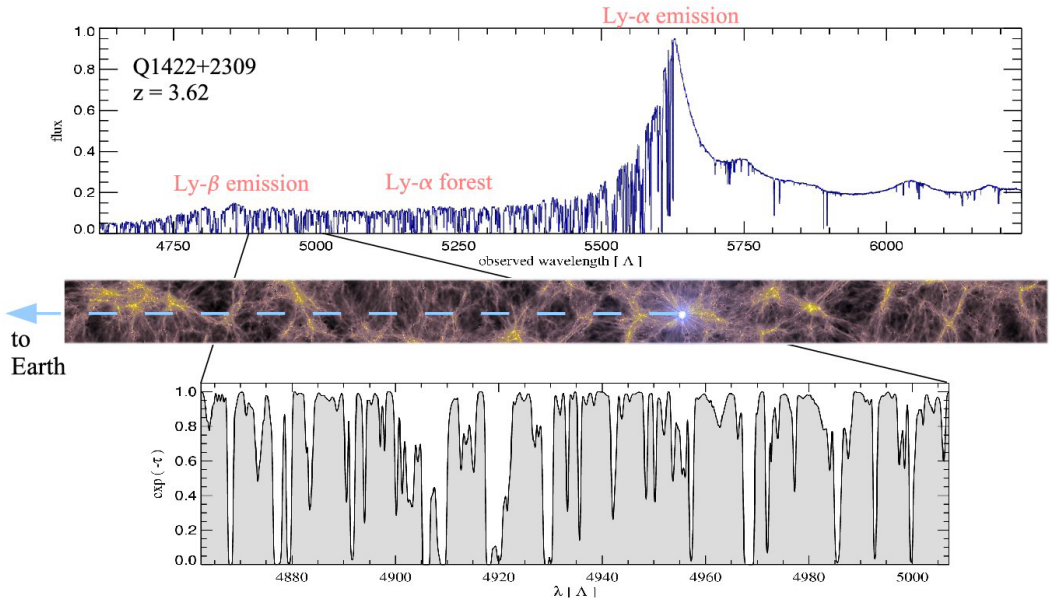

Figure 1.5: **The Lyman-α forest as a probe of the matter distribution.** The top panel shows a typical spectrum of a quasar at z = 3.62 with a redshifted Lyman-α emission line at 5616 angstrom. Towards shorter wavelengths, the spectrum shows the characteristic absorption lines arising from intervening neutral hydrogen along the line of sight from the quasar to the Earth. The middle panel shows a simulation of the intergalactic gas distribution assuming the ΛCDM model and how light-ray could travel from the quasar in the direction of the Earth. The bottom panel shows an example of a simulated absorption spectrum for such a scenario. This figure is taken from Springel et al. (2006), who consider the Lyman-α forest as a probe of large scale structure. For high-resolution spectra it can also be used to probe the small scale structure.

nature of dark matter as well as other cosmological quantities (Narayanan et al. 2000, Viel et al. 2005, Seljak et al. 2006, Viel et al. 2006, 2008).

The origin of the Lyman-α forest can be summarized as follows. A quasar emits light across a wide spectrum of wavelengths. As this light travels towards the observer it is redshifted by the expansion of space, so that Lyman-α absorption from neutral hydrogen clouds along the line of sight suppresses different parts of the original quasar spectrum at each redshift (as shown in Figure 1.5). As the matter power spectrum depends on the dark matter model, the comparison between mock spectra obtained from hydrodynamical simulations with different dark matter models and those derived from spectroscopic observations provides constraints on the nature of dark matter.

From the analysis of high-quality, high-resolution quasar absorption spectra at redshifts up to $z \approx 5.4$, Iršič et al. (2017) constrained the lower limit of the thermal relic particle mass to be 4.1 and 2.1 keV at the 95 per cent confidence level, assuming a smooth power-law and a free-form temperature evolution of the IGM, respectively. Recently, Murgia et al. (2017, 2018)

developed a broader approach to constrain generalised dark matter models (e.g. Archidiacono et al. 2019, Miller et al. 2019, Baldes et al. 2020) and obtained lower limits on the thermal relic particle mass of 3.6 and 2.2 keV at the 95 per cent confidence level for the two thermal histories above, respectively.

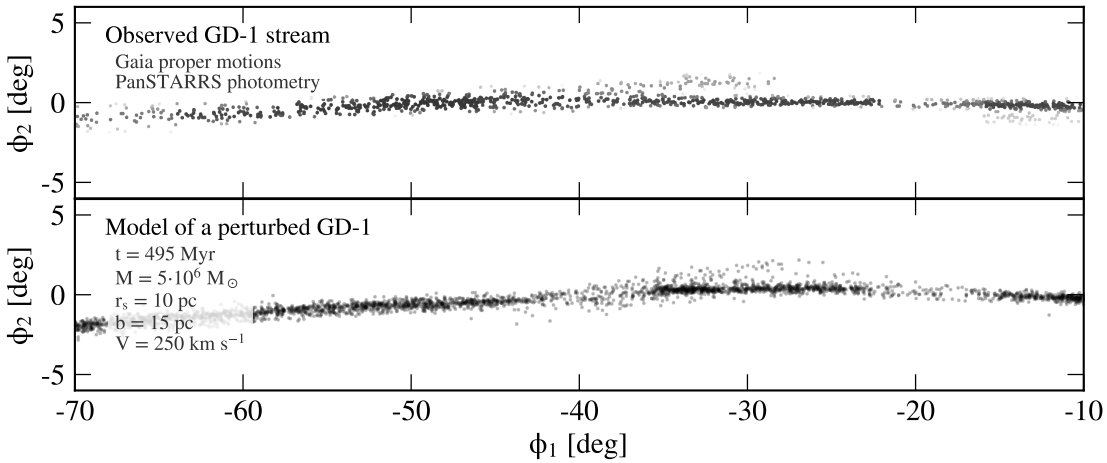

Figure 1.6: **The stellar stream GD-1 as identified with Gaia proper motions and PanSTARRS photometry and a reconstructed model.** Shown is the observed stream (top panel) and a model reconstruction (bottom panel). There exist two gaps in the stellar stream: at $\Phi \approx -20.0$, where its progenitor was disrupted, and at $\Phi \approx -40.0$, where a dark substructure interacts with the stream. This figure is taken from Bonaca et al. (2019).

Satellite galaxies

The number of luminous satellite galaxies in the Milky Way and other galaxies can also be used to constrain the properties of dark matter (e.g. Moore et al. 1999, Nierenberg et al. 2013). For example, Lovell et al. (2016) compared the luminosity function of the Milky Way satellites with predictions from semi-analytical models and found a lower limit on the sterile neutrino particle mass of 2 keV (for a lepton asymmetry in the range of $0.0 \lesssim L_6 \lesssim 700.0$). More recently, by comparing the luminosity function of Milky Way satellite galaxies to simulations and incorporating observational incompleteness in their model, Jethwa et al. (2017) obtained a lower limit of 2.9 keV on the thermal relic particle mass at the 95 per cent confidence level. Nadler et al. (2019) derived a more stringent lower limit of 3.26 keV from the analysis of the classical Milky Way satellites and those discovered by the Sloan Digital Sky Survey (SDSS). Combining data from the Dark Energy Survey (DES, Abbott et al. 2018) and the Pan-STARRS1 surveys (Chambers et al. 2019), Nadler et al. (2020b) derived a lower limit on the mass of thermal relic dark matter of 6.5 keV at the 95 per cent confidence level. Recently, Newton et al. (2020) derived a lower limit of $2.02 - 3.99$ keV by comparing the predicted number of subhaloes, calculated for a WDM Milky Way-mass halo using the Extended Press-

Schechter formalism, with the observed population, reconstructed by Newton et al. (2018a,b) from a combination of DES and SDSS data (Ahn et al. 2012).

Stellar streams

As globular clusters orbit within the potential of their host galaxy, they experience strong tidal forces which lead to the stripping of their stars and the formation of stellar streams (see e.g. Grillmair & Dionatos 2006, Ibata et al. 2016, Necib et al. 2019). If a stream encounters dark or luminous massive structures within the host galaxy its shape is perturbed and gaps in its stellar distribution become visible (Bovy et al. 2016, Erkal et al. 2016). Figure 1.6 shows, as an example, the effect of dark matter subhaloes in the Milky Way on the stellar stream GD-1 (Bonaca et al. 2019). The strength and amount of these perturbations can be used to constrain the number of subhalaoes within a galaxy and, therefore, the nature of dark matter (Banik et al. 2018). Using data from the Gaia data release 2 and the Pan-STARRS survey of the Milky Way streams GD-1 and Pal5, Banik et al. (2019a,b) derived a lower limit on the mass of thermal relic warm dark matter particles of 4.6 keV.

Strong gravitational lensing

Strong gravitational lensing, being sensitive only to gravity, allows the detection of low-mass haloes independently of their baryonic content. Therefore, it provides a direct method to quantify the dark matter distribution on subgalactic scales, where most of the structures are expected to be non-luminous. In practice, these low-mass haloes are detected because of their effect on the flux ratios of multiply imaged compact sources (flux-ratio anomalies; Mao & Schneider 1998) or on the surface brightness distribution of magnified arcs and Einstein rings from lensed galaxies (surface brightness anomalies or gravitational imaging; Koopmans 2005, Vegetti & Koopmans 2009a). So far, both approaches have led to the detection of individual low-mass haloes (Vegetti et al. 2010, 2012, Nierenberg et al. 2014, Hezaveh et al. 2016), as well as statistical constraints on the halo and subhalo mass functions and the related particle mass of sterile neutrinos and thermal relic warm dark matter models (Dalal & Kochanek 2002, Vegetti et al. 2014, 2018, Ritondale et al. 2019, Gilman et al. 2019a). In particular, the most recent analyses by Hsueh et al. (2019) and Gilman et al. (2019a) have derived a lower limit on the mass of a thermal relic dark matter particle of 5.6 and 5.2 keV at the 95 per cent confidence level, respectively. As strong lensing is the main focus of this thesis, we will give a brief overview of this phenomenon in Section 1.2.

1.2 Gravitational lensing

In this section, we give a short summary of the theory of gravitational lensing. We first describe the different types of gravitational lensing and then further focus on the regime of strong lensing. A more detailed introduction to this astrophysical phenomenon can be found in Schneider et al. (1992).

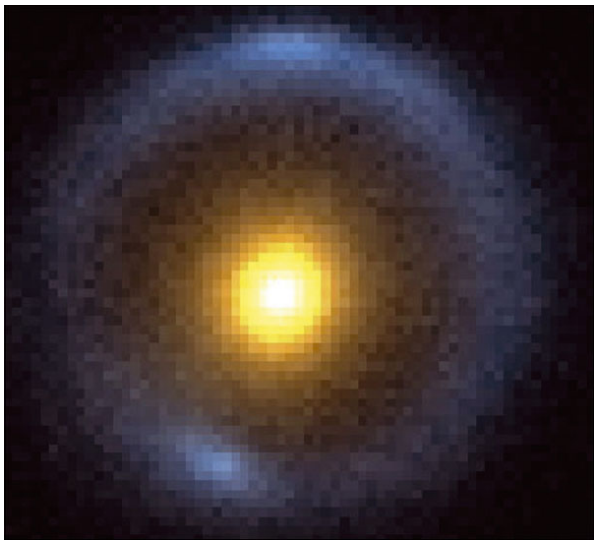

Figure 1.7: **Observation of the gravitational lens system B1938+666.** Shown are in false colours the lens galaxy (yellow) and the background source galaxy (blue). Image credit: David Lagattuta, Chris Fassnacht / UC Davis graphic.

1.2.1 The main types of gravitational lensing

Gravitational lensing is a key prediction of general relativity (Einstein 1936). It occurs when the presence of a massive object (e.g. a galaxy or a cluster of galaxies) deflects the paths of light rays originating from background sources (e.g. another galaxy). Traditionally, one distinguishes three main types of gravitational lensing: micro, weak and strong. In the micro lensing regime, the magnification through the deflection of light rays is so small that it is only detectable as a change in the light curve of the observed source (Wambsganss 2006). Micro lensing can be observed, for example, when the relative magnification of the images of strongly lensed quasars is altered by individual stars in the lens galaxy (Irwin et al. 1989, Vanderriest et al. 1989). In the weak gravitational lensing regime, the deflection of light rays is strong enough to produce a distorted and magnified image of the background source (Mellier 1999). Examples of weak lensing can be observed in the outer regions of galaxies and galaxy clusters (see e.g. Metzler et al. 1999, Giocoli et al. 2014), as well as in the effect of the large scale structure on the CMB (see. e.g. Metcalf & Silk 1997, Calabrese et al. 2008). In the strong gravitational lensing regime, the deflection of light leads to the creation of multiple highly distorted and magnified images of the background source. When the massive object and the background source align a so-called Einstein Ring, i.e. a ring-like distorted image of the background source is observed (see Figure 1.7). The radius of this ring is related to the mass of the deflector. The first observation of strong gravitational lensing was reported by Walsh et al. (1979).

1.2.2 Strong gravitational lensing

The deflection of light due to massive objects can be understood within the framework of general relativity. Under the influence of a Newtonian gravitational potential $\phi(\vec{r})$, the metric of space-time approximately follows:

$$ds^2 \approx \left(1 + \frac{2\phi(\vec{r})}{c^2}\right)c^2dt^2 - \left(1 - \frac{2\phi(\vec{r})}{c^2}\right)d\vec{r}^2\,. \tag{1.13}$$

Given that light travels along trajectories with a space time intervals of $ds^2 = 0$, the effective speed of light within such a potential is given by:

$$\left|\frac{d\vec{r}}{dt}\right| = c\sqrt{\left(1 + \frac{2\phi(\vec{r})}{c^2}\right)\Big/\left(1 - \frac{2\phi(\vec{r})}{c^2}\right)} := c/n(\vec{r})\,. \tag{1.14}$$

This effective speed of light is related to the Shapiro time delay (Shapiro 1964), and it can be interpreted as the result of a change in the refraction index. Assuming a weak gravitational field, i.e. $\phi(\vec{r})/c^2 \ll 1$, and considering Fermat's principle, i.e. that light rays always follow the paths that minimize the travel time, one can derive the equations of motion of a light-ray within a gravitational potential (e.g. Schneider et al. 1992):

$$-\vec{\nabla}_{\perp\dot{\vec{r}}}\frac{2\phi(\vec{r}(\lambda))}{c^2} = \ddot{\vec{r}}_{\perp\dot{\vec{r}}}(\lambda)\,, \tag{1.15}$$

where λ describes the position $\vec{r}(\lambda)$ along the path of the light-ray. The above relation assumes that $|\dot{\vec{r}}(\lambda)| = 1$ without any loss of generality, i.e. the parameter λ corresponds to the distance that the light has travelled. For a general trajectory the total deflection angle can then be calculated by integrating the above equation over λ:

$$\vec{\hat{\alpha}} = -\int_{\lambda_a}^{\lambda_b} d\lambda\, \ddot{\vec{r}}_{\perp\dot{\vec{r}}}(\lambda) = \frac{2}{c^2}\int_{\lambda_a}^{\lambda_b} d\lambda \vec{\nabla}_{\perp\dot{\vec{r}}}\phi(\vec{r}(\lambda))\,. \tag{1.16}$$

For simplicity, in the following we discuss the special case of a single lens and a single background source (see Figure 1.8 for a schematic representation). This is the typical scenario in galaxy-galaxy strong lensing observations. The lens is located at a distance D_l from the observer, while the distance between the source and the lens, and the source and the observer are referred to as D_s and D_ls, respectively. For any given position on the lens plane, $\vec{x} = D_\mathrm{l}\vec{\theta}$, and its corresponding position on the source plane, $\vec{y} = D_\mathrm{s}\vec{\beta}$, the displacement due to lensing is defined as $\|\vec{x} - \vec{y}\|$. As the latter is much smaller than D_s, we can assume that: $\tan(\theta) \approx \sin(\theta) \approx \theta$. Equation (1.16) then simplifies to an integral over an unperturbed path parametrized by its redshift z:

$$\vec{\hat{\alpha}}(\vec{\theta}) = \frac{2}{c^2}\int_{-\infty}^{\infty} dz \vec{\nabla}_{\vec{\theta}}\phi(D_\mathrm{l}\vec{\theta}, z)\,. \tag{1.17}$$

The angle β on the source plane is related to the corresponding angle θ on the lens plane by:

$$\vec{\beta}(\vec{\theta}) = \vec{\theta} - \frac{D_\mathrm{ls}}{D_\mathrm{s}}\vec{\hat{\alpha}}(D_\mathrm{l}\vec{\theta}) := \vec{\theta} - \vec{\alpha}(\vec{\theta})\,. \tag{1.18}$$

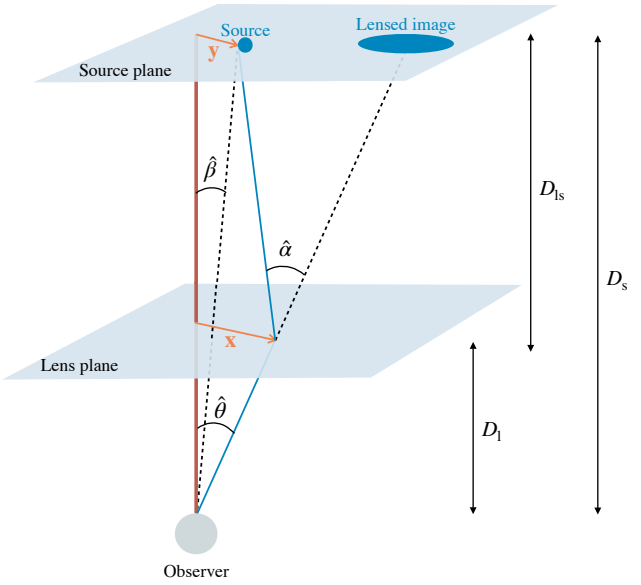

Figure 1.8: **Schematic representation of the gravitational lens geometry.** The blue lines illustrate the path of a light ray from the source to the observer. The light path is bent in the presence of a massive object (e.g., a galaxy or a cluster of galaxies). The red line shows the optical axis. Image credit: Francesca Rizzo.

The scaled deflection angle $\vec{\alpha}(\vec{\theta})$ can further be related to the derivative of the lensing potential, i.e. $\vec{\alpha}(\vec{\theta}) = \vec{\nabla}_{\vec{\theta}}\Psi(\vec{\theta})$, with:

$$\Psi(\vec{\theta}) = \frac{D_{\rm ls}}{D_{\rm l}D_{\rm s}}\frac{2}{c^2}\int_{-\infty}^{\infty} dz\phi(D_{\rm l}\vec{\theta}, z)\,. \tag{1.19}$$

The second derivative of $\Psi(\vec{\theta})$ is the convergence $\kappa(\theta)$ of the lens and corresponds to the projected mass surface density $\Sigma(D_{\rm l}\vec{\theta})$ normalised to the critical density $\Sigma_{\rm crit} = \frac{c^2}{4\pi G}\frac{D_{\rm S}}{D_{\rm L}D_{\rm LS}}$, i.e.:

$$\kappa(\vec{\theta}) = \frac{\Sigma(D_{\rm l}\vec{\theta})}{\Sigma_{\rm crit}} = \frac{1}{2}\vec{\nabla}_{\vec{\theta}}\vec{\alpha}(\vec{\theta}) = \frac{1}{2}\Delta_{\vec{\theta}}\Psi(\vec{\theta})\,. \tag{1.20}$$

Strong lensing only occurs if the projected surface density exceeds the critical density over some region of the sky. We notice that the critical density is a geometric property, which depends on the distance of the lens and the source with respect to the observer.

Taking the derivative of the source position with respect to the observed angle on the sky θ, we obtain the inverse of the differential magnification matrix $M(\vec{\theta})$ and its dependency on the lensing potential:

$$M^{-1}(\vec{\theta}) = \vec{\nabla}_{\vec{\theta}}\vec{\beta}^T(\vec{\theta}) = \mathbb{I} - \vec{\nabla}_{\vec{\theta}}\vec{\nabla}^T_{\vec{\theta}}\Psi(\vec{\theta})\,. \tag{1.21}$$

We notice that the magnification, $\mu = \det(M(\vec{\theta}))$, is different for every point $\vec{\theta}$ on the sky. The points where $\det(M^{-1}(\vec{\theta})) = \mu^{-1} = 0$ are referred to as the critical curves and the

corresponding curves on the source plane define the caustics. Point sources that fall on the caustics are infinitely magnified by the lens. When a source crosses the caustics, multiple images are created in pairs. Notice that if the determinant of $M(\vec{\theta})$ changes its sign, the parity of that image is changed. The magnification matrix can be decomposed further into components that relate to the convergence and to the shear $\vec{\gamma}(\vec{\theta})$:

$$M^{-1}(\vec{\theta}) = \begin{pmatrix} 1 - \kappa(\vec{\theta}) - \gamma_1(\vec{\theta}) & -\gamma_2(\vec{\theta}) \\ -\gamma_2(\vec{\theta}) & 1 - \kappa(\vec{\theta}) + \gamma_1(\vec{\theta}) \end{pmatrix} . \tag{1.22}$$

1.3 Bayesian statistics

Most of the analysis methods used in this thesis are based Bayesian statistics. In this section, we give a brief summary of Bayes theorem and present a small selection of Bayesian inference algorithms. For a more detailed description of Bayesian inference methods in the context of cosmology, we refer to Trotta (2017).

1.3.1 Bayes theorem

The methods used in this thesis often follow one of two goals: the inference of parameters of a given model or the comparison between different models.

The former goal requires the calculation of posterior probability distributions , $\mathcal{P}(\vec{x}|\vec{d}, \mathcal{M})$ of some target parameters $\vec{x}$ within the framework of a model $\mathcal{M}$ given the observed data $\vec{d}$. Bayes theorem states that such distributions can be calculated using the following equality:

$$\mathcal{P}(\vec{x}|\vec{d}, \mathcal{M}) = \mathcal{P}(\vec{d}|\vec{x}, \mathcal{M}) \times \frac{\mathcal{P}(\vec{x}|\mathcal{M})}{\mathcal{P}(\vec{d}|\mathcal{M})} , \tag{1.23}$$

with $\mathcal{P}(\vec{d}|\vec{x}, \mathcal{M})$ being the likelihood, $\mathcal{P}(\vec{x}|\mathcal{M})$ the prior probability distribution assigned to the parameters in this model, and $\mathcal{P}(\vec{d}|\mathcal{M}) = \int d\vec{x}\, \mathcal{P}(\vec{d}|\vec{x}, \mathcal{M}) \times \mathcal{P}(\vec{x}|\mathcal{M})$ the evidence of the model $\mathcal{M}$.

Considering two models $\mathcal{M}_\mathrm{A}$ and $\mathcal{M}_\mathrm{B}$, the ratio of evidences defines the Bayes factor:

$$K_{ij} = \frac{\mathcal{P}(\vec{d}|\mathcal{M}_\mathrm{A})}{\mathcal{P}(\vec{d}|\mathcal{M}_\mathrm{B})} = \frac{\mathcal{P}(\mathcal{M}_\mathrm{A}|\vec{d})\mathcal{P}(\mathcal{M}_\mathrm{B})}{\mathcal{P}(\mathcal{M}_\mathrm{B}|\vec{d})\mathcal{P}(\mathcal{M}_\mathrm{A})} . \tag{1.24}$$

In the context of this work, the Bayes factor is, for example, used as a criterion to compare different parametrisations of the mass distribution of lensing galaxies. The challenge of Bayesian inference is to find efficient algorithms to evaluate the posterior and the evidence in ways that are as accurate and precise as possible.

1.3.2 A selection of Bayesian inference algorithms

We now present a list of Bayesian inference algorithms, which were selected for their relevance to this thesis.

Markov Chain Monte Carlo algorithms

Markov Chain Monte Carlo (MCMC) algorithms aim to generate samples from a target probability distribution $\Pi(\vec{x})$, for example, the posterior of a Bayesian inference problem. This goal is achieved by creating a Markov chain, i.e. a sequence of random variables $\vec{x}_i$ for $i = 0, ..., k$, where each object in the sequence is drawn randomly, such that it only depends on the previous step $\vec{x}_{i-1}$.

The Metropolis-Hastings algorithm is the simplest version of an MCMC method. In order to create a Markov chain, one starts with a suitable starting value for the target parameters $\vec{x}_0$, which can be drawn from their prior. This random variable is then updated according to the following set of rules:

1. Draw a new test sample $\vec{y}$ from a proposal proposal distribution $Q(\vec{y}|\vec{x}_i)$.
2. Accept this new sample with a probability of $\min\left(1, \frac{\Pi(\vec{y})Q(\vec{x}_i,\vec{y})}{\Pi(\vec{x})Q(\vec{y},\vec{x}_i)}\right)$.
3. If accepted set $\vec{x}_{i+1} = \vec{y}$, else $\vec{x}_{i+1} = \vec{x}_i$.
4. Repeat until convergence.

The convergence properties of MCMC algorithms has been studied extensively (see e.g. Roy 2019, Betancourt 2018). We further notice, that there exist several other commonly used MCMC methods, for example, the Gibbs sampling algorithm (see e.g. George et al. 1992) and the Hamiltonian Monte Carlo method (see e.g. Neal 2012). While MCMC chains contain, in principle, all information required to calculate the evidence (see e.g. van Haasteren 2009), alternative methods, such as nested sampling, usually provide more reliable results.

Nested Sampling

Nested sampling is an algorithm used to calculate the evidence of a model $\mathcal{P}(\vec{d}|\mathcal{M})$ while producing a set of (weighted) random samples $\vec{x}_i$, that resemble the posterior distribution (Skilling 2004). It does so by integrating the likelihood while moving through thin nested surfaces $\mathcal{P}(\vec{d}|X, \mathcal{M})$ that are a function of the total prior mass X each of them encloses:

$$\mathcal{P}(\vec{d}|\mathcal{M}) = \int_0^1 dX\, \mathcal{P}(\vec{d}|X, \mathcal{M})\,. \tag{1.25}$$

Assuming that $\mathcal{P}(\vec{d}|X_j, \mathcal{M})$ can be evaluated for each of a given set of decreasing points, $0 < X_j < X_{j-1} < ... < X_1 < 1$, this integral can be expressed using the trapezoid rule:

$$\mathcal{P}(\vec{d}|\mathcal{M}) = \sum_{j=1} \mathcal{P}_j(\vec{d}|\mathcal{M}) = \sum_{j=1}^{j} \mathcal{P}(\vec{d}|X_j, \mathcal{M}) \frac{(X_{j+1} - X_{j-1})}{2}\,. \tag{1.26}$$

The total prior mass, X_j, attributed to a value of the likelihood value, $\mathcal{P}(\vec{d}|X_j)$, can be estimated from the number of drawn samples.

Nested sampling can be performed in the following way. First, one sets the evidence to $\mathcal{P}(\vec{d}|\mathcal{M}) = 0$ and the considered prior volume to $X_0 = 1$. Then, one draws N "live" points of the target parameters $\vec{x}_i$ from their prior probability distribution $\mathcal{P}(\vec{x}_i|\mathcal{M})$ and computes their likelihood $\mathcal{P}(\vec{d}|\vec{x}_i, \mathcal{M})$. New samples are then generated according to the following algorithm:

1. Find the lowest likelihood point $\vec{x}_j$.
2. Estimate its prior volume X_j, e.g. by assuming that $X_j = tX_{j-1}$ (with t being uniformly drawn between 0 and 1).
3. Add $\frac{\mathcal{P}(\vec{d}|X_j,\mathcal{M})}{2}(X_{j+1} - X_{j-1})$ to the current value of the evidence.
4. Turn $\vec{x}_j$ into a "dead" point, and replace it by a new point within the remaining prior volume, i.e. a new point with higher likelihood than $\mathcal{P}(\vec{x}_j|\mathcal{M})$.
5. Repeat until convergence.

A review on the convergence behaviour of nested sampling is given by Chopin & Robert (2010). It is then possible to assign weights to the accepted samples, such that the corresponding distribution matches the posterior distribution. For each sample $\vec{x}_j$ this weight is given by $w_j = \mathcal{P}_j(\vec{d}|\mathcal{M})/\mathcal{P}(\vec{d}|\mathcal{M})$.

Approximate Bayesian Computation

Approximate Bayesian Computation (ABC) provides an alternative inference approach in cases where the likelihood $\mathcal{P}(\vec{d}|\vec{x}, \mathcal{M})$ is not available or its evaluation is computationally expensive. The only requirement for such method is that it is possible to create mock data $\vec{d}^*$, such that $\vec{d}^* = \vec{f}(\vec{x}, \vec{v})$, with $\vec{v}$ describing the nuisance parameters of the considered inference problem. The simplest ABC approach is a rejection-sampling algorithm (see e.g. Akeret et al. 2015, for an application to cosmology). Its main steps can be summarized as follows:

1. Draw a large number of target and nuisance parameters $\vec{x}_i, \vec{v}_i$ from their prior $\mathcal{P}(\vec{x}, \vec{v})$.
2. Generate artificial data from these samples $\vec{d}_i^* = \vec{f}(\vec{x}_i, \vec{v}_i)$.
3. Accept $\vec{x}_i$ if the distance between the mocks and the observed data, $\rho(\vec{d}, \vec{d}^*)$, is smaller than some threshold ϵ.

In order to improve the acceptance rate one often considers only the distance between summary statistics ϕ of the data with a lower dimension, i.e. $\rho(\phi(\vec{d}), \phi(\vec{d}^*))$. For a suitable threshold, the accepted samples are approximately distributed like the posterior. The threshold

is sometimes chosen such that a certain fraction of the closest samples are accepted. Under normal circumstances the chosen threshold does do not give rise to any bias in the inference. We emphasize that in the case that all samples are accepted one will recover the prior distribution.

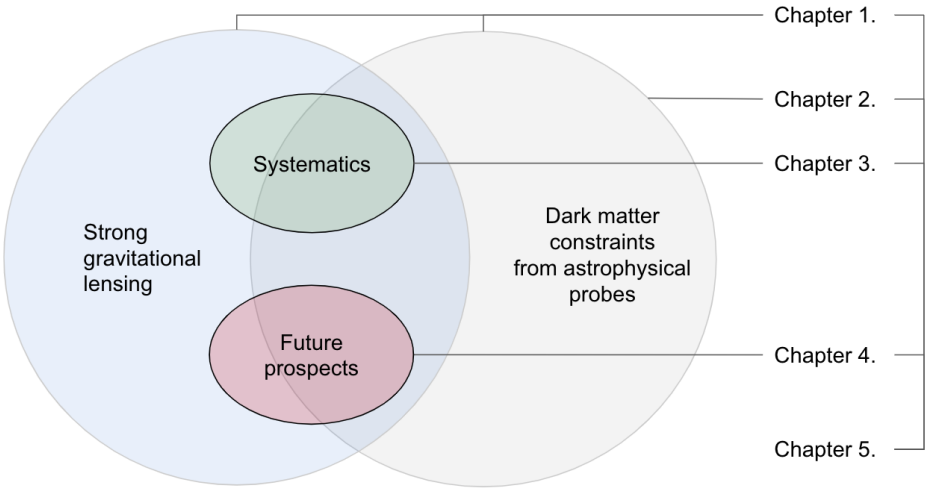

Figure 1.9: **Venn diagram illustrating the topics covered in this thesis.** Each Chapter focuses on a particular aspect of strong gravitational lensing and the inference of dark matter properties from astrophysical probes.

1.4 This thesis

Figure 1.9 illustrates the topics that are covered by this thesis, the remainder of which is structured as follows:

In Chapter 2, we derive constraints on the half-mode scale of thermal relic warm dark matter from the analyses of a selection of astrophysical observations. In particular, the probes that we consider are the Lyman-α forest, the observed amount of luminous Milky Way satellites, and galaxy-galaxy strong gravitational lensing observations. We compare the individual constraints with each other and derive joint constraints. We end the chapter by discussing the most important systematic challenges of each probe and give prospects on their future development.

In Chapter 3, we focus more closely on strong gravitational lensing. The goal of this chapter is to quantify degeneracies in strong lensing reconstructions and how they may affect the inference of the halo and subhalo mass function. To this end, we analyse a sample of twenty-four galaxy-galaxy strong gravitational lens systems with realistic lens galaxies and a background source, taken from the Illustris-1 simulation. We compare the amount of substructures inferred from this data with the one obtained in a control sample, for which we have perfect knowledge about its parameterized lensing potential. We determine important

biases in the lens modelling, which affect both the reconstruction of source galaxies and the reconstruction of the Hubble constant even though the data is fitted to the noise level.

In Chapter 4, we consider multiple observational strategies to determine whether they can strengthen the current constraints of strong gravitational lensing on the nature of dark matter. In particular, we examine if (three-dimensional) observations with ALMA and the JWST can provide the necessary data quality to improve upon the sensitivity of benchmark lens samples to detect low-mass haloes. We consider ALMA observations of the dust continuum emission and the $158\mu m$ fine-structure line of ionized carbon [CII], with both multiple velocity channels and their total average. We compare this sensitivity with the one of synthetic JWST data emulating observations with the onboard Nircam instrument. We then determine how many observations with the data quality of the considered observations are required to find at least ten detectable low-mass haloes, which would provide a significant test of the CDM model.

Finally, in Chapter 5, we provide a summary of the main findings of this thesis, along with a discussion of the main conclusions and how they fit into the bigger picture of astrophysical constraints on the nature of dark matter.

Chapter 2

Constraints on thermal relic dark matter from a selection of astrophysical probes

The content of this chapter is based on work that has been submitted for publication in the Monthly Notices of the Royal Astronomical Society as ***Enzi W.*** *et al. 2020, eprint arXiv:2010.13802. The presented results have been updated to those of the accepted version of the paper, eprint arXiv:2010.13802v3.*

2.1 Introduction

The nature of dark matter is one of the most important open questions in cosmology and astrophysics. While the standard cold dark matter paradigm successfully explains observations of structures larger than $\sim$ 1Mpc, it remains unclear whether observations on smaller (galactic and subgalactic) scales are consistent with this model (e.g. Bullock & Boylan-Kolchin 2017).

Possible alternatives include Warm Dark Matter models (e.g. Bode et al. 2001), in which dark matter particles have higher velocities in the early Universe than in the CDM model. This characteristic leads to the suppression of gravitationally bound structures at scales proportional to the mean free path of the particles at the epoch of matter-radiation equality (e.g. Schneider et al. 2012, Lovell et al. 2014a). Until now, several complementary approaches have been used to test CDM and WDM on these scales. Among these are methods based on observations of strong gravitational lens systems, the Lyman-α forest, and the satellite galaxies of the Milky Way.

In this chapter, we extend and combine the latest constraints on thermal relic WDM from the three methods above (Vegetti et al. 2018, Ritondale et al. 2019, Murgia et al. 2018, Newton et al. 2020) and derive joint constraints on the particle mass of a thermal relic dark matter model. Throughout this chapter we will consider the dark matter model described in Section 1.1.4. This chapter is structured as follows. In Section 2.2, we describe the method with which the different probes are analysed and combined. In Section 2.3, we discuss the results obtained from the individual probes and their joint analysis. We discuss the different sources

Table 2.1: **Main model parameters constrained by strong lensing observations and their relative prior ranges.** From top to bottom: the virial mass of subhaloes and field haloes, the half-mode mass, and the fraction of mass in subhaloes (note that $f_{\rm sub}$ is defined differently in the original analyses of V18 and R19).

Parameter	Original	this chapter
$m_{\rm sub}\,[M_\odot h^{-1}]$	$[\approx 10^5, 2\times10^{11}]$	$[10^6, 10^9]$
$m_{\rm los}\,[M_\odot h^{-1}]$	$[\approx 10^5, 2\times10^{11}]$	$[10^{5.26}, 10^{10.88}]$
$M_{\rm hm}\,[M_\odot h^{-1}]$	$[10^6, 2\times10^{12}]$	$[10^{-6}, 10^{14}]$
$f_{\rm sub}^{\rm CDM}$ [%]	[0.0,4.0]	[0.01, 10.0]

of systematic errors and the future prospects of each individual probe in Sections 2.4 and 2.5, respectively. Finally, we summarize the main results of this chapter in Section 2.6.

2.2 Methods and data

In this section, we provide details on the data, models, and analyses that are the main focus of this chapter. For each probe we rerun and/or extrapolate the analysis in order to match important model assumptions and to guarantee overlap in the prior range of $\lambda_{\rm hm}$.

2.2.1 Strong gravitational lensing

Galaxy-galaxy strong gravitational lensing occurs when the light from a background galaxy is deflected by the gravitational potential of another intervening galaxy. As a result, one observes multiple images of the background galaxy that are highly distorted and magnified. Substructures within the foreground lensing galaxy and low-mass haloes along the line-of-sight to the background object can produce additional perturbations to the lensed images, with a strength that depends on the mass of these (sub)haloes. Therefore, strong gravitational lensing provides a means to constrain the halo and subhalo mass functions directly.

In this chapter, to describe the CDM field halo mass function we assume the formulation introduced by Sheth & Tormen (1999). For the subhalo mass function we assume a power-law,

$$\frac{d}{dm} n_{\rm CDM}^{\rm sub} = A \times m^{-\xi}\,, \tag{2.1}$$

with a logarithmic slope of $\xi = 1.9$ (Springel et al. 2008). We determine the normalization constant A by considering the average fraction of projected mass contained in subhaloes, $f_{\rm sub}$, relative to the total projected mass of the main halo of the lens within two Einstein Radii,

$M(< 2R_{\rm E})$. For a general dark matter model $f_{\rm sub}$ is equal to:

$$f_{\rm sub} = \frac{4\pi R_{\rm E}^2 \int dm\ m \times A \times m^{-\xi} \times \left(1 + \frac{M_{\rm hm}}{m}\right)^{\beta}}{M(< 2R_{\rm E})}\,. \tag{2.2}$$

We derive the value of A by requiring a specific value of $f_{\rm sub}^{\rm CDM}$. Solving the above equation for A and setting $M_{\rm hm} = 0$ as is the case for ideal CDM, we find that the normalization is determined to be:

$$A = \frac{M(< 2R_{\rm E}) f_{\rm sub}^{\rm CDM}}{4\pi R_{\rm E}^2 \int dm\ m \times m^{-\xi}}\,. \tag{2.3}$$

Notice that while A is independent of the dark matter model being warm or cold, according to equation (2.2), the value of $f_{\rm sub}$ in WDM models is related to its CDM counterpart according to:

$$f_{\rm sub}^{\rm WDM} = f_{\rm sub}^{\rm CDM} \times \frac{\int dm\ m \times m^{-\xi} \times \left(1 + \frac{M_{\rm hm}}{m}\right)^{\beta}}{\int dm\ m \times m^{-\xi}}\,. \tag{2.4}$$

The target parameters in our inference process are, therefore, $f_{\rm sub}^{\rm CDM}$ and $M_{\rm hm}$, since they fully describe the mass function of subhaloes. We assume that $f_{\rm sub}^{\rm CDM} \in [0.01, 10]$ per cent with a uniform prior, which covers a wide range of previously inferred values of $f_{\rm sub}^{\rm CDM}$ with their uncertainties (see e.g. Hsueh et al. 2019). Using this parameterization we can enforce that for each lens system the range of normalizations A is the same in WDM and CDM cases and that it scales with the projected mass of the lens galaxy.

Notice, that this is a significant difference to previous works, such as Vegetti et al. (2018) and Ritondale et al. (2019), where the range of $f_{\rm sub}$ was chosen the same for all dark matter models, independently of how warm they are. This assumption is, however, in contradiction with equation (2.4). In particular, this choice would require much higher amplitudes for WDM models in order to obtain a large enough number of high mass objects that can reproduce the same $f_{\rm sub}$ as a CDM model. Since the only change between WDM and CDM models should arise from the conversion function in equation (1.11), the number of subhaloes at the high mass end should not increase in WDM and thus such a parametrization can potentially bias the inference results.

Another difference to previous works is our choice of the mass range for haloes. We chose a mass range of subhaloes of $m_{\rm sub} \in [10^6, 10^9] M_\odot h^{-1}$. The upper limit is chosen so that this range includes only masses corresponding to objects that are not expected to be visible, because they are either non-luminous or too faint to be observed (see e.g. Moster et al. 2010). We choose the mass range of line-of-sight haloes such that it contains the masses that show the most similar lensing effects to the lightest and heaviest substructures according to the mass-redshift relationship derived by Despali et al. (2018), $m_{\rm los} \in [10^{5.26}, 10^{10.88}] M_\odot h^{-1}$.

We neglect the dependence of the subhalo mass function on the lens redshift and mass. For both populations of haloes, the suppression in the number density at the low-mass end is calculated using equation (1.11). While this suppression relative to the CDM case tends to be stronger in the case of field haloes than in subhaloes, we ignore this effect in this chapter for simplicity (Lovell 2020).

Following previous works, we neglect the effect of tidal stripping and dark matter physics on the density profile and concentration of subhaloes and field haloes, as these have only a small effect on the number of detectable objects (of the order of 3 to 10 per cent; Despali et al. 2018). Both halo populations are assumed to have spherical NFW profiles (Navarro et al. 1996). We discuss the effect of these assumptions in Section 2.4. In Table 2.1, we summarize all of the relevant parameters, together with the corresponding priors used in this chapter and previous analyses. For the lensing analyses, we adopt the cosmology inferred by the Planck mission (Ade et al. 2014).

We consider the re-analyses of Vegetti et al. (2010, 2014), who studied a subsample of eleven gravitational lens systems from the SLACS survey (Bolton et al. 2008). Using the Bayesian gravitational imaging technique developed by Vegetti & Koopmans (2009b), only one low-mass subhalo was detected in the sample. Assuming a Pseudo-Jaffe profile, this subhalo shows an inferred mass of $3.5 \times 10^9 M_\odot$ ($\sim 10^{10} M_\odot$ for an NFW profile). Taking this detection and the non-detections into account, they constrained the subhalo mass function to be CDM. The lenses in this sample have a mean redshift of $z = 0.2$, while the background sources have a mean redshift of $z = 0.6$. In the remaining part of this chapter, we refer to this sample as the low-redshift sample.

The background source galaxies were modelled in a free-form fashion with a Delaunay mesh, while the foreground lenses were assumed to have an elliptical power-law mass density profile plus the contribution of an external shear component. Additional forms of complexity in the lenses not captured by the smooth power-law (including subhaloes) were identified using linear free-form corrections to the lensing potential. The statistical relevance of both detections and non-detections is determined via the sensitivity function. This function considers the Bayes factor between models with no substructure and those with a single substructure. A logarithmic Bayes factor of 50 provided a robust criterion to discriminate between reliable and non-reliable detections. Originally this description assumed a Pseudo-Jaffe parametric profile for the perturber (Vegetti et al. 2014). The analysis by Vegetti et al. (2014) was then extended by Vegetti et al. (2018, hereafter V18) to include the contribution of low-mass field haloes (i.e. haloes located along the line of sight; see Despali et al. 2018, Li et al. 2017) assuming an NFW profile for both populations of low-mass haloes. Furthermore, the effects of dark matter free streaming on the halo and subhalo mass functions were included via equation (1.11).

Ritondale et al. (2019, hereafter R19) have modelled a sample of seventeen gravitational lens systems from the BELLS-GALLERY survey (Shu et al. 2016), and reported zero detections of subhaloes and line-of-sight haloes. The mean redshift of the foreground lenses is $z \sim 0.5$, while the background source redshifts vary from $z = 2.1$ to 2.8. We refer to this sample as the high-redshift sample. The analysis by R19 used a more recent version of the Vegetti & Koopmans (2009b) lens modelling code that allows for a simultaneous inference

on the lens galaxy mass and light distribution. As in the original method, the source surface brightness distribution and low-mass haloes are defined on a grid of pixels. The calculation of the sensitivity function and the inference on the mass functions were performed in terms of the spherical NFW virial mass.

Here, we rerun the analyses of V18 and R19, while extending their prior ranges on the half-mode mass to $M_{\rm hm} \in [10^{-6}, 10^{14}]$ $\mathrm{M}_\odot\ \mathrm{h}^{-1}$.

2.2.2 Lyman-α forest

The second astrophysical probe that we consider is the analysis of high-quality optical spectra from Lyman-α emitting quasars at high redshifts (see Chapter 1 , Section 1.1.5).

We follow the approach of Murgia et al. (2018, hereafter M18), who obtain constraints on dark matter models by performing a Monte Carlo Markov Chain (MCMC) analysis of the full parameter space affecting the flux power spectrum $P_{\rm f}(k)$ reconstructed from high-redshift Lyman-α forest observations. We rerun the analysis of M18 changing two main elements; first, the results presented here are restricted to thermal relic warm dark matter, for which we choose a log-uniform prior on the particle mass, $m_{\rm WDM}$. Second, we extend the ranges of some model parameters, which are discussed below.

The data is provided by the high-resolution and high-redshift quasar samples from the HIRES/Keck and the MIKE/Magellan spectrographs (Viel et al. 2013). These samples include redshift bins of $z = 4.2$, 4.6, 5.0 and 5.4 over 10 wavenumber bins in the interval $k \in [0.001, 0.08]$ s km^{-1} (the range relevant for Lyman-α forest data). The spectral resolution of the HIRES and MIKE spectrographic data are 6.7 and 13.6 km s^{-1}, respectively. As in previous works, such as Viel et al. (2013), only the measurements with $k > 0.005$ s km^{-1} have been used to avoid systematic uncertainties on large scales due to continuum fitting. The highest redshift bin for the MIKE data has been excluded due to the large uncertainties in the spectra at that epoch (see Viel et al. 2013 for further details). A total of 49 data points in wavenumber k and redshift z are used in the analysis.

M18 determined the changes in the flux power spectra as a function of different model parameters by interpolation of (computationally expensive) realistically simulated mock spectra, which are generated for different astrophysical and cosmological parameters defined on a grid. This procedure allowed M18 to define a likelihood as a function of these parameters. The grid of mock simulations considers several values of the cosmological parameters and follows the approach of Iršič et al. (2017) to recover their effects on the likelihood. M18 considered five different values for the normalization of the linear matter power spectrum $\sigma_8 \in [0.754, 0.904]$ and its slope $n_{\rm eff} \in [-2.3474, -2.2674]$ (both defined on the typical scale probed by the Lyman-α forest of 0.009 s km^{-1}). For the rerun of this analysis, we consider ten ΛWDM simulations that correspond to thermal WDM masses of $m_{\rm wdm} \in [1, 10]$ keV, linearly spaced in steps of 1 keV.

Concerning the astrophysical parameters, the thermal history of the IGM is varied in the form of the amplitude T_0 and the slope γ of its temperature-density relation. This relation is parametrized as $T = T_0(1 + \delta_{\rm IGM})^{\gamma-1}$, with $\delta_{\rm IGM}$ being the overdensity of the IGM (Hui &

Gnedin 1997). Three different temperatures at mean density, $T_0(z = 4.2) = 6000$, 9200 and 12600 K, and three values for the slope of the temperature-density relation, $\gamma(z = 4.2) = 0.88$, 1.24 and 1.47, are considered here. The reference thermal history is defined by $T_0(z = 4.2) = 9200$ K and $\gamma(z = 4.2) = 1.47$ (see Bolton et al. 2016). The redshift evolution of $\gamma(z)$ is assumed to be a power law, that is, $\gamma(z) = \gamma^A[(1+z)/(1+z_p)]^{\gamma^S}$, where the pivot redshift z_p is the redshift at which most of the Lyman-α forest pixels originate ($z_p = 4.5$ for the MIKE and HIRES datasets).

M18 also considered three different redshift values of instantaneous reionization at $z_{\rm reion} \in \{7, 9, 15\}$, as well as ultraviolet (UV) fluctuations of the ionizing background of $f_{\rm UV} \in \{0, 0.5, 1\}$, where the case of $f_{\rm UV} = 0$ corresponds to a spatially uniform UV background. Nine values of the relative mean flux were considered, that is, $< F(z) > / < F_{\rm REF} > \in [0.6, 1.4]$ in linearly spaced intervals with steps of 0.1. The reference values $< F_{\rm REF} >$ are taken from the Sloan Digital Sky Survey (SDSS), i.e. the Baryon Oscillation Spectroscopic Survey (BOSS) measurements, which are part of SDSS-III (Anderson et al. 2014). Eight additional values of $< F(z) > / < F_{\rm REF} >$ are obtained by rescaling the optical depth $\tau = -\ln < F >$ (see M18).

For each of the resulting grid points in parameter space, hydrodynamical simulations are used to generate the mock spectra. All simulations are performed with `GADGET-3`, a modified version of the publicly available `GADGET-2` code (Springel 2005, Springel et al. 2001b). As in Iršič et al. (2017), the reference model simulation has a box length of 20 Mpc h^{-1} (comoving) with 2×768^3 gas and CDM particles (with gravitational softening lengths of 1.04 kpc h^{-1} comoving) in a flat ΛCDM Universe. The cosmological parameters are set to $\Omega_{\rm m} = 0.301$, $\Omega_{\rm b} = 0.0457$, $n_s = 0.961$, $H_0 = 70.2$ km s^{-1} Mpc^{-1} and $\sigma_8 = 0.829$ (Ade et al. 2016).

An Ordinary-Kriging scheme is used for the interpolation between different grid points and linearly extrapolated when necessary (Webster & Oliver 2007). The interpolation with respect to all the parameters happens in consecutive steps, first over the astrophysical and cosmological parameters, then over the different WDM models. This interpolation is then used to define a likelihood, which in return produces a posterior (e.g. Archidiacono et al. 2019). Table 2.2 gives a short overview of the model parameters, their ranges and (prior) probabilities. In this chapter, we replace the original prior on the WDM particle mass with a log-uniform prior in order to match the priors of the other probes considered here.

2.2.3 Milky Way luminous satellites

Our final astrophysical probe comes from the observed luminosity function of the satellite galaxies of the Milky Way. This method was the first to identify a potential challenge to the CDM model through the paucity of observed dwarf galaxies around our own galaxy (Kauffmann & White 1993, Klypin et al. 1999, Moore et al. 1999). Observational solutions, such as the lack of sufficient sky coverage/completeness (Koposov et al. 2008, Tollerud et al. 2008, Hargis et al. 2014, Jethwa et al. 2017), more realistic galaxy formation models (e.g. stellar feedback, Bullock et al. 2000, Benson et al. 2002, Somerville 2002) or revisions to the dark matter model (Bode et al. 2001, Green et al. 2004, Schneider et al. 2012) have since been proposed to reconcile this discrepancy. Here, we consider a new analysis of the number

Table 2.2: **The model parameters of the Lyman-α forest analysis.** For each parameter, we provide the ranges and its prior probability distributions, as used in this chapter. $^{(*)}$This prior is the same as the one described by Iršič et al. (2017).

Parameter	Range	Probability
$1/m_{\rm wdm}$ [keV^{-1}]	[0, 1]	Log-flat
$< F(z) > / < F_{\rm ref} >$	$(-\infty, \infty)$	Gaussian$^{(*)}$
$T_0(z)$ [10^4K]	[0, 2]	Flat
$\widetilde{\gamma}(z)$	[1, 1.7]	Flat
σ_8	[0.5, 1.5]	Flat
$z_{\rm reion}$	[7, 15]	Flat
$n_{\rm eff}$	[-2.6,-2.0]	Flat
$f_{\rm UV}$	[0,1]	Flat

density of luminous satellites that has been carried out by Newton et al. (2020).

Their approach assesses the viability of a given WDM model by comparing the predictions on the abundance of satellite galaxies within a Milky Way-mass halo for various dark matter models with the total satellite galaxy population inferred from those observed in the Milky Way. WDM scenarios that do not produce enough faint galaxies to be consistent with the Milky Way satellite population are ruled out with high confidence. As the current census of Milky Way satellites is incomplete, Newton et al. (2020) infer the total satellite galaxy population from observations, using a Bayesian formalism that was developed and tested robustly by Newton et al. (2018b, N18, hereafter). They use data from the SDSS and DES, as summarized in Table A1 of N18 (compiled from Watkins et al. 2009, McConnachie 2012, Drlica-Wagner et al. 2015, Kim et al. 2015, Koposov et al. 2015, Jethwa et al. 2016, Kim et al. 2016, Walker et al. 2016, Carlin et al. 2017, Li et al. 2017). The results obtained in a subsequent follow-up study of some of these objects and more recent discoveries of dwarf galaxy candidates are not incorporated into our analysis. However, it is unlikely that their inclusion would change the inferred population outside the uncertainties quoted in N18.

For each survey, the assumed observational selection function significantly affects the size of the total satellite population inferred from the observations. In particular, if the selection function overpredicts the completeness of faint objects in the survey, then the size of the inferred satellite population will be too small. While the SDSS selection function has been studied extensively and is now well-characterized (e.g. Walsh et al. 2009), no such study had been carried out for the DES before 2019. Therefore, N18 used the approximation proposed by Jethwa et al. (2016), one of the few estimates available for the DES at the time. Recently, the DES selection function was characterized in detail by Drlica-Wagner et al. (2020), and this was used in a follow-up study by Nadler et al. (2020b) to infer the total satellite population. Their results are consistent with N18, which suggests that the Jethwa et al. (2016) approximation of the selection function was reasonable.

The second ingredient of the analysis by Newton et al. (2020) is a set of estimates of the number of satellite galaxies formed in Milky Way-mass WDM haloes. They explore two ap-

proaches to obtain these predictions. In the first, they use the Extended Press-Schechter (EPS) formalism (Press & Schechter 1974, Bond et al. 1991, Bower 1991, Lacey & Cole 1993, Parkinson et al. 2008) and follow the approach of Kennedy et al. (2014), Schneider (2015), Lovell et al. (2016). Implicit in this technique is the assumption that all dark matter haloes form a galaxy, which allows Newton et al. to place a highly robust lower bound on the mass of the warm dark matter particle independently of assumptions about galaxy formation physics. However, the faint end of the satellite galaxy luminosity function is extremely sensitive to these processes, which prevent galaxies from forming in low-mass dark matter haloes. They are also complex and their details remain uncertain, permitting a large parameter space of viable descriptions of galaxy formation. In their second approach, Newton et al. (2020) use the GALFORM semi-analytic model of galaxy formation (Cole et al. 1994, 2000) to explore this space to understand how different parametrizations can affect the WDM constraints.

The main process affecting the Milky Way satellite galaxy luminosity function is the reionization of the Universe. In GALFORM, this is described by $z_{\rm reion}$, the redshift at which the intergalactic medium is fully ionized, after which the parameter $V_{\rm cut}$ prevents cooling of gas into haloes with circular velocities, $v_{\rm vir} < V_{\rm cut}$. Newton et al. (2020) assume the Lacey et al. (2016) version of galaxy formation and vary the reionization parameters in the ranges $6 \leq z_{\rm reion} \leq 8$ and $25{\rm km\,s^{-1}} \leq V_{\rm cut} \leq 35{\rm km\,s^{-1}}$, choosing a fiducial model with $z_{\rm reion} = 7$ and $V_{\rm cut} = 30{\rm km\,s^{-1}}$. From each GALFORM model they obtain Milky Way satellite galaxy luminosity functions which they compare with the luminosity functions inferred from observations (described above). They calculate the relative likelihood of a given model compared to the CDM case by convolving the probability density function of the number of satellites brighter than $M_{\rm V} = 0$ in a GALFORM WDM Milky Way halo with the cumulative distribution function of the inferred population of Milky Way satellites, which, according to N18, numbers 124^{+40}_{-27}. We use this approach in the analysis that follows.

2.2.4 Model consistency

The goal of this chapter is to derive joint constraints on the particle mass of a thermal relic dark matter model. However, $m_{\rm th}$ is not directly observable but is inferred from the different probes under different assumptions, as described above. In this section, we discuss the main differences between the three methods and how these can be treated to derive a meaningful joint inference on $m_{\rm th}$.

All the methods employed here constrain parameters describing the halo mass function; however, these parameters may differ in their meaning. The main parameter constrained by strong gravitational lensing observations is the half-mode mass $M_{\rm hm}$, while the analysis of the Lyman-α forest and the luminosity function of the Milky Way satellites are expressed directly in terms of $m_{\rm th}$. Converting from one to the other requires some assumptions about the physics of the dark matter particles (e.g. their type and production mechanism; see Section 1.1) and the cosmological parameters. As each of the considered analyses has adopted different cosmologies, we first express our inference in terms of the half-mode scale $\lambda_{\rm hm}$, which is less sensitive to the specific values of the cosmological parameters.

For all the different probes, we assume a uniform prior in $\lambda_{\rm hm}$ with the lowest possible value corresponding to the WIMP CDM model, that is, $M_{\rm hm}(\lambda_{\rm hm}^{\rm min}) = 10^{-6}$ $\mathrm{M}_\odot\,\mathrm{h}^{-1}$, and the upper limit $M_{\rm hm}(\lambda_{\rm hm}^{\rm max}) = 10^{14}$ $\mathrm{M}_\odot\,\mathrm{h}^{-1}$, corresponding to the lower limit of a thermal relic WDM particle mass of $m_{\rm th} = 0.07$ keV as constrained by Kunz et al. (2016) using observations of the CMB (Aghanim et al. 2016). We then express our results in terms of the half-mode and thermal relic particle masses, converting all results so that they adopt Planck cosmology (i.e. $\Omega_{\rm th} = 0.26$ and $\mathrm{h} = 0.68$; Ade et al. 2016) and the assumptions on the dark matter particles described in Section 1.1.4. Notice that logarithmic quantities $\log_{10}(M_{\rm hm})$ and $\log_{10}(m_{\rm th})$ are related to $\log_{10}(\lambda_{\rm hm})$ via linear transformations, so that the prior is flat in all of these parameters.

Another source of discrepancy is related to the different treatments of low-mass haloes. The strong gravitational lensing analyses by V18 and R19 made direct use of the subhalo and field halo mass functions (see Section 2.2.1) and expressed in terms of the virial mass of a spherical NFW profile. The Milky Way satellite analysis of N18 considers the number of subhaloes independent of their mass. Differences in the mass functions with respect to other astrophysical probes will, therefore, not lead to a systematic bias, as long as the analyses of these other probes consider mass functions that are in agreement with cosmological simulations.

Despali et al. (2018) have shown that the lensing effect of an NFW subhalo of a given mass $M_{\rm vir}^{\rm NFW}$ is a good approximation to a subhalo of equivalent mass found with the SUBFIND algorithm (Springel et al. 2001a) in cosmological simulations, indicating that our treatment of subhalo masses is consistent. Moreover, as the expressions (1.11) and (2.1) were also derived from numerical simulations, we expect only minor systematic uncertainty arising from the limited accuracy of fitting functions to non-CDM mass functions. The Lyman-α forest constraints obtained by M18 are expressed in terms of the matter power spectrum and are therefore not affected by any specific definition of halo masses. Moreover, as they have also been calibrated on numerical simulations, we assume that there is no strong discrepancy between the mass function that would be associated to their power spectrum and that used by the other two probes. Given these considerations, we conclude that any discrepancies in the definition of masses and mass functions are negligible and can be ignored.

2.3 Results

In this section, we present our main results on the half-mode scale and mass, and the thermal relic particle mass. We present the constraints from each of the individual probes as well as those of the joint statistical analysis. Our statistical summaries are presented in Table 2.3, which can be compared with previous results that we report in Table 2.4.

Table 2.3: **The posterior limits according to the 95 percentile criterium (section 2.3.2) and the Bayes factor (BF, section 2.3.2), probabilities of (in-)sensitivity, odds of being sensitive, ratio between the likelihoods of the maximum a likelihood half-mode mass $\lambda_{\rm hm}^{\rm ML}$ and that of the insensitive region.**

Reference	$\frac{m_{\rm th}}{\rm keV}$		$\frac{M_{\rm hm}}{10^{10}\ {\rm M}_\odot\ {\rm h}^{-1}}$		$\frac{\lambda_{\rm hm}}{\rm Mpc\ h^{-1}}$		$\mathcal{P}(\bar{S}\|d)$	$\mathcal{P}(S\|d)$	$\frac{\mathcal{P}(\bar{S}\|d)}{\mathcal{P}(S\|d)}$	$\frac{\mathcal{P}(d\|\lambda_{\rm hm}^{\rm ML})}{\mathcal{P}(d\|\lambda_{\rm hm}\in\bar{S})}$	$\frac{\lambda_{\rm hm}^{\rm ML}}{\rm Mpc\ h^{-1}}$
	BF	95% c.l.	BF	95% c.l.	BF	95% c.l.	%	%	1	1	1
V18	0.216	0.576	178.366	6.780	3.607	1.214	47.52	52.48	0.91	3.35	0.470
R19	-	0.121	-	1219.752	-	6.842	53.06	46.94	1.13	1.14	4.538
M18	1.197	3.571	0.594	0.016	0.540	0.160	74.30	25.70	2.89	1.04	0.029
N20	2.678	6.989	0.041	0.002	0.221	0.076	79.46	20.54	3.87	1.01	0.016
Joint	2.552	6.048	0.048	0.003	0.233	0.089	77.68	22.32	3.48	1.08	0.027

Table 2.4: **A summary of the lower limits reported on the thermal dark matter particle masses for a selection of past studies.** Note that additional model assumptions and assumed parameter ranges can widely differ. We include multiple limits in this table when they are derived for different assumptions.

Reference	Probe	$\frac{m_{\rm dm}}{\rm keV}$ 95% c.l.
this chapter	see Section 2.2	6.048
Birrer et al. (2017)	Grav. Imaging	2.0
V18 (Original)	Grav. Imaging	0.3
R19 (Original)	Grav. Imaging	0.26
Gilman et al. (2019b)	Flux Ratios	3.1, 4.4
Gilman et al. (2019a)	Flux Ratios	5.2
Hsueh et al. (2019)	Flux Ratios	5.6
Banik et al. (2018, 2019b)	Stellar streams	4.6, 6.3
Viel et al. (2005)	Lyman-α	0.55
Viel et al. (2006)	Lyman-α	2.0
Viel et al. (2013)	Lyman-α	3.3, 4.5
Seljak et al. (2006)	Lyman-α	2.5
Baur et al. (2016)	Lyman-α	4.09
Iršič et al. (2017)	Lyman-α	3.5, 5.3
M18 (Original)	Lyman-α	2.7, 3.6
Polisensky & Ricotti (2011)	Milky Way satellites	2.3
Kennedy et al. (2014)	Milky Way satellites	1.3, 5.0
Jethwa et al. (2017)	Milky Way satellites	2.9
Nadler et al. (2019)	Milky Way satellites	3.26
Nadler et al. (2020a)	Milky Way satellites	6.5
Newton et al. (2020) (Original)	Milky Way satellites	2.02, 3.99

2.3.1 Posterior distributions

Figure 2.1 shows the individual and the joint posterior on the half-mode scale, half-mode mass and thermal relic particle mass for each of the astrophysical probes considered here. Each of the posteriors is scaled so that its maximum value is equal to 1.

The shape of the joint posterior is mostly determined by the constraints coming from the Lyman-α forest (M18) and the luminous satellites in the Milky Way galaxy (N18) analyses, which - although being completely independent measurements - find a similar upper limit on the half-mode scale. Further data and more rigorous analyses may reveal differences between their respective constraining power in the future. Both posteriors are roughly shaped like sigmoid functions, reflecting that these probes reject those dark matter models that are too warm to explain the respective observations. As the analysis of V18 includes the detection of a relatively massive subhalo (Vegetti et al. 2010), it rules out slightly higher

values of $\lambda_{\rm hm}$ (as well as $M_{\rm hm}$), than the other probes. In contrast, the constraints from the BELLS-GALLERY sample turn out to be rather weak. As R19 reported no significant detections, the resulting posterior slightly prefers warmer dark matter models that predict a smaller number of (sub-) haloes. This may also be related to the sensitivities and the source redshifts of these lenses. The low-redshift SLACS sample has a higher sensitivity than the high-redshift BELLS-GALLERY. While this higher sensitivity allows the SLACS sample to detect objects with smaller masses, the high redshift sources of the BELLS-GALLERY allow to probe larger volumes, which should increase the expected number of line-of-sight objects. In combination with the high number of non-detections of low-mass haloes, the analysis of the BELLS-GALLERY provides slightly stronger statistical weight to large half-mode scales.

2.3.2 Statistical summaries

It is a common practice to report summary statistics of the posterior functions to characterise the strength of constraints on warm dark matter. One of the most reported quantities is the 95 percentile. However this comes with a caveat: the values of percentiles are strongly dependent on the specific choice of the lower limit of the model parameter range, since the likelihood (and posterior) functions become essentially flat for $\lambda_{\rm hm} < 0.013$ Mpc $\rm h^{-1}$ (corresponding to $M_{\rm hm} < 10^{5.0}$ $\rm M_\odot$ $\rm h^{-1}$). This flattening reflects a lack of sensitivity on these scales, i.e. that the analyses considered in this chapter are no longer capable of distinguishing between models of different half-mode scales.

In the posteriors shown in Figure 2.1, we choose a lower-limit of $\lambda_{\rm hm} = 3 \times 10^{-6}$ Mpc $\rm h^{-1}$ which corresponds to a WIMP CDM model (Schneider et al. 2013). We chose this limit mainly because a log uniform prior gives rise to a diverging posterior if we extend the inference to the idealised CDM case of $\lambda_{\rm hm} = 0$. However, it could be argued that even though our choice of lower limit in the parameter range is physically motivated, it arbitrarily excludes models that lie between the WIMP and the idealised case.

To account for some of the uncertainties in these a priori choices, we report two statistical summaries: one equivalent to the 95 percentiles within a rephrased version of the inference problem; the other based on the ratio of likelihoods and therefore, more independent of the chosen lower limit for $\lambda_{\rm hm}$ (and its prior). Notice that this does not affect our main conclusions, but only accommodates for different preferences in the way that posteriors are summarized.

95 percentiles

For the first summary, we rephrase our inference problem in terms of a hyper-model scenario with two models corresponding to an insensitive ($\bar{S}$) and a sensitive (S) region, respectively. In particular, we define the former as the range of half-mode masses $M_{\rm hm} \in \left[0, 10^5\right]$ $\rm M_\odot$ $\rm h^{-1}$ and the latter as $M_{\rm hm} \in \left[10^5, 10^{12}\right]$ $\rm M_\odot$ $\rm h^{-1}$. We know that the likelihood in the two regions is then defined as follows:

$$\mathcal{P}\left(d\middle|X\right) = \begin{cases} constant & \text{if } X = \bar{S}\,, \\ \int_S dM_{\rm hm} \mathcal{P}\left(d\middle|M_{\rm hm}\right) \times \mathcal{P}\left(M_{\rm hm}\middle|S\right) & \text{if } X = S\,. \end{cases} \tag{2.5}$$

We choose a log uniform prior distribution $\mathcal{P}\left(M_{\rm hm}\middle|S\right)$ on $M_{\rm hm}$ within S, which corresponds to a prior that is non-informative about the order of magnitude of the half-mode mass. We obtain the *constant* and $\mathcal{P}\left(d\middle|M_{\rm hm}\right)$ by dividing the posterior of the original analysis by its prior. We further enforce that all probabilities add up to 1 in the posterior in order to obtain the correct normalization.

This framework allows us to include the idealised CDM case while maintaining the log uniform prior regarding the sensitive region. It comes, however, at the small cost that we can only report an upper limit in the case that it happens to fall within the sensitive region. Our first summary is the 95 percentile of the posterior in this hyper model scenario, $M_{\rm hm}^{\rm CL}$, which is defined according to:

$$0.95 = \mathcal{P}(\bar{S}|d) + \int_{10^5\ {\rm M_\odot h^{-1}}}^{M_{\rm hm}^{\rm CL}} dM_{\rm hm}\, \mathcal{P}\left(d\middle|M_{\rm hm}\right) \times \mathcal{P}\left(M_{\rm hm}\middle|S\right)\frac{\mathcal{P}(S)}{\mathcal{P}(d)}\,. \tag{2.6}$$

In the equation above, $\mathcal{P}(S)$ is the prior probability of the sensitive case. The original parameter range contains all half-mode masses between the one corresponding to the coldest WIMP model and the constraints from the CMB. For a log uniform prior on the half-mode masses, this corresponds to $\mathcal{P}(S) = 1 - \mathcal{P}(\bar{S}) = 0.45$. We use this prior when reporting upper limits in this section for simplicity, but in general, one could choose different prior values. In Figure 2.2, we show how the 95 percentiles on the half-mode scale change as a function of prior mass attributed to the sensitive region. We find that the order of magnitude of these 95 percentiles is stable for values of $\mathcal{P}(S)$ between 0.5 and 1.0.

Following this approach we find a joint upper limit of $\lambda_{\rm hm}^{\rm CL} = 0.089$ Mpc h^{-1}. This rules out that haloes with a mass of $M_{\rm hm}^{\rm CL} = 3 \times 10^7$ M$_\odot$ h^{-1} are significantly suppressed with respect to the CDM scenario at the 2-σ level. Under the assumptions discussed in Section 1.1.4, we can express our constraints in terms of a lower limit on the thermal relic particle mass, i.e. $m_{\rm th}^{\rm CL} = 6.048$ keV at the 95 per cent confidence level. We mark these limits with dashed vertical lines in Figure 2.1. These constraints are in agreement with those derived by previous studies, as summarized in Table 2.4. We find that we require a higher sensitivity towards lower halo masses in order to rule out or confirm CDM models. Notice that our model assumptions, for example, on the IGM priors in the Lyman-α analysis (see Section 2.4.2), are rather conservative. While we obtain mildly weaker limits with respect to past literature, our limits are expected to be more robust.

Bayes factors

In order to be less dependent on the chosen parameter range and prior assumptions, the second summary statistic considers the ratio of likelihood with a model $\lambda_{\rm hm}$ and the model that maximises the likelihood $\lambda_{\rm hm}^{\rm ML}$ (corresponding to the Bayes factor between these two models, when each parameter value is considered to be different model). The value $\lambda_{\rm hm}^{\rm BF}$, above which the ratio of all models fullfill $\frac{\mathcal{P}(d|\lambda_{\rm hm}>\lambda_{\rm hm}^{\rm BF})}{\mathcal{P}(d|\lambda_{\rm hm}^{\rm ML})} \leq \frac{1}{20}$ gives then an upper limit in the

sense that all these models are strongly disfavoured (i.e. ruled out at 95% confidence limit) in comparison to the maximum likelihood case. We mark these upper limits with solid vertical lines in Figure 2.1.

We find for the joint posterior an upper limits of $\lambda_{\rm hm}^{\rm BF} = 0.233$ Mpc $\rm h^{-1}$, corresponding to $M_{\rm hm}^{\rm BF} = 4.8 \times 10^8$ $\rm M_\odot$ $\rm h^{-1}$ and $m_{\rm hm} = 2.552$ keV. This upper limit is mostly determined by the analysis of Milky Way satellites analysis, with $\lambda_{\rm hm}^{\rm BF} = 0.221$ Mpc $\rm h^{-1}$. The Lyman alpha forest, with $\lambda_{\rm hm}^{\rm BF} = 0.540$ Mpc $\rm h^{-1}$, turns out to be the second strongest constraint. We find that for the lensing probes only the SLACS sample exclude values according to this summary criterium, with $\lambda_{\rm hm} = 3.607$ Mpc $\rm h^{-1}$. In the case of the BELLS-GALLERY, the posteriors actually prefer warmer dark matter models. This is reflected in the ratio between the maximum likelihood value and the likelihood of the cold limit, which is 1/1.14 at $\lambda_{\rm hm}^{\rm ML} = 4.538$ Mpc $\rm h^{-1}$ for R19, respectively. We summarize the different results in Table 2.3, which furthermore gives additional information about the individual probes.

2.4 Systematic errors

In this section, we discuss the different sources of systematic errors that may affect each of the astrophysical probes considered here.

2.4.1 Strong gravitational lensing

The main sources of systematic errors that are common to strong gravitational lensing techniques are related to the assumptions on the mass density profile of the main lenses and their subhaloes, and the normalization of the halo mass function.

In the context of strong gravitational lensing by galaxies, the standard procedure is to parameterise the mass distribution of the lens with an elliptical power-law profile and a contribution of an external shear component. However, both numerical simulations (Xu et al. 2015, Hsueh et al. 2018) and observations (Gilman et al. 2017, Hsueh et al. 2016, 2017, Xu et al. 2013) demonstrated that for the analysis of lensed quasars, in some cases, important departures from this simplified model exist and have a non-negligible effect on the inference of low-mass haloes. For example, Hsueh et al. (2018) showed that the presence of an additional disc component could increase the probability of finding significant flux-ratio anomalies by 10 to 20 per cent, while baryonic structures in early-type galaxies lead to an increase of the order of 8 per cent. Similar effects are expected for departures in the mass distribution from a power-law in early-type galaxies in the analysis of extended sources (R19). However, for different conclusion see Enzi et al. (2020). Both V18 and R19 explicitly avoid this problem by including pixellated corrections to the lensing potential. These corrections are used to detect low-mass haloes and to distinguish them from other forms of complexity, i.e. they can also be used to account for large-scale deviations from the assumed elliptical power-law mass model (Vegetti et al. 2014).

Another common assumption is that the subhaloes are well described by spherical NFW

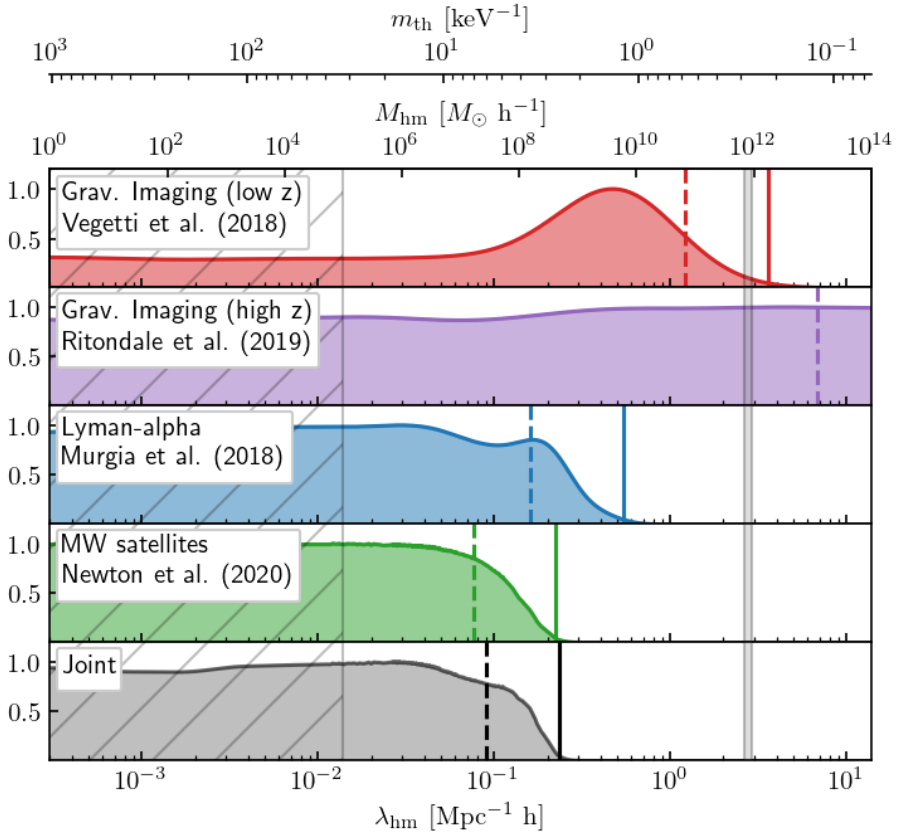

Figure 2.1: **Posterior probability distributions on the half-mode scale $\lambda_{\rm hm}$, half-mode mass $M_{\rm hm}$ and thermal relic particle mass $m_{\rm th}$.** Shown are the results of the analyses of gravitational lensing observations of extended arcs for the SLACS sample (red) and the BELLS sample (purple), the Lyman-α forest data (blue) and the luminous satellites of the Milky Way (green). All posteriors are scaled so that their maximum value is 1. The grey hatched area highlights the region in which all of the probes considered here become insensitive to the difference between the different models. The mass of the Milky Way within the 68 percent confidence interval, as inferred by Callingham et al. (2019), is shown with a grey line at $M_{\rm hm} \approx 10^{12}{\rm M}_\odot \approx M_{200}^{\rm MW}$. The vertical solid (dashed) lines indicate the upper limits determined from the Bayes factor (the 95 percentile) criterium.

profiles with the same mean concentration-mass relation derived from CDM simulations. However, due to non-linear effects, such as tidal stripping by the host halo, one expects their profile, and in particular, their concentration, to change as a function of distance from the host centre (Moliné et al. 2017). Moreover, as structure formation is delayed in WDM models, it is also expected that the concentration-redshift-mass relation will be different than in CDM. Despali et al. (2018) have explicitly quantified the effect of these assumptions and found it to be relatively small, underestimating the subhalo mass by, at most, 20 per cent. As a result, one sees a shift of the overall subhalo mass function that is small and within the intrinsic uncertainty of subhalo finding algorithms (Onions et al. 2012). A different conclusion is obtained by Gilman et al. (2019a), who claim that the changes in concentration have a significante impact on the lensing effect of small mass haloes, although only a limited number of studies have derived relations between concentrations in CDM and WDM models (Bose et al. 2015, Ludlow et al. 2016, Bohr et al. 2021).

Both studies assume that the line-of-sight halo mass functions are normalized according to the mean normalization value of CDM numerical simulations. However, it is unclear whether lens galaxies preferentially reside in lines of sight that are systematically over-dense, which could bias the results towards colder dark matter models. Moreover, the typical line of sight for different WDM models may not be the same as for the CDM case. This effect is potentially problematic for the analyses by V18 and R19, as their samples of lenses are homogeneous and consist of massive early-type galaxies. In the case of the gravitational lenses found by SLACS, Treu et al. (2009) showed that the line of sight density is comparable to those of non-lensing early-type galaxies with a similar redshift and mass.

2.4.2 Lyman-α forest

Here, we address the systematics affecting the analysis of the Lyman-α forest data by summarizing the discussion presented by Viel et al. (2013).

One of the potential systematics arises from the box size and particle number of the numerical simulations used for the model comparison, for which different setups usually show deviations at the 5 to 15 per cent level. M18 corrected for this effect in their analysis by comparing their simulations with those from standard cosmological simulations.

On small scales, the quasar spectra are influenced by the instrumental resolution, which in the case of the MIKE and HIRES data sets are at most on the level of 20 and 5 per cent, respectively. This uncertainty is independent of the redshift. The uncertainties arising from the signal-to-noise ratio of the spectra on the smallest scales vary from around 2 per cent at $z \leq 5$ to 7 per cent for the highest redshift bin. UV fluctuations in the spectra have been implemented using a rather extreme model that only takes into account the ionizing effect of the quasars. The systematic effect on flux power spectra is expected to be ≤ 10 per cent for the scales considered here and is scale-dependent (Croft 2004, McDonald et al. 2006). An additional systematic associated with the quasar spectra is the contamination with metal lines in the Lyman-α forest. However, this is expected to add less than 1 per cent to the uncertainty of flux power spectra on all of the scales considered in the analysis.

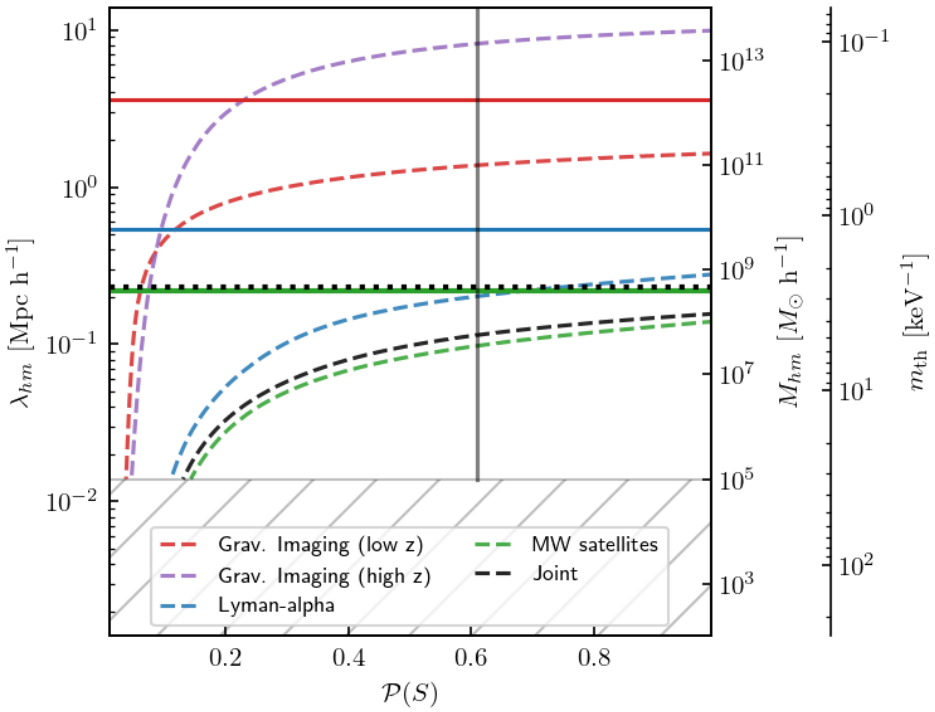

Figure 2.2: **The behaviour of 95 per cent upper c.l. (dashed curves) as a function of the prior mass attributed to the sensitive region.** The hatched area highlights the region in which none of the probes considered here is sensitive anymore. The vertical line shows the prior $\mathcal{P}(S)$ corresponding to the original box in which the analyses were performed (see Table 2.3). Notice that the order of magnitude of the 95 percentiles is stable over a large range of values. For reference we show the value of the upper limit according to the Bayes factor criterium for: the joint posterior (dotted black), the Lyman-α forest posterior (solid blue), the Milky Way satellites posterior (solid green), and the SLACS sample (solid red).

A well-known issue affecting Lyman-α forest analyses is the degeneracy between the small-scale impact of different WDM models, and the heating effects due to different thermal or reionization IGM histories (e.g., Garzilli et al. (2017, 2019). Unlike the IGM temperature, the WDM mass is a redshift-independent parameter. Thus, by simultaneously fitting power spectra at different redshift bins, one can partly break their degeneracy (M18). Furthermore, the limits presented in this chapter are obtained considering the IGM temperature $T_{\rm igm}(z)$ as a freely floating parameter, redshift bin by redshift bin. In other words, we did not make any assumption on its redshift-evolution, besides imposing $T_{\rm igm}(z) > 0$, and $\Delta_{\rm igm}(z) < 5000$ between adjacent redshift bins.

2.4.3 Milky Way luminous satellites

One of the major nuisance parameters affecting the constraints on dark matter obtained from the analysis of the luminous Milky Way satellites is the mass of the Milky Way, $M_{200}^{\rm MW}$. For this mass, N18 assume a value within the current observational constraints, such as those by Callingham et al. (2019) or Wang et al. (2019). Changes to the assumed Milky Way halo mass alter the number of subhaloes of a given mass that host a visible galaxy (see e.g. Sawala et al. 2017). For example, doubling the halo mass approximately doubles the number of subhaloes (Wang & White 2012, Cautun et al. 2014).

The analysis by Newton et al. (2020) considered here models baryonic processes assuming the Lacey et al. (2016) GALFORM model and a fiducial parametrization of reionization. Although the full details of these mechanisms have not been characterized fully, they have been constrained reasonably well. The Lacey et al. (2016) model was calibrated extensively by comparison with the properties of the local galaxy population. The fiducial reionization parameters that Newton et al. (2020) chose are at the centre of the ranges permitted by theoretical calculations and observational analyses. Imposing these choices produces more stringent constraints on the half-mode mass, $M_{\rm hm}$, than we could have obtained by considering only the abundance of dark matter subhaloes in Milky Way-mass haloes. However, the additional uncertainty this introduces, particularly in the value of $z_{\rm reion}$, weakens the robustness of the constraint. If the redshift at which the Universe is fully reionized is later than the fiducial choice, the upper limits on $M_{\rm hm}$ will move towards warmer models than we report here.

An important systematic associated with this astrophysical probe is the choice of the observed satellite population. Half of the non-classical satellites in the sample are drawn from the SDSS and have been spectroscopically confirmed as dark-matter-dominated dwarf galaxies. Newton et al. (2020) draw the other half from the DES, only 25 per cent of which are spectroscopically-confirmed. If later work reclassifies some of the DES objects to be globular clusters, then the inferred total satellite count will decrease for faint objects. However, this effect is likely to be small due to the good agreement in the inferred Milky Way satellite luminosity function when using only SDSS or only DES observations.

In their analysis, Newton et al. (2020) assume that the Milky way and its satellite system are typical examples of most dark matter haloes with similar masses. If this is not the case, for example, due to environmental effects, one expects that this would affect their results.

M31, for example, could introduce anisotropies into the Milky Way subhalo distribution. In general, if the radial distribution of subhaloes in simulations is different from the distribution within the Milky Way, it can lead to systematic uncertainties. Anisotropies would give rise to a correlation between satellites. N18 briefly studied the effects of anisotropy in the subhalo distribution and chose 300 kpc as their fiducial radius (smaller than the distance between Milky Way and M31) to minimize the effects from interactions with M31.

2.5 Future prospects

In this section, we discuss how the current constraints from the three different probes are likely to improve in the near future, and which steps will be necessary to obtain a more precise measurement on dark matter.

2.5.1 Strong gravitational lensing

The level of constraints currently obtainable with strong gravitational lensing is mainly determined by the low number of known systems, in particular at high-redshift (i.e. those for which the line of sight contribution is maximal). Moreover, in the particular case of extended sources, the lack of high-angular-resolution data strongly limits the possibility of detecting haloes with masses below 10^8 $M_\odot$. This hinders the exploration of the region of the parameter space, where the difference between different dark matter models is the largest.

Ongoing and upcoming surveys are expected to lead to the discovery of a large number of new gravitational lens systems. *Euclid*, for example, is expected to deliver as many as $O(10^5)$ new lensed galaxies (Collett 2015), while $O(10^3)$ lensed quasars are expected to be found in future surveys of the Vera C. Rubin Observatory, formerly known as the Large Synoptic Survey Telescope (LSST, Oguri & Marshall 2010). However, these new samples on their own will not be sufficient to significantly and robustly improve upon the present constraints. In particular, the gravitational imaging approach will require high-resolution follow-up observations to probe halo masses below the current limits. As the expected angular resolution of *Euclid* is about two times worse than currently available with the Hubble Space Telescope (HST) and about four times worse than what is already provided by current adaptive optics systems, these observations will only allow us to probe the halo mass function in a regime where predictions from different thermal relic dark matter models are essentially the same.

Follow-up observations will come from extremely large telescopes (ELTs), such as the Thirty Meter Telescope (TMT), the Giant Magellan Telescope (GMT), and the European Extremely Large Telescope (EELT), as well as VLBI observations at cm to mm-wavelengths, which will provide an angular resolution of the order of $\sim$ 0.2 to 5 milli-arcseconds. This will open up the possibility of detecting haloes with masses as low as 10^6 $M_\odot$ (McKean et al. 2015, Spingola et al. 2018). Furthermore, the JWST will not only provide an angular resolution of $\sim$ 0.02 to 0.1 arcsec, but will also allow us to maximise the contribution from the line of sight haloes by targeting high-redshift systems, and therefore, can potentially deliver

tighter constraints on the mass function in the mass ranges currently probed. In Chapter 4 of this thesis we will demonstrate, however, that JWST does not necessarily improve upon the current sensitivities to dark substructures.

We notice that flux-ratios of gravitationally lensed quasars also pose a very promising probe of dark matter. In order to take full advantage of their observations, deep follow-up imaging will be needed to quantify the frequency of galactic discs and other forms of complexity in the lens mass distribution, while long term monitoring will provide a robust measurement of the relative fluxes and possible variability in the lensed images (Harvey et al. 2019, Koopmans et al. 2003). It should also be considered that higher angular-resolution observations of such systems will allow one to resolve the extended source structure, and, therefore, permit an analysis using the gravitational imaging approach, which even though less sensitive is more robust.

With increasing resolution and sample sizes, fully understanding all sources of systematic errors will become increasingly important. To this end, high-resolution, realistic hydrodynamical simulations in different dark matter models will be required (e.g. Mukherjee et al. 2018, Enzi et al. 2020).

2.5.2 Lyman-α forest

In the near future, more accurate measurements of the IGM thermal history will provide stronger priors for this data analysis, allowing us to better constrain the small-scale cut-off in the linear power spectrum (see e.g. Boera et al. 2019).

Furthermore, the inclusion of the set of intermediate resolution and signal-to-noise quasar spectra observed by the XQ-100 survey (López et al. 2016), and of the new, high-resolution ones observed by the ESPRESSO spectrograph (Pepe et al. 2019), will improve the constraints presented in M18 and in this chapter, due both to improved large number statistics and to the complementary redshift and scale coverage, which will break some of the degeneracies among different parameters.

Another possible refinement might be achieved by including additional hydrodynamical simulations for which both astrophysical parameters, e.g., the IGM temperature, and WDM mass, are varied simultaneously. Constraints from the Lyman-α forest on small scales are indeed limited by the thermal cut-off in the flux power spectrum introduced by pressure and thermal motions of baryons in the ionised IGM. This makes the determination of accurate and independent constraints on the IGM thermal history essential in order to push current limits to even larger thermal relic masses. The 21 cm signal from neutral hydrogen gas before reionization could provide such an independent measurement (see e.g. Viel et al. 2013).

Concurrently with ongoing and future experimental efforts, further theoretical work is thus needed to interpret observations, accurately disentangle the impact of the various parameters, and combine outcomes from different observational methods.

2.5.3 Milky Way luminous satellites

There are two aspects of Local Group studies that are expected to improve in the future. The first relates to the theoretical predictions as simulations improve, and the second comes from the improving observational data as larger and deeper surveys are carried out, and new detection methods are developed.

Subhaloes in simulations can only be resolved above a certain particle number, which results in missing low-mass subhaloes. This issue can be approximately corrected for; however, future high-resolution simulations may lower the mass scale below which one needs to make these corrections. Also, next-generation simulations will assist attempts to understand better the relevant (baryonic) processes of satellite formation, potentially opening up the possibility to not just present an upper limit on $M_{\rm hm}$ from the abundance of Milky Way luminous satellites.

The method presented by N18 takes the observed satellites, which are found in surveys with various detectability limits, as a representative sample of the global population. However, there could be a population of faint and spatially extended dwarfs that are inaccessible to current surveys (see e.g. Torrealba et al. 2016a,b). The WDM constraints inferred from the satellite distribution could be improved further by deep observations of other nearby galaxies besides the Milky Way, such as M31, Centaurus A or the Virgo Cluster. Such external observations help to reduce uncertainties in the current analysis arising from the Milky Way halo mass and from the halo-to-halo scatter of the satellite luminosity function.

Finally, stronger limits on the halo mass of the Milky Way and especially the Large Magellanic Cloud (LMC) could help to provide a better model of the satellite number counts, as the LMC is known to have brought its own satellites (Kallivayalil et al. 2018, Patel et al. 2020) that need to be properly accounted for (Jethwa et al. 2016).

2.6 Conclusions

We have derived new constraints on thermal relic dark matter models from the joint statistical analysis of a set of different astrophysical probes. In particular, we extended two previous studies of strong gravitational lens systems and combined them with constraints from the Lyman-α forest and the luminous Milky Way satellites. Our results have interesting implications for the current status of dark matter studies, their limitations, as well as the most promising ways to improve upon them in the near future. We summarize them as follows:

1. We determined limits by considering the 95 percentiles of the parameters describing WDM models. From our joint posterior we find an limit on the half-mode scale of $\lambda_{\rm hm}^{\rm CL} = 0.089$ Mpc h^{-1}, corresponding to $M_{\rm hm}^{\rm CL} = 3 \times 10^7$ M$_\odot$ h^{-1} and $m_{\rm th}^{\rm CL} = 6.048$ keV under the assumption of Planck cosmology and a thermal relic dark matter model. These limits rule out the 7.1 keV sterile neutrino dark matter model for a lepton asymmetry $L_6 > 10$. If such sterile neutrino models aim to explain the observed 3.55 keV they are required to show a half-mode mass in the range of $\log_{10}(M_{\rm hm} \cdot {\rm M}_\odot^{-1}\,{\rm h}) \in [9, 11]$.

According to this summary statistics, we furthermore rule out the ETHOS-4 model of self-interacting dark matter, which shows a cutoff corresponding to a thermal relic with a mass $m_{\rm th} = 3.66$ keV (Vogelsberger et al. 2016). Amongst the considered probes, the Milky Way satellites and the Lyman-α forest provide the strongest constraints on the half-mode scale, i.e. $\lambda_{\rm hm}^{\rm CL} < 0.076$ Mpc h^{-1}, and $\lambda_{\rm hm}^{\rm CL} < 0.160$ Mpc h^{-1}, respectively. These values are followed by the strong gravitational lensing constraints of the SLACS sample $\lambda_{\rm hm}^{\rm CL} < 1.214$ Mpc h^{-1}, and the weakest constraints come from the high redshift BELLS-GALLERY with $\lambda_{\rm hm}^{\rm CL} < 6.842$ Mpc h^{-1}. The latter even shows a preference for warmer dark matter models, in contrast to the other probes. However, larger samples and higher-sensitivity lensing data are required to confirm such a trend.

2. We further considered the ratios of the joint likelihood and found that with respect to the maximum likelihood model, we rule out models above $\lambda_{\rm hm}^{\rm BF} = 0.233$ Mpc h^{-1} (corresponding to values above $M_{\rm hm}^{\rm BF} = 4.8 \times 10^8$ M$_\odot$ h^{-1} and below $m_{\rm th}^{\rm BF} = 2.552$ keV). Again, we find that the sterile neutrino dark matter models are ruled out. However, due to weaker constraints, the self-interacting dark matter model of ETHOS-4 is still allowed. In the case of Bayes factors, the limits are again mostly determined by the analysis of the Milky Way satellites (with $\lambda_{\rm hm}^{\rm BF} = 0.221$ Mpc h^{-1}). The Lyman-α analysis comes next, with $\lambda_{\rm hm}^{\rm BF} = 0.540$ Mpc h^{-1}). In the case of lensing, only the SLACS sample provides an upper limit under this criterium, with $\lambda_{\rm hm}^{\rm BF} < 3.607$ Mpc h^{-1}.

3. We highlight that the choice of a summary statistics is crucial for deciding which dark matter models are ruled out. In general, we find that the 95 percentiles provide stronger constraints, while the Bayes factor summary statistics provide more conservative limits (that are also more independent from prior assumptions).

4. None of the considered analyses are sensitive to half-mode masses below $M_{\rm hm} = 10^5$ M$_\odot$ h^{-1}, where the likelihood and posterior distributions flatten out. In the near future, we expect strong lensing observations with extended sources to increase their sensitivity towards these colder models thanks to the improvement in the angular resolution that will be provided by VLBI and the ELTs. High spectral resolution observations of quasars will provide Lyman-α forest constraints on smaller scales of the matter power spectrum and, therefore, smaller values of $\lambda_{\rm hm}$ (Iršič et al. 2017). For both probes, a larger sample of objects is expected to lead to more precise constraints. An analysis of the luminous Milky Way satellites, on the other hand, is by definition limited to satellites that are massive enough to host a galaxy. This restriction puts a limit on the lowest subhalo mass that can be detected, and the relative constraints will only improve with better control of systematic errors.

5. All probes are affected by their model assumptions (Section 2.2) and different sources of systematic errors (Section 2.4) that will need to be addressed to improve on the current level of accuracy. It is a well-known fact that current observations of the Lyman-α forest can be compatible with both CDM and WDM, depending on the

assumptions made on the thermal history of the IGM. The interpretation of the Milky Way satellite luminosity function is strongly affected by poorly constrained feedback and star formation processes, as well as the mass of the Milky Way (Lovell et al. 2012, 2014a and references therein). Inference on the halo and subhalo mass function from strong lensing observations can be significantly biased by assumptions made on the lens mass distribution (for both lensed galaxies and quasars, Vegetti et al. 2014, Hsueh et al. 2016, Gilman et al. 2017) and the size of the background sources (mainly for lensed quasars; Timerman et al., in prep.).

In this chapter, we have focused on three different astrophysical observations, that is, strong gravitational lensing, the Lyman-α forest and the luminosity function of the Milky Way satellites. However, other promising probes of the thermal relic particle mass exist, as for example, the number of non-luminous Milky Way subhaloes detectable with stellar streams (Banik et al. 2018, 2019b, Carlberg 2012, 2013, Erkal & Belokurov 2015, Yoon et al. 2011), the CMB (Ade et al. 2016) and the luminosity function of satellites in galaxies other than the Milky Way (Nierenberg et al. 2011, 2012, Corasaniti et al. 2017). In the future, we plan to extend our analysis to include other observational probes and dark matter models.

Chapter 3

Strong gravitational lensing: biases from realistic mass distributions

The content of this chapter is based on work that has been published in the Monthly Notices of the Royal Astronomical Society as ***Enzi W.*** *et al. 2019, eprint arXiv:1911.02581.*

3.1 Introduction

Galaxy-galaxy strong gravitational lensing is a powerful tool to investigate a large number of diverse astrophysical and cosmological inquiries (see Treu 2010, and references therein). For example, in the detailed analysis of the physical and kinematical properties of distant galaxies strong lensing provides the magnification that allows one to overcome the observational limitations of low signal-to-noise ratio and spatial resolution (e.g. Shirazi et al. 2014, Rybak et al. 2015b, Rizzo et al. 2018, Spingola et al. 2019). Measurements of the time-delay between the multiple lensed images of time-varying sources have been demonstrated to provide a promising constraint on the Hubble constant H_0 and weak constraints on other cosmological parameters (e.g. Suyu et al. 2010, Wong et al. 2019, Chen et al. 2019). The sensitivity of lensing to total mass has been used to quantify the amount of low-mass dark matter haloes and thereby the properties of dark matter, both with lensed quasars (Mao & Schneider 1998, Nierenberg et al. 2014, Gilman et al. 2019b, Hsueh et al. 2019) and lensed galaxies (Koopmans 2005, Vegetti & Koopmans 2009b, Vegetti et al. 2012, Hezaveh et al. 2016, Vegetti et al. 2010).

In order for these studies to be robust, a good understanding of the lenses gravitational potential is essential. For example, Ritondale et al. (2019) have shown from the analysis of the BELLS-GALLERY sample, that in certain cases it is necessary to extend the lensing mass density distribution beyond the standard assumption of a single or multiple power-laws in order to obtain a correct focusing of the reconstructed background source. Using hydrodynamical simulations, Xu et al. (2017) have shown that power-law mass models can bias the inference on the Hubble constant by up to 50 per cent, a deviation that is related to the mass sheet degeneracy (Falco et al. 1985). Using both numerical simulations and mock

observations Gilman et al. (2018) and Hsueh et al. (2018) have demonstrated that complex baryonic structures in elliptical galaxies can contribute to 8-10 percent of the strength of flux-ratio anomalies, and therefore constitute a potential bias in the inferred properties of dark matter (see also Möller et al. 2003, Xu et al. 2015). Similarly, Hsueh et al. (2016, 2017) have shown that lens galaxies with a significant edge-on disc are characterised by strong flux-ratio anomalies that can be explained by the presence of the disc without the need to include a significant population of substructures. Recently, from the analysis of extended lensed images from the BELLS-GALLERY sample, Ritondale et al. (2019) have concluded that complex mass distributions can emulate the effects of substructures and, therefore, lead to false-positive detections. More generally, Unruh et al. (2017) have shown that source-position transformations (Schneider & Sluse 2014) can explain why independent studies of the same lens system may result in different inferred lens parameters.

In this chapter, we quantify the systematic errors that may be induced by departures from simple power-law lensing mass distributions on the different lensing observables using galaxies taken from the cosmological hydrodynamical simulation Illustris-1 (Vogelsberger et al. 2014). We remove all substructures from the simulations and create mock lensing observations emulating the SLACS (Bolton et al. 2006) survey. We then model these data using a novel inference approach based on Approximate Bayesian Computation (ABC), under the classical assumption of a cored power-law mass model. We then infer the fraction of projected mass contained in substructures $f^{\rm sub}$ to quantify the degeneracy between substructures, time delays (which are relevant in the determination of H_0), other forms of complexity in the lensing potential, and the source surface brightness distribution. This chapter is structured as follows: in Section 3.2 we give a short overview of the Illustris-1 simulations and specify how we create the mock data. In Section 3.3 we describe the physical model as well as the statistical method that we use for the reconstruction of the mock data. In Section 3.4 we present and discuss our results. In Section 3.5 we conclude our study by summarising our main findings.

3.2 Lensing data generation

In order to study the lensing signal from galaxies with complex mass distributions, we focus on simulated galaxies from the hydrodynamical numerical simulation Illustris-1. In this section, we provide a short description of the simulation and how we generated mock lensing observations from it.

3.2.1 Lens and source galaxies selection

The Illustris-1 project is a series of hydrodynamical numerical simulations of cosmological volumes that follow the evolution of dark matter, cosmic gas, stars, and supermassive black holes from a starting redshift of $z = 127$ to the present time, all while accounting for realistic baryonic physics. In this chapter, we use the main run of Illustris-1, which has a box size of

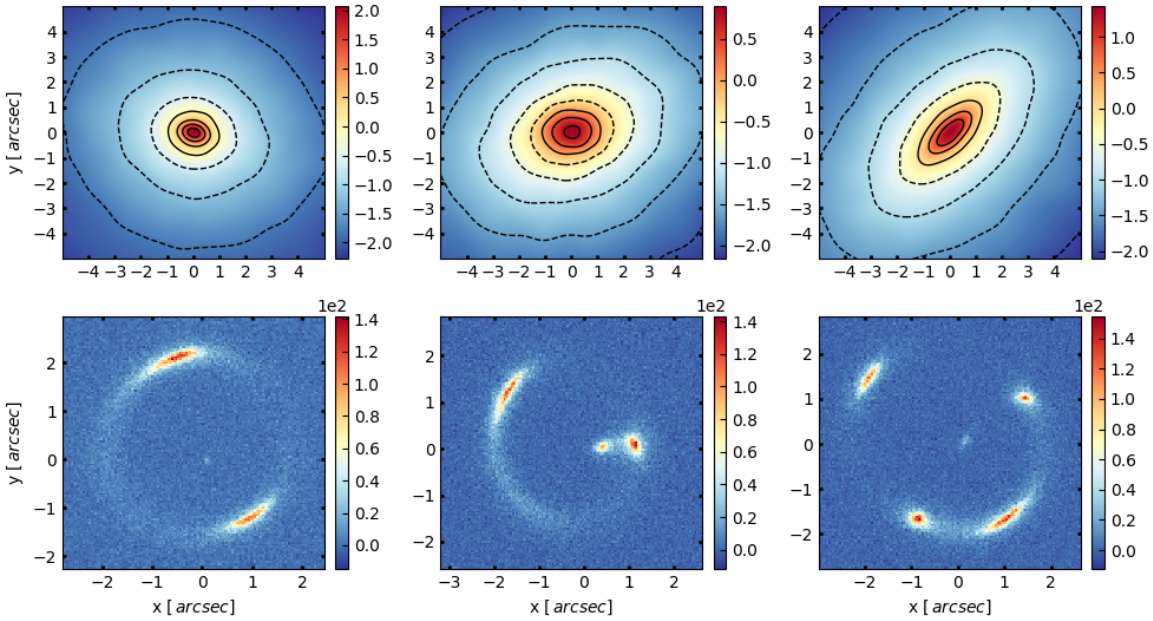

Figure 3.1: **The mass distribution of Illustris-1 galaxies and corresponding mock data.** Top panels: the logarithmic convergence maps of three representative galaxies selected from the Illustris-1 simulation. All substructures have been removed in these images. The dashed lines represent iso-convergence contours. Bottom panels: the corresponding mock lensing observations in our sample of lens systems.

106.5Mpc in all three dimensions and contains 1820^3 dark matter particles and (initial) gas cells. The simulations were run using the moving-mesh Arepo code (Springel 2010). The dark matter particle mass is $6.3\times10^6 M_\odot$ and the initial gas particle mass is $1.3\times10^6 M_\odot$, while the simulation uses softening lengths for dark matter and baryons that are $\epsilon_{\rm dm}$=1.4 kpc and $\epsilon_{\rm b}$=0.7 kpc, respectively. The simulations adopt the cosmological parameters $\Omega_m = 0.2726$, $\Omega_\Lambda = 0.7274$, $\Omega_b = 0.0456$, $\sigma_8 = 0.809$, $n_s = 0.963$, and $H_0 = 100\mathrm{h}\ \times \mathrm{km\ s^{-1}Mpc^{-1}}$ with $\mathrm{h} = 0.704$, which are consistent with the latest Wilkinson Microwave Anisotropy Probe 9 measurements (Hinshaw et al. 2013).

From the catalogue of Illustris-1 we choose ten galaxies at redshift $z = 0.2$ to be our sample of lenses. These are analogues of the SLACS lenses (Bolton et al. 2006) and were selected by Despali & Vegetti (2017) based on a number of properties: the total halo mass, stellar mass of the central galaxy, stellar effective radius, velocity dispersion, and dynamical properties compatible with an early-type morphology. In particular, they selected galaxies according to the fraction of stellar mass showing specific circularities $\epsilon = J_z/J(E)$, where J_z is the angular momentum of an individual stellar particle and $J(E)$ is the maximum angular momentum within its local environment. The fractions $f_{\epsilon<0.0}$ and $f_{\epsilon>0.7}$ are indicators of an early-type morphology (see e.g. Genel et al. 2015, Teklu et al. 2015), typical for lens galaxies at the considered redshift.

This selection ensures that we avoid galaxies with a significant disc component, for which the lensing properties have already been investigated by Hsueh et al. (2018). In Table 3.1 we list the main properties of the considered sample of lens galaxies. We also select a compact Illustris-1 galaxy at redshift $z = 0.6$ as our background source to be lensed. We use the same source galaxy for all created mocks.

3.2.2 Ray tracing

For each of the ten lens galaxies in our sample, we consider the projection of their mass density distribution onto the three principal planes and use the ray-tracing code GLAMER (Metcalf & Petkova 2014, Petkova et al. 2014)[1] to generate 3×10 maps of the lensed source surface brightness distribution. The focus of this chapter is on the lensing effect of small scale complexity in the lens mass distribution such as dense baryonic structures from tidally disrupted satellites and general departures from simple single power-law models. Therefore, we have removed all particles belonging to dark matter substructures and only considered the contribution from the main halo. However, it is worth mentioning that in the Illustris-1 simulation, the minimum resolved substructure mass is of the order of $\simeq 10^8 M_\odot$. For each main halo we take all its particles within $R_{200} \approx R_{\rm vir}$ as identified by the SUBFIND algorithm (Springel et al. 2001a).

For a given collection of particles (i.e. their positions, masses, and smoothing lengths as used in the original simulation), GLAMER determines the corresponding deflection angle maps with an efficient tree algorithm. In order to avoid unrealistic deflections of the light rays, each particle is represented by a B-spline in three dimensions as it is commonly done in smoothed particle hydrodynamics (SPH) simulations. We set the size of the smoothing length to the distance of the $N^{\rm smooth}$-th nearest neighbour. We choose $N^{\rm smooth} = 64$, which corresponds to scales $\lesssim O(1\text{kpc})$. This scheme provides smaller smoothing lengths where the particles are dense (for example at the centres of haloes), and larger ones where the particles are sparse and shot noise would otherwise pose a problem. We follow a similar approach for the baryonic cells by treating them as particles with a corresponding smoothing length in order to make them compatible with the implementation of GLAMER. The minimum and average smoothing lengths over the whole sample are $0.6 - 0.8$ kpc and 16 kpc respectively, which also taking into account the softening length of the simulation tell us that we can safely resolve inhomogeneities on scales $> O(1\text{kpc})$. Notice that the pixel size of the mock images is 0.04 arcsec.

As an example, Figure 3.1 shows the convergence maps (i.e. surface mass density normalised by the critical surface mass density) of three representative Illustris-1 galaxies. The presence of small scale features in the mass distribution, which are typically not captured by simplified parametric profiles, are highlighted by the irregular iso-convergence contours.

The galaxies in the Illustris-1 sample show unrealistically large cores. This effect is a well-known challenge for simulations and arises from finite smoothing lengths. Depending on the lensing configuration, these large cores lead to lensed images with morphologies

[1] http://glenco.github.io/glamer/

(see middle-bottom panel of Figure 3.1 for an example) which are rarely observed in real gravitational lens systems, where the central mass density distribution may be cuspier. For this reason we remove six realisations from our set and in the following we will only focus on the 24 systems with realistic lensed images. A similar problem was recently identified by Mukherjee et al. (2018) in a study of the EAGLE simulation (Schaye et al. 2015).

3.2.3 Observational effects

In order to generate mock HST observations we convolve the lensed surface brightness distribution $I(\vec{x})$ with the point spread function (PSF), $p(\vec{x})$, from the WFC3 camera in the F606W filter and add the contribution of observational noise. In particular, for each of the lensed images in our sample, we draw a realisation of Gaussian random noise $\vec{n}(\vec{x})$, with

$$P(\vec{n}|\mathbf{N}) = \mathcal{G}(\vec{n}, \mathbf{N}) := \exp\left(-\frac{1}{2}\vec{n}^t\mathbf{N}^{-1}\vec{n}\right)/\sqrt{\det(2\pi\mathbf{N})}\,, \tag{3.1}$$

which includes the contribution of a constant background b and another term proportional to the signal, which approximates the Poisson noise of photon counts. We assume the noise to be uncorrelated and therefore entries of the covariance matrix are multiplied by the Kronecker delta $\delta^K_{\rm ij}$:

$$\mathbf{N}_{\rm ij} = \left[b + (p * I)(\vec{x}_i)\right] \times \delta^K_{\rm ij}\,. \tag{3.2}$$

The level of noise is chosen in such a way to match the signal-to-noise ratio of the SLACS lenses analysed by Vegetti et al. (2014).

3.3 Lens modelling

In this section, we discuss the physical components that define our inference problem and present a novel inference approach based on Approximate Bayesian Computation (ABC). For an overview on ABC in cosmology, we refer the reader to Akeret et al. (2015). For examples of its application in the context of strong gravitational lensing we refer to Birrer et al. (2017), Gilman et al. (2018) and Gilman et al. (2019b).

3.3.1 Physical model

The physical model includes the following unknown components: the surface brightness distribution of the source galaxy, the mass distribution of the lens galaxy and the amount of lens-galaxy mass contained in substructures.

Source Surface Brightness distribution

We follow Vegetti & Koopmans (2009a) and define the source surface brightness distribution $s(\vec{x})$ on an adaptive grid, so that $\vec{s}_i = s(\vec{x}_i)$ corresponds to the brightness value at the position

Table 3.1: **The properties of the lens galaxies selected from the Illustris-1 simulation.** From left to right, the virial mass $M_{\rm vir}$, the mass of the dark matter component $M_{\rm dm}$, the stellar mass M_*, the gas mass $M_{\rm gas}$, the fraction of stellar mass showing a circularity higher than 0.7 $f_{\epsilon>0.7}$ and lower than 0.0 $f_{\epsilon<0.0}$, as well as the virial radius $r_{\rm vir}$ and stellar radius r_*.

$ID_{\rm fof}$	$M_{\rm vir}$	$M_{\rm dm}$	M_*	$M_{\rm gas}$	$f_{\epsilon>0.7}$	$f_{\epsilon<0.0}$	$r_{\rm vir}$	r_*
	$M_\odot$/h	$M_\odot$/h	$M_\odot$/h	$M_\odot$/h	1	1	kpc/h	kpc/h
28	2.314e+13	1.910e+13	4.107e+11	1.582e+11	1.580e-01	6.522e-01	6.319e+02	1.025e+01
40	2.789e+13	2.404e+13	5.945e+11	1.396e+11	1.662e-01	6.563e-01	6.725e+02	1.458e+01
51	2.094e+13	1.906e+13	4.691e+11	8.335e+10	7.534e-02	9.388e-01	6.112e+02	1.205e+01
55	2.327e+13	1.953e+13	4.341e+11	2.664e+11	1.323e-01	7.430e-01	6.330e+02	1.539e+01
65	1.749e+13	1.543e+13	6.486e+11	3.023e+10	1.284e-01	8.440e-01	5.756e+02	1.916e+01
84	1.179e+13	1.134e+13	3.427e+11	5.861e+10	1.664e-01	7.242e-01	5.047e+02	1.411e+01
91	1.355e+13	1.085e+13	2.147e+11	2.446e+10	6.730e-02	8.781e-01	5.286e+02	8.816e+00
95	1.094e+13	9.443e+12	3.164e+11	1.371e+11	5.936e-02	9.282e-01	4.922e+02	1.225e+01
121	6.297e+12	5.658e+12	1.389e+11	1.550e+11	1.085e-01	7.531e-01	4.095e+02	2.656e+00
140	8.737e+12	7.818e+12	2.268e+11	4.465e+10	1.077e-01	7.056e-01	4.567e+02	7.336e+00

$\vec{x}_i$ of the i-th vertex of a Delaunay mesh. The mesh vertices are obtained by ray tracing a subset of the pixels of the observed data $d(\vec{x})$ to the source plane. The source surface brightness at each vertex constitutes a free nuisance parameter of the model, while the source brightness within the triangles spanned by these vertices is determined via linear interpolation.

Lens projected mass distribution

We parametrise the projected mass distribution of the main deflector with an elliptical cored power-law profile of convergence

$$\kappa(\rho) = \frac{\kappa_0 \left(2 - \frac{\gamma}{2}\right) q^{-1/2}}{2(r_c^2 + \rho^2)^{(\gamma-1)/2}}, \tag{3.3}$$

where κ_0 is the amplitude, γ the 3D slope, q the projected minor-to-major axis ratio in projection, θ the angular orientation of the major axis (defined North to East), x_0 and y_0 the centre position coordinates, and r_c the core radius. After a rotation by θ and a translation of (x_0, y_0) the ellipse radius ρ is defined via $\rho^2 = x^2 + y^2/q^2$. Together with the strength Γ and position angle Γ_θ of an external shear component, these geometrical parameters constitute part of the key target parameters in the reconstruction process and we collectively refer to them as macro-model parameters or $\eta^{\rm macro} = (\kappa_0, q, \theta, x_0, y_0, r_c, \gamma, \Gamma, \Gamma_\theta)$. We use the FASTELL[2] library developed by Barkana (1998) to calculate deflection angles of non-isothermal ($\gamma \neq 2$) cored power laws.

Substructure

Each substructure is assumed to have a spherical NFW mass profile (Navarro et al. 1996) with a concentration given by the mass-concentration relation of Duffy et al. (2008) and a virial radius according to Bullock et al. (2001). We do not consider the effects of tidal truncation nor variations in the concentration-mass relation as a function of the distance from the host centre, as these are of secondary importance (see, e.g. Despali et al. 2018). We assume the number density of substructure in mass bins of width $[m, m + dm]$ (i.e. the substructure mass function) to be well described by a power-law (Springel et al. 2008)

$$\frac{dn^{\rm sub}(m)}{dm} = n_0 \times m^{-\xi}, \tag{3.4}$$

with a slope of $\xi \approx 1.9$ and a normalisation n_0 which depends in general on the redshift and the mass of the host galaxy (Gao et al. 2011, Xu et al. 2015). Here, we neglect both dependencies because the lens galaxies are all at the same redshift and span a narrow range in virial mass. The amplitude of the mass function is determined by the fraction of projected total mass of the host galaxy within the Einstein radius $R_{\rm E}$ (calculated according to Section 3.6.1) contained in substructure, $f^{\rm sub}(< R_{\rm E})$ (see Section 3.6.2 for a derivation). The number, masses and

[2] http://wise-obs.tau.ac.il/~barkana/codes.html

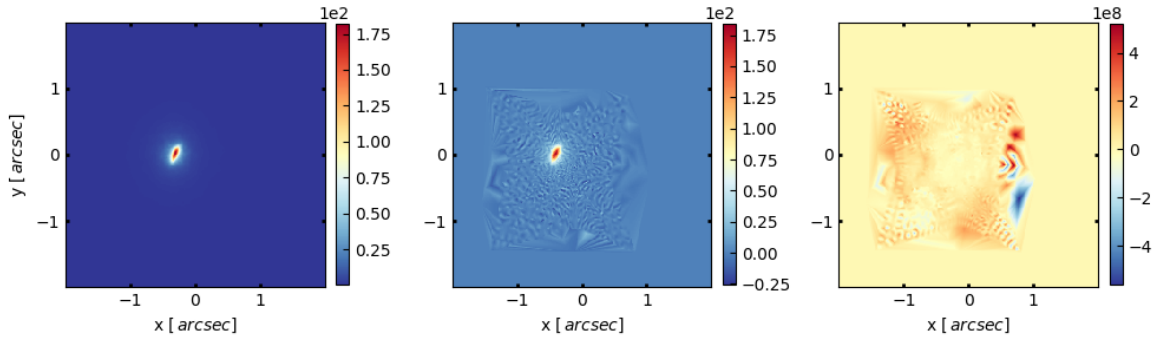

Figure 3.2: **Illustration of challenges of the source prior.** From left to right: the input source selected from the Illustris-1 simulation, the source inferred from the lens modelling of one of the lens systems in our sample, a prior sample of the surface brightness distribution. Notice that in the latter case the resulting field realization does not resemble a realistic source brightness distribution. The middle and right plots were generated using the same regularisation form $\mathbf{S}$ and strength λ.

projected positions of each substructure realisation $\eta^{\rm sub}$ are a free nuisance parameter, while $f^{\rm sub}(< R_{\rm E})$ is the main target parameter of the inference problem. In the following, we drop the argument of $f^{\rm sub}(< R_{\rm E})$ and write $f^{\rm sub}$, while always referring to the fraction within the Einstein radius unless stated otherwise.

3.3.2 Reconstruction approach

In this section, we describe the statistical method we use to infer the macro-model parameters $\eta^{\rm macro}$ and the substructure mass fraction $f^{\rm sub}$ while marginalising over the source surface brightness distribution and the substructure realisations.

Priors

For the parameters describing the lensing mass distribution we choose the following priors: a uniform prior for the amplitude κ_0, the axis ratio q, the angular orientation θ and the slope γ, a Gaussian prior for the position x_0, y_0 and a log normal distribution for the core radius r_c.

The tessellated source surface brightness distribution is drawn from a Gaussian regularising prior:

$$\mathcal{P}(\vec{s}|\mathbf{S},\lambda) = \mathcal{G}(\vec{s}, \lambda^{-1}\mathbf{S})\,, \tag{3.5}$$

where the matrix $\mathbf{S}$ determines the form of the regularisation and λ its strength. In particular, we choose a form of $\mathbf{S}$ that penalises either the gradient or the curvature at the vertices spanning the mesh (e.g. Suyu et al. 2006). We choose a log uniform prior on the regularisation strength (e.g. Vegetti & Koopmans 2009a).

The probability of a substructure having a mass m in $[m, m + dm]$ is related to the mass function introduced in Section 3.3.1 via the following relation:

$$\mathcal{P}(m) = \frac{dn^{\rm sub}}{dm} \Bigg/ \int_{10^6 {\rm M}_\odot}^{10^{10} {\rm M}_\odot} dm \frac{dn^{\rm sub}}{dm} \,. \tag{3.6}$$

We assume that the number of substructure $N^{\rm sub}$ follows a Poisson distribution,

$$\mathcal{P}(N^{\rm sub}|\mu) = \frac{e^{-\mu}\mu^{N^{\rm sub}}}{N^{\rm sub}!} \,, \tag{3.7}$$

with a mean value μ given by:

$$\mu = A^{\rm data} \times \frac{f^{\rm sub}(< R_{\rm E}) \times \Sigma_{\rm crit}}{< m >_{\mathcal{P}({\rm m})}} \,, \tag{3.8}$$

where $A^{\rm data}$ is the area covered by the data (see Section 3.6.2 for a derivation). We assume a prior probability for the substructure projected positions which is uniform within the extent of the considered lens plane (Xu et al. 2015). We choose a prior of $\mathcal{P}(f^{\rm sub}) \propto \left(f^{\rm sub}\right)^{-1/2}$, which is the Jeffreys prior of $\mathcal{P}(N^{\rm sub}|f^{\rm sub})$ (see Section 3.6.3).

Macro-model reconstruction

In order to find the optimal macro-model parameters, we propose a modified version of Approximate Bayesian Computation (ABC). ABC is usually motivated by inference problems, for which the likelihood is not accessible. In Section 3.6.4, we discuss why we follow this approach in this chapter and why the evaluation of the likelihood poses a challenge for our analysis. In its simplest form ABC is a rejection sampling algorithm in which prior samples of the target parameters - in this case samples of $\eta^{\rm macro}$ - are accepted/rejected according to the distance between the model realisation they generate $\vec{b}(\vec{x})$ and the observed data $\vec{d}(\vec{x})$. Accepting the subset of samples closest to the real observation yields a set of samples approximating the posterior distribution. Accepting only perfect matches would yield the exact posterior distribution, although these are extremely unlikely for a finite number of realisations. Therefore, it is customary to set a distance threshold or accept only the fraction of the drawn samples that are closest to the data. Alternatively, one can compare statistical summaries $\Phi(\vec{b}(\vec{x}, \eta_i^{\rm macro}))$ and $\Phi(\vec{d})$ rather than the mock data realisations and observations directly to each other. This approach improves the acceptance rate of the procedure, but only approximates the true posterior. As described below, the latter is the approach adopted in this chapter.

A fully forward modelling method requires that we draw a realisation of the source surface brightness distribution from a chosen prior (see Section 3.3.2). This requirement, together with the pixelated nature of our source, poses a challenge as every sampled source brightness distribution is almost guaranteed to be far away from the true one as illustrated in Figure

3.2. Searching through the parameter space of the target parameters $\eta^{\rm macro}$ while drawing new source realisations in every sample would, therefore, require a large number of samples, rendering the method infeasible. For this reason, we would like to obtain a distance measure that directly compares a lens configuration to the observations, without having to draw a source realisation explicitly. We propose, therefore, to use the following as a proxy to a conventional distance measure:

$$dist(\vec{d}, \eta^{\rm macro}) \approx -2 \log \mathcal{P}(\vec{d}|\mathbf{L}, \lambda) = -2 \log \left[\int d\vec{s}\, d\vec{n}\, \mathcal{P}(\vec{d}|\mathbf{L}, \vec{s}, \vec{n}) \mathcal{P}(\vec{s}|\mathbf{S}, \lambda) \mathcal{P}(\vec{n}|\mathbf{N}) \right] = \\ = -2 \log \left[\mathcal{G}(\vec{d}, \mathbf{N} + \lambda^{-1} \mathbf{PLSL}^T \mathbf{P}^T) \right], \quad (3.9)$$

where $\mathbf{L} = \mathbf{L}(\eta^{\rm macro})$ is the lensing operator for a macro-model $\eta^{\rm macro}$ and $\mathbf{P}$ a blurring matrix reproducing the effects of the point spread function $p(\vec{x})$. This expression corresponds to the evidence in the inference of the source for the given lens model. In particular, since one can evaluate the integral analytically, it will also include the more realistic source configurations that are of interest in the inference. Sampling individual source realisations will, in contrast, tend to be far away from what is expected for real galaxies (as discussed in the previous section). Since the priors on the source and the noise follow Gaussian distributions, the maximum-posterior source appears implicitly in the exponent of the Posterior probability distribution.

Heuristically, we expect those lens configurations with higher evidence to match the data more frequently than those with lower evidence. Another, more Bayesian viewpoint is that the evidence can be interpreted as a measure of credibility of the data $\vec{d}$ on the hypothesis of a $\eta_i^{\rm macro}$ (see e.g. Cox 1946). Both arguments indicate that we should prefer to accept higher evidence samples to obtain a posterior approximation. Following this consideration, we rank the drawn samples in order of their evidence values and accept the fraction of samples with the highest values. In order to improve the convergence behaviour further, we introduce a proposal distribution $\mathcal{Q}(\eta^{\rm macro})$, which we recursively update until it generates samples close to the posterior distribution.

Our inference on the macro-model parameters then proceeds in an iterative fashion as follows:

(i) We draw $N^{\rm samp}$ samples of macro-model parameters $\eta_i^{\rm macro}$ from the proposal distribution. At the first iteration we initialise the proposal distribution to match the prior introduced in Section 3.3.2, i.e. $\mathcal{Q}(\eta^{\rm macro}) = \mathcal{P}(\eta^{\rm macro})$.

(ii) For each sample we calculate its importance sampling weight $w_i = \mathcal{P}(\eta_i^{\rm macro})/\mathcal{Q}(\eta_i^{\rm macro})$ and evidence $\mathcal{P}_i = \mathcal{P}(\vec{d}|L(\eta_i^{\rm macro}))$. We accept the $N^{\rm acc}$ samples with the highest values of $P_i \times w_i$ (see Section 3.6.5 for further details). Whenever we find that the effective number of samples,

$$N^{\rm eff} = \left(\sum_{\rm i=1}^{N^{\rm samp}} \mathcal{P}_i \times w_i \right)^2 \Bigg/ \left(\sum_{\rm i=1}^{N^{\rm samp}} \mathcal{P}_i^2 \times w_i^2 \right), \quad (3.10)$$

is higher than $N^{\rm acc}$ we accept this effective number of samples instead. A chosen value of $N^{\rm acc}$ therefore determines the minimum relative acceptance threshold of $\epsilon = \frac{N^{\rm acc}}{N^{\rm samp}}$. The above acceptance strategy ensures that we do not underestimate the region within which we have to probe the parameter space in the following iteration.

(iii) We update the proposal distribution according to the following expression:

$$Q(\eta^{\rm macro}) = q \times \mathcal{G}(\eta^{\rm macro} - \mu, \Sigma) + (1 - q) \times \mathcal{P}(\eta^{\rm macro}) \,, \tag{3.11}$$

where $\mathcal{G}(\eta^{\rm macro} - \mu, \Sigma)$ is a Gaussian distribution estimated from the accepted samples. Above, q determines the fraction of samples drawn from the prior versus the fraction drawn from the estimated Gaussian distribution. We set $q = 0.25$, so that one fourth of all drawn samples are used to search through the initial parameter space thereby reducing the probability that we converge to a local minimum.

We repeat the steps (i) to (iii) $N^{\rm step}$ times, which results in a large number of samples, most of them drawn in the interesting regions of the parameter space. Following the discussion in Section 3.6.5 the approximate posterior is given by the distribution of the accepted parameters:

$$\mathcal{P}^{\rm ABC}(\eta^{\rm macro}|\vec{d}) = \frac{1}{K^{\rm acc}} \sum_{\rm i \in acc} \delta^D(\eta^{\rm macro} - \eta_i^{\rm macro}) \,, \tag{3.12}$$

with $K^{\rm acc}$ being the total number of accepted samples so far. The expectation value of some function $f(\eta^{\rm macro})$ is then given by the average of the accepted samples, i.e.:

$$\begin{aligned} \left\langle f(\eta^{\rm macro}) \right\rangle_{\mathcal{P}^{\rm ABC}(\eta^{\rm macro}|\vec{d})} = \\ = \int d\eta^{\rm macro} f(\eta^{\rm macro}) \mathcal{P}^{\rm ABC}(\eta^{\rm macro}|\vec{d}) = \frac{1}{K^{\rm acc}} \sum_{i \in \rm acc} f(\eta_i^{\rm macro}) \,. \end{aligned} \tag{3.13}$$

For every lens we perform the above iterative scheme a total of three times, inferring the macro-model parameters or the source regularisation level in an alternating fashion.

3.3.3 Substructure reconstruction

We now extend our inference to mass substructure within the lensing potential and the relative fraction $f^{\rm sub}$. To this end, we include two additional steps where we first draw values of $f^{\rm sub}$ from its prior probability distribution $\mathcal{P}(f^{\rm sub})$ and then draw explicit substructure realisations $\eta^{\rm sub}$ for each given value of $f^{\rm sub}$. We draw only one substructure realisation per macro-model. This is justified by the large number of drawn samples, which ensures that for each pair $(\eta_i^{\rm macro}, \eta_i^{\rm sub})$ there exist another one $(\eta_j^{\rm macro}, \eta_j^{\rm sub})$ with a different substructure realisation but similar macro-model parameters. For each macro-model and substructure combination we then calculate the evidence $\mathcal{P}_i = \mathcal{P}(\vec{d}|L(\eta_i^{\rm macro}, \eta_i^{\rm sub}))$ and follow the iterative scheme presented above.

In our approximate posterior we would like to take into account that each substructure realisation could also have been generated from a mass function with a different $f^{\rm sub}$ than the one it has actually been drawn from. To this end, for each of these realisations we also calculate the probability of having been drawn from a model with $\widetilde{f^{\rm sub}}$ and obtain the approximate posterior:

$$\mathcal{P}^{\rm ABC}(f^{\rm sub}|\vec{d}) = \frac{1}{K^{\rm acc}} \sum_{\rm i\in acc} v_i(f^{\rm sub})\,, \tag{3.14}$$

where the interpolating function v_i is defined as:

$$v_i(f^{\rm sub}) = \frac{\mathcal{P}(N_{\rm i}^{\rm sub}|f^{\rm sub})\mathcal{P}(f^{\rm sub})}{\int d\widetilde{f^{\rm sub}}\,\mathcal{P}\left(N_{\rm i}^{\rm sub}\middle|\widetilde{f^{\rm sub}}\right)\mathcal{P}\left(\widetilde{f^{\rm sub}}\right)}\,. \tag{3.15}$$

The generated substructure realisations will rarely match the one contained in the data. Rather than that, we expect that realisations drawn from the correct mass function will more often provide higher evidence than those drawn from a different mass function, as the latter have either too few or too many substructures. If a particular substructure realisation can reproduce the observed lensed images much better than any other, the weights introduced above will always ensure that more than a single value of $f^{\rm sub}$ will contribute to the posterior distribution. Using mock data with different level of $f^{\rm sub}$, we have thoroughly tested our inference approach, and found that we are able to recover the correct amount of mass in substructure with a precision that, as expected, depends on the number of lenses and the data quality.
Recently, Despali et al. (2018) have shown that low-mass haloes along the line of sight can significantly contribute to the number of detectable objects. For simplicity, here we focus only on the substructure contribution. However, our analysis can be easily extended to include both populations.

3.4 Results

In this section, we present the results of the inference analysis described in Section 3.3.2 when applied to the sample of twenty-four mock lensing data presented in Section 3.2. In particular, for each lens, we create $N^{\rm step} \times N^{\rm samp} = 20 \times 4000$ lens configurations and choose $N^{\rm acc} = 200$ at each step during the smooth modelling phase and $N^{\rm step} \times N^{\rm samp} = 40 \times 4000$ lens configurations and accept around half of them for the substructure inference.

3.4.1 Smooth modelling

Lensing mass distribution

A simple elliptical cored power-law mass distribution (with no substructure) can fit all of the mock observations to the noise level. Figures 3.3 and 3.4 show the best and the worst

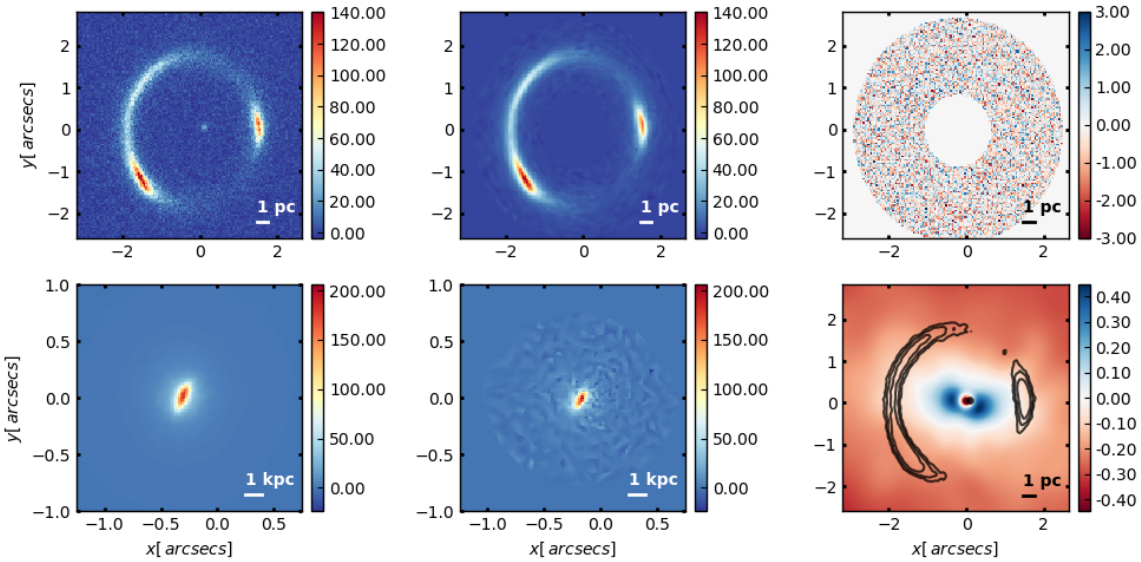

Figure 3.3: **Results from the best reconstruction of the Illustris-1 mock sample of galaxy-galaxy strong lenses.** From left to right, top to bottom: the mock data $d(\vec{x}) := I_{\rm gt}(\vec{x})$, the reconstructed lensed images $I_{\rm rec}(\vec{x})$, the image-plane residuals $r(\vec{x})$, the ground-truth source $s_{\rm gt}(\vec{x})$, the reconstructed source $s_{\rm rec}(\vec{x})$, and the relative difference in convergence $\Delta\kappa/\kappa(\vec{x}) = \kappa_{\rm gt}(\vec{x})/\kappa_{\rm rec}(\vec{x}) - 1$.

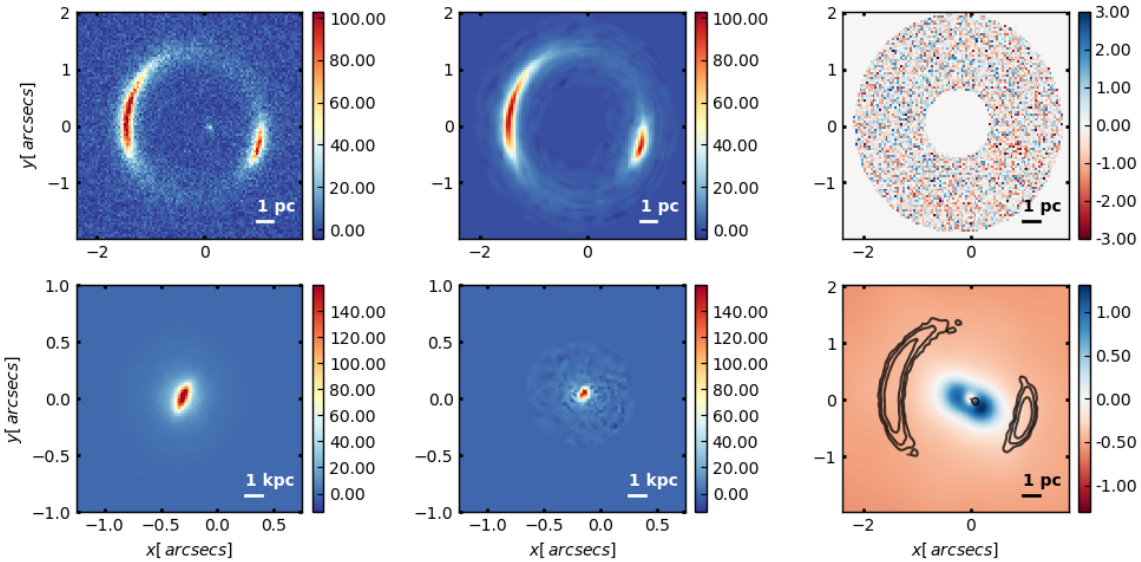

Figure 3.4: Same as Figure 3.3 but for the mock with the highest residuals.

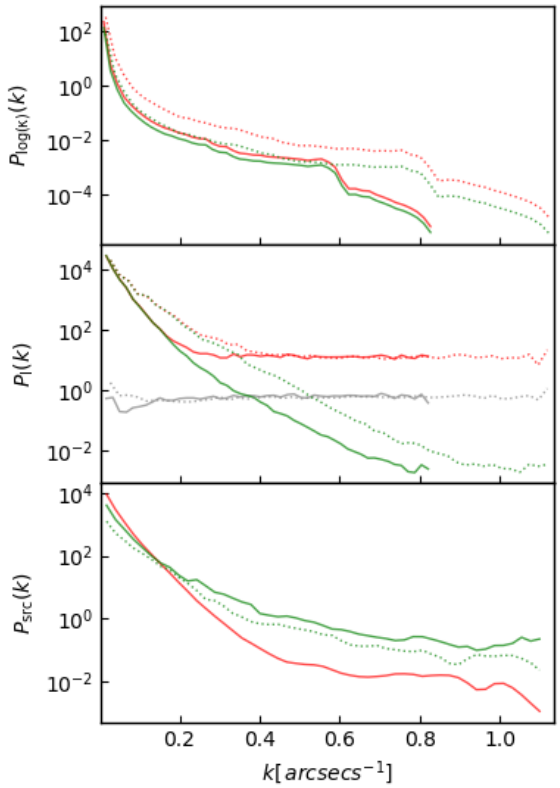

Figure 3.5: **Comparison between the ground-truth power spectra and the reconstructed ones.** The top panel shows the power spectrum of the logarithmic convergence maps $\log(\kappa(\vec{x}))$ for the ground-truth (red) and the model (green). The middle panel shows the power spectra of the data $d(\vec{x})$ (red), the reconstructed lensed images $I(\vec{x})$ (green), and the image plane residuals (gray). The bottom panel shows the power spectrum of the ground-truth source (red) and the reconstructed one (green). All panels show the spectra for the mock data with the best (solid curve) and worst (dotted curve) reconstruction.

reconstructions, respectively. For the worst reconstruction, we find correlated image residuals at the 1.3-σ level, which are related to the high dynamic range of the input lensed images which cannot be fully reproduced by the reconstructed source due to the effect of the regularisation.

Even though we can accurately reproduce the lensed images to the noise level, we infer a lensing mass density distribution with a core that is much smaller than the true value. Gravitational lensing is mainly sensitive to the total mass within the Einstein radius and not its specific distribution. This means that it is insensitive to the presence of a core, when it does not directly affect the lensed images. As a result, the Illustris-1 galaxies in our sample have a core size of $r_c = 0.15$-0.5 arcsec, while the reconstructed cores tend to be of the order of $r_c = 5 \times 10^{-4}$arcsec. As can be seen in the Figures 3.3 and 3.4, the relative difference in convergence between the simulated lenses and the power-law model stays constant over the extent of the lensed images, and hence can be compensated by a rescaling of the source size.

In order to further test the influence of the core we have also modelled one of the lens systems where the core directly affects the lensed images (e.g. middle-bottom panel of 3.1) and found that in this case we can recover the correct core size within one standard deviation and obtain a source that is of the same size as the original one.

Source and time delay reconstruction

In general, we recover sources that tend to be on average 50 per cent more compact than the original ones, as reflected by a lack of power in their power spectra at large scales of O(a few kpc). Figure 3.5 shows the power spectrum of the logarithmic convergence maps (top), the data (middle) and reconstructed model for the source (bottom); in all panels we can compare the ground-truth with the reconstructed model. This effect on the source size is in agreement with the findings by Ritondale et al. (2019) who have demonstrated, with the study of real gravitational lens systems, that a failure to fully capture the complexity of the lensing potential leads to a systematic error in the estimated morphological and physical properties of the reconstructed source by up to a factor of 22 per cent. In particular, we observe a correlation between the logarithm of $P^{\log(\kappa)}_{\rm gt}/P^{\log(\kappa)}_{\rm lens}(k)\ \forall\ k$ and the logarithm of $P^{\rm s}_{\rm gt}/P^{\rm s}_{\rm lens}(k)$ for $k = O($ kpc) with a correlation coefficient $\rho \approx 0.95$ in our sample. Moreover, we find that the reconstructed sources show increased power at small scales with respect to the ground-truth, which indicates a certain level of noise fitting, as also found by Bayer et al. (2018). A log-normal source prior may mitigate this issue significantly as it restricts the source brightness values to positive values. This potentially allows for reconstructing sources up to smaller scales and increase the sensitivity to substructures as a result.

We further consider the time delays of the reconstructed lens models. Following the discussion by Kochanek (2020) we calculate for each of the lens systems the expected fractional error on H_0 using the (mean) convergence at the Einstein radius $\kappa(R_{\rm E})$, i.e. $f_{H_0} = \frac{H_0^{gt}}{H_0^{rec}} - 1 = \frac{1-\kappa^{gt}(R_{\rm E})}{1-\kappa^{rec}(R_{\rm E})} - 1 = \frac{\Delta t_{ij}^{gt}}{\Delta t_{ij}^{rec}} - 1$, where κ_E^{gt} is the ground-truth convergence used to generate data, i.e. the Illustris galaxies (see Section 3.2) or the control sample (see Section 3.4.2), while κ^{rec} is the convergence at the Einstein radius obtained from the lens modelling of these two samples. We find a median and 1-σ percentiles of $f_{H_0} = 25^{+37}_{-19}$ per cent and $f_{H_0} = 12^{+6}_{-3}$ per cent, for the Illustris galaxies and the control sample, respectively. In both cases the Hubble constant is, therefore, systematically underestimated. Notice, that our analysis uses as constraints on the mass distribution the extended lensed images only, while cosmographic analyses (see e.g. Suyu et al. 2010, Chen et al. 2019), also make use of the additional information provided by the kinematics of the lens galaxy and the positions of quadruply-lensed quasars.

3.4.2 Substructure reconstruction

We now allow for the presence of substructure in the lens modelling to infer the amount of lens galaxy mass in substructures, $f^{\rm sub}$.

To validate this inference we generate a control sample of twenty-four lens systems with the same observational properties (i.e. signal-to-noise ratio and angular resolution) as our Illustris-1 sample. To this end, we use as input lensed images those corresponding to the highest evidence macro-model obtained in the previous section by modelling the Illustris-1 galaxies with a smooth power-law. This approach provides us with a sample for which we have a perfect knowledge of the lensing potential. We model both the mocks generated directly with the Illustris-1 galaxies and the control sample. As the input data in both samples do not contain any substructure, we expect the control sample to lead to an inferred value of $f^{\rm sub}$ consistent with zero. Any potential difference in $f^{\rm sub}$ between the two samples can be then attributed to the additional complexity of the Illustris-1 galaxies.

Figure 3.6 shows the joint posterior on $f^{\rm sub}$ for the Illustris-1 (solid line) and the control sample (dashed line). Furthermore, it shows the ABC likelihood of $f^{\rm sub}$ for one of the Illustris-1 galaxies in three projections (coloured lines). For small values of $f^{\rm sub}$ the likelihood is flat as the number of substructures affecting the lensed surface brightness approaches zero. This flattening may appear before the point where no substructures are placed due to the sensitivity of the data. Within the region where the likelihood is flat, the posterior inherits the shape of the prior. For large values of $f^{\rm sub}$, the likelihood decreases, reflecting an overabundance of substructures that is not compatible with the data.

From the analysis of all gravitational lens systems in each of the two samples, we find a posterior upper limit of $f^{\rm sub} \leq 1.7 \times 10^{-3}$ and $f^{\rm sub} \leq 2 \times 10^{-3}$ at the 68 per cent confidence level for the Illustris-1 and the control sample, respectively. The inferred $f^{\rm sub}$ from the two samples agree with each other within the uncertainty arising from the finite number of samples of $(\eta^{\rm macro}, \eta^{\rm sub})$ drawn in the analysis. In both cases, the posterior upper limits differ significantly from the prior upper limit of $f^{\rm sub} \leq 4.7 \times 10^{-2}$. Both samples are compatible with the null hypothesis that the lensing potential does not contain any substructures, i.e. $f^{\rm sub} = 0$. We conclude, therefore, that baryonic structure, as well as small scale complexity, do not have a significant impact on our capability of correctly inferring the amount of substructure. We find that the reconstructed sources do not change much with respect to the case without substructure. In order to test the effect of our choice for the input source galaxy, we repeat the analysis using HST observations of a local galaxy and find no significant change on the inferred upper limit on $f^{\rm sub}$.

Recently, from the analysis of flux-ratio anomalies in multiply imaged quasars lensed by elliptical galaxies from the Illustris-1 simulation with the same smoothing level, Hsueh et al. (2018) have found that baryonic components in the lensing galaxies are responsible for an increase in the probability of flux-ratio anomalies by 8 per cent. Similarly, Gilman et al. (2017) have concluded that 10 per cent of the flux-ratio anomalies in elliptical galaxies can be attributed to baryonic structures in their analysis of mock data based on real HST observations. Both results have important implications for the quantification of sub- and line-of-sight structure with gravitationally lensed quasars. We attribute the difference from our findings to the fact that flux ratios are sensitive to the second derivative of the lensing potential rather than the first, and are therefore more affected by general departures from simple smooth distributions. As the sensitivity of extended arcs and rings is strongly dependent on the

angular resolution of the data, we expect our results to change with increasing data quality.

On the other hand, our results are also in disagreement with Vegetti et al. (2014) and Ritondale et al. (2019), who from the analysis of real lens systems with a similar data quality as ours, have identified the presence of complex mass components which are degenerate with the presence of a large population of substructure and can lead to false detections. This discrepancy seems to indicate that galaxies in the Illustris-1 simulations are smoother than real galaxies and that their level of complexity at scales that we can resolve, i.e. scales $> \mathcal{O}(\mathrm{kpc})$, is such that it has a negligible lensing effect which is easily absorbed by changes in the macro-model and the source structure. It should also be considered that the simulated galaxies analysed here were chosen to match the global properties of known lenses, but unlike the latter they have not been selected based on their spectral properties. This may potentially lead to a sample of simulated lens galaxies which is more homogeneous and more regular than real ones. We therefore conclude that higher-resolution simulations together with more realistic selection criteria are key to properly characterise their central mass distribution in a way that is not affected by the presence of an un-physical core and investigate the level of complexity in more realistic galaxies and its effect on the detection of low-mass haloes.

3.5 Summary and conclusions

Using a novel Bayesian forward modelling technique based on ABC, we have analysed a sample of mock lensing observations generated from a simulated source and lensing potentials taken from the Illustris-1 simulation. Our results have interesting implications for a variety of measurements, from the study of lensed galaxies to constraints on cosmology and dark matter. We summarise them as follows:

1. Our capability of reproducing the lensed images down to the noise level without fully correctly modelling the cored central mass density distribution of the lenses indicates some form of the source-position transformation (SPT Schneider & Sluse 2014), in line with the previous findings by (Unruh et al. 2017). As a consequence, our reconstructions lead to a *systematic* fractional error on the Hubble constant of 25^{+37}_{-19} per cent (in comparison to a *statical* error of 12^{+6}_{-3} per cent when the shape of the lensing potential is perfectly known). This result is in agreement with the latest analysis of Blum et al. (2020), that shows that cored (dark matter) mass density distributions give rise to approximate Mass Sheet Degeneracies (MSDs), and an error on the inferred Hubble constant. The latest cosmographic analyses (see e.g. Wong et al. 2019) have attempted to break these degeneracies by including the information contained in the kinematic properties of the lens galaxies and the positions of the lensed quasar images. However, the validity of this approach has been recently debated by Kochanek (2020), who has demonstrated that departures from single power-law mass distributions are responsible for a fractional error on the Hubble constant of 30 per cent. While the cores in the simulations analysed in this chapter are artefacts related to the limited resolution, cored mass density distribution in real galaxies may be developed by the effect of baryonic

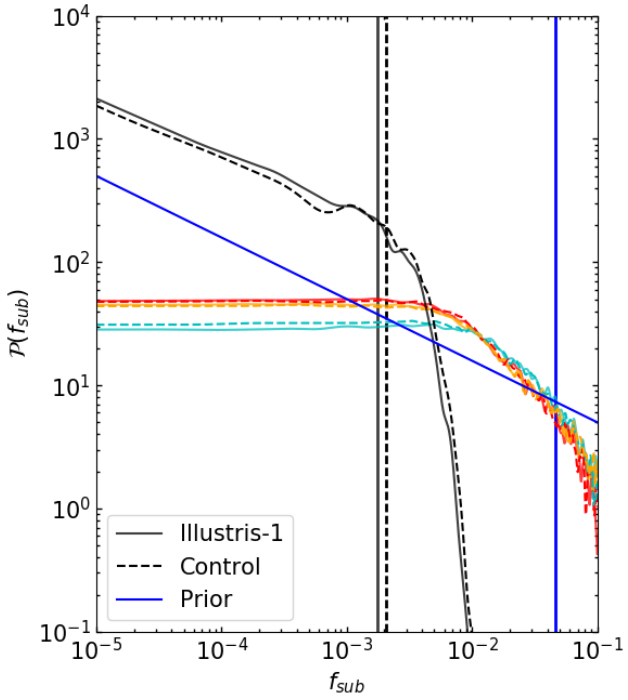

Figure 3.6: **The posterior distribution of f^{sub} for the Illustris-1 sample (continuous curve) and the control sample (dashed curve).** Furthermore, we show the likelihood for one of the Illustris-1 lens galaxies, and show the behaviour for all three projections (red, orange and cyan). We find that the upper confidence limits of f^{sub} do not change much between the mock and the control samples (in particular when comparing the reconstructions on the level of individual lens systems). The prior probability distribution of f^{sub} is shown in blue. The vertical lines indicate the 68 per cent confidence levels of the prior and posterior distributions for the mock and the control sample respectively.

processes (see e.g. Chan et al. 2015) or changes in the dark matter properties (Schive et al. 2014, Spergel & Steinhardt 2000). Moreover, similar additional complexities exist in real galaxies and are related, for example, to the presence of faint discs (Hsueh et al. 2018), bars, or other (baryonic) structures (Gilman et al. 2018, Xu et al. 2013). More generally, there exist many plausible deviations from a smooth power-law distribution, such as broken power laws (see e.g. Du et al. 2020) or multiple component models (see e.g. Nightingale et al. 2018), which can produce comparable degeneracies. Together with the findings of Blum et al. (2020), our results have important implications for the analysis of time delays and a potential solution to the H_0-tension (see e.g. Wong et al. 2019).

2. The MSD we encounter in this chapter further affects the reconstruction of the sources, for which we underestimate the size by 50 per cent on average. This finding is in agreement with Ritondale et al. (2018), who showed for a Lyman-α emitting galaxy

of the BELLS-GALLERY that a correct modelling of the complex lensing potential is essential to reconstruct the properties of the lensed galaxy. In particular, they find that mass components beyond the simple power-law model, in the form of pixelated potential corrections, are required for correct focusing of the source. The reconstructed sources also differ from the input source on small scales, reflecting the presence of noise in the data and the choice of a Gaussian regularising prior. We expect that a log-normal prior or a hyper-prior as in (Rizzo et al. 2018) may provide more stable results as well as a better sensitivity to low-mass haloes by reducing the freedom of the source structure to re-absorb the effect of the latter on the lensed images.

3. We find no degeneracy between complex structures in the lensing potential and low-mass substructures. After having removed all substructures from the simulations we have inferred a fraction of mass in substructure which is consistent with the absence of any substructure and with what was derived from a control sample for which we have a perfect knowledge of the lensing potential. This is in disagreement with Ritondale et al. (2019), Xu et al. (2015) and Gilman et al. (2019b), who have shown both with numerical simulations and observations that un-modelled small scale structures in the lensing mass distribution can lead to significant flux-ratio and surface-brightness anomalies. We believe the origin of this discrepancy to lie partly on the limited resolution of the simulations and partly on the selection criteria adopted in this chapter, as well as a different sensitivity of lensed quasars to changes in the potential.

It is important to notice that a simple power-law model provides a good fit to the lensing observable, implying that the quality of fit is not a good measure to exclude the potential presence of un-modelled mass components. In line with previous works, our results underlines the importance of complex lensing mass distributions that go beyond the standard power-law assumptions as well as the need of extra information (e.g. from stellar kinematics or time delays, Barnabè & Koopmans 2007, Schneider & Sluse 2014) to break existing degeneracies. From a numerical perspective, our work shows the importance of higher-resolution and more realistic numerical simulations to further quantify potential systematic effects in strong gravitational lensing reconstructions.

3.6 Supplementary material

3.6.1 The Einstein radius

The Einstein radius $R_{\rm E}$ is defined as the radius of a circle within which the average mass density is given by the critical density:

$$M(< R_{\rm E}) = A(< R_{\rm E}) \times \Sigma_{\rm crit} = \pi {R_{\rm E}}^2 \times \Sigma_{\rm crit}\,. \tag{3.16}$$

In this case, it is straightforward to calculate the mass within the Einstein radius by integrating the parameterised convergence within the circle and multiplying it by $\Sigma_{\rm crit}$. Solving

the resulting equation for $R_{\rm E}$ yields a definition of the Einstein radius in the chosen parameterisation of the convergence. In order to obtain the elliptical Einstein radius we consider the rescaled radius $R_{\rm E} \rightarrow \widetilde{R_{\rm E}} = R_{\rm E}/\sqrt{q}$, so that:

$$M\left(< \widetilde{R_{\rm E}}\right) = \pi \widetilde{R_{\rm E}}^2 \times \Sigma_{\rm crit}\,, \tag{3.17}$$

which is consistent with the previous definition for $q = 1$. For the parameterized version of the macro-model used in this chapter, we determine the mass within $R_{\rm E} = \widetilde{R_{\rm E}}/\sqrt{q}$ as:

$$\begin{aligned} M(< R_{\rm E}/\sqrt{q}) &= \Sigma_{\rm crit} \int_0^{2\pi} d\theta \int_0^{R_{\rm E}/\sqrt{q}} d\rho\; q\rho \times \kappa^{\rm macro}(\rho) = \\ &= \Sigma_{\rm crit} 2\pi \int_0^{R_{\rm E}/\sqrt{q}} d\rho\; q\rho \times \frac{\kappa_0 \left(2 - \frac{\gamma}{2}\right)\, q^{-1/2}}{2(r_c^2 + \rho^2)^{(\gamma-1)/2}} = \\ &= \Sigma_{\rm crit} \pi q \times \frac{\frac{1}{2}\left(2 - \frac{\gamma}{2}\right)}{(3-\gamma)/2} \frac{\kappa_0 q^{1/2}}{(r_c^2 + \rho^2)^{(\gamma-3)/2}} \Big|_{\rho=0}^{\rho=R_{\rm E}/\sqrt{q}} = \\ &= \Sigma_{\rm crit} \pi \times \kappa_0 \frac{\left(2 - \frac{\gamma}{2}\right)}{(3-\gamma)}\, q^{1/2} \times \left[\left(r_c^2 + {R_{\rm E}}^2/q\right)^{(3-\gamma)/2} - \left(r_c^2\right)^{(3-\gamma)/2}\right]. \end{aligned} \tag{3.18}$$

The Einstein radius $R_{\rm E}$ for a given set of macro-model parameters is then:

$${R_{\rm E}}^2 = \frac{\kappa_0}{2} \frac{(4-\gamma)}{(3-\gamma)}\, q^{1/2} \times \left[\left(r_c^2 + {R_{\rm E}}^2/q\right)^{(3-\gamma)/2} - \left(r_c^2\right)^{(3-\gamma)/2}\right]. \tag{3.19}$$

For $r_c = 0$ this definition of the Einstein radius matches the one described by Vegetti et al. (2014):

$${R_{\rm E}}^{\gamma-1} = \frac{\kappa_0}{2} \frac{(4-\gamma)}{(3-\gamma)}\, q^{(\gamma-2)/2}\,. \tag{3.20}$$

It further reduces to the definition by Kormann et al. (1994) in the isothermal case where $\gamma = 2$, i.e. $R_{\rm E} = \kappa_0$.

3.6.2 Mass function normalisation and mean number of substructures

In this section, we describe the derivation of the mass function normalisation constant from the condition that we have an average fraction of mass contained in substructures within the Einstein radius $f^{\rm sub}(< R_{\rm E})$. This condition can be written explicitly as:

$$f^{\rm sub}(< R_{\rm E}) = \left\langle \frac{\sum_{\rm k=1}^{N^{\rm sub}} m_k}{M(< R_{\rm E})} \right\rangle_{\mathcal{P}(\eta^{\rm sub}, < R_{\rm E})} = \frac{\mu(< R_{\rm E}) < m >_{\mathcal{P}(m)}}{M(< R_{\rm E})}, \tag{3.21}$$

where $\mu(< R_{\rm E})$ is the mean number of objects within the Einstein radius and $< m >_{\mathcal{P}(\rm m)}$ is the average mass of a substructure. $\mathcal{P}(N^{\rm sub}, \{m_k, \rho_k\})$ is the probability distribution associated with a substructure realisation with $N^{\rm sub}$ clumps, each with their respective mass m_k and radial position ρ_k.

The amplitude of the mass function n_0 can be determined by requiring that $\mu(< R_{\rm E})$ matches the definition given in equation (3.21), i.e.:

$$\mu(< R_{\rm E}) = A(< R_{\rm E}) \times \int_{10^6 {\rm M}_\odot}^{10^{10} {\rm M}_\odot} dm \, \frac{dn^{\rm sub}}{dm} \, . \tag{3.22}$$

Using the above definition for the Einstein radius, the proportionality constant in equation (3.4) is:

$$n_0 = \frac{f^{\rm sub}(< R_{\rm E}) \times \Sigma_{\rm crit}}{< m >_{\mathcal{P}(\rm m)} \times \int_{10^6 {\rm M}_\odot}^{10^{10} {\rm M}_\odot} dm \; m^{-\xi}} \, . \tag{3.23}$$

We assume that the density of substructures is constant throughout the extensions of the mock observation. In order to obtain the mean number of objects μ in the considered area, we integrate the mass function over m and multiply it with the area $A^{\rm data}$ covered by the mock observation:

$$\mu = A^{\rm data} \times \int_{10^6 {\rm M}_\odot}^{10^{10} {\rm M}_\odot} dm \frac{dn^{\rm sub}}{dm} = A^{\rm data} \times \frac{f^{\rm sub}(< R_{\rm E}) \times \Sigma_{\rm crit}}{< m >_{\mathcal{P}(\rm m)}} \, . \tag{3.24}$$

3.6.3 Jeffreys' prior

Jeffreys' prior for the inference of $f^{\rm sub}$ would require us to consider the full likelihood function $\mathcal{P}(\vec{d}|\mu)$ and marginalise it over the space data realisations (Jeffreys 1946). As this is computationally prohibitive, we consider the Jeffreys' prior of the probability density $\mathcal{P}(N^{\rm sub}|\mu)$, which being a Poisson distribution, leads to a prior of $\mathcal{P}(f^{\rm sub}) \propto \left(f^{\rm sub}\right)^{-1/2}$. We expect this simplified prior distribution to remain non-informative for our reconstruction process. Moreover, it has the advantage of being integrable under the assumption of a finite upper limit on $f^{\rm sub}$ and allows us to correctly probe a parameter which may vary by several orders of magnitude. We choose a range of $f^{\rm sub} \in [10^{-6}, 10^{-1}]$, which includes the case that $\mu(f^{\rm sub}) + \sqrt{\mu(f^{\rm sub})} < 1$, i.e. that there are no substructures within one standard deviation of the Poisson distribution $P(N^{\rm sub}|\mu)$.

3.6.4 Motivating an ABC approach

One of the main goals of this chapter is the reconstruction mass function parameters. The likelihood of this problem under the model assumptions specified in Section 3.3 is:

$$\mathcal{P}(\vec{d}|\{f_{\rm sub}, m_{\rm hm}\}) = \sum_{N^{\rm sub}=0}^{\infty} \left(\prod_{n=1}^{N^{\rm sub}} \int dm_n \int d\vec{x}_n \right) \mathcal{P}(m_n, \vec{x}_n|\{f_{\rm sub}, m_{\rm hm}\}) \times \\ \times \mathcal{P}\left(N \middle| \{f_{\rm sub}, m_{\rm hm}\}\right) \times \mathcal{P}\left(\vec{d} \middle| L(\eta^{\rm macro}, \{\vec{x}_n, m_n\}_{n=1}^{N^{\rm sub}})\right) . \quad (3.25)$$

While we can analytically calculate the last factor of this product, which corresponds to equation (3.9) once $\mathbf{L}(\eta^{\rm macro}, \eta^{\rm sub} := \{\vec{x}_n, m_n\}_{n=1}^{N^{\rm sub}})$ is determined, the combination with the integrals over positions $\vec{x}_n$ and masses m_n as well as the summation over $N^{\rm sub}$ pose a challenge for analytical evaluation.

In order to the evaluate expression (3.25) one may follow a reversible jump MCMC approach as done by Brewer et al. (2015) or Daylan et al. (2018). However, these methods usually require to reduce the parameters in the model, for example, by switching to a parametric source rather than a free form surface brightness field. Furthermore, these approaches tend to be computationally expensive.

Another way is to marginalize over the substructure realisations rigorously. Hsueh et al. (2019) follow this approach and explicitly integrate over $O(10^6)$ substructure realisations per $M_{\rm hm}$ value in their regular grid of mass function parameters. In this case, the flux-ratio analysis benefits from the small number of observables, i.e. $\dim(\vec{d}) \approx O(10)$. The evaluation of the analogous function to equation (3.9) is therefore much faster than in the case of extended arcs with $\dim(\vec{d}) \approx O(100^2)$. Notice that the computational cost of expression (3.25) would increase even further with the inclusion of line-of-sight haloes.

Given that our goal is to keep the physical model as general as possible, that we can generate substructure realisations with forward modelling, and that for a (somewhat) limited number of samples one will obtain conservative results with an ABC approach, we decided to follow this approach.

3.6.5 Importance sampling and weighting of samples

Standard importance sampling uses weights w_i and applies them to the samples that approximate the posterior distribution:

$$\mathcal{P}(\eta^{\rm macro}|\vec{d}) = \frac{1}{K^{\rm acc}} \sum_{\rm i \in acc} \delta^D(\eta^{\rm macro} - \eta_i^{\rm macro}) \times w_i \, . \quad (3.26)$$

However, in our implementation, as expressed by equation (3.12), we do not include any weight, as this is done implicitly by the acceptance strategy. In order to justify this, we consider the following relation:

$$\mathcal{P}(\vec{x}|\vec{d}) = \frac{\mathcal{P}(\vec{d}|\vec{x})}{\mathcal{P}(\vec{d})}\mathcal{P}(\vec{x}) = \frac{\mathcal{P}(\vec{d}|\vec{x})}{\mathcal{P}(\vec{d})}\frac{\mathcal{P}(\vec{x})}{Q(\vec{x})}Q(\vec{x}) = \frac{\widetilde{\mathcal{P}}(\vec{d}|\vec{x})}{\mathcal{P}(\vec{d})}Q(\vec{x}) \, , \quad (3.27)$$

with $\widetilde{\mathcal{P}}(\vec{d}|\vec{x}) = \mathcal{P}(\vec{d}|\vec{x})\frac{\mathcal{P}(\vec{x})}{Q(\vec{x})}$. This relation implies that one can interpret the importance sampling as a change in the likelihood[3] and the prior in such a way that their changes cancel each other out.

As we use the $P(\vec{d}|\vec{x})$ as our distance measure, we include the weight in the distance that decides whether or not a sample gets accepted rather than multiplying the weight to the corresponding sample. In doing so, we free ourselves from accounting for sample weights at the price of the caveat that the posterior distribution of the importance sampled parameters tend to resemble the shape of the proposal distribution.

In our case, we apply importance sampling on the macro-model parameters η^{macro}, for which previous tests have shown that their posteriors closely resemble Gaussian distributions, therefore justifying our choice of the proposal distribution. In order to improve the approximation, one could use Gaussian mixture models instead of a single Gaussian distribution to describe the proposition distribution. This form of the proposal distribution would allow one to account for more general shapes of posteriors.

[3]Which we refer to as the evidence in this chapter, as it is the evidence when reconstructing the source.

Chapter 4

Strong gravitational lensing and the nature of dark matter: future prospects

4.1 Introduction

The halo and subhalo mass functions of viable WDM models mainly differ from their CDM counterparts at masses $M \leq 10^8 M_\odot$ (see Figure 1.3). Strong constraints on the properties of dark matter with galaxy-galaxy strong lensing observations, therefore, require data that is sensitive below this limit.

Several aspects of the data determine the sensitivity of gravitational lensing observations to the presence of low-mass haloes. At a fixed signal-to-noise ratio and angular resolution, the sensitivity is higher for lensed images (and therefore background sources) with a higher degree of structure in their light distribution (see Koopmans 2005, Rau et al. 2013). The lowest detectable mass scales with both angular resolution and signal-to-noise ratio at a fixed amount of source structure. For subhaloes with an NFW profile, the size of the scale radius determines the maximum deflection angle. As a result, the mass for which this radius is equal to the FWHM of the PSF constitutes a rough lower limit for detectable substructures. At a given resolution, the lowest detectable mass scales linearly with the signal-to-noise ratio. A change in the latter also leads to a general increase in the fraction of image plane area that is sensitive to substructures, increasing the overall number of detectable objects as a result (Despali et al. in prep.).

Current lens samples provide more robust yet less stringent constraints on the nature of dark matter than other astrophysical probes (see Chapter 2). This limitation is related to the lack of large-enough samples of lens systems with high-enough data quality. In the future, large surveys (e.g., with Euclid and the Square Kilometer Array, Laureijs et al. 2011, Pawase et al. 2014, Collett 2015) coupled with high-resolution follow-up with the next generation of ground telescopes (e.g. the EELT, Neichel et al. 2018) will significantly increase the amount of high-quality data. In this chapter, we explore whether alternative observing strategies could improve upon the current sensitivity, now.

There exist mainly two ways to increase the number of detectable haloes: (i) probing

a larger cosmological volume with high-redshift lens systems (Ritondale et al. 2019); (ii) increasing the sensitivity with spatially resolved spectral observations (Hezaveh et al. 2013). It should be noticed that the first approach has also the advantage of increasing the sensitivity to lower-mass haloes, as those located in the foreground have a stronger lensing effect than a subhalo of the same mass (Despali et al. 2018). Here, we investigate these two observational strategies by focusing on real and simulated observations with ALMA and the JWST. In particular, the chapter is structured as follows: in Section 4.2 we present the adopted lens modelling technique. We provide a brief description of the data in Section 4.3. In Section 4.4 we present our main results and follow with a discussion in Section 4.5.

4.2 Lens modelling

In this section, we describe the lens modelling technique used throughout this chapter. We start by describing our smooth modelling approach (i.e. ignoring the presence of low-mass haloes), followed by a discussion on the sensitivity function.

4.2.1 Physical model

We assume that the data vector $\vec{d}$ is related to the model $\vec{s}$ of the surface brightness distribution of the background galaxy as follows:

$$\vec{d} = \mathbf{ML}\vec{s} + \vec{n}\,, \tag{4.1}$$

where $\mathbf{M}$ encodes the telescope response. In the case of optical observations, $\vec{d}$ is the surface brightness distribution of the lensed images, and $\mathbf{M}$ represents the effect of the PSF (see e.g. Vegetti & Koopmans 2009a). At radio wavelength $\vec{d}$ is a set of visibilities, and $\mathbf{M}$ includes a degridding, a Fast Fourier Transform operation, an apodization correction, and a zero-padding/masking operation Z (see Powell et al. 2020, for more details). In both cases, the noise $\vec{n}$ can be assumed to be Gaussian and uncorrelated between data points. The lensing operator $\mathbf{L}$ describes the lensing displacement field.

We follow the previous chapter and assume a cored elliptical power-law profile plus external shear for the main lens and a spherical NFW profile for subhaloes and field haloes. We refer to Sections 1.1.3 and 3.3.1 for more details on these models.

We reconstruct the source surface brightness distribution on an adaptive grid mesh and assume a regularization prior according to Suyu et al. (2006), which renders sources with less curvature more probable:

$$\mathcal{P}(\vec{s}|\lambda) = \mathcal{G}(\vec{s}, \lambda^{-1}\mathbf{S})\,, \tag{4.2}$$

where the covariance matrix is chosen such that $\mathbf{S}^{-1} = \mathbf{R}^T\mathbf{R}$ with $\mathbf{R}\vec{s}$ determining the curvature of a source. The strength of this regularisation λ is assumed to follow a log uniform prior. We note that when modelling data with spatially resolved spectra information, we use the latest version of the Vegetti & Koopmans (2009a) code as extended to the 3D domain by Rizzo et al.

(2018) and to the interferometric domain by (Powell et al. 2020). In this case, the source and data vectors are defined by the concatenation of the source and data in each spectral channel.

4.2.2 Smooth model reconstruction

We start by modelling the data assuming a smooth model $\mathcal{M}_\mathrm{s}$ that does not contain any subhalo. Using a three-level Bayesian inference approach, we determine the maximum-a-posteriori source $\vec{s}$ and lens parameters η^macro, as well as the Bayesian evidence of this model.

Linear optimization

At the first level of this inference, we infer the source surface brightness distribution for a set value of the lens parameters and source regularization. In particular, we maximise the following posterior:

$$\mathcal{P}(\vec{s}|\vec{d},\eta^\mathrm{macro},\lambda,\mathcal{M}_\mathrm{s}) = \mathcal{P}(\vec{d}|\vec{s},\eta^\mathrm{macro},\mathcal{M}_\mathrm{s}) \times \frac{\mathcal{P}(\vec{s}|\eta^\mathrm{macro},\lambda,\mathcal{M}_\mathrm{s})}{\mathcal{P}(\vec{d}|\eta^\mathrm{macro},\lambda,\mathcal{M}_\mathrm{s})}\,, \tag{4.3}$$

which we can rewrite under the assumptions presented in Section 4.2.1 as follows:

$$\mathcal{P}(\vec{s}|\vec{d},\eta^\mathrm{macro},\lambda,\mathcal{M}_\mathrm{s}) = \mathcal{G}(\vec{d}-\mathbf{ML}\vec{s},\mathbf{N}) \times \frac{\mathcal{G}(\vec{s},\lambda^{-1}\mathbf{S})}{\mathcal{P}(\vec{d}|\eta^\mathrm{macro},\lambda,\mathcal{M}_\mathrm{s})}\,. \tag{4.4}$$

We calculate the maximum posterior source $\vec{s}_\mathrm{mp}$ by requiring that $\nabla_{\vec{s}}\mathcal{P}(\vec{s}|\vec{d},\eta^\mathrm{macro}) = 0$, which leads to a linear system that can be solved numerically:

$$\left[\mathbf{L}^T\mathbf{M}^T\mathbf{N}^{-1}\mathbf{ML} + \lambda\mathbf{S}^{-1}\right]\vec{s}_\mathrm{mp} = \mathbf{L}^T\mathbf{M}^T\mathbf{N}^{-1}\vec{d}\,. \tag{4.5}$$

Non-linear optimization

At the second level of inference, the optimal lens parameters η^macro and source regularization λ are determined, by maximising the following posterior:

$$\begin{aligned}\mathcal{P}(\eta^\mathrm{macro},\lambda|\vec{d},\mathcal{M}_\mathrm{s}) &= \mathcal{P}(\vec{d}|\eta^\mathrm{macro},\lambda,\mathcal{M}_\mathrm{s}) \times \frac{\mathcal{P}(\eta^\mathrm{macro},\lambda|\mathcal{M}_\mathrm{s})}{\mathcal{P}(\vec{d}|\mathcal{M}_\mathrm{s})} = \\ &= \int d\vec{s}\,\mathcal{P}(\vec{d}|\vec{s},\eta^\mathrm{macro},\lambda,\mathcal{M}_\mathrm{s}) \times \mathcal{P}(\vec{s}|\eta^\mathrm{macro},\lambda,\mathcal{M}_\mathrm{s}) \times \frac{\mathcal{P}(\eta^\mathrm{macro},\lambda|\mathcal{M}_\mathrm{s})}{\mathcal{P}(\vec{d}|\mathcal{M}_\mathrm{s})} = \\ &= \mathcal{G}(\vec{d},\left[\mathbf{N}+\mathbf{MLSL}^T\mathbf{M}^T\lambda^{-1}\right]) \times \frac{\mathcal{P}(\eta^\mathrm{macro},\lambda|\mathcal{M}_\mathrm{s})}{\mathcal{P}(\vec{d}|\mathcal{M}_\mathrm{s})}\,.\end{aligned} \tag{4.6}$$

To this end, a non-linear optimization scheme is applied to λ and η^macro in an alternating fashion until the global maximum is found. We notice that the evaluation of the posterior above implicitly contains the inference of $\vec{s}_\mathrm{mp}$ as described in the previous section (see e.g. Rizzo et al. 2020).

Evidence calculation

At the third level of inference, we calculate the Bayesian evidence $\mathcal{P}(\vec{d}|\mathcal{M}_s)$ which is given by:

$$\mathcal{P}(\vec{d}|\mathcal{M}_s) = \int d\eta^{\rm macro} d\lambda\, \mathcal{P}(\vec{d}, \eta^{\rm macro}, \lambda|\mathcal{M}_s) = \\ = \int d\eta^{\rm macro} d\lambda\, \mathcal{P}(\vec{d}|\eta^{\rm macro}, \lambda, \mathcal{M}_s) \times \mathcal{P}(\eta^{\rm macro}, \lambda|\mathcal{M}_s)\,. \quad (4.7)$$

This multidimensional and highly non-linear integral is evaluated using the nested sampling algorithm MULTINEST (Feroz et al. 2009), which generates samples of the posterior distribution of λ and $\eta^{\rm macro}$ as a byproduct. These samples can be then used to determine the mean and uncertainty of individual parameters.

4.2.3 Sensitivity to low-mass haloes

In order to constrain the halo and subhalo mass functions, one needs to define and quantify the lowest mass that can be detected with a given observation. Extending on the work by Vegetti et al. (2014) and Ritondale et al. (2019), we define the lowest detectable substructure as the one that corresponds to a Bayes factor K, that fulfills $|\log K| = |\log \mathcal{P}(\vec{d}|\eta^{\rm sub}, \mathcal{M}_p) - \log \mathcal{P}(\vec{d}|\mathcal{M}_s)| > 50$. $\mathcal{P}(\vec{d}|\eta^{\rm sub}, \mathcal{M}_p)$ is the evidence of a perturbed model M_p that contains a substructure of a given mass and position defined by $\eta^{\rm sub}$. Under the assumption of Gaussian errors, this limit corresponds roughly to a 10-σ detection threshold. This limit was shown to minimise the number of both false positive and negative detections (see Ritondale et al. 2019 for more details). Using a less conservative threshold of 5σ (corresponding to a Bayes factor of 12.5), Hezaveh et al. (2016) have reported the detection of a substructure in the lens system SDP.81. Here, we make use of both thresholds to define what is the lowest-mass detectable subhalo. We then use the relations derived by (Despali et al. 2018) to derive the corresponding lowest-mass detectable for line-of-sight haloes as a function of their redshift.

As discussed in the introduction of this chapter, the sensitivity is tightly related to the structure of the lensed images. As such, it is expected to vary not only from one observation to the other but also between different pixels on the image plane of the same data set. We refer to the lowest detectable mass as a function of position on the image plane as the sensitivity function.

4.3 Data

The data analysed in this section consist of real and synthetic observations from ALMA and the JWST, respectively. Below, we provide further details on each data set.

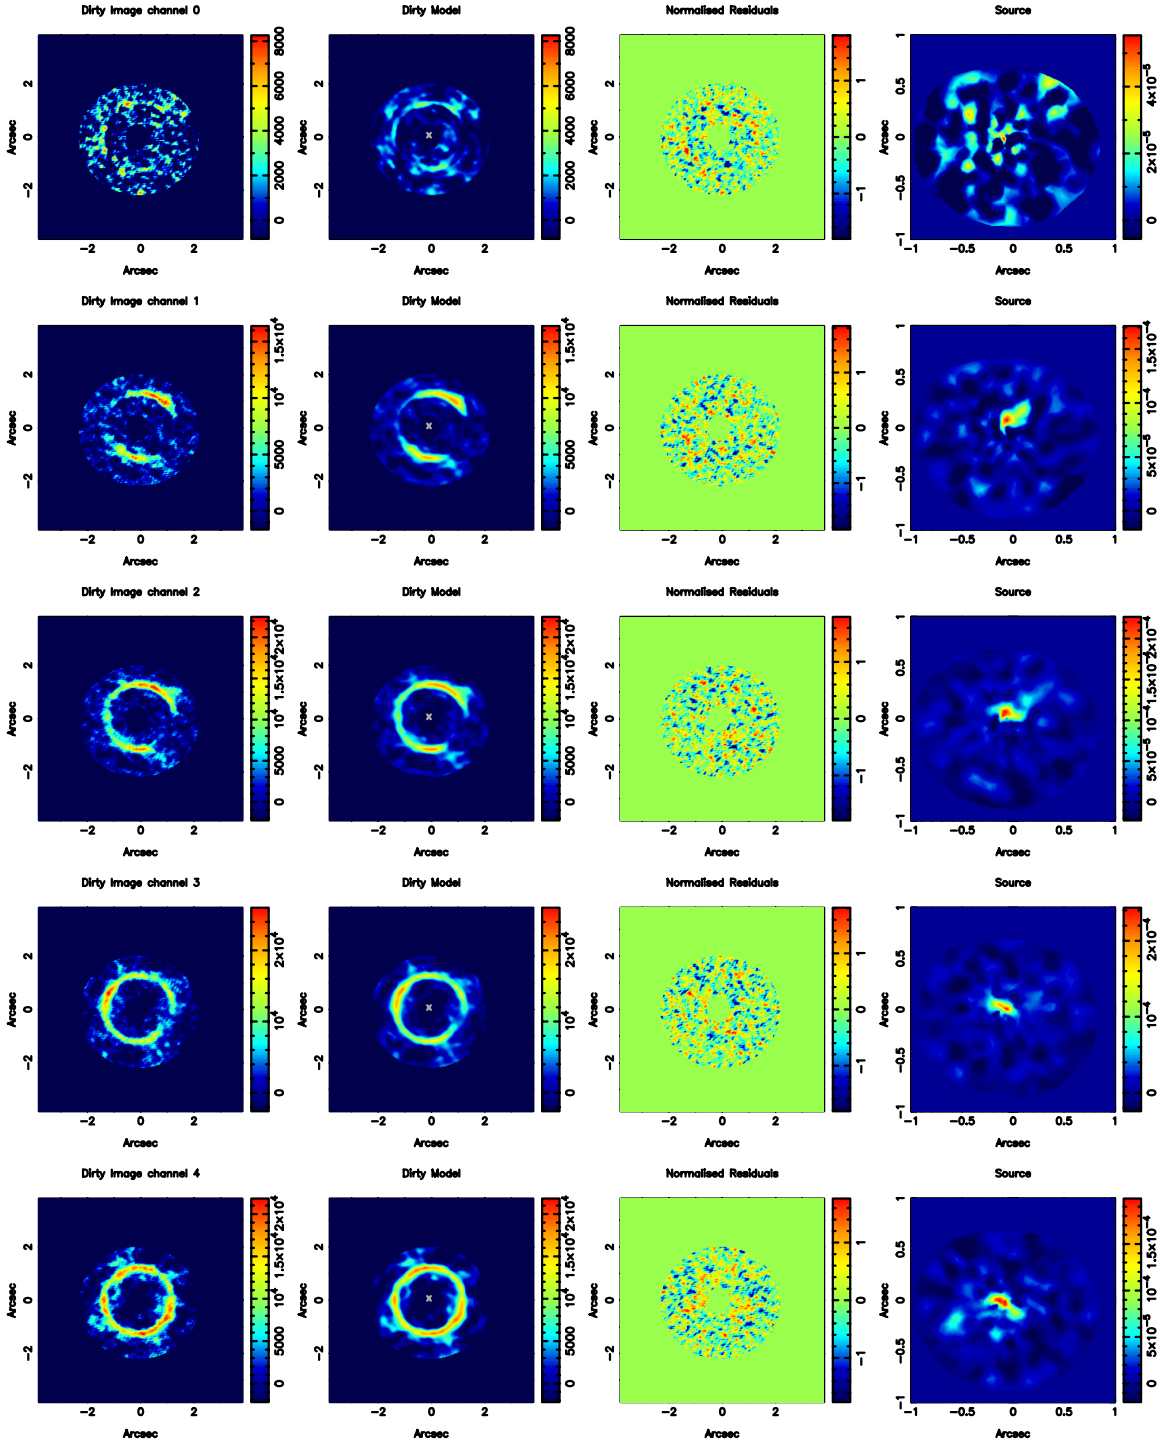

Figure 4.1: continues on next page.

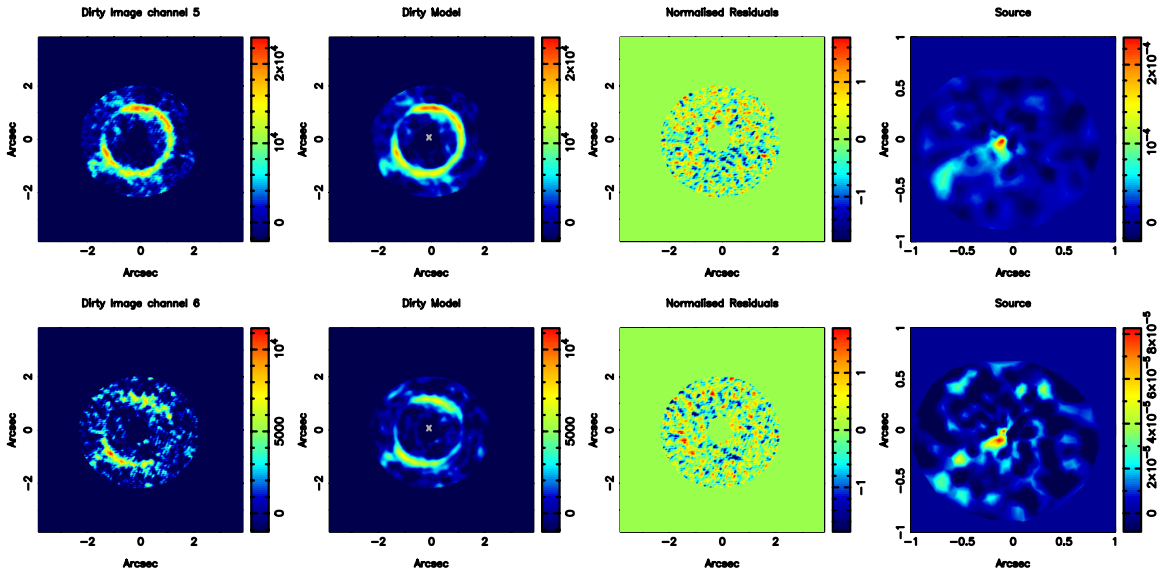

Figure 4.1: **Data and smooth model reconstruction of ALMA-D$_1$.** From left to right, the dirty image of the data set, the reconstructed dirty model, the residuals between model and data, and the reconstructed source brightness distribution. Every row corresponds to one of the velocity channels of ALMA-D$_1$.

4.3.1 ALMA data

As an example of a relatively good-quality ALMA data set, we consider the gravitationally lensed galaxy SPT–S J041839-4751.9 (Hezaveh et al. 2013, Vieira et al. 2013). The data were taken from the ALMA archive (project code 2016.1.01499.S, Principal investigator, K. Litke), calibrated and pre-processed by Rizzo et al. (2020). The source and the lens galaxy are located at $z_l = 0.3$ and $z_s = 4.2$, respectively.

We focus both on the thermal continuum emission from the dust and the emission from the 158μm fine-structure line of ionized carbon [CII]. The data were chosen for their angular and spectral resolutions of 0.3 arcsec and 40 km/s, respectively, and their high median signal-to-noise ratio of 10 over the spectral channels.

To improve the signal-to-noise ratio and the computational efficiency, we use the Common Astronomy Software Applications package (CASA, McMullin et al. 2007) to average the [CII] data into groups of seven (in the following data ALMA-D$_1$) and one (i.e. the zeroth-moment, in the following data ALMA-D$_2$) velocity channels. We average the continuum data into groups of eight (in the following data ALMA-D$_3$) and one channel (in the following data ALMA-D$_4$). The different data sets are shown in Figures 4.1 to 4.4, and their properties are summarised in Table 4.1.

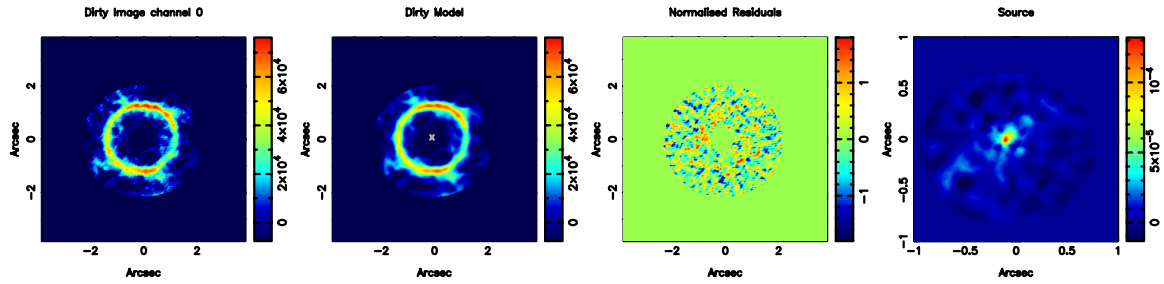

Figure 4.2: **Data and smooth model reconstruction of ALMA-D$_2$.** From left to right, the dirty image of the data set, the reconstructed dirty model, the residuals between model and data, and the reconstructed source brightness distribution.

4.3.2 JWST simulated data

To create a simulated JWST data set, we consider observations with the HST-F606W-filter of the gravitational lens system SL2SJ021325-074355 from the SL2L sample (Ruff et al. 2011, Gavazzi et al. 2012, Sonnenfeld et al. 2013a,b, 2015). The lens and source galaxy have a redshift of $z_l = 0.7$ and $z_s = 3.5$, respectively. We choose this system because it is the brightest one in the SL2L sample with a source redshift $z_s > 3$, making it potentially an interesting JWST target.

We then use the best lens and source model inferred by Sonnenfeld et al. (2013a) to create a simulated observation with the F070W filter (0.621-0.781 μm) of the Nircam instrument on board the JWST. Specifically, Sonnenfeld et al. (2013a) modelled SL2SJ021325-074355 assuming a single isothermal ellipsoidal profile for the lens mass distribution and a pixellated free-form model for the light distribution of the background source. For simplicity, we create our mock data by assuming for the source a single Sérsic component with a profile and amplitude matching those of the original HST observations. The flux is then extrapolated to the wavelength range of the F070W filter, assuming that the continuum emission follows a power-law with a logarithmic continuum slope of -1.5 (see e.g. Ritondale et al. 2018).

A realistic sky background is determined using the background calculator of the Euclid mission (Laureijs et al. 2011)[1] and added to the flux of the lensed images. The total sky brightness is then convolved with the PSF of the JWST-F070W-filter observations as calculated from the WEBBPSF python package (Perrin et al. 2014). The corresponding FWHM roughly matches the pixel size of $31\times31\text{mas}^2$. The noise in all pixels is generated by assuming Poisson noise with a standard deviation of $\sigma_n = \sqrt{s \cdot t + c}$, where c is the average readout error of the Nircam instrument taken from the JWST user documentation[2] and t is the exposure time, that was set to $t = 1800$s to match the original observation (Sonnenfeld et al. 2013b), resulting in a peak signal-to-noise ratio of 11.7. The corresponding data set, JWST-D, is shown in 4.5.

[1] https://irsa.ipac.caltech.edu/applications/BackgroundModel/
[2] https://jwst-docs.stsci.edu/

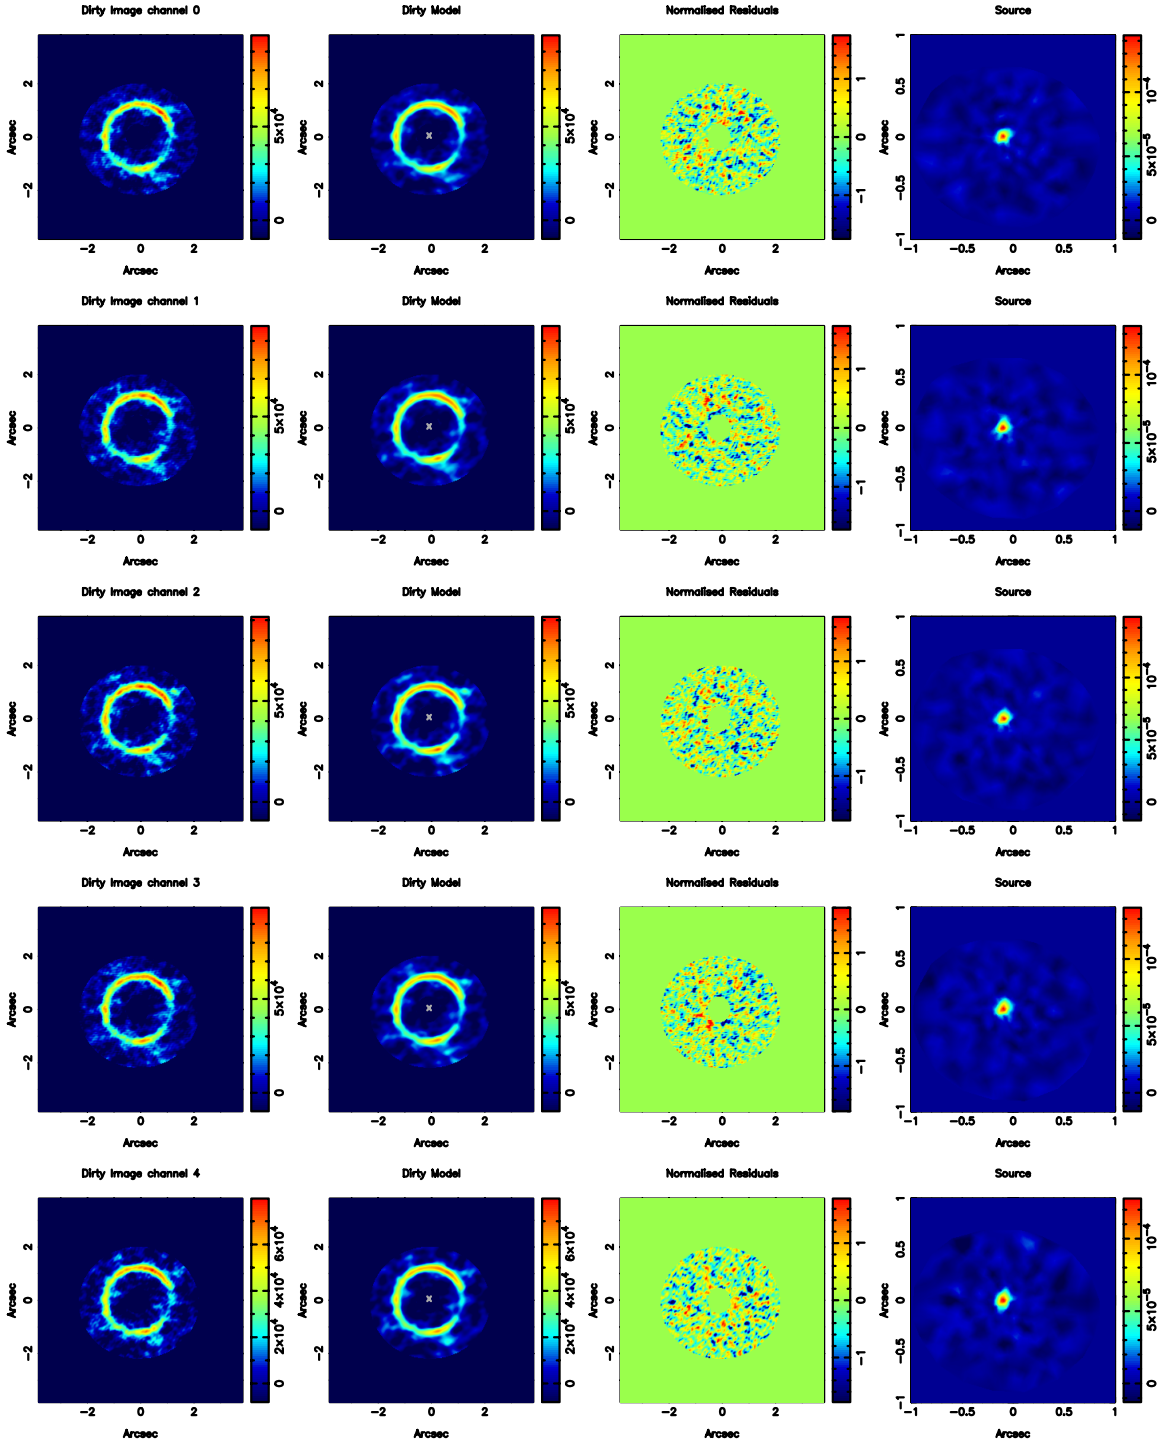

Figure 4.3: continues on next page.

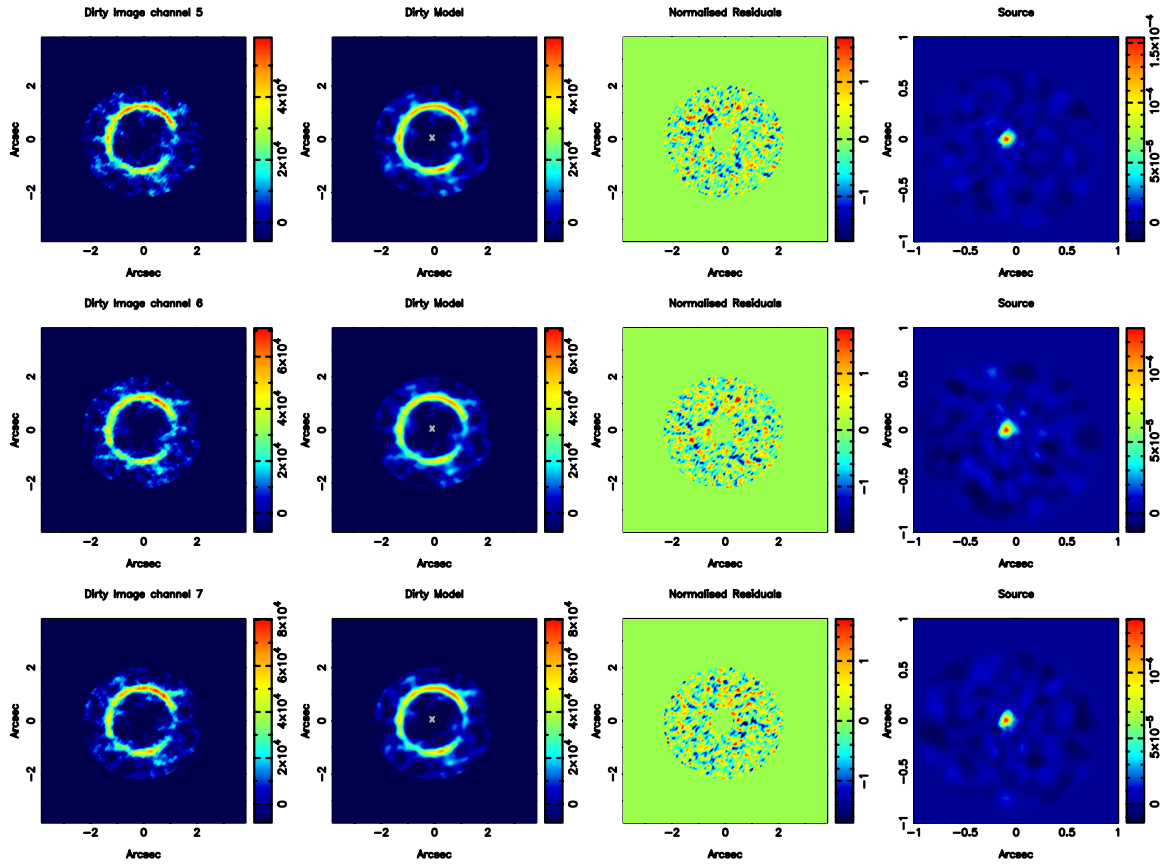

Figure 4.3: **Data and smooth model reconstruction of ALMA-D$_3$.** From left to right, the dirty image of the data set, the reconstructed dirty model, the residuals between model and data, and the reconstructed source brightness distribution. Every row corresponds to one of the velocity channels of ALMA-D$_3$.

4.4 Results

In this section, we present the smooth modelling results and the sensitivity function for each of the considered data sets. As discussed in Section 4.2.3, calculating the sensitivity function requires identifying for each pixel on the image plane the lowest substructure mass that has a strong enough effect to be supported or ruled out by the data at a chosen statistical threshold. In practice, one should calculate the Bayes factor for each pixel on the image plane by considering a set of perturbed models with a fixed substructure mass between $M_{\rm min}$ and $M_{\rm max}$. Here, we choose a lower limit of $10^8 M_\odot$, because we do not expect to be able to probe lower masses for the given angular resolution of the data. We do not consider substructures with a mass $M_{\rm sub} > 10^{10} M_\odot$, because we are only interested in the number of detectable subhaloes that are completely dark or too faint to be detected otherwise.

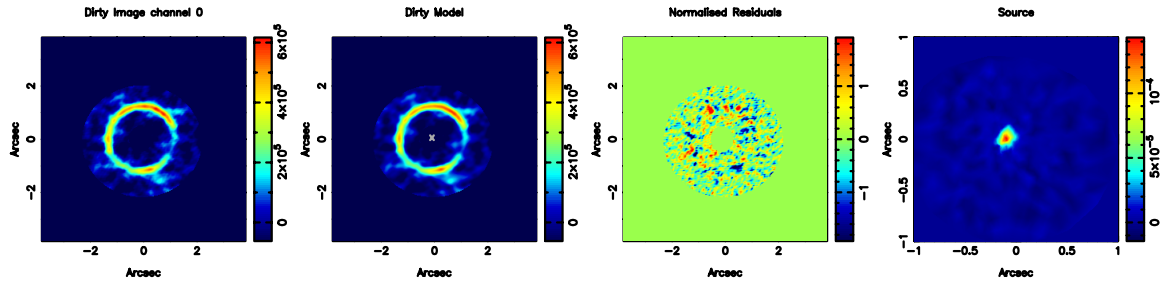

Figure 4.4: **Data and smooth model reconstruction of ALMA-D$_4$.** From left to right, the dirty image of the data set, the reconstructed dirty model, the residuals between model and data, and the reconstructed source brightness distribution.

Table 4.1: **Background information of the JWST and ALMA data sets considered in this chapter.** From left to right, the name of the data set, the number of channels N_{chan}, the peak signal-to-noise ratio $s_{\mathrm{max}}/\sigma_n$, and the beam width in both dimensions in terms of the full width half maxima $b_x \times b_y$ as obtained by fitting a 2D normal distribution to the dirty beam.

Name	Target	N_{chan}	$s_{\mathrm{max}}/\sigma_n$	$b_x \times b_y$
		1	1	mas^2
ALMA-D$_1$	[CII]	7	12.489	177×209
ALMA-D$_2$	[CII]	1	10.994	179×212
ALMA-D$_3$	Continuum	8	11.770	184×221
ALMA-D$_4$	Continuum	1	44.285	181×221
JWST-D	Continuum	1	11.700	31×31

4.4.1 ALMA data

Figures 4.1 to 4.4 show the results of the smooth modelling step for the data sets ALMA-D$_1$ to ALMA-D$_4$, respectively. We find that all data sets can be modelled almost to the noise level with residuals at the 2-σ level at most. The corresponding lens parameters are listed in Table 4.2. The values inferred from each data set are close but do exclude each other by a few standard deviations, which is partly related to a strong degeneracy between the amplitude and the logarithmic slope of the power-law mass profile, and partly to the tendency of nested sampling algorithms to underestimate the uncertainty of posterior parameters (Higson et al. 2018).

Figures 4.6 and 4.7 show the spatial dependency of $\log K$ for a $10^{10} M_\odot$ substructure for each of the four data sets. By inspecting these two figures we can see that:

1. a large fraction of pixels has a $|\log K|$ between zero and 12.5. This finding implies that irrespective of the data (continuum vs emission line) and the level of averaging, the lensing effect of a $10^{10} M_\odot$ substructure is not strong enough to induce an effect which is detectable at a significant statistical level, in most pixels.

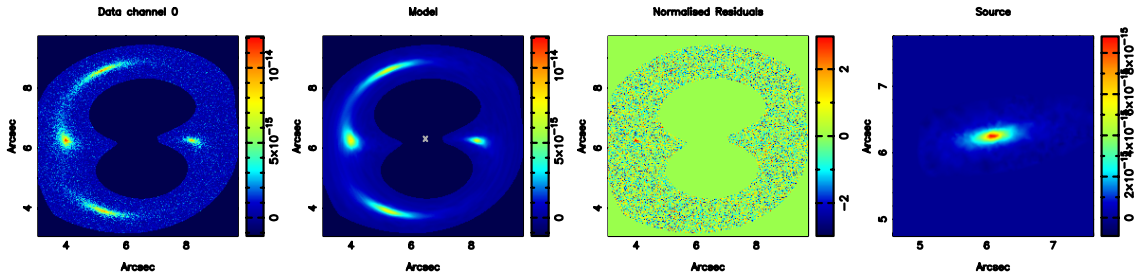

Figure 4.5: **Mock data and smooth model reconstruction of JWST-D.** From left to right, the original data set, the reconstructed dirty model, the residuals between model and data, and the reconstructed source brightness distribution.

2. By comparing the data set ALMA-D_1 with ALMA-D_2 and ALMA-D_3 with ALMA-D_4, we see that the range of $\log K$ values is larger for the less averaged data for both the thermal continuum and the [CII] line. This increase is related to the fact that the averaging process reduces the level of structure in the lensed images. This effect is stronger for the [CII] line, where the variation of the lensed images is greater between the different channels.

3. From the comparison of the data set ALMA-D_1 with ALMA-D_3 and ALMA-D_2 with ALMA-D_4, we find that the emission line observations are less sensitive, reflecting their overall lower signal-to-noise ratio.

   For reasons related to computational speed, we extend the calculation of the Bayes factor to the lower mass bins only for the data sets ALMA-D_3 and ALMA-D_4, which show pixels with $|\log K| > 12.5$ (5-σ detections). Specifically, we focus on those pixels where $|\log K(10^{10} M_\odot)| \gtrsim 5$. The results of this calculation are shown in Figure 4.6. We find that the values of $|\log K(10^8 M_\odot)|$ and $|\log K(10^9 M_\odot)|$ are even lower. This is not necessarily expected, as the presence of a $< 10^{10} M_\odot$ substructure in the data could lead to an increase of $|\log K|$ at lower masses. In particular, we find that none of the lower masses pass the 5-σ threshold.

4.4.2 JWST simulated data

In Figure 4.5, we illustrate the results of the smooth modelling step for the data set JWST-D. We find that a smooth elliptical power-law is able to fit the data down to the noise level (< 3-σ level) with no correlated residuals. The corresponding lens parameters and their uncertainties are listed in Table 4.2 and most of their ground-truth parameters fall within the 2-σ region.

For each pixel on the image plane, we compute the Bayes factor $\log K$ for a substructure with a mass of $10^8 M_\odot$, $10^9 M_\odot$ and $10^{10} M_\odot$ (see Figure 4.8). Notice that because this is a simulated data set with no substructures, it only makes sense to consider negative values of $\log K$. From Figure 4.8 we derive following:

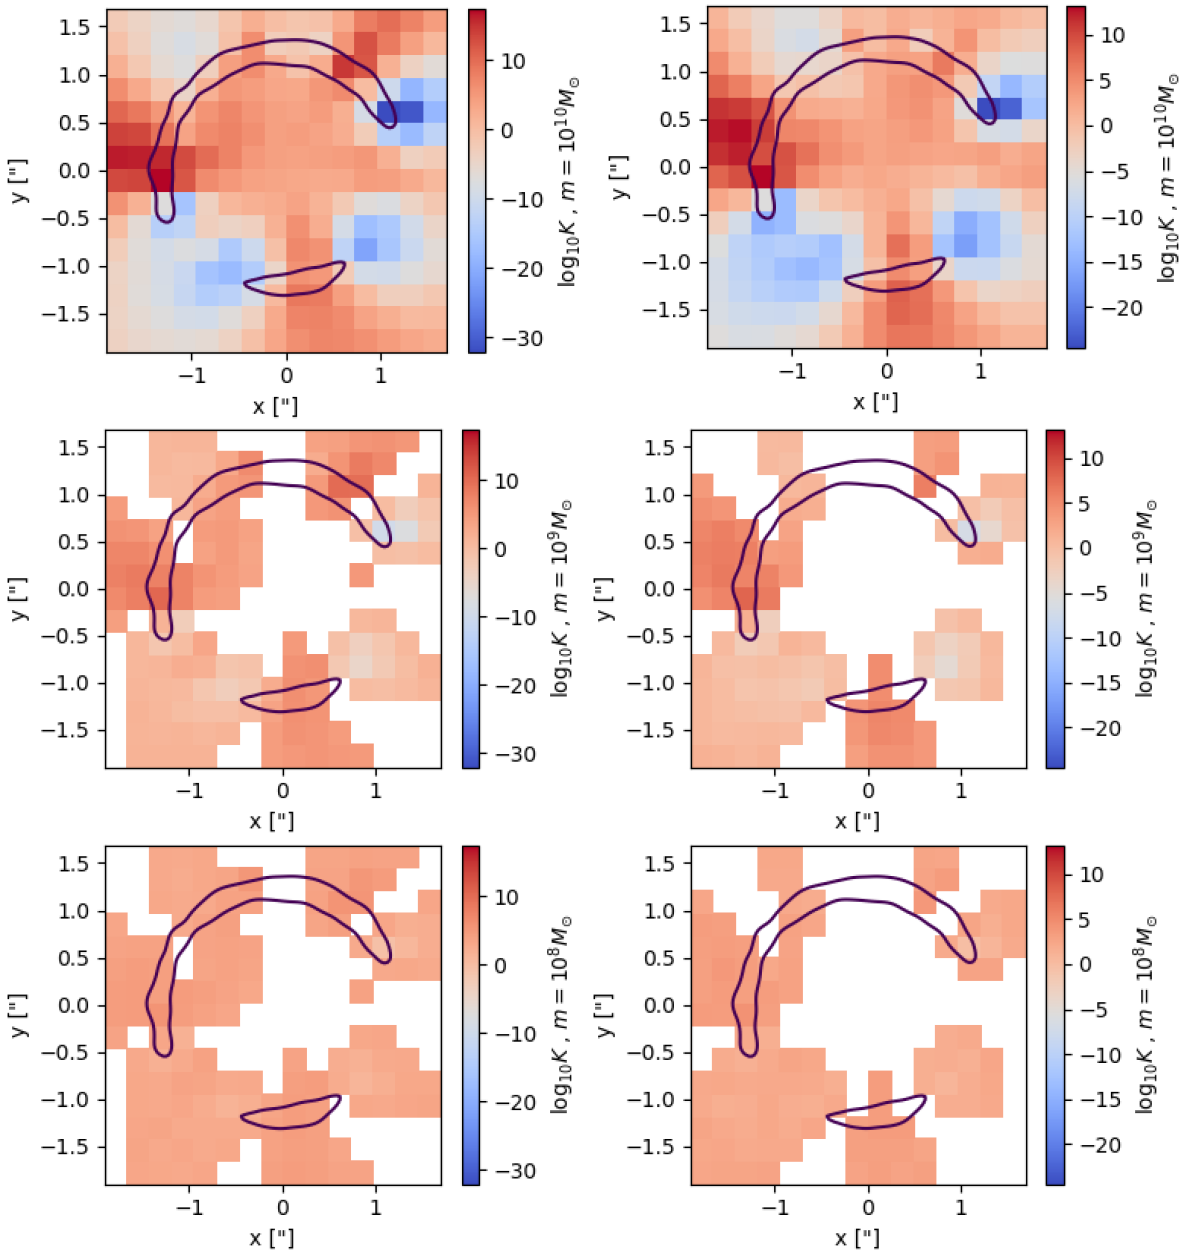

Figure 4.6: **The Bayes factor of a smooth model and a perturbed model containing a substructure for different pixels of the ALMA observation of the dust continuum.** Depicted are the maps for ALMA-D_3 (left) and ALMA-D_4 (right) for perturbed models with a substructure mass of $10^{10} M_\odot$ (top), $10^9 M_\odot$ (middle), and $10^8 M_\odot$ (bottom). The contours show the outline of the (channel averaged) lensed images.

1. despite the slightly improved angular resolution compared to the HST, our JWST observation does not provide an increase in sensitivity relative to benchmark lens samples,

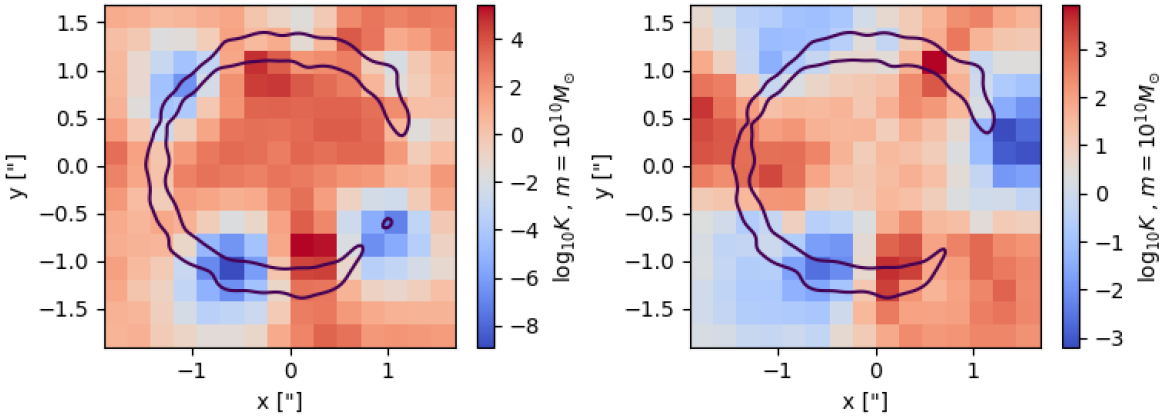

Figure 4.7: **The Bayes factors of a smooth model and a perturbed model containing a substructure for different pixels of the ALMA observation of the [CII] emission line.** Depicted are the maps for ALMA-D_1 (left) and ALMA-D_2 (right) for perturbed a model with a substructure mass of $10^{10}M_\odot$. The contours show the outline of the (channel averaged) lensed images.

such as the SLACS and BELLS-GALLERY sample. This result is related to our choice of a relatively smooth Sérsic source and the lower signal-to-ratio of the JWST simulated observation compared to the other two samples. We plan to investigate the effect of a more realistic source in the future. In the next section, we quantify the improvement in sensitivity obtainable with longer integration times.

2. Focusing on the $10^{10}M_\odot$ case, we find that our simulated JWST observation has a sensitivity that is comparable to the ALMA-D_1 and ALMA-D_2 data and lower than the ALMA-D_3 and ALMA-D_4 ones. We believe this to be related to the lower signal-to-noise ratio, which compensates for the increase in angular resolution.

3. We notice that a model with a substructure of $10^{10}M_\odot$ located in correspondence of the brightest pixel is preferred by the data with a Bayes factor given by 12.5, corresponding to a 5-σ detection. However, the data was created with a completely smooth potential. This result is consistent with Vegetti et al. (2014) and Ritondale et al. (2019) who have shown that for detection threshold progressively lower than $\log K < 50$, the rate of false positives significantly increases. In the specific case of the data considered here, we think this is related to the source regularization that does not allow us to reproduce the brightest pixel in the data properly (see Figure 4.5).

4. As expected, the effect of the lower-mass substructure is also below the detection threshold. In the following section, we use the Bayes factors from these mass bins to estimate the number of detectable substructures and line-of-sight haloes with improved data quality.

Table 4.2: **Reconstructed macro-models of the JWST and ALMA data sets.** From left to right, the Einstein radius, b, the position angle, θ, the minor-to-major axis ratio, f, the position coordinates, x and y, the logarithmic power law slope, q_h, the shear strength, Γ, and angle, Γ_θ.

data		b	θ	f	x	y	q_h	Γ	Γ_θ
ALMA-D_1	μ	1.255	11.427	0.922	-0.095	-0.016	0.477	0.008	66.414
	σ	0.010	1.909	0.011	0.003	0.003	0.017	0.001	12.614
	MAP	1.265	8.757	0.930	-0.096	-0.016	0.461	0.009	61.757
ALMA-D_2	μ	1.281	15.983	0.953	-0.099	-0.019	0.468	0.008	45.175
	σ	0.008	4.025	0.015	0.006	0.007	0.015	0.001	13.401
	MAP	1.291	20.505	0.971	-0.098	-0.020	0.451	0.008	36.361
ALMA-D_3	μ	1.242	13.525	0.910	-0.092	-0.033	0.468	0.010	61.407
	σ	0.014	2.253	0.007	0.002	0.002	0.023	0.002	6.873
	MAP	1.251	14.152	0.916	-0.092	-0.032	0.451	0.009	61.420
ALMA-D_4	μ	1.207	16.299	0.897	-0.092	-0.030	0.527	0.009	52.046
	σ	0.014	1.056	0.009	0.002	0.002	0.022	0.001	6.973
	MAP	1.195	16.035	0.899	-0.094	-0.032	0.545	0.010	49.223
JWST-D	μ	2.401	89.928	0.539	6.475	6.234	0.504	0.001	-0.691
	σ	0.011	0.071	0.004	0.003	0.001	0.003	0.000	105.337
	MAP	2.406	89.886	0.538	6.475	6.234	0.506	0.001	104.753

4.4.3 Testing the CDM paradigm

In this section, we quantify how many strong gravitational lens systems observed with ALMA and the JWST one would need, for a given exposure time, to rule out the CDM paradigm at a significant statistical level. In particular, we determine the observing strategy required to obtain an expected number of at least ten low-mass haloes (subhaloes plus field haloes) in the mass range between $10^8 M_\odot$ and $10^{10} M_\odot$ in a CDM cosmology. In the event of no detection, and assuming Poisson statistics, the Likelihood of this model would be $P(N = 0|\mu_{\rm CDM} = 10) = 4.5 \times 10^{-5}$. To this end, we calculate the mean number of detectable objects by integrating the subhalo and line-of-sight mass functions (see Chapter 2, Section 2.2.1) from the lowest detectable masses $M_{\rm sub}^{\rm min}$ (see Section 4.2.3 for a definition) and $M_{\rm los}^{\rm min}(z)$, respectively, to our upper limit $M_{\rm max} = 10^{10} M_\odot$. In the case of the ALMA data, we restrict this calculation to those pixels, where we have determined the Bayes factors for all three mass bins. This restriction leads to a conservative number of detectable objects, and therefore an upper limit on the number of gravitational lens systems required.

First, we start by noticing that given their current quality, none of the considered ALMA data sets would allow us to detect substructures in the mass range between 10^8 $M_\odot$ and 10^{10} $M_\odot$ for a detection threshold of $|\log K| > 50$. Assuming a less conservative cut of

$|\log K| > 12.5$, we find that the mean number of detectable objects within the same mass range is 0.0047 and 0.0012 for the ALMA-D_3 and ALMA-D_4 data sets, respectively. For either definition, no object is detectable with JWST observations matching the signal-to-noise ratio of our simulated data. Hence, we conclude that none of the considered observations is good enough to derive any constraint on the CDM paradigm.

Then, we determine the increase in the number of detectable objects in each data set as a function of the data signal-to-noise ratio. Despali et al. (in prep.) have shown that the Bayes factors of any given substructure mass scales by roughly a factor of 2.5 and 4, when the exposure time is increased by a factor of 4 and 10, respectively. In Table 4.3 we list the number of objects detectable with a single gravitational lens system, obtained by scaling the $\log K$ maps from the previous section accordingly.

Table 4.3: **The expected number of detectable objects for ALMA-D_3, ALMA-D_4 and JWST-D.** Presented are for different exposure times, t, the expected number of detectable objects, $\mu_{5\sigma}$ and $\mu_{10\sigma}$, and the corresponding number of observations required for ten detections, $N^{10}_{5\sigma}$ and $N^{10}_{10\sigma}$, at the 5-σ and 10-σ level respectively.

Name	N_{chan}	t	$\mu_{5\sigma}$	$\mu_{10\sigma}$	$N^{10}_{5\sigma}$	$N^{10}_{10\sigma}$
	1	*min*	1	1	1	1
ALMA-D_3	8	21	0.0047	0.0000	2128	-
		4×21	0.2047	0.0006	49	16666
		10×21	0.5333	0.0047	19	2128
ALMA-D_4	1	21	0.00125	0.00000	8016	-
		4×21	0.06248	0.00014	160	69629
		10×21	0.29067	0.00125	34	8333
JWST-D	1	30	0.00000	0.0000	-	-
		4×30	0.00074	0.0000	13450	-
		10×30	0.00871	0.0000	1149	-

Finally, we compute the number of gravitational lens systems required to detect, on average, ten low-mass haloes (see last two columns of Table 4.3). We find that a conservative detection threshold of $\log K > 50$ leads to a number of systems that is much larger than what is currently available for any of the considered integration times. Moreover, even though future surveys are expected to provide such a large number of targets, a prohibitively long total integration time would be necessary, making this observing strategy unfeasible. Assuming that one can safely lower the detection threshold to a value of $\log K > 12.5$, the number of necessary targets significantly decreases. For ALMA observations of the thermal continuum, this is well within the number of currently known gravitational lens systems. As a reference, the South Pole Telescope has discovered 100 gravitationally lensed galaxies, at least 47 of which have already been observed with ALMA at resolutions ranging from 0.5×0.6 mas^2 to 1.3×1.5 mas^2 (Spilker et al. 2016). However, it should be noticed that the minimum total observing time required to exclude the CDM model is as much as 69 hours.

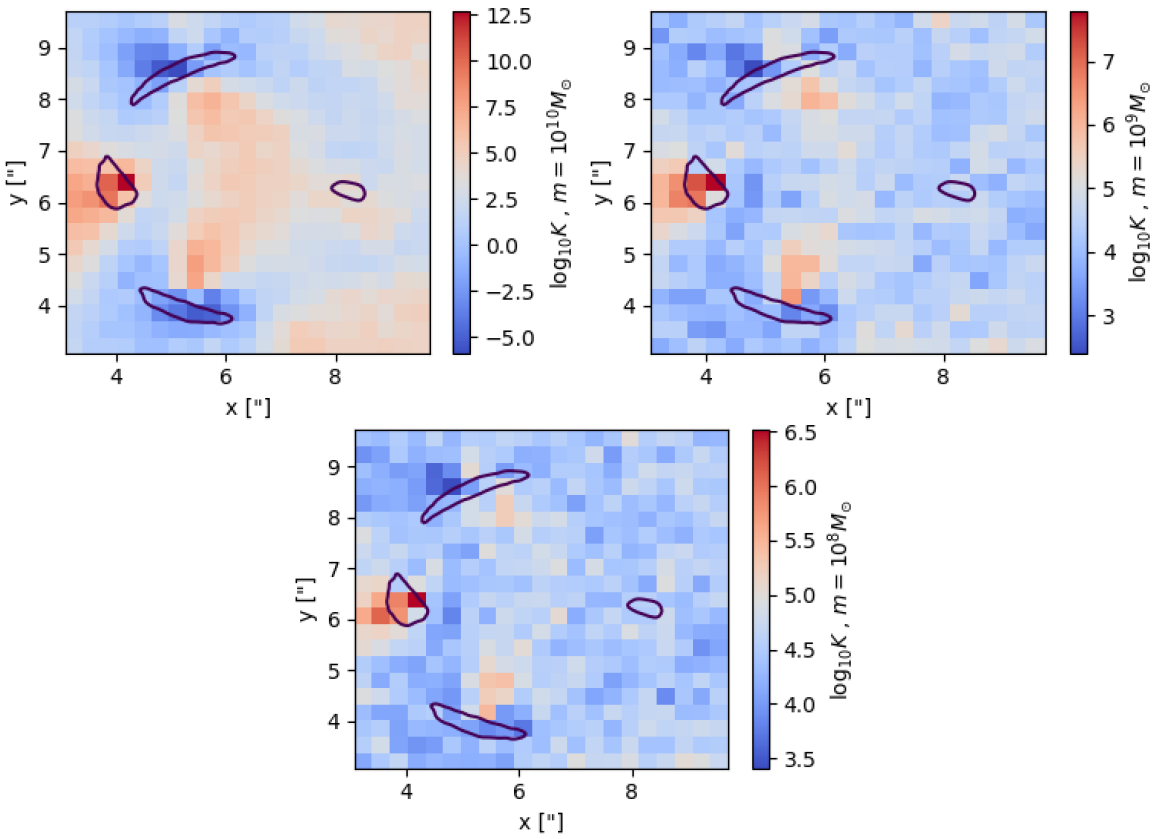

Figure 4.8: **The Bayes factor of a smooth model and a perturbed model containing a substructure for different pixels of the JWST mock data.** Shown are the Bayes factors with perturbed models containing substructures with masses $10^{10}M_\odot$, $10^9M_\odot$, and of $10^8M_\odot$.

4.5 Conclusions

Focusing on real and synthetic observations of strong gravitational lens systems with ALMA and the JWST, we quantified their sensitivity to low-mass haloes and identified possible observing strategies to rule out the CDM model at a high-significance level.

We found that publicly available [CII] and thermal continuum ALMA observations with data quality as reported in Table 4.1 are unlikely to provide stringent constraints on the nature of dark matter. Indeed, independently of the detection threshold adopted (10σ vs 5σ), the [CII] data are insensitive to the presence of subhaloes with masses $\leq 10^{10}M_\odot$, and the continuum observations only allow one to detect subhaloes with a mass of $10^{10}M_\odot$ for the less conservative threshold of $\log K \geq 12.5$. Similarly, JWST Nircam observations with a signal-to-noise ratio comparable to existing HST data from the SL2S survey provide only limited sensitivity to low-mass haloes, with no object being detectable at the 10-σ and 5-σ level.

Improving these limits will require either increasing the angular resolution, for example, with longer-baseline observations (as done for observations of SDP.81 with the ALMA long-

baseline campaign of SDP.81 Rybak et al. 2015a,b and Hezaveh et al. 2016) or the signal-to-noise ratio. At fixed angular resolution, we found that a minimum of 69 hours of ALMA observations would be required to rule out the CDM model for a 5-σ detection threshold. Unless brighter source galaxies can be identified, we rule out the JWST as a useful facility for constraining the CDM paradigm with strong gravitational lensing. In the near and long term future extensions to the ALMA array, with larger bandwidths and longer baselines (Wilson et al. 2019), will potentially increase the average signal-to-noise ratio and resolution of the data, improving the overall sensitivity to low-mass haloes, and lowering the number of targets and integration times required. In a follow-up study, we plan to extend our analysis to available ALMA observations with an improved angular resolution of $\sim 50\times50$ mas^2, as well as further investigate the issue of false detections affecting the 5-σ threshold.

Chapter 5

Conclusions

The goal of this thesis is two-fold: (i) to derive joint constraints on the mass of the thermal relic dark matter particle from a selected sample of astrophysical observations; (ii) with a deeper focus on strong gravitational lensing as a probe of dark matter, to quantify systematic errors in lensing reconstructions and derive projected constraints from ALMA and JWST observations. In the following sections, we summarise our main findings.

5.1 Astrophysical constraints on thermal relic dark matter

In Chapter 2 of this thesis, we derived new limits on thermal relic warm dark matter models by combining some of the latest constraints from the Lyman-α forest, galaxy-galaxy strong gravitational lensing, and the luminous satellites of the Milky Way. In particular, we derived a lower limit on the thermal relic particle mass of $m_{\rm th} = 2.552$ keV for a likelihood ratio of 1:20 and $m_{\rm th} = 6.048$ keV at the 95 per cent level. These limits also rule out the 7.1 keV sterile neutrino model with Lepton asymmetries $L_6 > 10$, as well as the ETHOS-4 model of self-interacting dark matter at the same confidence level.

Our results are consistent with the constraints from previous studies of strong gravitational lensing observations (e.g. Birrer et al. 2017, with $m_{\rm th} > 2.0$ keV), the Lyman-α forest (see e.g. Baur et al. 2016, with $m_{\rm th} > 4.09$ keV), and the Milky Way satellites (see e.g. Nadler et al. 2020a, with $m_{\rm th} > 6.5$ keV), and many of those presented in Table 2.4. Our constraints are somewhat weaker than those recently published by Nadler et al. (2021), who by combining studies of gravitationally lensed quasars with observations of the Milky Way satellites derived a limit of $m_{\rm th} > 9.7$ keV at the 95 per cent confidence level. This tighter limit is related to a more accurate analysis of the Milky Way satellites and to the higher sensitivity of lensed quasars.

A key ingredient to improve upon current constraints is removing all sources of systematic errors. In Chapter 3, we have extensively investigated the most dominant source of bias affecting the analysis of strong gravitational lensing observations. We refer to the next section for a summary of our main findings. The most critical uncertainties in the interpretation of the Lyman-α forest are related to the modelling of the ionizing effect of quasars, which affect

the UV fluctuations in the observed flux spectra, and the thermal and reionization history of the IGM. Significant uncertainties also arise from the limited spectral resolution of the data. The primary uncertainty in the analysis of the Milky Way satellites are the effects of baryonic processes that determine which Milky Way subhalo hosts a detectable satellite. (see e.g. Benitez-Llambay & Frenk 2020). Further uncertainties include the modelling of the Large Magellanic Cloud (Jethwa et al. 2016), the precision of the mass of the Milky Way (Wang & White 2012, Sawala et al. 2017), and anisotropies in the satellite distribution (see e.g. Newton et al. 2018b).

Further improvements can be obtained by increasing the size and quality of the observations. We refer to Chapter 4 and the last section of this chapter for a discussion on the shorter- and longer-term prospects of strong gravitational lensing. The XQ-100 survey (López et al. 2016) and the ESPRESSO spectrograph (Pepe et al. 2019) are expected to significantly increase the number of available quasar spectra with intermediate and high resolution, respectively. A larger signal-to-noise ratio will characterise these data, which provide a complementary redshift and scale coverage, and represent, therefore, an interesting way forward for the study of the Lyman-α forest and its relation to dark matter. Constraints from galaxy satellites are also expected to improve thanks to deeper observations of the Milky Way and other nearby galaxies such as M31 (Geha et al. 2017, Drlica-Wagner et al. 2019, The MSE Science Team et al. 2019, Mao et al. 2021). Overall, these new data will increase the constraining power of each probe. In combination with better hydrodynamical simulations, they will allow for a better understanding and treatment of systematic errors.

5.2 Strong gravitational lensing: biases from realistic mass distributions

In Chapter 3, we quantified the level of systematic errors in strong gravitational lensing reconstructions arising from realistic lensing potentials. To this end, we analysed a set of twenty-four gravitational lens systems generated from the Illustris-1 hydrodynamical simulation with a data quality matching the SLACS sample (Bolton et al. 2006).

State-of-the-art lens models assume that a single elliptical power-law mass profile well represents galaxy-scale lenses. However, SIDM models as well as baryonic effects in CDM can significantly affect the central mass distribution of galaxies (see e.g. Nelson et al. 2015, Schaye et al. 2015, Despali et al. 2019b). The aim of Chapter 3 was, therefore, to investigate the effect of more complex potentials on the quantities of main interest: the properties of the background sources, the Hubble constant, and the presence of substructures.

We found that the Illustris-1 lenses tend to display a shallow mass profile with a core of $0.51 - 1.7$kpc in size. Such a shallow central mass distribution is difficult to constrain with lensing observations alone, as its effect on the lensed images can be compensated by a change in the source size via the so-called source-position transformation (Schneider & Sluse 2014, Unruh et al. 2017). Indeed, we were able to model the data down to the noise level, without reproducing the core sizes correctly and at the price of underestimating the source sizes by 50

per cent. Moreover, we found that the presence of these cores is responsible for a systematic bias on the Hubble constant of 25 per cent. It should be noted that the cores in the simulations are an artifacts of the relative low resolution of the simulation. However, as mentioned above, other physical processes can lead to shallow density profiles in real galaxies. Our results confirm, therefore, the findings of Blum et al. (2020) and Kochanek (2020) and demonstrate the importance of additional observational constraints on the lensing potential, for example, from kinematic information.

Recently, Hsueh et al. (2016, 2017) and Gilman et al. (2017) have investigated the effect of complex mass distributions in ETGs on the relative fluxes of strongly lensed quasars. They found that baryonic structures can increase the level of flux-ratio anomalies by 8-10 per cent, with problematic implications for studies of dark matter with these observations. Our analysis of SLACS-like lens systems discloses no degeneracy between complex structures in the lensing potential and low-mass substructures. After removing all substructures from the simulations, we inferred a fraction of mass contained in substructure, which is consistent with their absence. Our results suggest that the origin of this difference lies partly on the limited resolution of the simulations and partly on the selection criteria adopted in this chapter and a higher sensitivity of lensed quasars to changes in the lensing potential. Higher resolution simulations and observations are required to test the validity of this result at smaller scales.

5.3 Strong gravitational lensing and the nature of dark matter: future prospects

In the next five to ten years, upcoming surveys with, for example, Euclid and the Vera Rubin Observatory (Collett 2015, Oguri & Marshall 2010) will lead to a dramatic increase in the number of known strong gravitational lens systems. Coupled with the next generation of 40m-class telescopes, this will allow us to probe the low-mass end of the halo mass function with significantly improved precision.

In Chapter 4 of this thesis, we focused on existing telescopes to identify successful observing strategies that may allow us to improve upon our current constraints in the short-time scale. To this end, we analysed publicly available ALMA observations of the lensed dusty starburst galaxy SPT–S J041839-4751. The data consisted of the thermal continuum emission from dust and the $158\mu m$ fine-structure line of ionized carbon [CII]. Furthermore, we considered synthetic observations in the F070W filter of the Nircam instrument a board of the JWST. These mock data were created from existing HST observations of the lens system SL2SJ021325-074355 from the SL2L lens sample (Ruff et al. 2011, Gavazzi et al. 2012, Sonnenfeld et al. 2013a,b, 2015).

To quantify their constraining power on the halo and subhalo mass functions, we determined, for each data set, the lowest detectable subhalo and field halo mass as a function of position on the image plane (i.e i.e. the sensitivity function, Vegetti et al. 2014, Ritondale et al. 2019), for a 10-σ and 5-σ detection threshold. We found that, due to their low signal-to-noise ratio, the JWST data are insensitive to substructures with masses $\leq 10^{10} M_{\odot}$, independently of

the adopted detection threshold. The same conclusion can be drawn for the ALMA emission line data, independently of the level of data averaging. Slightly more promising are our findings for the ALMA continuum data, which are sensitive to substructure masses between $10^9 M_\odot$ and $10^{10} M_\odot$, at least for the less conservative detection threshold of 5σ. However, as the mean total number of detectable objects within this mass range is less than zero, we concluded that neither ALMA nor JWST observations with the considered signal-to-noise ratio and angular resolution provide a viable avenue to probe the nature of dark matter.

As the sensitivity to low-mass haloes is tightly related to the data quality, we then quantified how many lens systems with ALMA and JSWT observations would be required to test the CDM paradigm as a function of their integration time. We found that, even for the less conservative detection threshold, more lens systems than are currently known would be required to measure the properties of dark matter with the JWST. Conversely, only 19 to 160 lens systems observed in the continuum with ALMA would be sufficient. However, this would require a non-negligible total exposure time of 69 hours, at least.

From the analysis presented in Chapter 4, we concluded, therefore, that the JWST does not provide a meaningful way forward for constraining the properties of dark matter with galaxy-galaxy strong gravitational lensing. Instead, continuum ALMA observations may deliver interesting results in the short term future at a moderate observational cost.

Bibliography

Abbott, T. M. C., Abdalla, F. B., Allam, S., et al. 2018, ApJS, 239, 18

Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., et al. 2014, Astronomy & Astrophysics, 571, A16

Ade, P. A. R., Aghanim, N., Arnaud, M., et al. 2016, Astronomy & Astrophysics, 594, A13

Aghanim, N., Akrami, Y., Arroja, F., et al. 2020, Astronomy & Astrophysics, 641, A1

Aghanim, N., Arnaud, M., Ashdown, M., et al. 2016, Astronomy & Astrophysics, 594, A11

Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2012, ApJS, 203, 21

Akeret, J., Refregier, A., Amara, A., Seehars, S., & Hasner, C. 2015, J. Cosmology Astropart. Phys., 2015, 043

Anderson, L., Aubourg, E., Bailey, S., et al. 2014, Monthly Notices of the Royal Astronomical Society, 441, 24–62

Angus, G., Diaferio, A., Famaey, B., Gentile, G., & van der Heyden, K. 2014, Journal of Cosmology and Astroparticle Physics, 2014, 079

Archidiacono, M., Hooper, D. C., Murgia, R., et al. 2019, Journal of Cosmology and Astroparticle Physics, 2019, 055–055

Baer, H., Barger, V., Salam, S., Sengupta, D., & Sinha, K. 2020, Midi-review: Status of weak scale supersymmetry after LHC Run 2 and ton-scale noble liquid WIMP searches

Baldes, I., Decant, Q., Hooper, D. C., & Lopez-Honorez, L. 2020, Non-Cold Dark Matter from Primordial Black Hole Evaporation

Baltz, E. A. 2004, Dark Matter Candidates

Banik, N., Bertone, G., Bovy, J., & Bozorgnia, N. 2018, Journal of Cosmology and Astroparticle Physics, 2018, 061–061

Banik, N., Bovy, J., Bertone, G., Erkal, D., & de Boer, T. J. L. 2019a, Evidence of a population of dark subhalos from Gaia and Pan-STARRS observations of the GD-1 stream

Banik, N., Bovy, J., Bertone, G., Erkal, D., & de Boer, T. J. L. 2019b, Novel constraints on the particle nature of dark matter from stellar streams

Barkana, R. 1998, ApJ, 502, 531

Barnabè, M. & Koopmans, L. V. E. 2007, ApJ, 666, 726

Baur, J., Palanque-Delabrouille, N., Yèche, C., Magneville, C., & Viel, M. 2016, Journal of Cosmology and Astroparticle Physics, 2016, 012–012

Bayer, D., Chatterjee, S., Koopmans, L. V. E., et al. 2018, arXiv e-prints, arXiv:1803.05952

Benitez-Llambay, A. & Frenk, C. 2020, Monthly Notices of the Royal Astronomical Society, 498, 4887–4900

Benson, A. J. 2010, Physics Reports, 495, 33–86

Benson, A. J., Lacey, C. G., Baugh, C. M., Cole, S., & Frenk, C. S. 2002, MNRAS, 333, 156

Bertone, G., Hooper, D., & Silk, J. 2005, Physics Reports, 405, 279–390

Bertone, G. & Tait, T. M. P. 2018, Nature, 562, 51–56

Betancourt, M. 2018, The Convergence of Markov chain Monte Carlo Methods: From the Metropolis method to Hamiltonian Monte Carlo

Birrer, S., Amara, A., & Refregier, A. 2017, Journal of Cosmology and Astroparticle Physics, 2017, 037–037

Blum, K., Castorina, E., & Simonović, M. 2020, The Astrophysical Journal, 892, L27

Blumenthal, G. R., Faber, S. M., Primack, J. R., & Rees, M. J. 1984, Nature, 311, 517

Bocquet, S., Saro, A., Dolag, K., & Mohr, J. J. 2016, MNRAS, 456, 2361

Bode, P., Ostriker, J. P., & Turok, N. 2001, The Astrophysical Journal, 556, 93–107

Boera, E., Becker, G. D., Bolton, J. S., & Nasir, F. 2019, Astrophys. J., 872, 101

Bohr, S., Zavala, J., Cyr-Racine, F.-Y., & Vogelsberger, M. 2021, The halo mass function and inner structure of ETHOS haloes at high redshift

Bolton, A. S., Burles, S., Koopmans, L. V. E., et al. 2008, The Astrophysical Journal, 682, 964–984

Bolton, A. S., Burles, S., Koopmans, L. V. E., Treu, T., & Moustakas, L. A. 2006, ApJ, 638, 703

Bolton, J. S., Puchwein, E., Sijacki, D., et al. 2016, Monthly Notices of the Royal Astronomical Society, 464, 897–914

Bonaca, A., Hogg, D. W., Price-Whelan, A. M., & Conroy, C. 2019, The Astrophysical Journal, 880, 38

Bond, J. R., Cole, S., Efstathiou, G., & Kaiser, N. 1991, ApJ, 379, 440

Bose, S., Hellwing, W. A., Frenk, C. S., et al. 2015, Monthly Notices of the Royal Astronomical Society, 455, 318–333

Bovy, J., Erkal, D., & Sanders, J. L. 2016, Monthly Notices of the Royal Astronomical Society, 466, 628

Bower, R. G. 1991, MNRAS, 248, 332

Boyarsky, A., Drewes, M., Lasserre, T., Mertens, S., & Ruchayskiy, O. 2019, Progress in Particle and Nuclear Physics, 104, 1–45

Boyarsky, A., Franse, J., Iakubovskyi, D., & Ruchayskiy, O. 2015, Phys. Rev. Lett., 115, 161301

Boyarsky, A., Ruchayskiy, O., Iakubovskyi, D., & Franse, J. 2014, Physical Review Letters, 113

Bozek, B., Boylan-Kolchin, M., Horiuchi, S., et al. 2016, Monthly Notices of the Royal Astronomical Society, 459, 1489–1504

Bozek, B., Fitts, A., Boylan-Kolchin, M., et al. 2018, Monthly Notices of the Royal Astronomical Society, 483, 4086–4099

Brewer, B. J., Huijser, D., & Lewis, G. F. 2015, Monthly Notices of the Royal Astronomical Society, 455, 1819

Bridle, S., Elvin-Poole, J., Evans, J., et al. 2017, Physics Letters B, 764, 322–327

Bulbul, E., Markevitch, M., Foster, A., et al. 2014, ApJ, 789, 13

Bullock, J. S. & Boylan-Kolchin, M. 2017, ARA&A, 55, 343

Bullock, J. S., Kolatt, T. S., Sigad, Y., et al. 2001, MNRAS, 321, 559

Bullock, J. S., Kravtsov, A. V., & Weinberg, D. H. 2000, ApJ, 539, 517

Burkert, A. 2000, ApJ, 534, L143

Calabrese, E., Slosar, A., Melchiorri, A., Smoot, G. F., & Zahn, O. 2008, Physical Review D, 77

Callingham, T. M., Cautun, M., Deason, A. J., et al. 2019, Monthly Notices of the Royal Astronomical Society, 484, 5453–5467

Carlberg, R. G. 2012, The Astrophysical Journal, 748, 20

Carlberg, R. G. 2013, The Astrophysical Journal, 775, 90

Carlin, J. L., Sand, D. J., Muñoz, R. R., et al. 2017, AJ, 154, 267

Cautun, M., Hellwing, W. A., van de Weygaert, R., et al. 2014, MNRAS, 445, 1820

Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2019, The Pan-STARRS1 Surveys

Chan, J. H. H., Schive, H.-Y., Wong, S.-K., Chiueh, T., & Broadhurst, T. 2020, Phys. Rev. Lett., 125, 111102

Chan, T. K., Kereš, D., Oñorbe, J., et al. 2015, Monthly Notices of the Royal Astronomical Society, 454, 2981

Chen, G. C.-F., Fassnacht, C. D., Suyu, S. H., et al. 2019, Monthly Notices of the Royal Astronomical Society, 490, 1743

Chopin, N. & Robert, C. P. 2010, Biometrika, 97, 741–755

Cirelli, M., Corcella, G., Hektor, A., et al. 2011, Journal of Cosmology and Astroparticle Physicas, 2011, 051–051

Clowe, D., Bradač, M., Gonzalez, A. H., et al. 2006, The Astrophysical Journal, 648, L109–L113

Clowe, D., Gonzalez, A., & Markevitch, M. 2004, The Astrophysical Journal, 604, 596–603

Cole, S., Aragón-Salamanca, A., Frenk, C. S., Navarro, J. F., & Zepf, S. E. 1994, MNRAS, 271, 781

Cole, S., Lacey, C. G., Baugh, C. M., & Frenk, C. S. 2000, MNRAS, 319, 168

Collett, T. E. 2015, The Astrophysical Journal, 811, 20

Corasaniti, P., Agarwal, S., Marsh, D., & Das, S. 2017, Physical Review D, 95

Cox, R. T. 1946, American Journal of Physics, 14, 1

Croft, R. A. C. 2004, The Astrophysical Journal, 610, 642–662

Csáki, C., Lombardo, S., & Telem, O. 2018, TASI Lectures on Non-Supersymmetric BSM Models

Cyr-Racine, F.-Y., Sigurdson, K., Zavala, J., et al. 2016, Physical Review D, 93

Dalal, N. & Kochanek, C. S. 2002, ApJ, 572, 25

Dar, A. 1995, The Astrophysical Journal, 449, 550

Davé, R., Spergel, D. N., Steinhardt, P. J., & Wandelt, B. D. 2001, ApJ, 547, 574

Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, ApJ, 292, 371

Daylan, T., Cyr-Racine, F.-Y., Diaz Rivero, A., Dvorkin, C., & Finkbeiner, D. P. 2018, ApJ, 854, 141

Despali, G., Lovell, M., Vegetti, S., Crain, R. A., & Oppenheimer, B. D. 2019a, Monthly Notices of the Royal Astronomical Society, 491, 1295–1310

Despali, G., Sparre, M., Vegetti, S., et al. 2019b, Monthly Notices of the Royal Astronomical Society, 484, 4563–4573

Despali, G. & Vegetti, S. 2017, MNRAS, 469, 1997

Despali, G., Vegetti, S., White, S. D. M., Giocoli, C., & van den Bosch, F. C. 2018, Monthly Notices of the Royal Astronomical Society, 475, 5424–5442

Diez-Tejedor, A. & Marsh, D. J. E. 2017, Cosmological production of ultralight dark matter axions

Dodelson, S. & Widrow, L. M. 1994, Phys. Rev. Lett., 72, 17

Drewes, M. & Garbrecht, B. 2017, Nuclear Physics B, 921, 250

Drlica-Wagner, A., Bechtol, K., Mau, S., et al. 2020, The Astrophysical Journal, 893, 47

Drlica-Wagner, A., Bechtol, K., Rykoff, E. S., et al. 2015, ApJ, 813, 109

Drlica-Wagner, A., Mao, Y.-Y., Adhikari, S., et al. 2019, arXiv e-prints, arXiv:1902.01055

Du, W., Zhao, G.-B., Fan, Z., et al. 2020, The Astrophysical Journal, 892, 62

Duffy, A. R., Schaye, J., Kay, S. T., & Dalla Vecchia, C. 2008, MNRAS, 390, L64

Eggemeier, B., Redondo, J., Dolag, K., Niemeyer, J. C., & Vaquero, A. 2020, Physical Review Letters, 125

Einasto, J. 1965, Trudy Astrofizicheskogo Instituta Alma-Ata, 5, 87

Einstein, A. 1936, Science, 84, 506

Eke, V. R., Cole, S., Frenk, C. S., & Navarro, J. F. 1996, MNRAS, 281, 703

Ellis, J., Linde, A., & Nanopoulos, D. 1982, Physics Letters B, 118, 59

Enzi, W., Murgia, R., Newton, O., et al. 2020, arXiv e-prints, arXiv:2010.13802

Erkal, D. & Belokurov, V. 2015, Monthly Notices of the Royal Astronomical Society, 450, 1136

Erkal, D., Belokurov, V., Bovy, J., & Sanders, J. L. 2016, Monthly Notices of the Royal Astronomical Society, 463, 102–119

Fahlman, G., Kaiser, N., Squires, G., & Woods, D. 1994, The Astrophysical Journal, 437, 56

Falco, E. E., Gorenstein, M. V., & Shapiro, I. I. 1985, ApJ, 289, L1

Feng, J. L. 2005, Annals of Physics, 315, 2–51

Feng, J. L. 2010, Annual Review of Astronomy and Astrophysics, 48, 495–545

Feroz, F., Hobson, M. P., & Bridges, M. 2009, Monthly Notices of the Royal Astronomical Society, 398, 1601–1614

Ferreira, E. G. M. 2020, Ultra-Light Dark Matter

Frenk, C. S. & White, S. D. M. 2012, Annalen der Physik, 524, 507

Fukuda, Y., Hayakawa, T., Ichihara, E., et al. 1998, Phys. Rev. Lett., 81, 1562

Gao, L., Frenk, C. S., Boylan-Kolchin, M., et al. 2011, MNRAS, 410, 2309

Gao, L., Yoshida, N., Abel, T., et al. 2007, MNRAS, 378, 449

Garzilli, A., Boyarsky, A., & Ruchayskiy, O. 2017, Physics Letters B, 773, 258–264

Garzilli, A., Ruchayskiy, O., Magalich, A., & Boyarsky, A. 2019, How warm is too warm? Towards robust Lyman-α forest bounds on warm dark matter

Gaskins, J. M. 2016, Contemporary Physics, 57, 496–525

Gavazzi, R., Treu, T., Marshall, P. J., Brault, F., & Ruff, A. 2012, ApJ, 761, 170

Geha, M., Wechsler, R. H., Mao, Y.-Y., et al. 2017, The Astrophysical Journal, 847, 4

Genel, S., Fall, S. M., Hernquist, L., et al. 2015, ApJ, 804, L40

George, E. I., Casella, G., & George, E. I. 1992, The American Statistician

Gilman, D., Agnello, A., Treu, T., Keeton, C. R., & Nierenberg, A. M. 2017, MNRAS, 467, 3970

Gilman, D., Birrer, S., Nierenberg, A., et al. 2019a, Monthly Notices of the Royal Astronomical Society, 491, 6077–6101

Gilman, D., Birrer, S., Treu, T., Keeton, C. R., & Nierenberg, A. 2018, MNRAS, 481, 819

Gilman, D., Birrer, S., Treu, T., Nierenberg, A., & Benson, A. 2019b, Monthly Notices of the Royal Astronomical Society, 487, 5721–5738

Giocoli, C., Meneghetti, M., Metcalf, R. B., Ettori, S., & Moscardini, L. 2014, Monthly Notices of the Royal Astronomical Society, 440, 1899

Graham, P. W., Irastorza, I. G., Lamoreaux, S. K., Lindner, A., & van Bibber, K. A. 2015, Annual Review of Nuclear and Particle Science, 65, 485–514

Green, A. M., Hofmann, S., & Schwarz, D. J. 2004, Monthly Notices of the Royal Astronomical Society, 353, L23

Grillmair, C. J. & Dionatos, O. 2006, ApJ, 643, L17

Hamann, J., Hannestad, S., Raffelt, G. G., Tamborra, I., & Wong, Y. Y. Y. 2012, Journal of Physics: Conference Series, 375, 032003

Hamilton, A. J. S., Kumar, P., Lu, E., & Matthews, A. 1991, ApJ, 374, L1

Hargis, J. R., Willman, B., & Peter, A. H. G. 2014, The Astrophysical Journal, 795, L13

Harvey, D., Valkenburg, W., Tamone, A., et al. 2019, Monthly Notices of the Royal Astronomical Society, 491, 4247

Heitmann, K., Lukić, Z., Fasel, P., et al. 2008, Computational Science and Discovery, 1, 015003

Hezaveh, Y. D., Dalal, N., Marrone, D. P., et al. 2016, The Astrophysical Journal, 823, 37

Hezaveh, Y. D., Marrone, D. P., Fassnacht, C. D., et al. 2013, The Astrophysical Journal, 767, 132

Higson, E., Handley, W., Hobson, M., & Lasenby, A. 2018, Bayesian Analysis, 13, 873

Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, ApJS, 208, 19

Hsueh, J.-W., Despali, G., Vegetti, S., et al. 2018, MNRAS, 475, 2438

Hsueh, J.-W., Enzi, W., Vegetti, S., et al. 2019, Monthly Notices of the Royal Astronomical Society

Hsueh, J. W., Fassnacht, C. D., Vegetti, S., et al. 2016, MNRAS, 463, L51

Hsueh, J. W., Oldham, L., Spingola, C., et al. 2017, MNRAS, 469, 3713

Hu, W., Barkana, R., & Gruzinov, A. 2000, prl, 85, 1158

Hui, L. & Gnedin, N. Y. 1997, Monthly Notices of the Royal Astronomical Society, 292, 27–42

Iakubovskyi, D. 2014, New emission line at 3.5 keV - observational status, connection with radiatively decaying dark matter and directions for future studies

Ibata, R. A., Lewis, G. F., & Martin, N. F. 2016, The Astrophysical Journal, 819, 1

Iršič, V., Viel, M., Haehnelt, M. G., Bolton, J. S., & Becker, G. D. 2017, Phys. Rev. Lett., 119, 031302

Iršič, V., Viel, M., Haehnelt, M. G., et al. 2017, Phys. Rev. D, 96, 023522

Irwin, M. J., Webster, R. L., Hewett, P. C., Corrigan, R. T., & Jedrzejewski, R. I. 1989, AJ, 98, 1989

Jasche, J., Leclercq, F., & Wandelt, B. 2015, Journal of Cosmology and Astroparticle Physics, 2015, 036–036

Jeffreys, H. 1946, Proceedings of the Royal Society of London Series A, 186, 453

Jenkins, A., Frenk, C. S., White, S. D. M., et al. 2001, Monthly Notices of the Royal Astronomical Society, 321, 372–384

Jethwa, P., Erkal, D., & Belokurov, V. 2016, MNRAS, 461, 2212

Jethwa, P., Erkal, D., & Belokurov, V. 2017, Monthly Notices of the Royal Astronomical Society, 473, 2060–2083

Kallivayalil, N., Sales, L. V., Zivick, P., et al. 2018, ApJ, 867, 19

Kauffmann, G. & White, S. D. M. 1993, MNRAS, 261, 921

Keeton, C. R. 2001, The Astrophysical Journal, 561, 46–60

Kennedy, R., Frenk, C., Cole, S., & Benson, A. 2014, MNRAS, 442, 2487

Kim, D., Jerjen, H., Geha, M., et al. 2016, ApJ, 833, 16

Kim, D., Jerjen, H., Mackey, D., Costa, G. S. D., & Milone, A. P. 2015, ApJ, 804, L44

Klypin, A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, ApJ, 522, 82

Klypin, A. A. & Shandarin, S. F. 1983, MNRAS, 204, 891

Kochanek, C. S. 1995, ApJ, 445, 559

Kochanek, C. S. 2020, Monthly Notices of the Royal Astronomical Society, 493, 1725–1735

Koopmans, L. V. E. 2005, MNRAS, 363, 1136

Koopmans, L. V. E., Biggs, A., Blandford, R. D., et al. 2003, ApJ, 595, 712

Koposov, S., Belokurov, V., Evans, N. W., et al. 2008, ApJ, 686, 279

Koposov, S. E., Casey, A. R., Belokurov, V., et al. 2015, ApJ, 811, 62

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

Kulkarni, M. & Ostriker, J. P. 2020, arXiv e-prints, arXiv:2011.02116

Kunz, M., Nesseris, S., & Sawicki, I. 2016, Phys. Rev. D, 94, 023510

Lacey, C. & Cole, S. 1993, Monthly Notices of the Royal Astronomical Society, 262, 627

Lacey, C. G., Baugh, C. M., Frenk, C. S., et al. 2016, MNRAS, 462, 3854

Lagos, C. d. P., Theuns, T., Stevens, A. R. H., et al. 2016, Monthly Notices of the Royal Astronomical Society, 464, 3850

Larson, R. B. 2003, Reports on Progress in Physics, 66, 1651–1697

Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, Euclid Definition Study Report

Lavaux, G. & Jasche, J. 2015, Monthly Notices of the Royal Astronomical Society, 455, 3169–3179

Li, R., Frenk, C. S., Cole, S., Wang, Q., & Gao, L. 2017, MNRAS, 468, 1426

Li, T. S., Simon, J. D., Drlica-Wagner, A., et al. 2017, ApJ, 838, 8

Li, X., Hui, L., & Bryan, G. L. 2019, Physical Review D, 99

López, S., D'Odorico, V., Ellison, S. L., et al. 2016, A&A, 594, A91

Lovell, M. R. 2020, Towards a general parametrization of the warm dark matter halo mass function

Lovell, M. R., Bose, S., Boyarsky, A., et al. 2016, Monthly Notices of the Royal Astronomical Society, 461, 60–72

Lovell, M. R., Bose, S., Boyarsky, A., et al. 2017, Monthly Notices of the Royal Astronomical Society, 468, 4285–4298

Lovell, M. R., Eke, V., Frenk, C. S., et al. 2012, Monthly Notices of the Royal Astronomical Society, 420, 2318–2324

Lovell, M. R., Frenk, C. S., Eke, V. R., et al. 2014a, MNRAS, 439, 300

Lovell, M. R., Frenk, C. S., Eke, V. R., et al. 2014b, Monthly Notices of the Royal Astronomical Society, 439, 300

Ludlow, A. D., Bose, S., Angulo, R. E., et al. 2016, Monthly Notices of the Royal Astronomical Society, 460, 1214–1232

Lukić, Z., Heitmann, K., Habib, S., Bashinsky, S., & Ricker, P. M. 2007, ApJ, 671, 1160

Macciò, A. V., Paduroiu, S., Anderhalden, D., Schneider, A., & Moore, B. 2012, Monthly Notices of the Royal Astronomical Society, 424, 1105–1112

Maleki, A., Baghram, S., & Rahvar, S. 2020, Physical Review D, 101

Mao, S. & Schneider, P. 1998, MNRAS, 295, 587

Mao, Y.-Y., Geha, M., Wechsler, R. H., et al. 2021, The Astrophysical Journal, 907, 85

Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, Monthly Notices of the Royal Astronomical Society

Marsh, D. J. 2016, Physics Reports, 643, 1–79

Marsh, D. J. & Niemeyer, J. C. 2019, Physical Review Letters, 123

Marsh, D. J. E. 2017, Axions and ALPs: a very short introduction

May, S. & Springel, V. 2021, arXiv e-prints, arXiv:2101.01828

McConnachie, A. W. 2012, AJ, 144, 4

McDonald, P., Seljak, U., Burles, S., et al. 2006, The Astrophysical Journal Supplement Series, 163, 80–109

McKean, J. P., Jackson, N., Vegetti, S., et al. 2015, Strong gravitational lensing with the SKA

McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in Astronomical Society of the Pacific Conference Series, Vol. 376, Astronomical Data Analysis Software and Systems XVI, ed. R. A. Shaw, F. Hill, & D. J. Bell, 127

McQuinn, M. 2016, Annual Review of Astronomy and Astrophysics, 54, 313–362

Meiksin, A. A. 2009, Rev. Mod. Phys., 81, 1943

Mellier, Y. 1999, Annual Review of Astronomy and Astrophysics, 37, 127–189

Metcalf, R. B. & Petkova, M. 2014, MNRAS, 445, 1942

Metcalf, R. B. & Silk, J. 1997, The Astrophysical Journal, 489, 1–6

Metzler, C. A., White, M., Norman, M., & Loken, C. 1999, The Astrophysical Journal, 520, L9

Miller, C., Erickcek, A. L., & Murgia, R. 2019, Physical Review D, 100

Miralda-Escude, J. 2002, The Astrophysical Journal, 564, 60–64

Mitsou, V. A. 2015, in Journal of Physics Conference Series, Vol. 651, Journal of Physics Conference Series, 012023

Mo, H., van den Bosch, F., & White, S. 2010, Galaxy Formation and Evolution (Cambridge University Press)

Mocz, P., Fialkov, A., Vogelsberger, M., et al. 2019, Phys. Rev. Lett., 123, 141301

Moliné, A., Sánchez-Conde, M. A., Palomares-Ruiz, S., & Prada, F. 2017, Monthly Notices of the Royal Astronomical Society, 466, 4974

Möller, O., Hewett, P., & Blain, A. W. 2003, MNRAS, 345, 1

Moore, B., Ghigna, S., Governato, F., et al. 1999, ApJ, 524, L19

Moster, B. P., Somerville, R. S., Maulbetsch, C., et al. 2010, The Astrophysical Journal, 710, 903–923

Mukherjee, S., Koopmans, L. V. E., Metcalf, R. B., et al. 2018, MNRAS, 479, 4108

Murgia, R., Iršič, V., & Viel, M. 2018, Physical Review D, 98

Murgia, R., Merle, A., Viel, M., Totzauer, M., & Schneider, A. 2017, Journal of Cosmology and Astroparticle Physics, 2017, 046–046

Naab, T. & Ostriker, J. P. 2017, Annual Review of Astronomy and Astrophysics, 55, 59–109

Nadler, E. O., Birrer, S., Gilman, D., et al. 2021, Dark Matter Constraints from a Unified Analysis of Strong Gravitational Lenses and Milky Way Satellite Galaxies

Nadler, E. O., Drlica-Wagner, A., Bechtol, K., et al. 2020a, Milky Way Satellite Census. III. Constraints on Dark Matter Properties from Observations of Milky Way Satellite Galaxies

Nadler, E. O., Gluscevic, V., Boddy, K. K., & Wechsler, R. H. 2019, ApJ, 878, L32

Nadler, E. O., Wechsler, R. H., Bechtol, K., et al. 2020b, The Astrophysical Journal, 893, 48

Narayanan, V. K., Spergel, D. N., Dave, R., & Ma, C.-P. 2000, Astrophys. J., 543, L103

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, The Astrophysical Journal, 462, 563

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, ApJ, 490, 493

Neal, R. M. 2012, MCMC using Hamiltonian dynamics

Necib, L., Ostdiek, B., Lisanti, M., et al. 2019, Evidence for a Vast Prograde Stellar Stream in the Solar Vicinity

Neichel, B., Mouillet, D., Gendron, E., et al. 2018, Overview of the European Extremely Large Telescope and its instrument suite

Nelson, D., Genel, S., Vogelsberger, M., et al. 2015, Monthly Notices of the Royal Astronomical Society, 448, 59

Nelson, D., Springel, V., Pillepich, A., et al. 2019, The IllustrisTNG Simulations: Public Data Release

Newton, O., Cautun, M., Jenkins, A., Frenk, C. S., & Helly, J. C. 2018a, Proceedings of the International Astronomical Union, 14, 109–113

Newton, O., Cautun, M., Jenkins, A., Frenk, C. S., & Helly, J. C. 2018b, MNRAS, 479, 2853

Newton, O., Leo, M., Cautun, M., et al. 2020, Constraints on the properties of warm dark matter using the satellite galaxies of the Milky Way

Nierenberg, A. M., Auger, M. W., Treu, T., Marshall, P. J., & Fassnacht, C. D. 2011, The Astrophysical Journal, 731, 44

Nierenberg, A. M., Auger, M. W., Treu, T., et al. 2012, The Astrophysical Journal, 752, 99

Nierenberg, A. M., Treu, T., Menci, N., Lu, Y., & Wang, W. 2013, The Astrophysical Journal, 772, 146

Nierenberg, A. M., Treu, T., Wright, S. A., Fassnacht, C. D., & Auger, M. W. 2014, Monthly Notices of the Royal Astronomical Society, 442, 2434–2445

Nightingale, J. W., Dye, S., & Massey, R. J. 2018, Monthly Notices of the Royal Astronomical Society, 478, 4738

Nori, M. & Baldi, M. 2018, Monthly Notices of the Royal Astronomical Society, 478, 3935

Nori, M. & Baldi, M. 2021, MNRAS, 501, 1539

Oguri, M. & Marshall, P. J. 2010, Monthly Notices of the Royal Astronomical Society, no–no

Onions, J., Knebe, A., Pearce, F. R., et al. 2012, Monthly Notices of the Royal Astronomical Society, 423, 1200–1214

Parkinson, H., Cole, S., & Helly, J. 2008, MNRAS, 383, 557

Patel, E., Kallivayalil, N., Garavito-Camargo, N., et al. 2020, ApJ, 893, 121

Pawase, R. S., Courbin, F., Faure, C., Kokotanekova, R., & Meylan, G. 2014, Monthly Notices of the Royal Astronomical Society, 439, 3392–3404

Peacock, J. A. 1996, arXiv e-prints, astro

Peacock, J. A. 2003, arXiv e-prints, astro

Peccei, R. D. & Quinn, H. R. 1977, Phys. Rev. Lett., 38, 1440

Peebles, P. J. E. 1993, Principles of Physical Cosmology

Pepe, F. et al. 2019, ESO Messenger, 153, 6

Perrin, M. D., Sivaramakrishnan, A., Lajoie, C.-P., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9143, Space Telescopes and Instrumentation 2014: Optical, Infrared, and Millimeter Wave, ed. J. Oschmann, Jacobus M., M. Clampin, G. G. Fazio, & H. A. MacEwen, 91433X

Peter, A. H. G. 2012, arXiv e-prints, arXiv:1201.3942

Peter, A. H. G., Gluscevic, V., Green, A. M., Kavanagh, B. J., & Lee, S. K. 2013, WIMP physics with ensembles of direct-detection experiments

Petkova, M., Metcalf, R. B., & Giocoli, C. 2014, Monthly Notices of the Royal Astronomical Society, 445, 1954

Polisensky, E. & Ricotti, M. 2011, Physical Review D, 83

Powell, D., Vegetti, S., McKean, J. P., et al. 2020, Monthly Notices of the Royal Astronomical Society, 501, 515–530

Press, W. H. & Schechter, P. 1974, Astrophys. J., 187, 425

Primack, J. R. & Blumenthal, G. R. 1984, What is the Dark Matter? Implications for Galaxy Formation and Particle Physics, ed. J. Audouze & J. Tran Thanh Van (Dordrecht: Springer Netherlands), 163–183

Raffelt, G. G. 1995, in Dark Matter in Cosmology, Clocks and Tests of Fundamental Laws, ed. B. Guiderdoni, G. Greene, D. Hinds, & J. Tran Thanh Van, 159

Randall, S. W., Markevitch, M., Clowe, D., Gonzalez, A. H., & Bradač, M. 2008, The Astrophysical Journal, 679, 1173–1180

Rau, S., Vegetti, S., & White, S. D. M. 2013, MNRAS, 430, 2232

Ritondale, E., Auger, M. W., Vegetti, S., & McKean, J. P. 2018, Monthly Notices of the Royal Astronomical Society, 482, 4744–4762

Ritondale, E., Vegetti, S., Despali, G., et al. 2019, Monthly Notices of the Royal Astronomical Society, 485, 2179–2193

Rizzo, F., Vegetti, S., Fraternali, F., & Di Teodoro, E. 2018, MNRAS, 481, 5606

Rizzo, F., Vegetti, S., Powell, D., et al. 2020, Nature, 584, 201–204

Robertson, A., Harvey, D., Massey, R., et al. 2019, MNRAS, 488, 3646

Robertson, A., Massey, R., Eke, V., et al. 2018, MNRAS, 476, L20

Ross, A. J., Samushia, L., Howlett, C., et al. 2015, Monthly Notices of the Royal Astronomical Society, 449, 835–847

Roszkowski, L., Sessolo, E. M., & Trojanowski, S. 2018, Reports on Progress in Physics, 81, 066201

Roy, V. 2019, Convergence diagnostics for Markov chain Monte Carlo

Rubin, V. C., Burstein, D., Ford, W. K., J., & Thonnard, N. 1985, ApJ, 289, 81

Rubin, V. C., Ford, W. K., J., & Thonnard, N. 1978, ApJ, 225, L107

Rubin, V. C. & Ford, W. Kent, J. 1970, ApJ, 159, 379

Ruff, A. J., Gavazzi, R., Marshall, P. J., et al. 2011, ApJ, 727, 96

Rybak, M., McKean, J. P., Vegetti, S., Andreani, P., & White, S. D. M. 2015a, MNRAS, 451, L40

Rybak, M., Vegetti, S., McKean, J. P., Andreani, P., & White, S. D. M. 2015b, MNRAS, 453, L26

Sanders, R. H., van Albada, T. S., & Oosterloo, T. A. 1984, ApJ, 278, L91

Sawala, T., Pihajoki, P., Johansson, P. H., et al. 2017, MNRAS, 467, 4383

Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, MNRAS, 446, 521

Scherrer, R. J. & Turner, M. S. 1986a, Phys. Rev. D, 34, 3263

Scherrer, R. J. & Turner, M. S. 1986b, Phys. Rev. D, 33, 1585, [Erratum: Phys.Rev.D 34, 3263 (1986)]

Schive, H.-Y., Chiueh, T., & Broadhurst, T. 2014, Nature Physics, 10, 496–499

Schneider, A. 2015, MNRAS, 451, 3117

Schneider, A., Smith, R. E., Macciò, A. V., & Moore, B. 2012, Monthly Notices of the Royal Astronomical Society, 424, 684

Schneider, A., Smith, R. E., & Reed, D. 2013, Monthly Notices of the Royal Astronomical Society, 433, 1573

Schneider, P., Ehlers, J., & Falco, E. E. 1992, Gravitational Lenses

Schneider, P. & Sluse, D. 2014, A&A, 564, A103

Sedrakian, A. 2016, Physical Review D, 93

Seljak, U. c. v., Makarov, A., McDonald, P., & Trac, H. 2006, Phys. Rev. Lett., 97, 191303

Shapiro, I. I. 1964, Phys. Rev. Lett., 13, 789

Sheth, R. K., Mo, H. J., & Tormen, G. 2001, MNRAS, 323, 1

Sheth, R. K. & Tormen, G. 1999, MNRAS, 308, 119

Shirazi, M., Vegetti, S., Nesvadba, N., et al. 2014, Monthly Notices of the Royal Astronomical Society, 440, 2201

Shirman, Y. 2010, The Dawn of the LHC Era

Shu, Y., Bolton, A. S., Kochanek, C. S., et al. 2016, The Astrophysical Journal, 824, 86

Sikivie, P. 2008, Axions, 19–50

Skilling, J. 2004, in American Institute of Physics Conference Series, Vol. 735, American Institute of Physics Conference Series, ed. R. Fischer, R. Preuss, & U. V. Toussaint, 395–405

Somerville, R. S. 2002, ApJ, 572, L23

Sonnenfeld, A., Gavazzi, R., Suyu, S. H., Treu, T., & Marshall, P. J. 2013a, ApJ, 777, 97

Sonnenfeld, A., Treu, T., Gavazzi, R., et al. 2013b, ApJ, 777, 98

Sonnenfeld, A., Treu, T., Marshall, P. J., et al. 2015, ApJ, 800, 94

Spergel, D. N. & Steinhardt, P. J. 2000, Physical Review Letters, 84, 3760–3763

Spilker, J. S., Marrone, D. P., Aravena, M., et al. 2016, The Astrophysical Journal, 826, 112

Spingola, C., McKean, J. P., Auger, M. W., et al. 2018, Monthly Notices of the Royal Astronomical Society, 478, 4816–4829

Spingola, C., McKean, J. P., Massari, D., & Koopmans, L. V. E. 2019, A&A, 630, A108

Springel, V. 2005, MNRAS, 364, 1105

Springel, V. 2010, MNRAS, 401, 791

Springel, V., Frenk, C. S., & White, S. D. M. 2006, Nature, 440, 1137–1144

Springel, V., Wang, J., Vogelsberger, M., et al. 2008, MNRAS, 391, 1685

Springel, V., White, S. D. M., Jenkins, A., et al. 2005, Nature, 435, 629

Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001a, Monthly Notices of the Royal Astronomical Society, 328, 726–750

Springel, V., Yoshida, N., & White, S. D. 2001b, New Astronomy, 6, 79–117

Stoughton, C., Lupton, R. H., Bernardi, M., et al. 2002, AJ, 123, 485

Suyu, S. H., Marshall, P. J., Auger, M. W., et al. 2010, ApJ, 711, 201

Suyu, S. H., Marshall, P. J., Hobson, M. P., & Bland ford, R. D. 2006, MNRAS, 371, 983

Tegmark, M., Silk, J., Rees, M. J., et al. 1997, ApJ, 474, 1

Teklu, A. F., Remus, R.-S., Dolag, K., et al. 2015, ApJ, 812, 29

The MSE Science Team, Babusiaux, C., Bergemann, M., et al. 2019, arXiv e-prints, arXiv:1904.04907

Thoul, A. A. & Weinberg, D. H. 1995, ApJ, 442, 480

Tollerud, E. J., Bullock, J. S., Strigari, L. E., & Willman, B. 2008, ApJ, 688, 277

Torrealba, G., Koposov, S. E., Belokurov, V., & Irwin, M. 2016a, Monthly Notices of the Royal Astronomical Society, 459, 2370–2378

Torrealba, G., Koposov, S. E., Belokurov, V., et al. 2016b, Monthly Notices of the Royal Astronomical Society, 463, 712–722

Treu, T. 2010, Annual Review of Astronomy and Astrophysics, 48, 87

Treu, T., Gavazzi, R., Gorecki, A., et al. 2009, ApJ, 690, 670

Trotta, R. 2017, Bayesian Methods in Cosmology

Undagoitia, T. M. & Rauch, L. 2015, Journal of Physics G: Nuclear and Particle Physics, 43, 013001

Unruh, S., Schneider, P., & Sluse, D. 2017, A&A, 601, A77

van Haasteren, R. 2009, Bayesian evidence: can we beat MultiNest using traditional MCMC methods?

Vanderriest, C., Schneider, J., Herpe, G., et al. 1989, Optical Determinations of the Time Delay in 0957+561, ed. J. M. Moran, J. N. Hewitt, & K.-Y. Lo, Vol. 330, 90

Vegetti, S., Despali, G., Lovell, M. R., & Enzi, W. 2018, Monthly Notices of the Royal Astronomical Society, 481, 3661–3669

Vegetti, S. & Koopmans, L. V. E. 2009a, MNRAS, 392, 945

Vegetti, S. & Koopmans, L. V. E. 2009b, MNRAS, 400, 1583

Vegetti, S., Koopmans, L. V. E., Auger, M. W., Treu, T., & Bolton, A. S. 2014, Monthly Notices of the Royal Astronomical Society, 442, 2017–2035

Vegetti, S., Koopmans, L. V. E., Bolton, A., Treu, T., & Gavazzi, R. 2010, Monthly Notices of the Royal Astronomical Society, 408, 1969

Vegetti, S., Lagattuta, D. J., McKean, J. P., et al. 2012, Nature, 481, 341–343

Verde, L., Heavens, A. F., Percival, W. J., et al. 2002, Monthly Notices of the Royal Astronomical Society, 335, 432–440

Vieira, J. D., Marrone, D. P., Chapman, S. C., et al. 2013, Nature, 495, 344

Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013, Physical Review D, 88

Viel, M., Becker, G. D., Bolton, J. S., et al. 2008, Phys. Rev. Lett., 100, 041304

Viel, M., Lesgourgues, J., Haehnelt, M. G., Matarrese, S., & Riotto, A. 2005, Physical Review D, 71

Viel, M., Lesgourgues, J., Haehnelt, M. G., Matarrese, S., & Riotto, A. 2006, Phys. Rev. Lett., 97, 071301

Vogelsberger, M., Genel, S., Springel, V., et al. 2014, MNRAS, 444, 1518

Vogelsberger, M., Zavala, J., Cyr-Racine, F.-Y., et al. 2016, Monthly Notices of the Royal Astronomical Society, 460, 1399–1416

Vogelsberger, M., Zavala, J., Simpson, C., & Jenkins, A. 2014, Monthly Notices of the Royal Astronomical Society, 444, 3684

Walker, M. G., Mateo, M., Olszewski, E. W., et al. 2016, ApJ, 819, 53

Walsh, D., Carswell, R. F., & Weymann, R. J. 1979, Nature, 279, 381

Walsh, S. M., Willman, B., & Jerjen, H. 2009, AJ, 137, 450

Wambsganss, J. 2006, arXiv e-prints, astro

Wang, J., Bose, S., Frenk, C. S., et al. 2020, Nature, 585, 39–42

Wang, W., Han, J., Cautun, M., Li, Z., & Ishigaki, M. N. 2019, arXiv e-prints, arXiv:1912.02599

Wang, W. & White, S. D. M. 2012, MNRAS, 424, 2574

Watkins, L. L., Evans, N. W., Belokurov, V., et al. 2009, MNRAS, 398, 1757

Webster, R. & Oliver, M. 2007, Geostatistics for Environmental Scientists, Statistics in Practice (Wiley)

Wechsler, R. H., Bullock, J. S., Primack, J. R., Kravtsov, A. V., & Dekel, A. 2002, The Astrophysical Journal, 568, 52–70

Weinberg, S. 1978, Phys. Rev. Lett., 40, 223

Wilczek, F. 1978, Phys. Rev. Lett., 40, 279

Wilson, C., Dong, R., di Francesco, J., et al. 2019, in Canadian Long Range Plan for Astronony and Astrophysics White Papers, Vol. 2020, 19

Wong, K. C., Suyu, S. H., Chen, G. C.-F., et al. 2019, Monthly Notices of the Royal Astronomical Society, 498, 1420–1439

Xu, D., Sluse, D., Gao, L., et al. 2015, MNRAS, 447, 3189

Xu, D., Springel, V., Sluse, D., et al. 2017, MNRAS, 469, 1824

Xu, D. D., Sluse, D., Gao, L., et al. 2013, How well can cold-dark-matter substructures account for the observed lensing flux-ratio anomalies?

Yoon, J. H., Johnston, K. V., & Hogg, D. W. 2011, The Astrophysical Journal, 731, 58

Yoshida, N., Springel, V., White, S. D. M., & Tormen, G. 2000, The Astrophysical Journal, 544, L87–L90

Yèche, C., Palanque-Delabrouille, N., Baur, J., & Bourboux, H. d. M. d. 2017, Journal of Cosmology and Astroparticle Physics, 2017, 047–047

Zhang, J., Liu, H., & Chu, M.-C. 2018, Frontiers in Astronomy and Space Sciences, 5, 48

Zwicky, F. 1933, Helvetica Physica Acta, 6, 110

Acknowledgements

I want to express my deepest gratitude to my supervisor Simona Vegetti for all her support, encouragement, and guidance throughout the last years. This thesis would not have been possible without the time she dedicated to me and my project.

I further want to thank everybody from my group for sharing their expertise and unique perspectives with me, in particular, Giulia Despali, Elisa Ritondale, Francesca Rizzo, Devon Powell, and Hannah Stacey.

It has been a privilege to work with the people from the Max Planck Institute for Astrophysics. I thank Torsten Enßlin for being my official supervisor at the Ludwig-Maximilians-University Munich and many stimulating conversations about Bayesian statistics. I thank Simon White for the many opportunities he gave me to discuss various aspects of this thesis. Likewise, I am very grateful for my colleague Mohammadreza Ayromlou and his most considerate and collegial nature: It has been a pleasure sharing the best office with you throughout this time! Additionally, I would like to thank Benedetta Ciardi, Matteo Frigo, Stefan Hilbert, Chris Byrohl, Ulrich Steinwandel, Reimar Leike, Jakob Knollmüller, Maria Depner, Gabriele Kratschmann, Hannelore Hämmerle, Andreas Weiss, Stefan Schuldt, Sherry Suyu, Sonia Ayari, and Thorsten Naab.

I further thank all the collaborators with whom I worked over the past years: Chris Fassnacht, John McKean, Matt Auger, Mark Lovell, Leon Koopmans, Jens Jasche, Florent Leclercq, R. Benton Metcalf, Jen-Wei Hsueh, Matteo Viel, Oliver Newton, Riccardo Murgia, Carlos Frenk, Marius Cautun. Thank you for the extremely valuable feedback over the course of this dissertation.

Finally, I thank Bastian, Rudi, and Regina Mörtl, Timo Kluth, Markus Thoma, Peter Rogler, Tamara Roth, Paul Mitbach, Sabine Zips, Ulrike Zweck, Gabrielle Peterson, Bernhard Springer, Philipp Klose, Mayra Munoz, Achim Suchowierksi, and my family: Peter, Rosemarie, Bernhard, and Elisabeth, who always supported me throughout my academic career.