



Universiteit
Leiden
The Netherlands

Statistical strong lensing: III. Inferences with complete samples of lenses

Sonnenfeld, A.

Citation

Sonnenfeld, A. (2022). Statistical strong lensing: III. Inferences with complete samples of lenses. *Astronomy & Astrophysics*, 659. doi:10.1051/0004-6361/202142301

Version: Accepted Manuscript

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/3561976>

Note: To cite this publication please use the final published version (if applicable).

Statistical strong lensing. III. Inferences with complete samples of lenses

Alessandro Sonnenfeld¹

Leiden Observatory, Leiden University, Niels Bohrweg 2, 2333 CA Leiden, the Netherlands
e-mail: sonnenfeld@strw.leidenuniv.nl

ABSTRACT

Context. Existing samples of strong lenses have been assembled by giving priority to sample size, at the cost of having a complex selection function. With the advent of the next generation of wide-field photometric surveys, however, it might become possible to identify subsets of the lens population with well-defined selection criteria, trading sample size for completeness.

Aims. There are two main advantages of working with a complete sample of lenses. First, it is possible to recover the properties of the general population of galaxies, of which strong lenses are a biased subset. Second, the relative number of lenses and non-detections can be used to further constrain models of galaxy structure. This work illustrates how to carry out a statistical strong lensing analysis that takes advantage of these features.

Methods. I introduced a general formalism for the statistical analysis of a sample of strong lenses with known selection function, then tested it on simulated data. The simulation consists of a population of 10^5 galaxies with an axisymmetric power-law density profile, a population of background point sources, and a subset of $\sim 10^3$ strong lenses, complete above an observational cut.

Results. The method allows to recover the distribution in Einstein radius and mass density slope of the galaxy population in an unbiased way. The number of non-lenses helps to constrain the model when magnification data are not available.

Conclusions. Complete samples of lenses are a powerful asset to turn precise strong lensing measurements into accurate statements on the properties of the general galaxy population.

Key words. Gravitational lensing: strong – Galaxies: fundamental parameters

1. Introduction

Strong gravitational lenses have been used extensively to study the structure and late ($z < 1$) evolution of massive galaxies, both through the detailed study of a small number of systems (see for example Sonnenfeld et al. 2012; Barnabè et al. 2013; Smith et al. 2015a; Collett et al. 2018; Schuldt et al. 2019), and with the statistical combination of tens of objects (Auger et al. 2010; Bolton et al. 2012; Schechter et al. 2014; Oguri et al. 2014; Sonnenfeld et al. 2015; Oldham & Auger 2018). Strong lenses, however, are rare systems, and they usually require dedicated follow-up observations to confirm their nature and to obtain all of the necessary information for a successful analysis. For these reasons, past lens finding campaigns have been carried out with the primary goal of maximising sample size while minimising costs, in terms of observational time. This approach has made it possible, in the past decade, to carry out invaluable investigations of the structure of galaxies. At the same time, however, it has made the selection function of the existing samples of lenses a highly non-trivial one.

Let us consider, for example, the Strong Lensing Legacy Survey (SL2S, Ruff et al. 2011; Gavazzi et al. 2012; Sonnenfeld et al. 2013a,b). SL2S lens candidates were first identified in ground-based imaging data, with the automatic lens-finding tool RINGFINDER (Gavazzi et al. 2014), the performance of which was quantified on simulated lenses. Subsequently, lens candidates were visually inspected and ranked on the basis of their likelihood of being strong lenses, according to the judgement of strong lensing experts. A few tens of the most promising lens candidates were followed-up with spectroscopic observa-

tions from $8m$ -class telescopes and with high-resolution imaging data from the Hubble Space Telescope, with a higher priority given to lenses with larger photometric redshift. Finally, the systems for which follow-up observations yielded measurements of the lens and source redshift were used for the science analyses.

The final sample was then the result of a combination of algorithms, human decisions, and the efficiency of obtaining source redshifts from spectroscopic observations. The probability of a strong lens to be included in the SL2S sample is, at the very least, a function of Einstein radius, source surface brightness, image configuration, lens-source contrast, source redshift and source emission line strength. Clearly, it is very difficult to obtain an accurate description of such a high-dimensional probability distribution.

The main problem of working with lens samples with an unknown selection function, is that it is difficult to generalise the measurements obtained with them to the population of galaxies from which they are drawn. Strong lens samples are a biased subset of the general galaxy population, for two reasons. First, galaxies that are more massive and more compact have a higher cross-section for strong lensing, and therefore are more likely to be strong lenses. Second, the probability of a lens to be included in a given survey depends on the efficiency of the survey at discovering it and obtaining all of the necessary data for its analysis.

In the next few years, upcoming wide-field surveys such as Euclid¹, the LSST² and the Chinese Space Station Telescope

¹ <https://www.euclid-ec.org/>

² <https://www.lsst.org/>

(CSST) will enable the discovery of order 10^5 new strong lenses (Collett 2015). The first paper of this series (Sonnenfeld & Cauntun 2021, hereafter Paper I), however, showed how a sample of 10^3 strong lenses can be sufficient to obtain precise and accurate measurements of several important parameters of the structure of galaxies, such as the stellar mass-to-light ratio and the dark matter density profile. This means that, using the full set of 10^5 Euclid-LSST-CSST lenses as starting point, there will be ample freedom in determining how such a sample is defined. Depending on the efficiency of the lens finding algorithms that will be used, it might be possible to define subsamples of the lens population that are complete within a well-defined region in observational space. The goal of this paper is to illustrate the advantages of working with such a complete subsample of lenses.

When the properties of the background source population are known, there are two benefits in having a complete lens sample. The first is that it is possible to model the mapping between the general population of galaxies and that of the lens sample, thus removing the bias introduced by strong lensing. The second is that the number of non-detections can be used along with the lensing data to further constrain models of galaxy structure. In this work, I illustrate how to take advantage of these features in a statistical strong lensing inference.

Building on past work (Sonnenfeld et al. 2015, 2019), I first introduce a general formalism for inferring the properties of the galaxy population from the analysis of a strong lensing survey with known selection function. I then extend it to incorporate information from the number of non-lenses in the survey. I test this method on a simulated complete set of lenses, generated under a series of simplifying assumptions. Although additional work is needed to apply this formalism to real samples of lenses, this is the necessary first step towards enabling accurate inferences of galaxy properties with samples of strong lenses.

The structure of this work is the following. In Section 2 I introduce the theoretical formalism on which this work is based. In Section 3 I describe the simulations used to test the analysis method. In Section 4 I describe the model that I used to fit the simulated data. In Section 5 I show the results of the experiment. I discuss the results in Section 6 and draw conclusions in Section 7.

The Python code used for the simulation and analysis of the lens sample can be found in a dedicated section of a GitHub repository³.

2. Theory

In this section I introduce the theoretical foundation for the analysis method that I use in this work. In section 2.1 I define the quantities on which this method is based. In section 2.2 I explain how to infer the properties of a population of galaxies from the observation of a subset of strong lenses drawn from it. In section 2.3 I explain how to combine information from the strong lenses with non-detections.

2.1. Definitions

2.1.1. Parameters and distributions

Let us consider a strong lens survey that searches for lenses among a well-defined population of foreground galaxies and background sources, with a lens selection criterion S . Let us assume that all of the properties of a galaxy that are relevant for

determining its probability of being a lens can be summarised with a set of parameters ψ_g . This set of parameters will include the redshift of the galaxy and quantities describing its mass distribution. Let ψ_s be a set of parameters that are relevant for determining the probability of a background source to be strongly lensed. These parameters will include the redshift of the source, and quantities describing its surface brightness distribution.

Given a lens-source pair, (ψ_g, ψ_s) , let $P_{\text{det}}(\psi_g, \psi_s|S)$ be the probability of the two objects to produce a lens that is detected by our survey, given the selection criterion S . Finally, let $P_g(\psi_g)$ and $P_s(\psi_s)$ be the probability distribution of the foreground galaxy and background source population, respectively. With these definitions, the (unnormalised) probability distribution of the population of strong lenses resulting from the survey is

$$P_{\text{SL}}(\psi_g, \psi_s|S) \propto P_g(\psi_g)P_s(\psi_s)P_{\text{det}}(\psi_g, \psi_s|S). \quad (1)$$

The main quantity of interest is the distribution of the properties of the lens galaxy population, $P_g(\psi_g)$. Strong lensing data, however, can only directly constrain the left-hand side term, $P_{\text{SL}}(\psi_g, \psi_s|S)$. This is, in a broad sense, an inverse problem. As I show in this work, if the background source population $P_s(\psi_s)$ and the lens detection probability $P_{\text{det}}(\psi_g, \psi_s|S)$ are known, it is possible to effectively invert Equation 1 and solve for $P_g(\psi_g)$. Table 1 lists the definitions of all of the quantities introduced in this section, as well as those that are defined later in this paper.

2.1.2. The lensing cross-section

A very important parameter, that is necessary to evaluate the likelihood of strong lensing data, is the source position relative to the lens, β . This is a source parameter, therefore it is included in ψ_s . However, it is convenient to split ψ_s as follows:

$$\psi_s = (\psi_s^{(-\beta)}, \beta), \quad (2)$$

where $\psi_s^{(-\beta)}$ includes all of the remaining source parameters.

Let us assume that the angular position in the sky does not correlate with any other source property. Given a strong lens with parameters $(\psi_g, \psi_s^{(-\beta)})$, the probability of the source position being β is proportional to

$$P_{\text{SL}}(\beta|\psi_g, \psi_s^{(-\beta)}, S) \propto U_{\mathbb{R}^2}(\beta)P_{\text{det}}(\psi_g, \psi_s^{(-\beta)}, \beta|S), \quad (3)$$

where $U_{\mathbb{R}^2}(\beta)$ is a uniform distribution over the sky. The inverse of the normalisation constant of this distribution is

$$\sigma_{\text{SL}}(\psi_g, \psi_s^{(-\beta)}|S) = \int_{\mathbb{R}^2} d\beta P_{\text{det}}(\psi_g, \psi_s^{(-\beta)}, \beta|S). \quad (4)$$

I refer to this quantity as the strong lensing cross-section.

I point out that σ_{SL} does not necessarily correspond to a well-defined region in the source plane, within which a source is always lensed. This is because the detection probability P_{det} can take any value between 0 and 1, meaning that certain lens-source pairs only give rise to detected lenses with some probability.

2.1.3. Selection criteria

In this work, I focus on a particular class of lens selection criteria, based on a cut in observational space. Let \mathbf{d} be an array that summarises the observational data that are relevant for the detection and modelling of a strong lens. This array will include, for instance, the observed lens and source redshift (formally, these are different from the true redshifts that are included in ψ_g and

³ https://github.com/astrosonnen/strong_lensing_tools

Table 1. Notations used in this work

Quantity	Description
ψ_g	Set of parameters describing a foreground galaxy.
ψ_s	Set of parameters describing a background source.
$P_g(\psi_g)$	Probability distribution of the foreground galaxy population.
$P_s(\psi_s)$	Probability distribution of the background source population.
S	Selection criteria of the strong lens survey.
$P_{SL}(\psi_g, \psi_s S)$	Probability distribution of the strong lens population, given the survey selection criteria.
β	Angular position of the source.
$\psi_s^{(-\beta)}$	Set of background source parameters, excluding the angular position.
$\sigma_{SL}(\psi_g, \psi_s^{(-\beta)} S)$	Strong lensing cross-section of a galaxy with parameters ψ_g , relative to a source with parameters $\psi_s^{(-\beta)}$.
\mathbf{d}	Array of observational data (of the entire survey, or relative to a generic system).
Ω	Volume of the data space defined by the selection criteria of the survey.
η	Set of parameters describing the population of foreground galaxies.
N_s	Number of strongly lensed sources associated to a foreground galaxy.
λ	Expectation value of the number of strongly lensed sources associated to a galaxy.
n_{bkg}	Projected number density of background sources.
$P(\psi_s \text{Lens}, S)$	Probability distribution of the source parameters, given that the source is strongly lensed.
θ	Angular coordinate in the image plane of a strong lens.
$\alpha(\theta)$	Deflection angle of the lens.
θ_{Ein}	Einstein radius of a lens, in angular units.
γ	Slope of the density profile of a power-law lens.
θ_1, θ_2	Coordinates of image 1 and 2 (the main image of a strongly lensed source and the counter image).
$\mathcal{N}(\mu, \sigma^2)$	Normal distribution with mean μ and variance σ^2 .
m_s	Source intrinsic (unlensed) magnitude.
m_s^*	Knee of the source magnitude distribution function.
α_s	Faint-end slope of the source magnitude distribution function.
$\mu(\theta)$	Lensing magnification at image position θ .
m_1, m_2	True magnitudes of images 1 and 2.
$m_1^{\text{obs}}, m_2^{\text{obs}}$	Observed magnitudes of images 1 and 2.
ϵ_m	Observational uncertainty (1σ) on the image magnitudes.
m_{max}	Maximum allowed magnitude of image 2, for a system to be defined a strong lens.
μ_{Ein}	Parameter describing the mean $\log \theta_{\text{Ein}}$ of the galaxy population.
σ_{Ein}	Parameter describing the scatter in $\log \theta_{\text{Ein}}$ of the galaxy population.
μ_γ	Parameter describing the mean γ of the galaxy population.
σ_γ	Parameter describing the scatter in γ of the galaxy population.
ξ_{asymm}	Image configuration asymmetry, defined in Equation 30.
θ_{Ein}^*	Value of θ_{Ein} needed to reproduce the observed image positions, for a given γ .
β^*	Value of β needed to reproduce the observed image positions, for a given γ .
$\det J$	Jacobian determinant of the variable change $(\log \theta_{\text{Ein}}, \gamma) \rightarrow (\theta_1, \theta_2)$.

ψ_s), the surface brightness distribution of the lens galaxy, the number of images of the strongly lensed source, their position and surface brightness distribution. I then consider, as selection criterion, one that only selects lenses within a volume Ω of the multi-dimensional space spanned by the data array:

$$S = \mathbf{d} \in \Omega. \quad (5)$$

An example of such a criterion is the following: only lens galaxies brighter than a minimum luminosity, within a given redshift range, with at least two strongly lensed images brighter than a given threshold are included in the sample.

Finally, I assume that it is possible to define a volume Ω in observational space and design a lens survey that retrieves all of the existing lenses in Ω . I refer to that as a complete lens sample over the volume Ω .

2.2. Inference with strong lenses alone

Let us assume that the galaxy distribution can be described with an analytical model, summarised by a set of population-level pa-

rameters η :

$$P_g(\psi_g) = P_g(\psi_g|\eta). \quad (6)$$

The goal of this method is to infer the posterior probability distribution of η given the data from the ensemble of strong lenses, $\{\mathbf{d}\}$. From Bayes' theorem, this is given by the product between the prior probability of η and the likelihood of observing the data given η :

$$P(\eta|\{\mathbf{d}\}) \propto P(\eta)P(\{\mathbf{d}\}|\eta). \quad (7)$$

Measurements obtained on different lenses are independent from each other. Therefore, the likelihood term can be written as

$$P(\{\mathbf{d}\}|\eta) = \prod_i P(\mathbf{d}_i|\eta), \quad (8)$$

where \mathbf{d}_i is the data array relative to the i -th lens. The observed data depends only indirectly on the lens population parameters η : only after specifying the parameters describing the lens and the source properties, ψ_g and ψ_s , it is possible to evaluate the

likelihood. Since these are not univocally set by $\boldsymbol{\eta}$, it is necessary to marginalise over all possible values taken by them. In other words, each factor in the product of Equation 8 becomes

$$P(\mathbf{d}_i|\boldsymbol{\eta}) = \int d\boldsymbol{\psi}_g d\boldsymbol{\psi}_s P(\mathbf{d}_i|\boldsymbol{\psi}_g, \boldsymbol{\psi}_s) P_{\text{SL}}(\boldsymbol{\psi}_g, \boldsymbol{\psi}_s|\boldsymbol{\eta}, S), \quad (9)$$

where $P_{\text{SL}}(\boldsymbol{\psi}_g, \boldsymbol{\psi}_s|\boldsymbol{\eta}, S)$ is the normalised probability distribution of the strong lenses, given the galaxy population parameters $\boldsymbol{\eta}$ and the lens selection criterion S . With these definitions, it is possible, in principle, to explore the parameter space of $\boldsymbol{\eta}$ and compute its posterior probability distribution.

Equation 9 is very similar in form to Equation (30) of Paper I. The main difference is that the formalism adopted in that case was aimed purely at inferring the properties of the lens galaxies. Here, instead, the goal is to recover the distribution of the general galaxy population, $P_g(\boldsymbol{\psi}_g|\boldsymbol{\eta})$, which enters the integrand function of Equation 9 through the product of Equation 1.

Evaluating integrals of the kind of Equation 9 in an accurate way requires a good knowledge of the source population and the lens detection probability. Having a complete lens set is not a strict requirement for this purpose, but, as I discuss later, simplifies the problem and guarantees robustness.

2.3. Inference with lenses and non-detections

Section 2.2 explained how to infer the properties of the galaxy population, summarised by the parameters $\boldsymbol{\eta}$, when using only data from detected strong lenses. When considering also non-detections, we need to adopt a slightly different formalism.

First of all, the data array \mathbf{d} is modified to include an additional observable: the number of strongly lensed sources, N_s . This is $N_s = 0$ for the non-detections and $N_s = 1$ (or larger, in the rare cases in which more than one source is strongly lensed by the same galaxy) for the lenses. Second, the product of Equation 8 is now calculated over the entire galaxy sample. Third, only a fraction of the galaxies in the sample have a background source associated with them. Given these considerations, it is useful to first expand the integrands of Equation 8 in terms of the galaxy parameters only:

$$P(\mathbf{d}_i|\boldsymbol{\eta}) = \int d\boldsymbol{\psi}_g P(\mathbf{d}_i|\boldsymbol{\psi}_g, S) P(\boldsymbol{\psi}_g|\boldsymbol{\eta}). \quad (10)$$

In this case, the galaxy probability distribution is simply described by $P(\boldsymbol{\psi}_g|\boldsymbol{\eta})$, which is the distribution of the general galaxy population. This is because galaxies are included in the sample independently of whether they produce a detectable strong lens or not.

Let us consider the case in which a given galaxy is not a strong lens. The only data available in this situation are the number of strongly lensed sources ($N_s = 0$) and observables related to the lens galaxy, which I indicate with \mathbf{d}_g . The likelihood term $P(\mathbf{d}_i|\boldsymbol{\psi}_g, S)$ can then be written as

$$P(\mathbf{d}_i|\boldsymbol{\psi}_g, S)^{(\text{Non-lens})} = P(N_s = 0|\boldsymbol{\psi}_g, S) P(\mathbf{d}_g|\boldsymbol{\psi}_g). \quad (11)$$

The first factor on the right-hand side of the above equation is the likelihood of observing no strongly lensed sources, given the galaxy parameters and the selection criteria of the survey. The probability of detecting N_s lensed sources is given by Poisson statistics:

$$P(N_s|\boldsymbol{\psi}_g, S) = \frac{\lambda(\boldsymbol{\psi}_g, S)^{N_s} e^{-\lambda(\boldsymbol{\psi}_g, S)}}{N_s!}, \quad (12)$$

where $\lambda(\boldsymbol{\psi}_g, S)$ is the average number of strongly lensed sources that are expected for a galaxy with parameters $\boldsymbol{\psi}_g$ and a survey with selection criteria S . In the limit $\lambda \ll 1$, usually satisfied in the galaxy-scale strong lensing regime, we have

$$\begin{aligned} P(0|\boldsymbol{\psi}_g, \boldsymbol{\psi}_s, S) &\approx 1 - \lambda \\ P(1|\boldsymbol{\psi}_g, \boldsymbol{\psi}_s, S) &\approx \lambda \\ P(N_s > 1|\boldsymbol{\psi}_g, \boldsymbol{\psi}_s, S) &\approx 0, \end{aligned} \quad (13)$$

to first order in λ . Assuming that the background sources are uniformly distributed in space with average projected number density n_{bkg} , then $\lambda(\boldsymbol{\psi}_g, S)$ is given by

$$\lambda(\boldsymbol{\psi}_g, S) = n_{\text{bkg}} \int d\boldsymbol{\psi}_s^{(-\beta)} P(\boldsymbol{\psi}_s^{(-\beta)}) \sigma_{\text{SL}}(\boldsymbol{\psi}_g, \boldsymbol{\psi}_s^{(-\beta)}|S), \quad (14)$$

which is the product between n_{bkg} and the strong lensing cross-section, averaged over the distribution of source parameters other than position, $P(\boldsymbol{\psi}_s^{(-\beta)})$.

If a galaxy is a strong lens, then the data vector includes also observables related to the lensed source(s), \mathbf{d}_s . In order to compute the likelihood term $P(\mathbf{d}_i|\boldsymbol{\eta})$, then, one must expand Equation 10 by including also the source parameters. In the case of a single source, this becomes

$$\begin{aligned} P(\mathbf{d}_i|\boldsymbol{\eta})^{(\text{Lens})} &= \int d\boldsymbol{\psi}_g d\boldsymbol{\psi}_s P(N_s|\boldsymbol{\psi}_g, S) P(\mathbf{d}_g|\boldsymbol{\psi}_g) P(\mathbf{d}_s|\boldsymbol{\psi}_g, \boldsymbol{\psi}_s) \times \\ &P(\boldsymbol{\psi}_g|\boldsymbol{\eta}) P(\boldsymbol{\psi}_s|\text{Lens}, \boldsymbol{\psi}_g, S). \end{aligned} \quad (15)$$

In the integrand function, the probability of observing the data given the lens and source parameters and the selection criteria is written as the product of three factors. The first factor is the probability of observing N_s lensed sources given the lens parameters and selection criteria, which is computed in the same way as in the non-detection case: first by calculating λ via Equation 14, then by taking the Poisson probability of Equation 12. The second factor in the integrand is the probability of observing galaxy-related data \mathbf{d}_g given the parameters $\boldsymbol{\psi}_g$, while the third factor is the probability of observing the lensed image configuration \mathbf{d}_s given the lens and source parameters.

Finally, the last factor in the integrand of Equation 15 is the probability distribution of the source parameters. This factor must take into account the fact that the source is strongly lensed. Therefore, its unnormalised probability distribution is given by

$$P(\boldsymbol{\psi}_s|\text{Lens}, \boldsymbol{\psi}_g, S) \propto P_s(\boldsymbol{\psi}_s) P_{\text{det}}(\boldsymbol{\psi}_g, \boldsymbol{\psi}_s|S). \quad (16)$$

There is a subtle difference between the source parameters $\boldsymbol{\psi}_s$ of Equation 15 and those entering the estimation of λ in Equation 14. The former describe a well-defined background source that corresponds to the observed strongly lensed images. The latter refer generically to the source population, and not to a specific source. For this reason, the probability of observing N_s sources depends only on the lens parameters and not on those of its associated source.

On a related note, there need not be a strict correspondence between the number of detected sources and the number of sources that are taken into account in a lens model. For example, if more than a single source is detected, one may decide to only explicitly model one of them, using Equation 15, all of them, by expanding Equation 15 over new dimensions $\boldsymbol{\psi}_{s,2}, \dots, \boldsymbol{\psi}_{s,N_s}$, or even none of them. In this last case, the likelihood can be evaluated simply with Equation 10.

3. Simulations

In this section I describe the simulations created to test the method introduced in Section 2. The main purpose of these simulations is to show concrete examples of how to carry out an inference with a complete set of lenses. In order to make these examples easy to follow, I created the simulations by making the following series of simplifying assumptions:

- All of the foreground galaxies have an axisymmetric mass distribution.
- All of the background sources are point-like.
- All of the foreground galaxies are at the same redshift, z_g .
- All of the background sources are at the same redshift, z_s .

In section 3.1 I describe the properties of the foreground galaxy population. In section 3.2 I describe the background source population. In section 3.3 I describe the selection criterion of the resulting strong lens sample.

3.1. Foreground galaxy population

Each foreground galaxy is described by an axisymmetric power-law mass distribution. Let θ be the angular coordinate in the image plane along an axis with origin on the lens centre. The lensing deflection angle of such a galaxy as a function of θ is

$$\alpha(\theta) = \theta_{\text{Ein}} \frac{\theta}{|\theta|} \left(\frac{|\theta|}{\theta_{\text{Ein}}} \right)^{2-\gamma}, \quad (17)$$

where θ_{Ein} is the Einstein radius and γ is the density slope (the spherical deprojection of a power-law lens has a density profile $\rho(r) \propto r^{-\gamma}$).

As section 2.1 of Paper I shows, the types of lensed image configurations that such a lens can produce depend on the value of γ . When $\gamma < 2$, the lens has a non-degenerate radial caustic: sources located inside the caustic are strongly lensed into three images. Assuming without loss of generality that the source position β is positive, the three images ($\theta_1, \theta_2, \theta_3$) are located at $\theta_1 > \theta_{\text{Ein}}$, $-\theta_{\text{Ein}} < \theta_2 < -\theta_{\text{rad}}$ and $-\theta_{\text{rad}} < \theta_3 < 0$, where θ_{rad} is the radius of the radial critical curve.

Lenses with $\gamma > 2$ always produce two multiple images, at $\theta_1 > \theta_{\text{Ein}}$ and $-\theta_{\text{Ein}} < \theta_2 < 0$. However, the magnification of image 2 goes to zero when the source position is large. Finally, when $\gamma = 2$ (not shown in Paper I), sources located within $\beta < \theta_{\text{Ein}}$ are strongly lensed into two images. For more details of the lensing properties of axisymmetric power-law lenses, I refer to section 2.1 of Paper I.

I generated 10^5 such galaxies, drawing $\log \theta_{\text{Ein}}$ and γ from the following distributions:

$$\log \theta_{\text{Ein}} \sim \mathcal{N}(-0.3, 0.2^2) \quad (18)$$

$$\gamma \sim \mathcal{N}(2.0, 0.2^2), \quad (19)$$

where $\mathcal{N}(\mu, \sigma^2)$ indicates a Gaussian with mean μ and variance σ^2 . The Einstein radius is in units of arcsec, and contains implicitly information on the lens and source redshift and on the cosmology.

3.2. Background source population

I obtained the distribution in apparent magnitude m_s of the background sources as follows. First, I considered the rest-frame 1500\AA luminosity function of galaxies at $1.5 < z < 2.5$ measured by Parsa et al. (2016). Then, I converted it to the g -band

apparent magnitude assuming a fixed redshift of $z = 1.5$, which is a typical value for a galaxy-scale strong lens, and a fixed spectral energy distribution, the ‘starb1’ template from the Kinney-Calzetti Spectral Atlas of Galaxies (Calzetti et al. 1994; Kinney et al. 1996). The resulting apparent magnitude distribution is the following:

$$m_s \sim \left(10^{-0.4(m_s - m_s^*)} \right)^{\alpha_s + 1} \exp \left\{ -10^{-0.4(m_s - m_s^*)} \right\}, \quad (20)$$

with $m_s^* = 24.47$ (also referred to as the knee of the magnitude distribution) and $\alpha_s = -1.32$ (the faint-end slope). For the sake of computational efficiency, I truncated the distribution outside of the range $m_s \in [23.0, 28.0]$. This cut does not affect the results: the number of galaxies brighter than 23 in the distribution of Equation 20 is highly suppressed, and, because of the lens selection criterion that is introduced in section 3.3, the fraction of galaxies fainter than 28 that results in detected strong lenses is also very small.

I generated, behind each lens, a random realisation of background source positions, drawn from a uniform distribution with average number density $n_{\text{bkg}} = 100 \text{ arcmin}^{-2}$. Most of these sources are very faint, by the standard of current wide-field photometric surveys. When considering only sources brighter than 26, n_{bkg} is a factor of four smaller and matches closely the number density of sources in the Hyper Suprime-Cam Subaru Strategic Program (HSC survey Aihara et al. 2018) weak lensing shear catalogue (approximately 25 arcmin^{-2} , Mandelbaum et al. 2018).

3.3. Selection criteria

For each lens, I defined a search region in the source plane to identify background sources that are multiply imaged. For lenses with $\gamma < 2$, this search region was simply the area within the radial caustic. For lenses with $\gamma > 2$, I set the limit to a very large value of β . For each multiply-imaged source, I solved the lens equation,

$$\beta = \theta - \alpha(\theta), \quad (21)$$

to find the image positions. I assumed that only image 1 and 2 are observable, and discarded image 3, which, when present, is typically highly de-magnified. I calculated the magnification μ of each image and used it to compute its magnitude:

$$m_1 = m_s - 2.5 \log(|\mu(\theta_1)|) \quad (22)$$

$$m_2 = m_s - 2.5 \log(|\mu(\theta_2)|). \quad (23)$$

With m_1 and m_2 I indicate the true magnitudes of the two images. Then, I added a random observational error of $\epsilon_m = 0.1$ mag to m_1 and m_2 to generate observed magnitudes:

$$m_1^{(\text{obs})} \sim \mathcal{N}(m_1, \epsilon_m^2) \quad (24)$$

$$m_2^{(\text{obs})} \sim \mathcal{N}(m_2, \epsilon_m^2). \quad (25)$$

Given these data, I defined as a strong lens any lens-source pair for which the observed magnitude of image 2 is brighter than a limiting magnitude $m_{\text{max}} = 26.6$. The choice of this value is based on the lens search of Sonnenfeld et al. (2018) among data from the HSC survey: it corresponds to the 5σ detection limit for a point source in the g -band, which is the band used by Sonnenfeld et al. (2018) for the detection of lensed images. The selection criterion is then

$$S = \{m_2^{(\text{obs})} < m_{\text{max}}\}. \quad (26)$$

The rationale for this choice is the following. The process of finding lenses involves always a visual inspection step. During this step, lensing experts estimate to what degree a lens candidate can be considered a real strong lens. Only lenses that are identified with high confidence are used for a scientific analysis. One very important element that experts look for in their classification is the presence of multiple images (see for example Sonnenfeld et al. 2020). In all practical cases, image 1 is always brighter than image 2. Therefore, the condition of Equation 26 ensures that both images are detected.

Finally, I assumed that the strong lens survey is complete above the limit imposed with Equation 26: all of the lenses with an observed magnitude brighter than image 2 are included in the sample; all of the remaining lenses are discarded and treated as non-detections. Under these assumptions, the detection probability of a lens-source pair that produces two images is

$$P_{\text{det}}(\theta_{\text{Ein}}, \gamma, m_s, \beta | S) = \frac{1}{2} \left[1 - \text{erf} \left(\frac{m_2 - m_{\text{max}}}{\sqrt{2} \epsilon_m} \right) \right]. \quad (27)$$

The generated sample consists of 867 lenses, out of the initial 10^5 galaxies. This is a typical value of the ratio between the number of lenses and the number of inspected galaxies in the HSC lens search (Sonnenfeld et al. 2018).

Figure 1 shows the distribution in Einstein radius and density slope of the lens sample, compared to that of the parent population of galaxies. The lenses have on average a larger Einstein radius. This can be easily understood: at fixed density slope, a galaxy with a larger Einstein radius has a larger lensing cross-section (to be precise, $\sigma_{\text{SL}} \propto \theta_{\text{Ein}}^2$), and therefore is more likely to be a lens. Additionally, the average value of γ of the lens sample is also larger than that of all galaxies (this can be clearly seen in the histogram at the bottom right of Figure 1). This suggests that the lensing cross-section is also an increasing function of γ .

To better illustrate this feature, Figure 2 shows σ_{SL} as a function of γ , for lenses at fixed $\theta_{\text{Ein}} = 1''$ and sources of different magnitudes. The dotted line shows the area of the source-plane region within which a source is lensed into multiple images. This can be considered as the lensing cross-section of an infinitely bright source, which is an upper limit on σ_{SL} .

For low values of γ , σ_{SL} equals this upper limit, independently of the source magnitude: this means that, for any source position within the radial caustic, the magnification of image 2 is sufficiently large to guarantee its detection. As γ increases, we observe a different behaviour, depending on m_s . At faint magnitudes, σ_{SL} decreases: this means that, with increasing γ , a larger area of the source plane is mapped into sets of images with a low magnification of image 2. At bright magnitudes, however, image 2 can be detected even when it is de-magnified, and therefore σ_{SL} increases with γ .

To better understand this behaviour, Figure 3 shows the magnification of image 2 as a function of source position, for a few values of γ . For $\gamma < 2$, the area of the source plane that produces multiple images (limited by the dashed line) decreases with decreasing γ , but the magnification of image 2 increases at fixed β . For $\gamma > 2$, the $\mu - \beta$ curves of different lenses cross each other. This crossing is responsible for the different trends of σ_{SL} with γ of sources with different magnitudes.

The net effect, when averaging over the source population, is that galaxies with a higher γ have a larger cross-section, and therefore the mean γ of the lens population is larger than that of the parent galaxy population. This, however, is a result of the particular choice of the source magnitude distribution, and is not a

general property of power-law lenses. For example, with a fainter knee of the source magnitude distribution (a higher value of parameter m_s^* in Equation 20), the difference in the average density slope between the general population of galaxies and the strong lenses becomes smaller.

4. Model

The simulation produced the following data:

- 867 strong lenses, each with measurements of the two image positions, θ_1^{obs} and θ_2^{obs} , and related magnitudes, m_1^{obs} and m_2^{obs} .
- 99133 non-detections.

I fitted these data with a population of galaxies with power-law density profile, and a population of background sources described by the same magnitude distribution as that used to generate the sample, Equation 20. I described each galaxy with the pair of parameters $(\log \theta_{\text{Ein}}, \gamma)$, and assumed that their population distribution is the following product of Gaussians:

$$P_{\text{g}}(\log \theta_{\text{Ein}}, \gamma | \boldsymbol{\eta}) = \mathcal{N}(\mu_{\text{Ein}}, \sigma_{\text{Ein}}^2) \mathcal{N}(\mu_{\gamma}, \sigma_{\gamma}^2). \quad (28)$$

The population parameters of the model are the mean and standard deviations of these two Gaussians:

$$\boldsymbol{\eta} \equiv \{\mu_{\text{Ein}}, \sigma_{\text{Ein}}, \mu_{\gamma}, \sigma_{\gamma}\}. \quad (29)$$

The distribution is expressed in terms of the logarithm of the Einstein radius, rather than θ_{Ein} , as this was the form used in section 3.1 to generate the galaxy population. In general, a log-normal distribution in θ_{Ein} is preferred over a simple Gaussian in θ_{Ein} when the Einstein radius distribution of the sample is not limited to a well-defined order of magnitude. This is the case, for example, when both galaxy-scale and group- or cluster-scale lenses are included in the population.

This model reproduces exactly the true distribution that was assumed when creating the mock, for $\boldsymbol{\eta} = \{-0.3, 0.2, 2.0, 0.2\}$. In other words, the model is accurate by construction. This makes it possible to test for the accuracy of the inference method of Section 2.

Each lens has four degrees of freedom, θ_{Ein} , γ , β and m_s , and four observational constraints: the two image positions and magnitudes. Therefore, unlike in Paper I, the model is fully constrained on an individual lens basis. The two magnitudes, however, are only measured with a precision of 0.1 mag. When fitting a power-law model to an individual lens, this translates into an uncertainty of on the density slope γ in the range 0.05 – 0.3, depending on the image configuration. This is illustrated in Figure 4, which shows the 68% credible region on the Einstein radius and density slope of a subset of lenses, obtained by assuming flat priors on all of the parameters. The circles in Figure 4 show the true values of θ_{Ein} and γ of the lenses, colour-coded by the image asymmetry parameter, which is defined as

$$\xi_{\text{asymm}} = \frac{\theta_1 + \theta_2}{\theta_1 - \theta_2}. \quad (30)$$

Lenses with more asymmetric configurations (darker points) allow for a more precise inference on γ , though with a stronger degeneracy with θ_{Ein} . These individual posteriors are just shown for illustrative purposes and are not used in the analysis.

I fitted the model first to the sample of strong lenses, and then to the whole sample, including the non-detections. I assumed a flat prior on each parameter, over the range reported in Table 2.

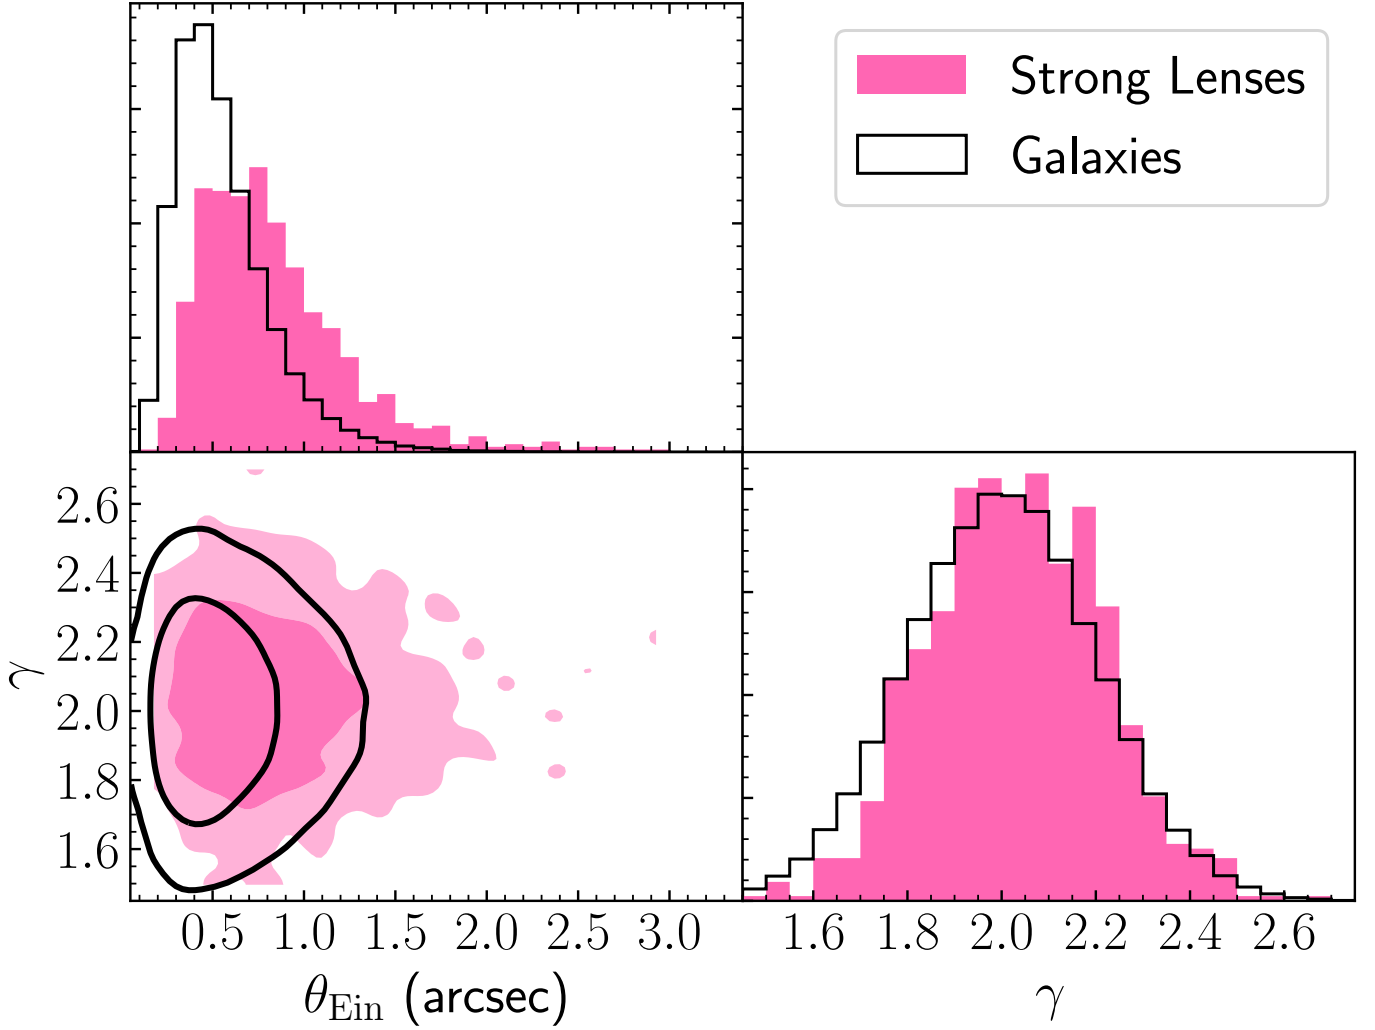


Fig. 1. Distribution in Einstein radius and density slope of the strong lens sample (filled contours and histograms) and of the parent population of galaxies (solid lines). Contour levels enclose 68% and 95% of the sample. The histograms are normalised by the number of objects.

I sampled the posterior probability distribution $P(\boldsymbol{\eta}|\{\mathbf{d}\})$ with a Markov Chain Monte Carlo (MCMC) method.

The main technical challenge in this experiment is the computation of the likelihood. Each system in the sample contributes to the likelihood with a multi-dimensional integral: Equation 9 (when using only lenses) or Equation 15 (when using the whole sample). In the next two sections, I explain the procedure that I adopted to compute these integrals.

4.1. Likelihood computation: lenses-only case

When fitting the observed data with the power-law model introduced in this section, Equation 9 becomes the following 4-dimensional integral,

$$\begin{aligned}
 P(\mathbf{d}|\boldsymbol{\eta}) = A(\boldsymbol{\eta}) \int d\gamma dm_s d \log \theta_{\text{Ein}} d\beta & P(\theta_1^{\text{obs}}, \theta_2^{\text{obs}}|\theta_{\text{Ein}}, \gamma, \beta) \times \\
 & \times P(m_1^{\text{obs}}, m_2^{\text{obs}}|\theta_{\text{Ein}}, \gamma, \beta, m_s) P(\theta_{\text{Ein}}, \gamma|\boldsymbol{\eta}) \times \\
 & \times P_{\text{det}}(\theta_{\text{Ein}}, \gamma, m_s, \beta|S) P(m_s), \quad (31)
 \end{aligned}$$

where $A(\boldsymbol{\eta})$ is a normalisation constant defined as follows:

$$\begin{aligned}
 A(\boldsymbol{\eta}) \int d\gamma dm_s d \log \theta_{\text{Ein}} d\beta & P(\theta_{\text{Ein}}, \gamma|\boldsymbol{\eta}) P_{\text{det}}(\theta_{\text{Ein}}, \gamma, m_s, \beta|S) \times \\
 & \times P(m_s) = 1. \quad (32)
 \end{aligned}$$

At fixed θ_{Ein} , γ and m_s , the integral over β in Equation 32 gives the strong-lensing cross-section defined in Equation 4. The full integral of Equation 32 is therefore the population-averaged σ_{SL} . This integral separates into a part that depends only on θ_{Ein} , and a part that depends on (γ, m_s) . I computed the former with a general purpose 1D integration method, and the latter via Monte Carlo integration.

Let us now consider the image position likelihood term in Equation 31. Assuming, as in Paper I, that the image positions are measured exactly, we have

$$\begin{aligned}
 P(\theta_1^{\text{obs}}, \theta_2^{\text{obs}}|\theta_{\text{Ein}}, \gamma, \beta) = & \delta(\theta_1(\theta_{\text{Ein}}, \gamma, \beta) - \theta_1^{\text{obs}}) \times \\
 & \delta(\theta_2(\theta_{\text{Ein}}, \gamma, \beta) - \theta_2^{\text{obs}}), \quad (33)
 \end{aligned}$$

where θ_1 and θ_2 are the values of the image positions predicted by the model. At fixed γ and m_s , the above product of Dirac delta functions can be easily eliminated after the following variable change,

$$(\log \theta_{\text{Ein}}, \beta) \rightarrow (\theta_1, \theta_2), \quad (34)$$




Fig. 2. Strong lensing cross-section, defined according to Equation 4, as a function of density slope, for a power-law lens with Einstein radius $\theta_{\text{Ein}} = 1''$. Coloured lines correspond to different source magnitudes, according to the legend. The dotted line indicates the area in the source plane that is lensed into multiple images (infinite for $\gamma > 2$).




Fig. 3. Negative magnification of image 2 as a function of source position, for power-law lenses with $\theta_{\text{Ein}} = 1''$. Coloured lines correspond to different values of the density slope γ , according to the legend. The dashed lines indicate the radial caustic position for the lenses of the corresponding colour.

and subsequent integration over θ_1 and θ_2 . The result of this operation is

$$P(\mathbf{d}|\boldsymbol{\eta}) = A(\boldsymbol{\eta}) \int d\gamma dm_s P(m_1^{\text{obs}}, m_2^{\text{obs}} | \theta_{\text{Ein}}^*, \beta^*, m_s) P(\theta_{\text{Ein}}^*, \gamma | \boldsymbol{\eta}) \times P_{\det}(\theta_{\text{Ein}}^*, \gamma, m_s, \beta^* | S) P(m_s) |\det J|_{(\theta_{\text{Ein}}, \beta) = (\theta_{\text{Ein}}^*, \beta^*)}, \quad (35)$$

where θ_{Ein}^* and β^* are the values of the Einstein radius and source position needed to reproduce the observed image positions, and $\det J$ is the Jacobian determinant of the variable change of Equation 34.

At fixed γ , the integral over m_s is independent of the population parameters $\boldsymbol{\eta}$. I computed this integral for each lens, on a grid of values of γ , with a spline approximation and integration, prior to running the MCMC. Finally, I computed the inte-




Fig. 4. Inference on the Einstein radius and density slope of a subset of lenses, obtained assuming flat priors on all of the lens model parameters. The black bars indicate the 68% credible region. Circles are true values of θ_{Ein} and γ of the lenses, colour-coded by the image asymmetry parameter ξ_{asymm} , defined in Equation 30. The grey dotted lines connect the true values to the corresponding measurement.

gral over γ via spline integration, for each lens at each draw of parameters $\boldsymbol{\eta}$.

4.2. Likelihood computation: full sample case

The calculation of the likelihood terms of Equation 8 takes different avenues, depending on whether the object is a lens or not. In the case of a non-detection, the likelihood of observing the data given the population parameters is obtained by combining Equation 10, Equation 11 and Equation 12. The result is

$$P(\mathbf{d}|\boldsymbol{\eta}) = \int d\gamma d \log \theta_{\text{Ein}} e^{-\lambda} P(\theta_{\text{Ein}}, \gamma | \boldsymbol{\eta}), \quad (36)$$

where λ is the expected number of lensed sources given the galaxy parameters. For a power-law lens, this quantity can be written as

$$\lambda = n_{\text{bkg}} \left(\frac{\theta_{\text{Ein}}}{1''} \right)^2 \int dm_s P(m_s) \sigma_{\text{SL}}(\theta_{\text{Ein}} = 1'', \gamma, m_s | S), \quad (37)$$

where the lensing cross-section in the integrand is evaluated at a unit Einstein radius. In other words, the dependence of λ on the Einstein radius is analytical. To obtain λ as a function both θ_{Ein} and γ , I computed the integral of Equation 36 on a grid of values of γ and then interpolated the result with a spline polynomial. This calculation was done once, before running the MCMC on the population parameters. I then computed the integral of Equation 36 via Monte Carlo integration. Since the non-lenses of the simulation are indistinguishable from one another, this integral is the same for all non-detections. Therefore it only needs to be computed once (but with high precision), for each draw of $\boldsymbol{\eta}$.

In the case of a strong lens, the likelihood is given by Equation 15, which for a power-law lens model becomes

$$P(\mathbf{d}|\boldsymbol{\eta}) = \int d\gamma d \log \theta_{\text{Ein}} dm_s d\beta \frac{\lambda^{N_s} e^{-\lambda}}{N_s!} P(\theta_1^{\text{obs}}, \theta_2^{\text{obs}} | \theta_{\text{Ein}}, \gamma, \beta) \times P(m_1^{\text{obs}}, m_2^{\text{obs}} | \theta_{\text{Ein}}, \gamma, \beta, m_s) P(\theta_{\text{Ein}}, \gamma | \boldsymbol{\eta}) \quad (38)$$

$$\times P(m_s, \beta | \text{Lens}, \theta_{\text{Ein}}, \gamma, S). \quad (39)$$

Using the same method as in section 4.1, I reduced this to a two-dimensional integral over γ and m_s , which I then computed via spline integration.

Out of the 867 strong lenses, 13 have $N_s = 2$. For the purpose of evaluating the Poisson likelihood term of Equation 12, I considered these as double-source lenses. However, I only used image positions and magnitudes for one source, picked randomly. In a real case scenario, all of the available lensed sources would be used for the analysis, as multiple sources allow for more precise constraints on the lens model parameters. That, however, would result in a few lenses having extremely precise inferences of γ , which would then dominate the inference on the population parameters. I chose to only use a single source in order to give equal weight to each lens on the inference, which makes the interpretation of the results easier.

5. Results

Figure 5 shows the posterior probability distribution of the population model parameters η , obtained by fitting the strong lenses only (green filled contours) and the entire sample of galaxies (red solid lines). The median and 68% credible region of the marginal posterior probability of each parameter is reported in Table 2. In both cases, the model is able to recover the true distribution of θ_{Ein} and γ of the galaxy sample (dashed black lines).

Figure 5 also shows the values of the model parameters of the population of lenses (dotted red lines), which I obtained by fitting the model of Equation 28 directly to the true values of θ_{Ein} and γ of the 867 lenses. The average value of γ of the lens population is 2.03. This is the value of μ_γ that one would infer by fitting the lens sample without taking into account selection effects (that is, following the method of Paper I). Thanks to the formalism introduced in this paper, this solution is ruled out at the 3σ level, and the true value $\mu_\gamma = 2.0$ is recovered.

The posterior probability obtained when using the full galaxy sample (red solid lines in Figure 5) is almost identical to the one based purely on the sample of lenses. The full sample brings in additional information: the number of non-detections and the number density of background sources, n_{bkg} . This result, however, shows that this information contributes little to the posterior, which is dominated by the strong lenses, with the exception of the parameter μ_{Ein} .

The statement above is conditional on the type and quality of the data that the model is fitted to. For the analysis presented so far, I assumed that, for each lens, image positions and magnitudes are available. The magnitudes of the two images provide a constraint on the ratio of the magnifications at positions θ_1 and θ_2 , which in turn can be used to constrain the density slope (see O’Riordan et al. 2019, 2020). Often, however, statistical strong lensing analyses do not make use of magnification information (see for example Auger et al. 2010; Bolton et al. 2012; Sonnenfeld et al. 2015, 2019), which usually requires high-quality imaging data to be extracted.

I investigated an alternative scenario, in which the only lensing data available are the image positions and the knowledge of the lens detection probability, P_{det} , as a function of the model parameters. Since the selection criterion of the simulated sample, Equation 26, relies on the magnitude of the second image, it is difficult to imagine a situation in which P_{det} is known but the image magnitudes are not. However, the rationale of this experiment is to simulate a more general scenario in which the selection function is known, but no magnification information is available.

Figure 6 shows the posterior probability distribution obtained by fitting the same model of Section 4 to lensing data, while removing the term related to the observed magnitudes, $P(m_1^{\text{obs}}, m_2^{\text{obs}} | \theta_{\text{Ein}}, \gamma, \beta, m_s)$, from the likelihood. The inference on the parameters describing the γ distribution is much less precise, as a result of the lack of direct constraints on the magnification ratios. When using only the strong lenses, I obtained $\mu_\gamma = 1.97 \pm 0.10$, which is an order of magnitude more uncertain than what the full dataset allows. In this case, the addition of the non-detections reduces this uncertainty by a factor of two: $\mu_\gamma = 2.01 \pm 0.05$.

Interestingly, image position data are still able to constrain μ_γ , to some extent, even when non-detections are not considered. This is because certain image configurations can only be produced by lenses with certain values of γ : for example, in lenses with $\gamma > 2$, the position of image 2 can be arbitrarily close to the lens centre. If $\gamma < 2$, however, there is a lower limit to the value of θ_2 , set by the position of the radial critical curve. This explains the shape of the $\mu_\gamma - \sigma_\gamma$ degeneracy observed in the fit to the lens sample (see the panel in the third column, bottom row of Figure 6): if the average γ is smaller than 2, the scatter in γ must be large, in order to allow for the presence of lenses with $\gamma > 2$ in the sample, which are needed to reproduce highly asymmetric image configurations.

6. Discussion

6.1. Advantages of a complete lens sample

I introduced a formalism for the analysis of strong lensing data from a survey with known selection function, then applied it to a simulated sample of lenses that is complete above a well-defined observational cut. Sample completeness is not a strict requirement in order to apply the method of Section 2: in principle, it is sufficient to know the lens detection probability P_{det} of the survey. However, if a fraction of lenses is missed by the survey, in a way that correlates with some lens property and that is not taken into account by the model for P_{det} , biases will be introduced.

For example, let us suppose that the lens detection efficiency of a survey like the one considered in Section 3 depends not only on the magnitude of image 2, but also on the surface brightness contrast between the lens and the source light. Let us suppose that lenses with a small value of θ_2 are, for this reason, missed by the survey. If this selection is not taken into account, the model might be biased against solutions with large values of γ , which are instead needed to reproduce lenses with highly asymmetric configurations. Having a complete lens sample guarantees against such a possibility. In this example case, this requires ensuring that no lenses are missed, regardless of their image configuration.

6.2. Technical challenges

In order to apply this method to real samples of lenses, several challenges need to be overcome. The most important one is obtaining an accurate description of the lens detection probability, in order to avoid biases like the one in the example above. The experiment that I carried out assumed point-like sources, but virtually all strongly-lensed sources have non-zero size. Lensed galaxies are usually well-resolved, which means that their detection depends on their surface brightness distribution. The source parameter set, ψ_s , will then need to include quantities such as the half-light radius and the Sérsic index, in addition to the magnitude.




Fig. 5. Posterior probability distribution of the galaxy population model parameters, η , obtained by fitting only the strong lenses (green filled contours) and the full sample (red solid contour lines). All of the available lensing data (image positions and magnitudes) were used for the fit. Black dashed lines mark the true values of the parameters. Red dotted lines indicate the values of the parameters obtained by fitting the model of Equation 28 to the subsample of strong lenses (the mean $\log \theta_{\text{Ein}}$ parameter value is $\mu_{\text{Ein}} = -0.12$, which is outside of the plot range).

Table 2. Inference on the model parameters.

Parameter	Truth	Lenses	Prior	Fit to full dataset		Fit to image positions only		Description
				Lenses only	Full sample	Lenses only	Full sample	
μ_{Ein}	-0.30	-0.12	$U(-0.5, 0.5)$	$-0.297^{+0.010}_{-0.010}$	$-0.298^{+0.007}_{-0.008}$	$-0.298^{+0.010}_{-0.011}$	$-0.295^{+0.011}_{-0.011}$	Mean $\log \theta_{\text{Ein}}$
σ_{Ein}	0.20	0.20	$U(0.0, 0.3)$	$0.199^{+0.004}_{-0.004}$	$0.201^{+0.004}_{-0.004}$	$0.200^{+0.004}_{-0.004}$	$0.201^{+0.004}_{-0.004}$	Scatter in $\log \theta_{\text{Ein}}$
μ_{γ}	2.00	2.04	$U(1.8, 2.2)$	$1.997^{+0.009}_{-0.009}$	$2.000^{+0.009}_{-0.009}$	$1.97^{+0.09}_{-0.10}$	$1.97^{+0.04}_{-0.04}$	Mean γ
σ_{γ}	0.20	0.18	$U(0.0, 0.5)$	$0.201^{+0.008}_{-0.008}$	$0.201^{+0.008}_{-0.008}$	$0.20^{+0.06}_{-0.07}$	$0.19^{+0.03}_{-0.03}$	Scatter in γ

Notes. Column (2): true values of the population parameters. Column (3): values of the lens population parameters, obtained by fitting the model of Equation 28 to the sample of detected strong lenses. Column (4): priors on the parameters. Column (5): median, 16th and 84th percentile of the marginal posterior probability distribution of each parameter obtained by fitting the sample of strong lenses with the model of Section 4. Column (6): inference on the full sample (lenses and non-detections). Column (7): inference using image positions only (no magnitudes), lens-only sample. Column (8): inference using image positions only, full sample.

As new parameters are introduced, the model describing the source population, $P(\psi_s)$, must be updated accordingly. While, in principle, ψ_s can be measured on the large population of galaxies that are not strongly lensed, in practice the lens detection probability could depend on features that are below the resolution limit of non-lensed galaxies. For example, the presence of bright clumps can boost the detectability of a lens, but the appearance

of a clumpy galaxy changes when it is strongly lensed, due to the magnification effect (see for example Faure et al. 2021). It is possible, in principle, to generalise the model of Section 4 and infer the source property distribution $P(\psi_s)$ simultaneously with the lens distribution. In practice, this would increase the computational burden of the inference procedure.

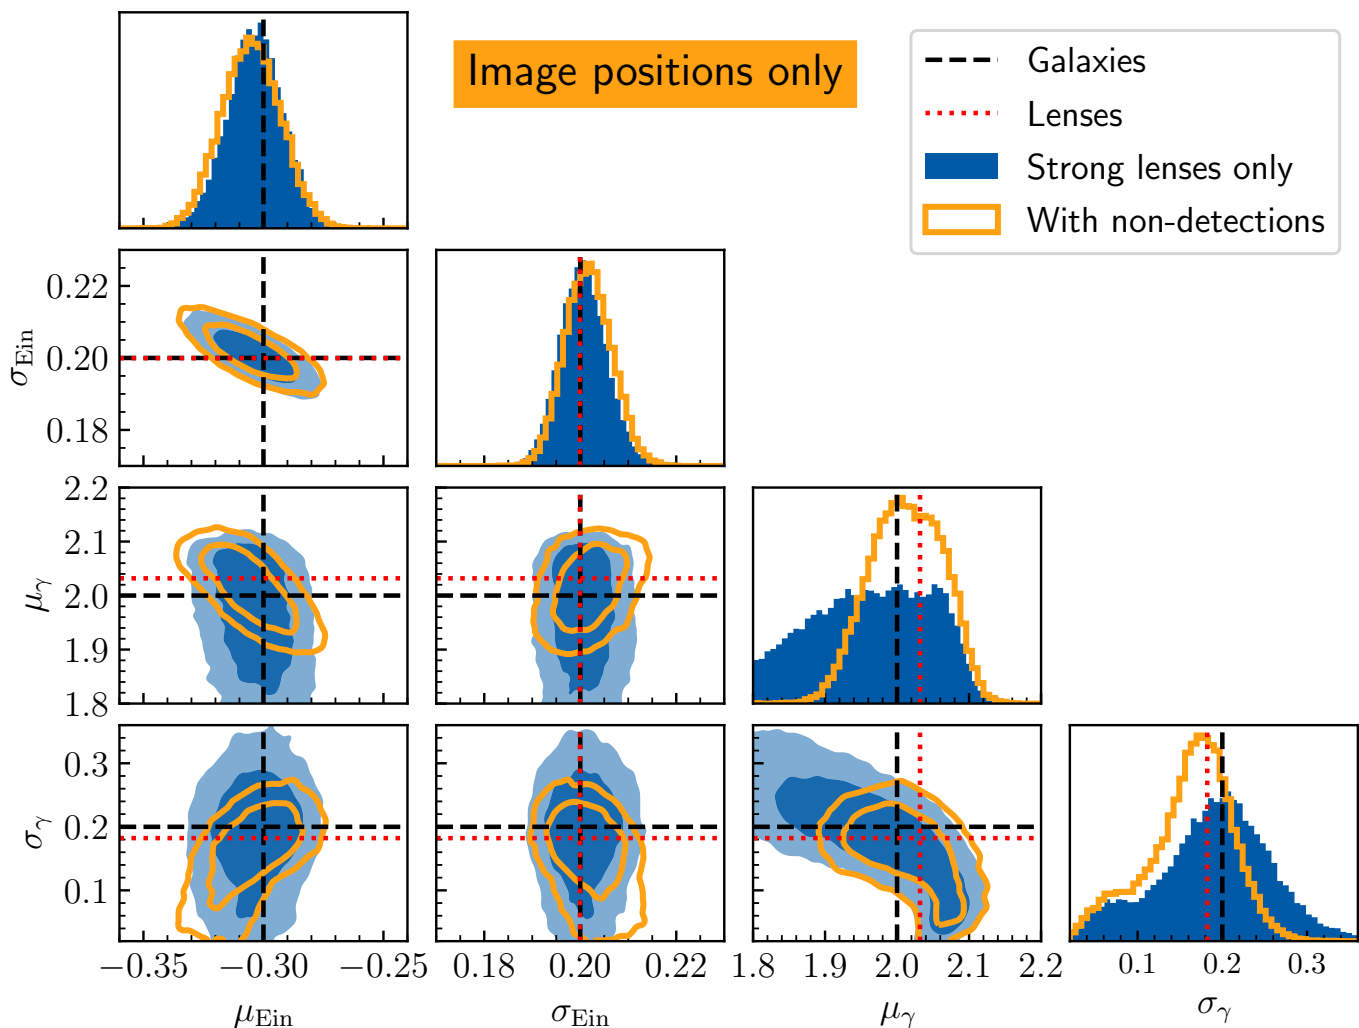


Fig. 6. Same as Figure 5, but obtained by fitting only the image positions (no magnitudes) of the lenses.

In case of a lens sample selected on the basis of the availability of spectroscopic redshifts, then the selection function needs to also take into account the efficiency of obtaining such measurements. On the one hand, this can complicate the selection function even further, as in the case of the SL2S survey discussed in Section 1. On the other hand, if the primary means of the lens search is spectroscopic, meaning that the presence of a lensed source is assessed with the same observational data used to measure its redshift, then the image detection and source redshift measurement aspects of the problem can be merged.

This is the case, for example, of the Sloan ACS Lens Survey (SLACS, Bolton et al. 2006), and the SINFONI Nearby Elliptical Lens Locator Survey (SNELLS, Smith et al. 2015b). Arneson et al. (2012) carried out a study of the selection function of SLACS, which, if extended, could be used to carry out an analysis like the one presented in this paper. SNELLS, and the more recent MNELLS (Collier et al. 2020), are particularly suited for the application of this method, in virtue of their clean selection function.

One class of systems for which it is easier to obtain the selection function, compared to lensed galaxies, is that of strongly lensed quasars. Current lensed quasar searches are already able to produce highly complete lens samples (see for example Agnello & Spiniello 2019). It could be possible, in the near future, to identify a complete set of strongly-lensed quasars: for instance,

by carrying out a systematic search on a well-defined sample of massive galaxies, applying a selection criterion such as that of Equation 26, and obtaining spectroscopic follow-up observations to obtain redshifts and rule out contaminants (as was done by Agnello et al. 2018).

Lensed quasars, however, are subject to microlensing, which introduces a source of scatter in the relation between the intrinsic source magnitude and the observed one. Microlensing depends on the fraction of mass in compact objects (that is, stars) and on the size of the light-emitting region of the quasar, which is the accretion disk. A rigorous analysis would require modelling microlensing, in order to obtain an accurate selection function of lensed quasars. This would also provide an opportunity to further constrain the stellar mass-to-light ratio of the lens galaxies (Vernardos 2019). There is, however, an associated cost in terms of model complexity.

If the lens detection probability is a function of image configuration, whether directly, or, as in the simulation of Section 3, indirectly (highly asymmetric configuration are rarely detected, because of the de-magnification of image 2), then the lens population model, $P(\psi_g)$, needs to be extended. This is because, unlike in the experiment carried out in this work, real lenses are not axisymmetric, and the ellipticity of a lens can change the number, the shape and the magnification of the strongly-lensed images. At the very least, it is necessary to add two degrees of

freedom to the model: the axis ratio and orientation of the mass component. Moreover, lens models often require to account for external shear in order to correctly reproduce the observed images. Galaxy ellipticity and external shear could be incorporated in the hierarchical model used to describe the lens population, but this would make computations more challenging (especially the evaluation of integrals of the kind of Equation 9).

In summary, most practical applications of the method proposed in this paper will likely require an extension of the model considered so far. This will pose important computational challenges, analogous to those discussed in section 6.5 of Paper I.

6.3. Other applications

In addition to providing unbiased descriptions of the structural parameter distribution of galaxies, this method can be used for combining information from different samples of lenses. One important case in which this is needed is the analysis of time-delay lenses for the study of cosmology. The current approach of the leading team in this field, the TDCOSMO⁴ collaboration, is that of using information obtained from samples of strong lenses with no time delays as a prior on a smaller set of time-delay lenses (Birrer et al. 2020). In Paper II (Sonnenfeld 2021), I showed how, with such an approach, it is possible to obtain 1%-level measurements of the Hubble constant.

Equation 1, however, shows that two different lens samples, $P_{SL,1}$ and $P_{SL,2}$, with different source populations, $P_{s,1}$ and $P_{s,2}$, and selection functions $P_{det,1}$, $P_{det,2}$, are biased in different ways with respect to the general population of galaxies P_g . Therefore, using information relative to one lens sample as a prior on the other might lead to biases (the amplitude of which needs to be quantified with dedicated simulations).

The formalism introduced in this work offers a solution to this problem: if the parent population of galaxies, among which strong lenses are searched, is selected in the same way in the two surveys, then P_g is by construction the same distribution for the two samples. It is relatively easy to achieve this: usually, galaxy populations are selected by means of cuts applied at the catalogue level (for example, by imposing that galaxies lie in a given redshift range and are brighter than a luminosity threshold), so it is sufficient to apply the same cuts consistently to the two samples. In that case, it is possible, with the method of Section 2, to infer P_g from one sample and then use it as a prior on the other sample.

Another possible application is the analysis of strong lenses with no spectroscopic measurement of the source redshift. Although for the experiment carried out in this work I assumed that the source redshifts were known, this is not a strict requirement. The main requirement is for the source parameter distribution, $P(\psi_s)$, which includes the source redshift, to be known. If that is the case, then the source redshift can be treated as a nuisance parameter and marginalised over. This is the subject of Paper IV (Sonnenfeld 2022).

Finally, information from the relative number of lenses and non-detections can in principle be used to break the mass-sheet degeneracy. This is because, in general, a mass-sheet transformation (see section 2.1 of Paper II for a definition) changes the strong lensing cross-section of a galaxy. I tried to explore this possibility on the sample generated in Section 3. The change in σ_{SL} with the mass-sheet transformation parameter, however, varies with source magnitude, and any signal is washed out after averaging over the source population.

7. Conclusions

This paper introduced the formalism necessary for inferring the structural properties of a sample of galaxies from the observations of a subset of strong lenses drawn from it. Although the strong lenses are a biased sample, when the properties of the background source population and the selection function of the sample are known, it is possible, with this method, to remove this bias.

There are two variations of the method: one that uses only the strong lenses, and one that combines the lens sample with non-detections. I tested both of these methods on a simulated sample of galaxies that included a complete set of lenses. In both cases, the true properties of the general galaxy population were accurately recovered.

The experiment also revealed that, when the lensing data provides constraints on the relative magnification between the images, in addition to the image positions, then the non-detections add little constraining power. Nevertheless, in such cases, the number of non-detections could be used as a consistency check, to test the model assumptions. However, when magnification constraints are not available, non-detections reduce model uncertainties significantly.

There are several technical challenges to be overcome, in order to apply this method to real samples of lenses. Nevertheless, the formalism developed in this work constitutes the first necessary step in order to use strong lenses to make unbiased statements on the galaxy population.

References

- Agnello, A., Schechter, P. L., Morgan, N. D., et al. 2018, *MNRAS*, 475, 2086
Agnello, A. & Spiniello, C. 2019, *MNRAS*, 489, 2525
Aihara, H., Arimoto, N., Armstrong, R., et al. 2018, *PASJ*, 70, S4
Arneson, R. A., Brownstein, J. R., & Bolton, A. S. 2012, *ApJ*, 753, 4
Auger, M. W., Treu, T., Bolton, A. S., et al. 2010, *ApJ*, 724, 511
Barnabè, M., Spiniello, C., Koopmans, L. V. E., et al. 2013, *MNRAS*, 436, 253
Birrer, S., Shajib, A. J., Galan, A., et al. 2020, *A&A*, 643, A165
Bolton, A. S., Brownstein, J. R., Kochanek, C. S., et al. 2012, *ApJ*, 757, 82
Bolton, A. S., Burles, S., Koopmans, L. V. E., Treu, T., & Moustakas, L. A. 2006, *ApJ*, 638, 703
Calzetti, D., Kinney, A. L., & Storchi-Bergmann, T. 1994, *ApJ*, 429, 582
Collett, T. E. 2015, *ApJ*, 811, 20
Collett, T. E., Oldham, L. J., Smith, R. J., et al. 2018, *Science*, 360, 1342
Collier, W. P., Smith, R. J., & Lucey, J. R. 2020, *MNRAS*, 494, 271
Faure, B., Bournaud, F., Fensch, J., et al. 2021, *MNRAS*, 502, 4641
Gavazzi, R., Marshall, P. J., Treu, T., & Sonnenfeld, A. 2014, *ApJ*, 785, 144
Gavazzi, R., Treu, T., Marshall, P. J., Brault, F., & Ruff, A. 2012, *ApJ*, 761, 170
Kinney, A. L., Calzetti, D., Bohlin, R. C., et al. 1996, *ApJ*, 467, 38
Mandelbaum, R., Miyatake, H., Hamana, T., et al. 2018, *PASJ*, 70, S25
Oguri, M., Rusu, C. E., & Falco, E. E. 2014, *MNRAS*, 439, 2494
Oldham, L. J. & Auger, M. W. 2018, *MNRAS*, 476, 133
O’Riordan, C. M., Warren, S. J., & Mortlock, D. J. 2019, *MNRAS*, 487, 5143
O’Riordan, C. M., Warren, S. J., & Mortlock, D. J. 2020, *MNRAS*, 496, 3424
Parsa, S., Dunlop, J. S., McLure, R. J., & Mortlock, A. 2016, *MNRAS*, 456, 3194
Ruff, A. J., Gavazzi, R., Marshall, P. J., et al. 2011, *ApJ*, 727, 96
Schechter, P. L., Pooley, D., Blackburne, J. A., & Wambsganss, J. 2014, *ApJ*, 793, 96
Schuldt, S., Chirivì, G., Suyu, S. H., et al. 2019, *A&A*, 631, A40
Smith, R. J., Lucey, J. R., & Conroy, C. 2015a, *MNRAS*, 449, 3441
Smith, R. J., Lucey, J. R., & Conroy, C. 2015b, *MNRAS*, 449, 3441
Sonnenfeld, A. 2021, *A&A*, 656, A153, (Paper II)
Sonnenfeld, A. 2022, *A&A*, 659, A133, (Paper IV)
Sonnenfeld, A. & Cautun, M. 2021, *A&A*, 651, A18, (Paper I)
Sonnenfeld, A., Chan, J. H. H., Shu, Y., et al. 2018, *PASJ*, 70, S29
Sonnenfeld, A., Gavazzi, R., Suyu, S. H., Treu, T., & Marshall, P. J. 2013a, *ApJ*, 777, 97
Sonnenfeld, A., Jaelani, A. T., Chan, J., et al. 2019, *A&A*, 630, A71
Sonnenfeld, A., Treu, T., Gavazzi, R., et al. 2012, *ApJ*, 752, 163
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
Sonnenfeld, A., Verma, A., More, A., et al. 2020, *A&A*, 642, A148
Vernardos, G. 2019, *MNRAS*, 483, 5583

⁴ <http://tdcosmo.org>