



HAL
open science

Measuring line-of-sight shear with Einstein rings: a proof of concept

Natalie B. Hogg, Pierre Fleury, Julien Larena, Matteo Martinelli

► **To cite this version:**

Natalie B. Hogg, Pierre Fleury, Julien Larena, Matteo Martinelli. Measuring line-of-sight shear with Einstein rings: a proof of concept. *Monthly Notices of the Royal Astronomical Society*, 2023, 520 (4), pp.5982-6000. 10.1093/mnras/stad512 . hal-03833837

HAL Id: hal-03833837

<https://hal.science/hal-03833837>

Submitted on 19 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Measuring line-of-sight shear with Einstein rings: a proof of concept

Natalie B. Hogg,¹★ Pierre Fleury,¹ Julien Larena,² and Matteo Martinelli^{3,4}

¹Université Paris-Saclay, CNRS, CEA, Institut de physique théorique, 91191, Gif-sur-Yvette, France

²Laboratoire Univers et Particules de Montpellier, Université de Montpellier, CNRS, Montpellier, France

³INAF - Osservatorio Astronomico di Roma, via Frascati 33, 00040 Monteporzio Catone (Roma), Italy

⁴INFN - Sezione di Roma, Piazzale Aldo Moro, 2 - c/o Dipartimento di Fisica, Edificio G. Marconi, I-00185 Roma, Italy

Accepted 2023 February 13. Received 2023 February 13; in original form 2022 October 21

ABSTRACT

Line-of-sight effects in strong gravitational lensing have long been treated as a nuisance. However, it was recently proposed that the line-of-sight shear could be a cosmological observable in its own right, if it is not degenerate with lens model parameters. We firstly demonstrate that the line-of-sight shear can be accurately measured from a simple simulated strong lensing image with percent precision. We then extend our analysis to more complex simulated images and stress test the recovery of the line-of-sight shear when using deficient fitting models, finding that it escapes from degeneracies with lens model parameters, albeit at the expense of the precision. Lastly, we check the validity of the tidal approximation by simulating and fitting an image generated in the presence of many line-of-sight dark matter haloes, finding that an explicit violation of the tidal approximation does not necessarily prevent one from measuring the line-of-sight shear.

Key words: gravitational lensing: strong – cosmology: theory – software: development

1 INTRODUCTION

Strong gravitational lensing is one of the most elegant probes of the late-time Universe. Massive objects such as galaxies and galaxy clusters bend the local spacetime geometry and produce multiple, distorted images of distant background sources. When the source–lens–observer alignment is good, the image forms an Einstein ring. In an otherwise perfectly homogeneous and isotropic Universe, the morphology of that ring is a direct probe of the gravitational field generated by the strong lens.

However, on the scales probed by strong gravitational lensing, the Universe is very inhomogeneous. Therefore, we expect that inhomogeneities – typically dark matter haloes – near the line of sight (LOS) will generate weak lensing distortions that will affect the images produced by the strong lensing process. We generically refer to the effects induced by this weak lensing of strong lensing images as *LOS effects*. These effects have been studied for their impact on, for example, the determination of the Hubble parameter via time-delay cosmography (Gilman et al. 2020), in which case they are considered to be a nuisance effect. But when LOS effects are observable, they also provide some insight into the distribution of dark matter. For example, the effects of individual haloes of dark matter, whether in the main lens or along the LOS, can lead to the detection of dark matter structures on small scales (Vegetti & Koopmans 2009; Vegetti et al. 2012; Despali et al. 2018; Sengül et al. 2022; Nightingale et al. 2022) and the study of their properties (Vegetti & Koopmans 2009; Vegetti et al. 2014; Ostdiek et al. 2022; Zhang et al. 2022). Line-of-sight haloes and main-lens subhaloes lead to different signatures on strong lensing images, and their effects might be disentangled (Dhanasingham et al. 2022), although the detectability of individual

haloes remains an open question (He et al. 2022). Finally, LOS effects in strong lensing may constitute a novel probe of weak lensing in their own right, which is the motivation of the present work.

Line-of-sight perturbations can be modelled in two complementary ways. One option is to add thin perturbers along the LOS in a multi-plane lensing formalism (Blandford & Narayan 1986), a methodology which is particularly suitable for taking into account the presence of several subdominant yet clearly identifiable mass concentrations near the LOS. However, this approach is not suitable for modelling the collective effect of a large number of perturbers, due to an excessive number of extra parameters thereby added to the lens model.

This difficulty can be overcome by encapsulating the collective effects of LOS inhomogeneities into an effective, external convergence and shear perturbing the LOS (Kovner 1987; Bar-Kana 1996; Schneider 1997; Birrer et al. 2017; Fleury et al. 2021b). Although this neglects higher-order effects such as flexion, this parametric method is usually considered good enough on large scales. An analysis of both approaches and how to connect them is presented in Fleury et al. (2021a), where one can also find a discussion of higher-order effects.

The idea that the LOS shear can, in principle, be extracted from strong lensing images in galaxy lensing was proposed by Birrer et al. (2017, 2018) as an alternative probe of cosmic shear. Based on this approach, a first attempt at correlating this LOS shear with standard weak lensing was made by Kuhn et al. (2021). However, the set-up developed in this programme relied on a non-optimal treatment of LOS effects that overlooked some degeneracies between foreground shear and the main lens modelling. It also employed a somewhat optimistic and simple model for the main lens.

In this paper, we re-analyse the possibility of measuring the LOS shear with Einstein rings generated by galaxy lensing using the minimal model introduced by Fleury et al. (2021a), in lieu of the parametrisation used by Birrer et al. (2017). This model fully accounts

★ E-mail: natalie.hogg@ipht.fr

for degeneracies between the main lens and its LOS corrections, thereby providing an effective description that depends on only one additional complex parameter. This parameter, denoted γ_{LOS} , is expected to be mostly independent of the main lens model. In our analysis, we take great care in studying the level of complexity required when modelling the main lens, showing that too simplistic a model can result in biases and a loss of accuracy in the reconstruction of the LOS shear.

We present the first demonstration that the LOS shear *can* be observed using a series of simulated strong lensing images of increasing complexity. In particular, we show that the LOS shear is generally measurable without systematic biases, and that the next-order LOS effects, while present, do not introduce further biases. These results are obtained thanks to our modifications of the `lenstronomy` software¹ (Birrer & Amara 2018; Birrer et al. 2021), to which we contributed a new subpackage, `LineOfSight`, included in the software from version 1.11.0.

The paper is organised as follows: in Section 2 we recapitulate the LOS theory and demonstrate how degeneracies between the LOS shears and other parameters are avoided in the so-called minimal model; in Section 3, we show how the LOS shear is systematically measurable; in Section 4 we explore the validity of the tidal regime. In Section 5 we present our conclusions. In Appendix A we describe our modifications to the strong lensing software `lenstronomy`, while Appendix B contains details of the parameter choices used to create the mock image data analysed in this work.

Throughout this work, we employ the spatially flat Λ CDM cosmological model described by the *Planck* 2018 data release (Aghanim et al. 2020), with $H_0 = 67.4 \text{ kms}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.315$.

2 THE MINIMAL LINE-OF-SIGHT MODEL AND ITS ADVANTAGES

In this section, we review the theoretical underpinnings of LOS effects as presented in Fleury et al. (2021a) and demonstrate how the inherent degeneracies can be absorbed in the so-called *minimal model*, paving the way for the LOS shear to be measured.

2.1 Tidal perturbations to a strong lens

As a ray of light passes through the inhomogeneous Universe from source to observer, its path is perturbed by the presence of massive objects. In the simplest strong lensing situation, the deflection of light is assumed to be caused by a single lens. In this case, the whole problem is encapsulated in the lens equation,

$$\beta = \theta - \alpha(\theta), \quad (1)$$

where β is the angular position of the source, θ is the angular position of the observed image and $\alpha(\theta)$ is the displacement angle, which is in turn determined by the gravitational potential associated with the surface density of the lens, projected along the LOS.

A more elaborate description can be made by including the effects of other nearby objects on the image. The most general approach to deal with this situation is the multi-plane lensing set-up (Blandford & Narayan 1986), in which each additional lens mass is considered to reside in its own plane along the LOS. However, the recursive nature of the multi-plane lens equation makes analytical treatments of complicated lensing situations intractable.




Figure 1. When LOS perturbations to a strong lens can be treated in the tidal regime, they intervene in the lens equation (2) via three distortion matrices: \mathcal{A}_{od} , \mathcal{A}_{os} , \mathcal{A}_{ds} . For each couple ab , \mathcal{A}_{ab} represents the distortions to a small source placed at (b) as observed from (a) due to the perturbers only.

A simpler approach is to treat the inhomogeneities along the LOS as tidal perturbations which induce a shear and convergence on the image produced by the main lens. This has the advantage of a far simpler theoretical description than the multi-plane case (Kovner 1987; Seitz & Schneider 1994; Bar-Kana 1996; Schneider 1997; Fleury et al. 2021b); see McCully et al. (2014); Fleury et al. (2021a) for explicit derivations of the perturbed lens equation in the tidal regime, starting from the general multi-plane framework.

In this simplified description, the lens equation (1) becomes

$$\beta = \mathcal{A}_{\text{os}}\theta - \mathcal{A}_{\text{ds}}\alpha(\mathcal{A}_{\text{od}}\theta), \quad (2)$$

where the amplification matrices \mathcal{A}_{os} , \mathcal{A}_{ds} , \mathcal{A}_{od} encapsulate all the LOS perturbations. They are commonly parameterised as

$$\mathcal{A}_{ab} = \mathbf{1} - \Gamma_{ab}, \quad \Gamma_{ab} = \begin{bmatrix} \kappa_{ab} + \text{Re}(\gamma_{ab}) & \text{Im}(\gamma_{ab}) - \omega_{ab} \\ \text{Im}(\gamma_{ab}) + \omega_{ab} & \kappa_{ab} - \text{Re}(\gamma_{ab}) \end{bmatrix}, \quad (3)$$

for $ab \in \{\text{os}, \text{ds}, \text{od}\}$. The quantities κ_{ab} , γ_{ab} , ω_{ab} represent the convergence, the complex shear, and the rotation of an image as seen from the position (a) with respect to its source at the position (b), due to the inhomogeneities between (a) and (b) – see Fig. 1. For example, if a circular source were placed in the source plane (s), and observed from the main-lens plane (d), then γ_{ds} would be the complex ellipticity of its image due to the perturbations between (d) and (s). When all the perturbations are set to zero, $\Gamma_{\text{os}} = \Gamma_{\text{ds}} = \Gamma_{\text{od}} = \mathbf{0}$, we recover the standard single-plane lens equation (1).

It is expected that the tidal regime, and hence Equation (2), is valid when all the perturbers besides the main lens are either a smoothly distributed mass on the LOS, which would induce convergence, or compact masses lying far from the LOS, which would induce shear; rotation is caused by successive shears along different axes. Throughout this paper, we shall assume that κ , γ , ω are all very small, although Equation (2) would remain valid even for large values of those parameters. This description of a perturbed strong lensing system is simple enough to be easily implemented in and rigorously tested by the current state-of-the-art strong lensing software, as we will describe in Section 2 and Section 3. It is, in addition, easily extended beyond the tidal regime following the dominant-lens approach of Fleury et al. (2021a), as we will touch on in Section 4.

¹ <https://github.com/lenstronomy/lenstronomy>.

2.2 Absorbing degeneracies in a minimal model

In this work, we aim to show that the LOS shear can be measured from strong lensing images. However, as shown by Fleury et al. (2021a), the three shears involved in Equation (2), $\gamma_{os}, \gamma_{ds}, \gamma_{od}$, cannot be measured independently, because they are degenerate with the main-lens properties and with one another. This degeneracy is a generalisation of the well-known mass-sheet degeneracy (Falco et al. 1985), and is called the source-position transformation (Schneider & Sluse 2014), since it stems from our ignorance of the exact position and morphology of the source that is lensed. For example, one may multiply Equation (2) with any matrix \mathcal{A} and redefine the unknown source position as $\beta \rightarrow \mathcal{A}\beta$; the resulting lens equation is formally identical to Equation (2), but with $\mathcal{A}_{os} \rightarrow \mathcal{A}\mathcal{A}_{os}$ and $\mathcal{A}_{ds} \rightarrow \mathcal{A}\mathcal{A}_{ds}$.

Degeneracies may be removed by arbitrarily fixing a definition for the unknown source β . This is analogous to lifting a gauge freedom by fixing the gauge. As shown in Fleury et al. (2021a), there exists a particularly meaningful choice, leading to a *minimal model* for the perturbed lens equation. The first step consists in rewriting the displacement angle as the gradient of a potential,

$$\alpha(\theta) = \frac{d\psi}{d\theta}, \quad (4)$$

where

$$\psi(\theta) \equiv \frac{D_{ds}}{D_{od}D_{os}} \hat{\psi}(D_{od}\theta), \quad (5)$$

and $\hat{\psi}(\mathbf{x})$ is twice the projected gravitational potential produced by the surface density of the main lens, $\Sigma(\mathbf{x})$,

$$\hat{\psi}(\mathbf{x}) \equiv \frac{4G}{c^2} \int d^2\mathbf{y} \Sigma(\mathbf{y}) \ln |\mathbf{x} - \mathbf{y}|, \quad (6)$$

where G is Newton's constant and c the speed of light.

Then, multiplying Equation (2) with the combination $\mathcal{A}_{od}\mathcal{A}_{ds}^{-1}$, we obtain the *minimal model* for the lens equation,

$$\tilde{\beta} = \mathcal{A}_{LOS}\theta - \frac{d\psi_{\text{eff}}}{d\theta}, \quad (7)$$

with the transformed source position $\tilde{\beta} \equiv \mathcal{A}_{od}\mathcal{A}_{ds}^{-1}\beta$, and

$$\mathcal{A}_{LOS} \equiv \mathcal{A}_{od}\mathcal{A}_{ds}^{-1}\mathcal{A}_{os}, \quad (8)$$

$$\psi_{\text{eff}}(\theta) \equiv \psi(\mathcal{A}_{od}\theta). \quad (9)$$

Equation (7) effectively describes a main lens with potential ψ_{eff} and external tidal perturbations, \mathcal{A}_{LOS} , located in the same plane.

Since Equations (2) and (7) are equivalent up to a source position transformation, the corresponding lens models would succeed or fail equally when fitting a strong lensing image. Therefore, any two distinct situations with the same \mathcal{A}_{LOS} and ψ_{eff} are observationally indistinguishable. This is why the (od), (os) and (ds) perturbations cannot be independently measured. However, when the perturbations are small, $|\Gamma_{od}|, |\Gamma_{ds}|, |\Gamma_{os}| \ll 1$, $\Gamma_{LOS} = \mathbf{1} - \mathcal{A}_{LOS}$ is linearised as

$$\Gamma_{LOS} \approx \Gamma_{od} + \Gamma_{os} - \Gamma_{ds}. \quad (10)$$

This means that the potentially measurable shear component is

$$\gamma_{LOS} \approx \gamma_{od} + \gamma_{os} - \gamma_{ds}. \quad (11)$$

We shall see in Section 2.3 that γ_{LOS} can indeed be measured from a mock strong lensing image with a much higher precision than any one of γ_{od}, γ_{os} or γ_{ds} taken independently.

Another degeneracy revealed by the minimal model (7) is between the foreground perturbations (od) and the properties of the main lens, through the effective potential ψ_{eff} of Equation (9). Specifically, the $\psi_{\text{eff}} = \text{cst}$ contours are the images of the $\psi = \text{cst}$ contours by the

foreground perturbations encoded in \mathcal{A}_{od} . This generally makes it hard to distinguish between, for example, the effect of the ellipticity of the main lens and a foreground shear γ_{od} . Note, however, that in many widely used elliptical lens models, such as the singular isothermal ellipsoid (SIE, Kormann et al. 1994), or its generalisation, the elliptical power law (EPL, Tessore & Metcalf 2015), the foreground shear–ellipticity degeneracy is only approximate, because in those models the lens ellipticity is implemented at the level of the iso-density contours, not in the iso-potential contours. That is why, in their pioneering numerical experiments, Birrer et al. (2017) could measure both γ_{od} and $\gamma_{os} - \gamma_{ds} = \gamma_{LOS} - \gamma_{od}$. Since this only applies to very specific lens models that are not necessarily well motivated from the astrophysical point of view, such an operation should nevertheless be considered very risky; we thus recommend not to try measuring γ_{od} in general, unless there is a compelling reason to trust the result.

We have chosen to focus the above discussion on the shear. The other components of the tidal matrix Γ_{LOS} are the convergence κ_{LOS} and the rotation ω_{LOS} . It is well known that the convergence and source position can be transformed in such a way that the observed lensing image is left invariant – this is precisely the mass-sheet degeneracy – which means that any attempt to measure or constrain the LOS convergence from a strong lensing image would be a pointless undertaking. For this reason, we shall keep the convergence parameters fixed to zero throughout this paper, both when generating mock images and when fitting them.

The rotation is induced by lens–lens coupling, and is generally small when the secondary lenses are individually weak, $\omega = \mathcal{O}(\gamma^2)$. However, it can play a significant role when the shear on the image is large – say, 10% or more. Shears this large would be unusual in real strong lensing data so when generating our mock images we restrict ourselves to smaller values, and keep the rotation parameters fixed to zero. Nevertheless, when fitting images with the minimal model, we include ω_{LOS} as a free (and in this case nuisance) parameter. Its presence accounts for any slight non-linear effect that may be induced by larger shears, and thus ensures a better recovery of γ_{LOS} from a given mock image.

2.3 A simple example with lenstronomy

To demonstrate the advantage of the minimal model over the full model with a concrete example, we simulate a strong lensing image using the software `lenstronomy`, which we modified to include the LOS formalism described above. Our modifications are now part of the public version of this code. For full details of our modifications of `lenstronomy`, see Appendix A.

We create a simple lens using an EPL profile, with additional shears $\gamma_{od}, \gamma_{os}, \gamma_{ds}$. The ellipticity of the profile, along with its centre and the LOS shear components, are all drawn randomly. We model the source and the lens light with elliptical Sérsic light profiles. The ellipticity of the source is also generated at random. Throughout this work we simulate the noise on our mock images according to the Hubble Space Telescope Wide Field Camera 3 F160W noise settings in `lenstronomy` (Windhorst et al. 2011). The resulting simulated image is shown in Fig. 2 and the specific model parameters associated with this image are listed in Table 1.

Using the affine-invariant ensemble sampling Markov chain Monte Carlo (MCMC) parameter inference method available in the `emcee` package² (Foreman-Mackey et al. 2013), we then fit the simulated image shown in Fig. 2 with two different models: a model comprised

² <https://github.com/dm/emcee>.




Figure 2. The image produced by a simple lens model consisting of an elliptical power law mass profile and a elliptical Sérsic lens light profile, with an elliptical Sérsic source (model parameters listed in Table 1).

of an EPL plus the full LOS model (2), i.e. the same model as was used to generate the mock image; and a model comprised of an EPL plus the minimal model (7). We sample all the parameters listed in Table 1 except for the lens and lens light central positions, which are kept fixed at the centre of the image. We impose uniform priors on all the parameters and run chains containing ten walkers per parameter. As discussed above, we expect that the individual shear components in the full model will be degenerate with each other, while the γ_{LOS} combination of the minimal model will be accurately recovered. In both cases (full and minimal model), we expect the foreground shear γ_{od} to be degenerate with the ellipticity of the main lens.

In Fig. 3, we show the one- and two-dimensional marginalised posterior distributions for the (os) and (ds) shears (left panel) and (od) shear and lens ellipticity (right panel) which result from fitting the simple simulated image with the full LOS model. Note that all the contour plots presented in this paper were produced using the `chainconsumer` package³ (Hinton 2016).

From these plots, we can see that while the shear parameters are correctly recovered at the level of 1σ , there are very strong degeneracies between the (os) and (ds) shears. This results in large uncertainties for γ_{os} and γ_{ds} , on the order of 10 %, i.e. comparable or larger than the expected signal. Furthermore, the expected degeneracy between the (od) shear and the lens ellipticity is clearly present. As discussed in Section 2.2, the simplicity of the EPL model artificially alleviates that degeneracy, giving the impression that the main lens’s ellipticity can be measured independently of the foreground shear.

In Fig. 4, we show the one- and two-dimensional marginalised posterior distributions for the (od) shear and lens ellipticity (left panel) and LOS shear and lens ellipticity (right panel) which result from fitting the simple simulated image with the minimal model. From these plots, we can see that, as expected, the minimal LOS shear γ_{LOS} evades any of the previous degeneracies. The input value is perfectly recovered with, in this simple example, a remarkable precision on the order of a few parts in 10^4 .

The results presented in this section demonstrate the clear advantage of using the minimal model over the full model. This advantage can

Table 1. The parameters used to simulate the image shown in Fig. 2, along with the priors on the parameters sampled in the MCMC to produce the results shown in Fig. 3 and Fig. 4. The ellipticities, positions and shear parameters were all drawn at random. The lens light position and ellipticity are fixed to be the same as those of the EPL profile. The magnitudes quoted are apparent magnitudes. The expected values of the minimal LOS quantities γ_{LOS} , ω_{LOS} are derived from Equation (8) using the decomposition of Equation (3).

Component	Parameter	Value	Prior
Lens (EPL)	θ_{E}	1.2''	[0.01'', 10'']
	γ	2.6	[0, 4.0]
	e_1	−0.053	[−0.5, 0.5]
	e_2	0.077	[−0.5, 0.5]
	x	0.014''	–
	y	−0.048''	–
Lens light (Sérsic ellipse)	Magnitude	20	–
	$R_{\text{Sérsic}}$	0.5''	[0.001, 10.0]
	$n_{\text{Sérsic}}$	4.0	[0.5, 5.0]
Source light (Sérsic ellipse)	Magnitude	24	–
	$R_{\text{Sérsic}}$	0.3''	[0'', 10'']
	$n_{\text{Sérsic}}$	1.0	[0.5, 5]
	e_1	0.104	[−0.5, 0.5]
	e_2	−0.038	[−0.5, 0.5]
	x	−0.003''	[−0.1'', 0.1'']
LOS	y	0.048''	[−0.1'', 0.1'']
	γ_1^{od}	−0.0027	[−0.5, 0.5]
	γ_2^{od}	−0.0060	[−0.5, 0.5]
	γ_1^{os}	−0.0091	[−0.5, 0.5]
	γ_2^{os}	−0.028	[−0.5, 0.5]
	γ_1^{ds}	0.019	[−0.5, 0.5]
Derived	γ_2^{ds}	0.018	[−0.5, 0.5]
	γ_1^{LOS}	−0.031	[−0.5, 0.5]
	γ_2^{LOS}	−0.052	[−0.5, 0.5]
	ω_{LOS}	0.00033	[−0.5, 0.5]

be understood as being due to a more efficient parameterisation of the problem at hand. It is comparable to the use of $S_8 \propto \sigma_8 \sqrt{\Omega_m}$ in standard cosmic shear, which is better constrained than σ_8 (see Hall 2021 for a thorough explanation as to why); or to the use of $\omega = \Omega h^2$ in cosmic microwave background constraints.

³ <https://github.com/Samreay/ChainConsumer>.




Figure 3. Inference with the full LOS model (2), from fitting the simulated image in Fig. 2. *Left panel:* one- and two-dimensional marginalised posterior distributions of the (os) and (ds) shears. *Right panel:* Same for the (od) shear and lens ellipticity. The dashed lines show the input value of each parameter.

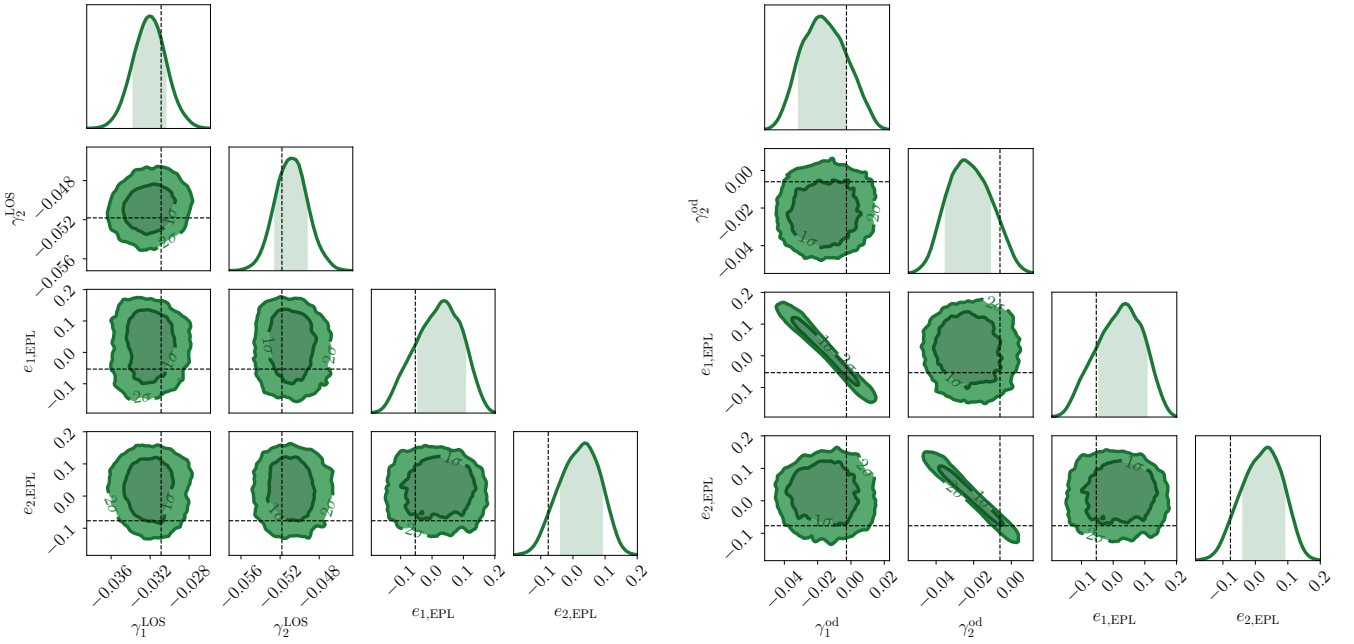


Figure 4. Inference with the minimal model (7), from fitting the simulated image in Fig. 2 with the minimal model. *Left panel:* one- and two-dimensional marginalised posterior distributions of the effective LOS shear and lens ellipticity. *Right panel:* Same for the (od) shear and lens ellipticity. The dashed lines show the input value of each parameter.

3 LINE-OF-SIGHT SHEAR MEASUREMENTS FROM REALISTIC MOCK IMAGES

In this section, we show how the LOS shear is systematically measurable from a larger set of more complicated images than the simple example shown in the previous section. We also assess how the measurement of the LOS shear is affected if the images are fit with

lens models that are slightly to drastically different to those used to create the images.

3.1 Methodology

We generate a catalogue of 64 mock images, depicted in Fig. 6, following a protocol which we will describe in Section 3.1.1. We




Figure 5. Schematic of our composite lens model on the celestial sphere. The baryonic component (orange) is described by an elliptical Sérsic model whose centre is placed at $(0, 0)$. The dark-matter halo (blue) has an elliptical NFW profile and its centre (x_h, y_h) is offset with respect to the baryons. This composite main lens is then perturbed by the LOS shears γ_{od} , γ_{os} , γ_{ds} . The source (green) is a perturbed elliptical Sérsic placed at (x_s, y_s) .

then fit those images with four different models which we present in Section 3.1.2, focusing on the inference of the shear γ_{LOS} . The precision and accuracy of the recovery γ_{LOS} for each model is assessed using statistical tools we will describe in Section 3.1.3.

3.1.1 Complex mock images from a composite lens model

We increase the complexity of the lens modelling in a number of ways compared to Section 2, so as to mimic real strong lensing images. We now model the lens with two constituent parts: a baryonic component modelled with an elliptical Sérsic profile, and a dark matter halo modelled with an elliptical Navarro–Frenk–White (NFW) profile (Navarro et al. 1996). Note that in both profiles, the ellipticity is implemented at the level of the *potential* rather than at the level of the convergence. This means that their iso-potential contours are elliptical, but not their iso-density contours. The halo is offset from the centre of the baryons and the ellipticities of each component are not constrained to be of the same magnitude or direction. The set-up is depicted in Fig. 5.

We model the lens and source light with elliptical Sérsic profiles. Furthermore, we perturb the source light by adding three additional elliptical Sérsic profiles on top of the original, each with an amplitude of 1% of the main source. Such perturbations will not be included in the fitting model, and thus aim to account for our general ignorance of the details of the source light. Finally, we shear each image using the full LOS shear model.

We simulate 64 such systems, each with every model parameter drawn at random from physically motivated distributions. The full parameter ranges and distributions are described in detail in Appendix B. The resulting simulated images are shown in Fig. 6, where their broad diversity in morphology can be clearly seen.

We quantify the quality of an image with a notion of cumulated signal-to-noise ratio. Specifically, for an image i composed of pixels p , we define

$$Q_i \equiv \sum_{p \in i} \text{SNR}(p) \Theta[\text{SNR}(p) - 1], \quad (12)$$

where $\text{SNR}(p)$ denotes the signal-to-noise ratio of pixel p . The lens

light is not included in the signal because it is unrelated to the quality of the lensed image, and the noise accounts for both the background and shot noise. The Heaviside function Θ in the sum of Equation (12) implies that only the pixels with a signal-to-noise ratio larger than one are accounted for in the computation of the quality Q_i . The advantage of this definition of quality is that it accounts for both the traditional notion of SNR (the brighter, the better) and for the number of pixels that are covered by the image (the larger and more resolved, the better).

3.1.2 Four fitting models

We employ four models to fit the simulated images in our catalogue, each with a different level of complexity compared to the mocks. We describe them below; see Table 2 for a summary of the parameters included in or excluded from each model.

(i) The *comprehensive model* is identical to the set-up used to generate the mock images, except that LOS effects are described by the minimal model (7) with parameters γ_{od} , γ_{LOS} , ω_{LOS} ; the rotation term ω_{LOS} is a nuisance parameter included to absorb the small non-linear effects arising from shear–shear couplings. As mentioned earlier, perturbations to the source light are not included.

(ii) The *no foreground shear model* is identical to the comprehensive model, except that the foreground shear γ_{od} involved in the effective potential of Equation (9) is set to zero, thereby implying $\psi_{\text{eff}} = \psi$ here. The motivation for this model is twofold. On the one hand, the presence of γ_{od} in the minimal model is precisely what makes it different from a more traditional model comprising a composite lens plus external shear, such as those used for some of the HOLiCOW lenses (Wong et al. 2020). It is thus essential to evaluate the role of this new parameter. On the other hand, we expect a degeneracy to exist between the ellipticity of the main lens and the foreground shear. If this degeneracy were perfect, then only two of the three complex parameters γ_{od} , $e^{\text{Sérsic}}$, e^{NFW} would suffice to describe an image, allowing us to set one of them to zero. The performance of the no foreground shear model will thus assess the validity of this degeneracy in a composite lens scenario.

(iii) In the *aligned halo model*, the centres of the baryon and dark matter components are both fixed to $(0, 0)$, thus removing the possibility for the halo to be offset from the baryons. The purpose of this model is to check if neglecting this offset leads to systematic biases in the inference of the other parameters, as was noticed by Gomer & Williams (2021). The aligned model is otherwise identical to the comprehensive one.

(iv) In the *power law model*, finally, the entire main lens is modelled using a single EPL profile, and with no additional NFW halo, as we used in our simple example in Section 2. This model aims to quantify the impact of neglecting the composite character of a lens on the recovery of the LOS shear.

3.1.3 Evaluating the precision and accuracy of the models

We now present the quantitative criteria that we use to assess the precision and accuracy of the inference of γ_{LOS} in each model. Firstly, we note that when fitting the images in our mock catalogue, the MCMC parameter inference does not always succeed in obtaining either an upper or lower bound on one or both components of γ_{LOS} , particularly in the less informative models. We therefore discard the cases where one or both bounds were not obtained from the quantitative comparison of the models that will follow. We refer to




Figure 6. Our catalogue of 64 simulated images produced according to the protocol presented in Section 3.1.1.

the remaining images in the mock catalogue as the *productive images*, \mathcal{P} , denoting their number as $N_{\mathcal{P}}$.

Furthermore, for an image $i \in \mathcal{P}$ with a given model, the posterior distribution of γ_{LOS} is generally non-Gaussian; in particular, the 68%-confidence domain is not always symmetric, which translates into asymmetric error bars. We call $\sigma_{a,i}^{\pm}$ the corresponding upper and lower uncertainties at 68%-confidence level on the component $a \in \{1, 2\}$ of γ_{LOS} for the i^{th} image. This implies that for each image and model, the uncertainty on the LOS shear is a priori described by four numbers, which is somewhat cumbersome. For further simplicity, we summarise them all by the geometric mean of the symmetrised

uncertainty of each component of the LOS shear:

$$\sigma_i \equiv \sqrt{\left(\frac{\sigma_{1,i}^+ + \sigma_{1,i}^-}{2}\right) \left(\frac{\sigma_{2,i}^+ + \sigma_{2,i}^-}{2}\right)}. \quad (13)$$

The general *precision* of a model is then quantified from the mean uncertainty on the LOS shear over all its productive images,

$$\bar{\sigma} \equiv \frac{1}{N_{\mathcal{P}}} \sum_{i \in \mathcal{P}} \sigma_i. \quad (14)$$

The *accuracy* of a model, i.e. its ability to recover the true value of

Table 2. Summary of the properties and results of the four models used to fit the mock images of Fig. 6.

	Comprehensive model	No foreground shear model	Aligned model	Power law model
Features				
Foreground shear	✓	✗	✓	✓
LOS shear	✓	✓	✓	✓
Baryon ellipticity	✓	✓	✓	–
Halo ellipticity	✓	✓	✓	–
Baryon–halo offset	✓	✓	✗	–
Results				
Productive images: $N_{\mathcal{P}}$	63 (98 %)	62 (97 %)	56 (88 %)	28 (44 %)
Precision: $\bar{\sigma}$	0.011	0.011	0.010	0.010
Accuracy: χ^2	1.0	3.4	62.9	92.2
Outliers: $N_{>2\sigma}$	0	2	36	20

γ_{LOS} , may be evaluated with a quantity inspired by the reduced χ^2 , accounting for the asymmetry of the uncertainties. Namely, for each component $a \in \{1, 2\}$ of the LOS shear, we define

$$\chi_a^2 \equiv \frac{1}{N_{\mathcal{P}}} \sum_{i \in \mathcal{P}} \frac{1}{\sigma_{a,i}^2} \left(\gamma_{a,i}^{\text{LOS,out}} - \gamma_{a,i}^{\text{LOS,in}} \right)^2 \quad (15)$$

with

$$\sigma_{a,i} \equiv \begin{cases} \sigma_{a,i}^+ & \text{if } \gamma_{a,i}^{\text{LOS,out}} \leq \gamma_{a,i}^{\text{LOS,in}} \\ \sigma_{a,i}^- & \text{if } \gamma_{a,i}^{\text{LOS,out}} > \gamma_{a,i}^{\text{LOS,in}} \end{cases} \quad (16)$$

where $\gamma_{a,i}^{\text{LOS,in}}$ and $\gamma_{a,i}^{\text{LOS,out}}$ are respectively the true (input) value and the best fit (output) value of the component a of the LOS shear of the i^{th} image. Again, for further simplicity and to avoid describing the accuracy of a model with two numbers χ_1^2, χ_2^2 , we shall summarise them by their arithmetic mean,

$$\chi^2 \equiv \frac{1}{2} \left(\chi_1^2 + \chi_2^2 \right). \quad (17)$$

Lastly, we can examine the number $N_{>n\sigma}$ of outliers, i.e. the images for which the best fit output for γ_{LOS} lies further than a certain number n of σ from the true input value. Specifically, we consider that an image i is an $n\sigma$ outlier if

$$\left| \gamma_i^{\text{LOS,out}} - \gamma_i^{\text{LOS,in}} \right|^2 > n^2 \left(\sigma_{1,i}^2 + \sigma_{2,i}^2 \right), \quad (18)$$

where $\sigma_{a,i}$ was defined in Equation (16). This definition thus accounts for the asymmetry of the error bars.

3.2 Results

3.2.1 Examining the fits of a single image

We firstly present the results of the fits to a single image as a representative example of the results in the rest of this section. In Fig. 7 we show the 44th image in our mock catalogue, along with the one- and two-dimensional marginalised posterior distributions for γ_1^{LOS} and γ_2^{LOS} obtained when fitting with all four models described above. We chose this particular image to focus on as it is of roughly median quality and the shear parameters were well-sampled in all of the models used to fit it.

From the contour plot, we can see that in both the comprehensive model and the no foreground shear model, the LOS shear is well recovered. On the other hand, in both the aligned model and in the power law model, the posteriors are fairly biased away from the expected values. This indicates that the less informative the model is, the worse the recovery of the shear is. However, the precision of the recovered shears is roughly the same in all four models, implying that the error bar on γ_{LOS} is insensitive to the model used. This is in line with the results we presented in Section 2, where we showed how γ_{LOS} is not degenerate with lens model parameters. We will explore this point further in the following.

We now show the result of fitting our entire simulated image catalogue with the four models described above. In Fig. 8, we show the input LOS shear components as computed from the three individual shears of the full model versus the output LOS shear components as recovered from the images by the MCMC parameter inference. The colour bars show the log of the quality criterion Q . Darker colours thus represent larger values of Q and hence larger SNR values.

Only the results from the productive images in each model are shown in this figure.

3.2.2 Results of the comprehensive model

In the top two panels of Fig. 8, we show the result of fitting all the images in our catalogue with the comprehensive model. The number of productive images in this case is $N_{\mathcal{P}} = 63$, that is 98 % of the complete sample. We can see that in this model, the LOS shear is always very well recovered from the images.

The mean absolute precision of the recovered shear is $\bar{\sigma} = 0.01$ and does not show any correlation with the amplitude of γ_{LOS} . In comparison, the precision of the inference of γ_{LOS} found in Section 2 when demonstrating the advantage of the minimal model was about two orders of magnitude better. The better precision obtained in that case can be attributed to the relative simplicity of the lens and source model being used.

Our measure for the accuracy of the LOS shear recovery yields $\chi^2 = 1.0$, which shows that uncertainties are well estimated for both components. As for the number of outliers, we find $N_{>2\sigma} = 0$ points away from the true value by more than 2σ . This is well within the expected proportion of 2σ outliers in the Gaussian case, 5 %.




Figure 7. *Left panel:* the 44th image in our catalogue. *Right panel:* the one- and two-dimensional marginalised posterior distributions for γ_1^{LOS} and γ_2^{LOS} obtained by fitting this image with the power law model (blue), aligned halo model (green), no foreground shear model (magenta) and comprehensive model (orange). The dashed lines represent the expected values.

Furthermore, the image-quality criterion Q defined in Equation (12) is well correlated with the size of the error bars on each point, with the most precise constraints generally coming from the best images (the darkest coloured points).

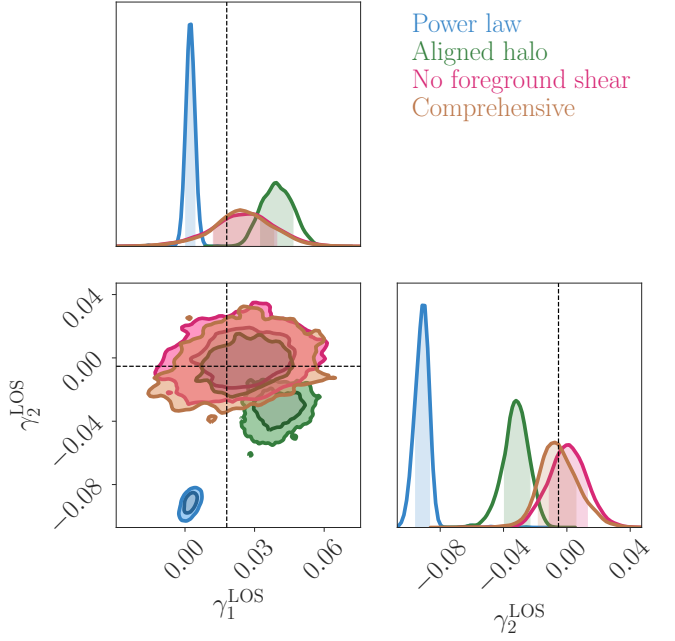
3.2.3 Results of the model with no foreground shear

Next, we fit our catalogue of mock images with the model containing no foreground shear, $\gamma_{\text{od}} = 0$. The input–output plot for this case is shown in the second two panels of Fig. 8, with the points plotted in magenta. The number of productive images in this case is one fewer than in the comprehensive model, $N_{\mathcal{P}} = 62$, that is 97% of the catalogue.

With the exception of two outliers (which are not visible due to the scale of the plot), the two LOS shear components are generally well recovered in this case too. The precision measure is unchanged compared to the comprehensive model, with $\bar{\sigma} = 0.01$. The accuracy is slightly reduced overall, $\chi^2 = 3.4$, while the number of outliers is increased, with $N_{>2\sigma} = 2$ (3% of the productive images), though still in line with the expectation for a normal distribution. This implies that the inclusion of foreground shear in the comprehensive model is not strictly necessary for a better fit to an image. We will discuss this point further below.

3.2.4 Results of the aligned model

Thirdly, we fit our catalogue of images with the aligned halo model. The input–output plot for this case is shown in the third two panels of Fig. 8, with the points plotted in green. We now see how worsening the model can degrade the fit; there is a larger number of outliers and unconstrained points, leading to the number of productive images shrinking to $N_{\mathcal{P}} = 56$, that is 88% of the catalogue. While the precision of the recovery of the two components of γ_{LOS} is unaffected,



$\bar{\sigma} = 0.01$, the accuracy is now very bad, $\chi^2 = 62.9$, with more than two thirds of the productive images being outliers, $N_{>2\sigma} = 36$.

3.2.5 Results of the power law model

Lastly, we fit our catalogue of images with the power law model. The input–output plot for this case is shown in the bottom two panels of Fig. 8, with the points plotted in blue. Of the four models used to fit the simulated images, this is clearly the worst, with a large number of unconstrained points that result from the difficulty the MCMC has when trying to fit complex images with an overly simplistic model: $N_{\mathcal{P}} = 28$. Of the remaining points, the accuracy measure is $\chi^2 = 92.2$, with again over two thirds of the productive images leading to a measurement of γ_{LOS} that is biased by more than 2σ , $N_{>2\sigma} = 20$. The mean precision of the recovered shear components remains in line with that of the other models, $\bar{\sigma} = 0.01$.

3.3 Discussion

From these results, summarised in Table 2, we can clearly see the importance of using a sufficiently feature-rich model to fit complicated images, if a reliable and precise measurement of the LOS shear is desired, with well-sampled posterior distributions that yield upper and lower bounds on the recovered shear parameters. This is emphasised in Fig. 9, where we plot the histogram of the mean error on the shear in each model (left panel) and the histogram of the difference between the input and output values of the shear (right panel). From this plot, we can see how the distribution of the size of the error is extremely similar across all the models. However, in the right-hand panel, we can see how the comprehensive model yields consistently more accurate shear measurements than the rest, with the greatest difference between input and output shear being less than 2σ ; this is in stark contrast to the aligned halo and power law models, where

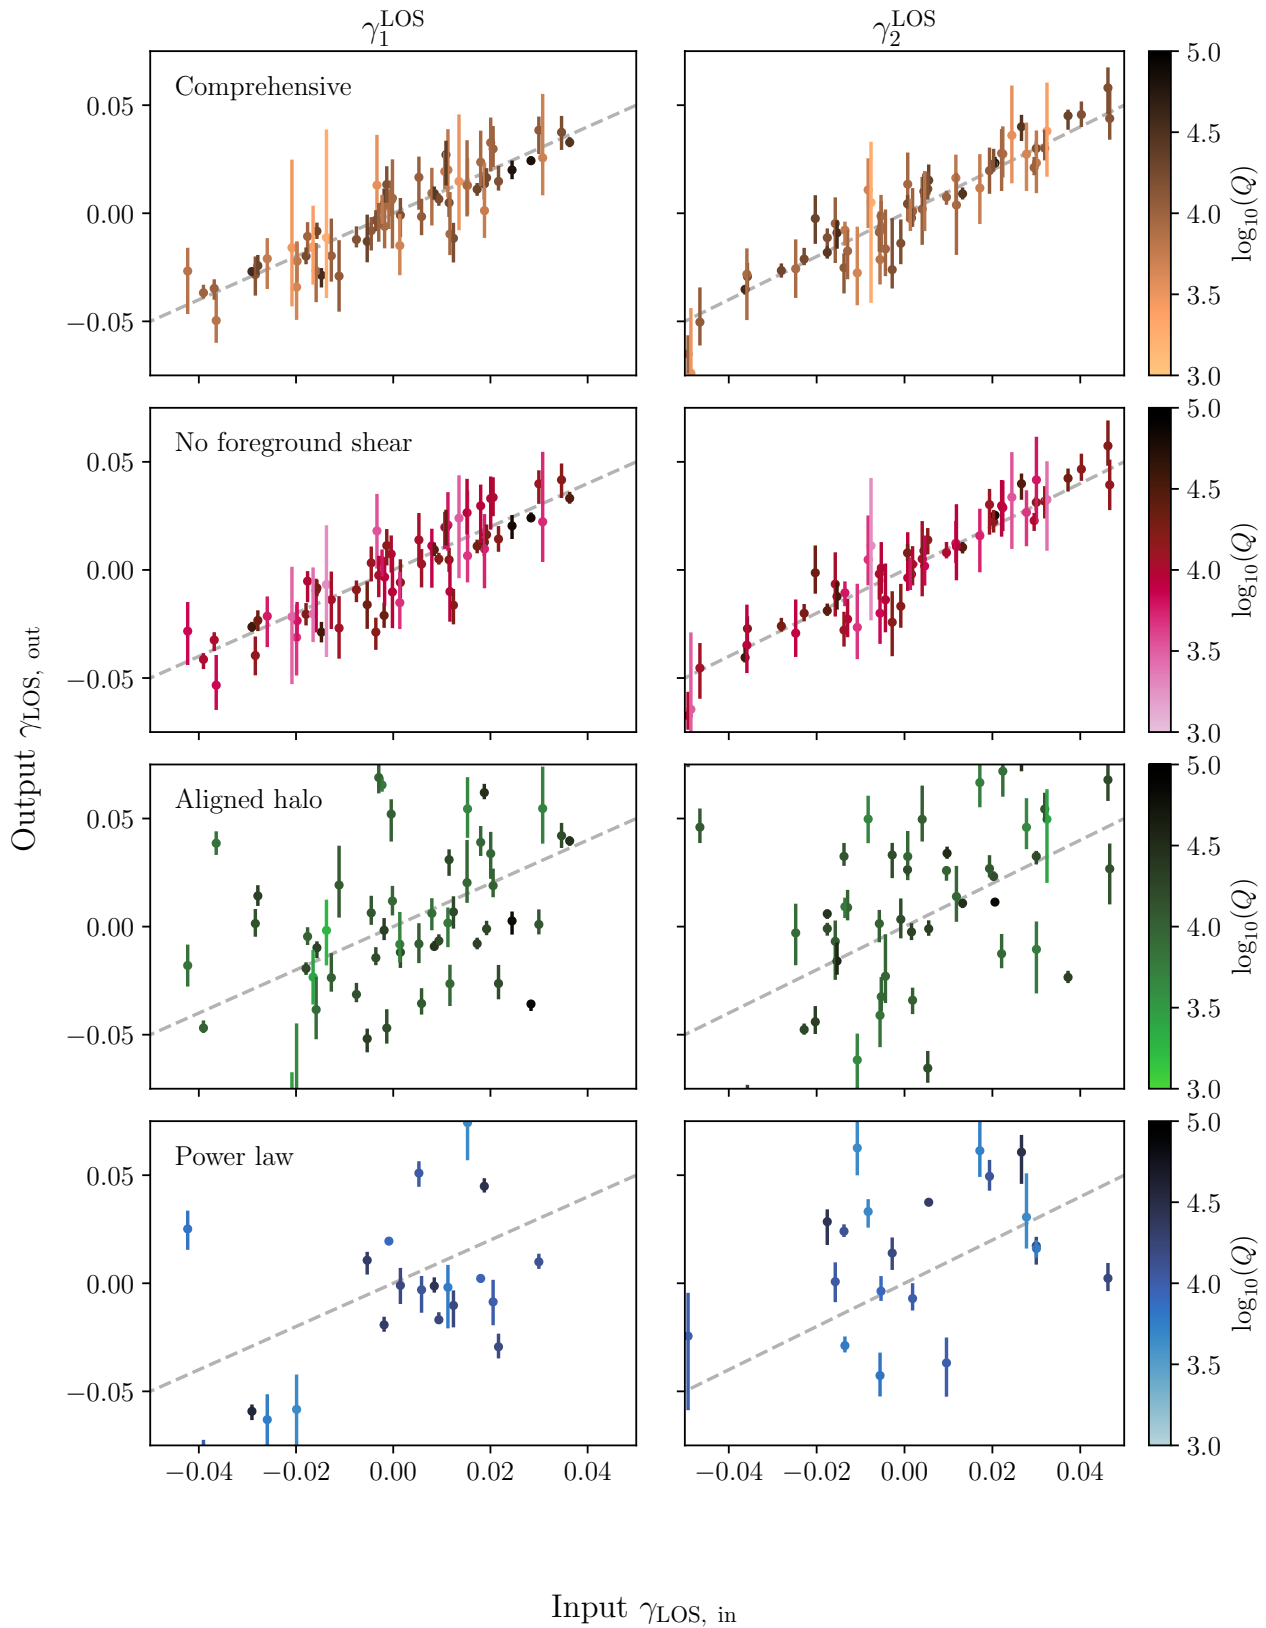


Figure 8. Input $\gamma_{\text{LOS, in}}$ versus output $\gamma_{\text{LOS, out}}$ in the comprehensive model (orange), no foreground shear model (magenta), aligned halo model (green) and power law model (blue). The left-hand column shows the results for γ_1^{LOS} and the right-hand column shows γ_2^{LOS} . The dashed line in each panel is where $\gamma_{\text{LOS, in}} = \gamma_{\text{LOS, out}}$. The brightness of each point indicates the log of the quality criterion Q for the associated system. A larger Q (darker colour) corresponds to a larger SNR, or in other words, a better quality image.

the greatest difference grows as large as 20σ . The accuracy of the no foreground shear model is very close to that of the comprehensive model, barring the two outliers previously mentioned above.

The MCMC sampling is also less successful if a poor model is used: when using the no foreground shear model, we found two-sided bounds on both components the LOS shear for nearly 100% of the catalogue, whereas when using the power law model, we found two-sided bounds on both the shear parameters for less than half of the catalogue.

We can further examine the correlation between the quality measure Q_i defined in Equation (12) and the uncertainty σ_i on γ_{LOS} defined in Equation (13) for each image i , as plotted in the left-hand panel of Fig. 10. From this plot, we can see a strong negative correlation between these quantities, with larger average errors coming from images of worse quality (smaller Q). A simple linear fit yields $\log_{10} \sigma \approx -0.76 \log_{10} Q + 1.01$. This implies that for precise measurements of γ_{LOS} to be obtained, poor quality images should be removed from the catalogue. On the other hand, in the right-hand panel of Fig. 10, we plot the difference between the input and output γ_{LOS} in units of σ against Q . We can see that there is virtually no correlation between these quantities; in other words, the *accuracy* of the shear recovery is not affected by the image quality. Note that we chose to plot these relationships for the comprehensive model as an example, but the trend remains the same for the no foreground shear model.

The excellent recovery of γ_{LOS} obtained from fitting with the no foreground shear model show that the degeneracy between the foreground shear and the properties of the main lens is indeed strong enough that the composite lens images can be satisfactorily obtained by the MCMC parameter inference even when γ_{od} is fixed to zero. This is another powerful argument in favour of the use of the minimal model and again strongly emphasises the fact that γ_{LOS} is the only notion of shear that can evade such degeneracies and thus be accurately measured from strong lensing images.

For test purposes, we tried to increase the amplitude of the source perturbations, which we recall are not accounted for in the fitting models, so as to mimic a poorer modelling of the source light. In that case, we found that the reduced χ^2 , which quantifies the accuracy of the measurement of γ_{LOS} , increases and becomes larger than unity even for the comprehensive model. This suggests that a too simple modelling of the lens light could lead to systematic biases in the measurements of the LOS shear. Such a systematic effect could be mitigated by using complex source models when fitting the images, but this approach may not translate easily when it comes to analysing real data (Nightingale & Dye 2015; Nightingale et al. 2018).

The results presented in this section show that the LOS shear is systematically measurable from simulated images if the lens model used to fit the images is of comparable complexity to that used to generate them. They also show that the foreground shear does not need to be included for good inference if the model for the main lens is sufficiently rich.

4 VALIDITY OF THE TIDAL APPROXIMATION

The analysis and results presented so far were based on the lens equation (2), which supposes that all LOS perturbations to a strong lens – typically dark matter haloes that would lie near the optical axis – can be modelled as tidal fields, whose effect is encoded in the amplification matrices \mathcal{A}_{od} , \mathcal{A}_{os} , \mathcal{A}_{ds} . In this section, we assess the validity of this tidal approximation by fully simulating the effect of randomly distributed haloes on an image.

4.1 The tidal regime and beyond

A LOS perturber is said to be in the tidal regime if the deflection field that it produces can be described by a quadratic potential across the strong lensing image under consideration. Suppose for instance that the observer sees a source that is lensed by a halo h in the absence of the main lens and any other perturber. The corresponding Fermat potential can be expanded at second order in the position θ in the vicinity of the LOS ($\theta = \mathbf{0}$),

$$\psi_h(\theta) = \psi_h(\mathbf{0}) + \alpha_h(\mathbf{0}) \cdot \theta + \frac{1}{2} \theta \cdot \mathcal{A}_h \theta + \dots \quad (19)$$

where $\alpha_h(\mathbf{0})$ is the displacement angle caused by h for an image observed along the optical axis, and \mathcal{A}_h is the contribution of h to the amplification matrix \mathcal{A}_{os} . A similar expansion could be performed for the (od) and (ds) cases; see Fleury et al. (2021a) for further details on the general theory of LOS perturbations in strong lensing. We have chosen to drop the (os) subscripts or superscripts when there is no ambiguity to alleviate notation.

The tidal regime for the halo h is satisfied as long as Equation (19), together with its (od) and (ds) analogues, remain valid up to an angle on the order of the main lens's Einstein radius, θ_E . This requirement is satisfied when h lies sufficiently far from the LOS; as it gets closer, the higher-order terms hidden in the suspension points of Equation (19) may become non-negligible. Among them, the first one that we may worry about is the cubic term, $\mathcal{O}(\theta^3)$, generally referred to as *flexion* because its effect is to produce arc-shaped images (Bacon et al. 2006). Flexion can be understood as gradients of \mathcal{A} , i.e. gradients of convergence and shear; it consists of two complex numbers⁴

$$\mathcal{F} \equiv \partial\gamma = \partial^* \kappa, \quad \mathcal{G} \equiv \partial^* \gamma, \quad (20)$$

where $\partial \equiv (\partial/\partial\theta_1 - i\partial/\partial\theta_2)/2$ is the standard complex derivative with respect to the complex counterpart $\theta_1 + i\theta_2$ of $\theta = (\theta_1, \theta_2)$. When the flexion due to LOS perturbers is properly accounted for in the lens equation, Equation (2) is significantly more complicated, with the addition of four extra complex parameters (Fleury et al. 2021a).

4.2 A line of sight populated with haloes

Since all the previous results we presented were obtained under the assumption that the tidal approximation is valid, we need to understand if it is a good description of the real Universe: how often do perturbers escape this regime, and do they jeopardise our ability to measure the LOS shear from strong lensing images?

In order to test this, we simulate a simple strong lens with a very large number of dark matter haloes in a volume around and along the LOS using the multi-plane lensing formalism. We compute the shear induced on the image by these LOS haloes and then fit the image using the same MCMC procedure followed in the previous sections, with a model again based on the minimal model. If the tidal approximation is valid, we expect the model to provide a good fit and an accurate recovery of the induced shear from the image. If the tidal approximation is significantly broken, we expect to see some sign of it, which could be a biased recovery of the shear, or a bad fit due to the non-negligible role of higher-order distortions.

The main lens and source light are modelled in the same way as in Section 2, except that, for the sake of simplicity, we do not include any lens light. The ellipticities of the lens and source are drawn at random, while the remaining parameters are chosen in a physically

⁴ These definitions follow the conventions of Fleury et al. (2021a), which differ by a factor two from those of Bacon et al. (2006).




Figure 9. *Left panel:* histogram of the uncertainty σ on γ_{LOS} , as defined in Equation (13) for the four models. *Right panel:* histogram of the difference between the input and best-fit output values of γ_{LOS} in units of σ .




Figure 10. *Left panel:* quality criterion Q versus the estimated uncertainty σ on γ_{LOS} [Equation (13)] in the comprehensive model. The grey dashed line represents a linear fit, $\log_{10} \sigma \approx -0.76 \log_{10} Q + 1.01$, highlighting the correlation between these quantities. *Right panel:* quality criterion Q versus the residuals of γ_{LOS} in units of σ in the no foreground shear model. There is no significant correlation between these quantities.

realistic way. Specifically, we set the Einstein radius of the main lens $\theta_E = 1.5''$, the slope of the elliptical power law potential $\gamma_{\text{EPL}} = 2.6$, the apparent magnitude of the source as 23, the half-light radius of the source $R_{\text{Sérsic}} = 0.3''$ and the Sérsic index $n_{\text{Sérsic}} = 6.0$. The main lens is placed at $z_d = 0.5$ and the source is placed at $z_s = 1.5$. The angular positions of the centre of the main lens and of the source are chosen so as to be observed at $\theta = \mathbf{0}$ in the presence of the haloes, $\beta_s = -\alpha_{\text{os}}(\mathbf{0})$ and $\beta_d = -\alpha_{\text{od}}(\mathbf{0})$. Conversions between distances

and redshifts are performed with *astropy* (Astropy Collaboration et al. 2022). As a reminder, we work in the spatially flat Λ CDM cosmology described by the *Planck* 2018 data (Aghanim et al. 2020), with $H_0 = 67.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.315$.

We generate the halo population as follows. We consider a region of space consisting of a comoving cylinder around the optical axis, with comoving radius $R = 11.6 \text{ Mpc}$. This number is chosen so that a point lens with mass $M = 10^{12} M_\odot$ located halfway between the




Figure 11. The Tinker et al. (2008) halo mass function between $10^8 M_\odot$ and $10^{15} M_\odot$.

observer and the source and with a transverse comoving distance R from the optical axis would produce an (os) shear of 10^{-6} on the latter. We then randomly populate this region with a uniform comoving number density

$$n = \int_{M_{\min}}^{M_{\max}} \frac{dM}{M} \frac{dn}{d \ln M}, \quad (21)$$

where $dn/d \ln M$ denotes the halo mass function (HMF) evaluated following Tinker et al. (2008) as implemented in the `colossus` package⁵ (Diemer 2018) at $z_d = 0.5$. We fix a lower mass cutoff to $M_{\min} = 10^8 M_\odot$ corresponding to the typical mass of galactic subhaloes in the cold dark matter scenario, and $M_{\max} = 10^{15} M_\odot$, which yields $n = 5.6 \text{ Mpc}^{-3}$. The mass of each halo is randomly drawn from the same HMF, shown in Fig. 11.

This procedure results in a very large number of haloes, $N_{\text{tot}} = 1.1 \times 10^7$ in the cylindrical region, among which many are practically too light or too far from the LOS to be relevant to the analysis. For computational efficiency, we thus remove the haloes for which a point lens with the same mass would individually produce an (os) shear with magnitude below 10^{-6} on the optical axis. After this operation, we are left with $N = 3.4 \times 10^4$ haloes, whose positions are depicted in Fig. 12. The points in red indicate the haloes that significantly evade the tidal regime (see Section 4.3).

The haloes themselves are modelled with NFW profiles. We compute their concentration using the Diemer & Joyce (2019) model for the mass–concentration relation, which is available in `colossus`. We then use the cosmology subpackage of `lenstronomy` to convert these into the angular scale radius and the deflection angle at the scale radius which constitute the remaining parameters of the NFW mass profile in `lenstronomy`. The left and middle panels of Figure 13 show the distribution of individual (os) shear and convergence produced by the N haloes on the LOS. We can see that the individual tidal effects are generally small, except for a handful of haloes producing convergences and shears of percent order.

Lastly, in each plane containing a halo, we add a uniform component of negative density in order to compensate for the added mass under the form of haloes. This procedure mostly aims to avoid producing a too strong net convergence along the LOS. In practice, we set




Figure 12. The cylindrical region populated with our sample of N haloes, shown in projection in blue (for clarity only a tenth of the population is plotted), of which around 14 % have a relative variation of shear Δ_h greater than 0.1 across θ_E (of which we again plot a tenth), shown in red.

the negative convergence in each of those planes h so as to exactly compensate the positive convergence $\kappa_h(\mathbf{0})$ produced by the halo on the optical axis. This choice is not strictly consistent from a physical point of view, but it has the virtue of simplicity.

4.3 Is the tidal regime a good approximation?

A first, qualitative, assessment of the tidal approximation can be made by examining the variations of the convergence and shear fields due to the haloes only, i.e. in the absence of the main lens. This is shown in Fig. 14, where we observe that the (os) convergence and shear due to the haloes display variations on the order of 100 % on the scale of the main lens’s Einstein radius. The convergence map in particular (left panel of Fig. 14) reveals the presence of a handful of haloes sitting very close to the optical axis, hence strongly evading the tidal regime, while remaining subcritical.

This diagnostic can be further quantified by computing the relative variation Δ of the (os) shear γ as we move away from the optical axis by one unit of the main lens’s Einstein radius, $\theta_E = 1''$,

$$\Delta^2 = \left| \frac{\gamma(\theta_E \hat{x}) - \gamma(\mathbf{0})}{\gamma(\mathbf{0})} \right|^2 + \left| \frac{\gamma(\theta_E \hat{y}) - \gamma(\mathbf{0})}{\gamma(\mathbf{0})} \right|^2, \quad (22)$$

where (\hat{x}, \hat{y}) is an orthonormal basis of the image plane. Note that if the gradients of the shear can be considered constant across the image (flexion regime), then Δ can be expressed in terms of the two complex flexions defined in Equation (20) as

$$\Delta^2 = \frac{1}{2} \theta_E^2 |\gamma|^{-2} (|\mathcal{F}|^2 + |\mathcal{G}|^2). \quad (23)$$

We first consider the value Δ_h of Δ for each halo h taken individually. The distribution of Δ_h for the N haloes is shown in the right panel of Fig. 13. We can see that a considerable number of haloes produce a shear with significant variations on the scale of θ_E . Figure 12 shows in red the haloes for which $\Delta_h \geq 0.1$; they are unsurprisingly clustered around the optical axis. We find that about 14 % of the total sample (4.8×10^3 haloes) have these especially large values of Δ_h .

As expected from the above consideration, the combined effect of all the haloes does escape the tidal regime. We indeed find $\Delta_{\text{tot}} = 0.44$ for the total (os) shear, whose central value is $\gamma_{\text{os}} = -2.4 \times 10^{-2} + 3.8i \times 10^{-3}$. This number drops to $\Delta_{\text{tot}} = 0.045$ if we remove all

⁵ <https://bitbucket.org/bdiemer/colossus>.




Figure 13. *Left panel:* the histogram of the convergence κ_h produced by individual haloes h on the optical axis. *Middle panel:* same for the modulus of the shear $|\gamma_h|$. *Right panel:* same for the relative variation of shear Δ_h on the angular scale θ_E . These are all the (os) quantities.




Figure 14. Convergence field $\kappa_{os}(\theta)$ (*left panel*) and shear field $\gamma_{os}(\theta)$ (*middle and right panels*) produced by the randomly placed haloes around the LOS. Dashed circles indicate the Einstein radius of the main lens.

the haloes h with $\Delta_h \geq 0.1$, in which case the cumulative shear becomes $\gamma_{os} = -1.8 \times 10^{-2} - 9.8i \times 10^{-4}$. This indicates that the haloes escaping the tidal regime are not negligible in the total shear produced along the LOS.

In order to quantify the impact of beyond-tidal effects on the measurement of the LOS shear, we generate two images: the first one is the image that would be observed from a system consisting of our main lens plus N haloes and the second one the image that would be observed from the same system but without the haloes h such that $\Delta_h \geq 0.1$. We do this using the multi-plane formalism in `lenstronomy`. The two images are shown in the top and bottom left-hand panels in Fig. 15.

We then fit these two images using the minimal model, plus an EPL mass profile for the main lens. For this part of the work, we changed our MCMC parameter inference method to use ensemble slice sampling (Karamanis & Beutler 2021) as implemented in the `zeus` package⁶ (Karamanis et al. 2021), which we again modified `lenstronomy` to include. We made this change due to the increased difficulty of the inference which required a faster and more powerful sampling method than the affine-invariant ensemble sampler provided in `emcee`.

In Fig. 16, we show the one- and two-dimensional marginalised posterior distributions of the (od) and LOS shears that result from fitting the images in Fig. 15. The dashed lines indicate expected values of $\gamma_{od}(\mathbf{0})$ and $\gamma_{LOS}(\mathbf{0}) = \gamma_{os}(\mathbf{0}) + \gamma_{od}(\mathbf{0}) - \gamma_{ds}(\mathbf{0})$ exactly along the optical axis, in the absence of the main lens. We firstly consider the

left-hand panel of this figure, which shows the results from fitting the image produced in the presence of the full population of haloes. The expected LOS shear is recovered with good accuracy, but with a precision degraded by about two orders of magnitude compared to the results of Section 2, whose image was generated from a strictly tidal LOS perturbation model. Furthermore, as we can see from the top right panel of Fig. 15, there is an obvious strong residual between the observed image and the reconstructed image. The dipolar morphology of these residuals is reminiscent of the \mathcal{F} -type flexion signal predicted by Fleury et al. (2021a) (see Fig. 7 therein). Further investigation is nevertheless required in order to assess the detectability of such a signal in a realistic set-up.

Taking now the contour plot in the right-hand panel of Fig. 16, which results from fitting the image with the strongly non-tidal haloes removed, we can see that once again the LOS shear is recovered with good accuracy, while the precision is not significantly improved. The residuals exhibited in the case with the full halo population are still present with a similar amplitude, as can be seen in the bottom right panel of Fig. 15.

These results indicate that, quite surprisingly, significant violations of the tidal regime do not entirely prevent one from measuring a notion of LOS shear from a strong lensing image. The tidal model is thus extremely resilient. Furthermore, the best-fit LOS shear appears to be unbiased with respect to the expected values on the optical axis.⁷ Beyond-tidal effects then significantly increase the uncertainty on the

⁶ <https://github.com/minaskar/zeus>.

⁷ This conclusion regarding the impact of higher-order LOS perturbations can be compared with the recent finding that the boxiness or diskiness of strong lenses do not significantly bias H_0 measurements (Van de Vyvere et al. 2022).




Figure 15. The image produced in the presence of the full population of haloes (top left) along with the reconstruction from the best-fit parameters found by the MCMC (top middle) and the normalised residuals of the difference between the two (top right). The bottom panels show the same but without the haloes h escaping the tidal regime ($\Delta_h \geq 0.1$).




Figure 16. The one- and two-dimensional marginalised posterior distributions of the (od) and LOS shears that result from fitting the image produced by the total population of haloes (left) and the population without the haloes h for which $\Delta_h \geq 0.1$ (right). Dashed lines indicate the expected values on the optical axis, $\theta = 0$.

shear measurement, which may be attributed to the fact that the shears are simply not single-valued across the image. They also produce visible patterns in the residuals of the best fit. A natural question is whether those patterns can be attributed to a combination of flexion signals; we shall address this point in future work.

We stress that the quantitative conclusions of this section should be taken with a pinch of salt. Indeed, in the scenario we constructed, haloes were initially randomly distributed in space without any clustering. However, in reality, halo positions are strongly correlated, with many small haloes found in the vicinity of larger ones, and following the cosmic two-point correlation function on large scales. Such correlations imply the presence of large coherent structures far from the LOS that could produce a rather uniform shear. In other words, we expect the role of beyond-tidal effects to be *overestimated* in the present analysis. An important next step for future work will be to simulate a distribution of haloes that follows a cosmologically motivated two-point correlation function, possibly by obtaining a distribution of haloes and their physical characteristics such as mass and concentration from an N -body simulation.

5 CONCLUSIONS

In this paper, we have provided a clear demonstration that the minimal model, based on the dominant lens approximation and the tidal regime, to describe line-of-sight effects in strong gravitational lensing, evades the degeneracies inherent in other formalisms which try to address the same problem. Armed with this powerful model, we showed how the line-of-sight shear can be measured with percent-level accuracy from complicated strong lensing images even when source perturbations are not modelled. Furthermore, we tested the validity of the tidal approximation by simulating the effect of around 3.4×10^4 dark matter haloes on a strong lensing image, finding that while the approximation may not be completely robust, LOS shear measurements are nevertheless possible.

Concretely, in Section 2, we laid out the theoretical background to this work, beginning with the treatment of LOS effects in the tidal regime. We argued that the degeneracies present between the three shear terms in the lens equation in this regime means that they cannot be measured independently of each other or of the model parameters describing the main lens, but that a so-called minimal model can be constructed in which these degeneracies are eliminated. Using *lenstronomy* – modified for this work to include this LOS formalism – we demonstrated the advantage that the minimal model has over the full model, by constructing a simple lens model with shear and fitting the resulting image with both the full and minimal model. As expected, the aforementioned degeneracies were present in the full model but not in the minimal model.

Equipped with this demonstration of the advantage of the minimal model, we applied it to the inference of the LOS shear from a set of 64 simulated strong lensing images, each produced by a rich composite lens model consisting of an elliptical dark matter halo offset from an elliptical baryonic core, along with LOS shear and perturbations to the source light, in Section 3. We fitted these images with a series of models, among which the optimal performance was obtained in the case where all the features of the simulated lenses are included save the foreground shear. This model is very similar to the composite lens models used to fit the H0LiCOW lens SDSS 1206+4332 by [Birrer et al. \(2019\)](#) and the STRIDES lens DES J0408-5354 by [Shajib et al. \(2020\)](#), both re-analysed by the TDCOSMO collaboration, again with a composite lens model ([Millon et al. 2020](#)), and is similar to the multi-baryonic component lens models used by [Williams & Zegeye](#)

(2020) and [Nightingale et al. \(2019\)](#), for example. In this case, we concluded that the LOS shear could be measured with an average absolute uncertainty of 1 %, without any systematic bias or degeneracy with the lens model parameters. We also showed how the inference of the LOS shear worsens with decreasing model complexity, as is to be expected.

Finally, in Section 4, we explored the validity of an important assumption: that LOS perturbers can be treated in the tidal regime. By simulating a scenario in which a volume around the LOS is populated with dark matter haloes, we analysed the images produced in the case where potentially beyond-tidal haloes were present and in the case where these haloes were removed. We observed a clear breakdown of the tidal regime, albeit with no effect on the accuracy of the measurement of the LOS shear. However, the precision was significantly worsened with respect to the best-case scenario of a very simple lens model in Section 2. As a result of the way we modelled the distribution of the dark matter haloes, it is likely that we have overestimated their beyond-tidal effects, making this a very conservative estimate of the achievable precision of the LOS shear.

Lastly, we would like to highlight the important debate in the literature regarding whether external shear is truly *external* or not – in other words, if the distortions to a strong lensing image which have been traditionally attributed to the presence of perturbers external to the main lens are in fact due to another, perhaps poorly understood or controlled-for effect. One example of this was brought to light by [Van de Vyvere et al. \(2020\)](#), who showed how incorrect truncation at the edges of a mass profile for a main lens can induce a spurious external shear. While this effect is important for reconstructions of lensing potentials from mass maps, it does not play a role in our current work due to our use of exact mass profiles which essentially extend to infinity. In that work, the spurious shear was also found to decrease to negligible magnitudes provided a large enough region for the mass map was used. The existence of external shear was recently further called into question by [Etherington et al. \(2023\)](#), who showed how the best-fit value of an external shear parameter does not faithfully represent the actual shear if the lens mass is modelled by a simplistic elliptical power law. This finding is completely in agreement with the results of the current work.

In conclusion, the prospects for the use of the LOS shear as a new cosmological probe are excellent – *provided that the main lens mass is modelled correctly*. The LOS shear, which as we have shown is free from the usual degeneracies that plague typical strong lensing observables, will be able to be measured not only in current strong lensing data but from the many thousands of new strong lenses which are expected to be observed in the next five years by LSST and *JWST*, and further afield, by *Euclid*. The LOS shear thus has the potential to be employed in synergy with weak lensing surveys, magnifying the information that we receive from the inhomogeneous sky and boosting the constraints on cosmological parameters that are already in our hands.

ACKNOWLEDGEMENTS

We thank all the participants of the LOS workshop held in Montpellier in June 2022 for the lively and enriching discussions about this work; we are particularly grateful to Simon Birrer for detailed discussions about *lenstronomy* and to Daniel Johnson for the implementation of LOS effects in the Fermat potential which allowed for the time delay comparison in Appendix A. We also thank Daniel Gilman for reviewing our *lenstronomy* pull request, Anowar Shajib for providing information about TDCOSMO lens modelling, Sherry

Suyu for suggesting that we consider a notion of signal-to-noise ratio to quantify the quality of our images, and Dominique Sluse and Liliya Williams for their detailed comments on the paper. We are also grateful to the anonymous referee for their comments which led to a substantial improvement of our results. Some of the colourmaps used in this work were provided by the CMasher package⁸ (Van der Velden 2020); other colour palettes were taken from [colorbrewer2](#). Some of the numerical computations for this work were done on the Sciama High Performance Compute (HPC) cluster, which is supported by the ICG, SEPNet and the University of Portsmouth. NBH is supported by a postdoctoral position funded by the French Commissariat à l'énergie atomique et aux énergies alternatives (CEA). JL's work was supported by a starting grant from the University of Montpellier and the French Ministry of Higher Education, Research and Innovation. MM acknowledges funding by the Agenzia Spaziale Italiana (ASI) under agreement no. 2018-23-HH.0 and support from INFN/Euclid Sezione di Roma.

DATA AVAILABILITY

The software `lenstronomy` was modified during this work, and the modifications are now part of the public version of that software. The full modification history of the code can be found in the pull request on GitHub: <https://github.com/lenstronomy/lenstronomy/pull/363>. Our wrapper for `lenstronomy` built to produce and analyse the mock image catalogue presented here is publicly available at <https://github.com/nataliehogg/analysis>, along with the image data itself, while Jupyter notebooks for the rest of the analysis in this work can be found at https://github.com/nataliehogg/los_proof_of_concept.

CREDIT AUTHORSHIP CONTRIBUTION STATEMENT

Natalie B. Hogg: Methodology, software, investigation, data curation, writing – original draft, visualisation. **Pierre Fleury:** Conceptualisation, methodology, software, validation, investigation, writing – original draft, visualisation. **Julien Larena:** Conceptualisation, methodology, writing – original draft. **Matteo Martinelli:** Resources, writing – review & editing.

References

Aghanim N., et al., 2020, *Astronomy & Astrophysics*, 641, A6
 Astropy Collaboration et al., 2022, *The Astrophysical Journal*, 935, 167
 Bacon D. J., Goldberg D. M., Rowe B. T. P., Taylor A. N., 2006, *Monthly Notices of the Royal Astronomical Society*, 365, 414
 Bar-Kana R., 1996, *The Astrophysical Journal*, 468, 17
 Birrer S., Amara A., 2018, *Physics of the Dark Universe*, 22, 189
 Birrer S., Welschen C., Amara A., Refregier A., 2017, *Journal of Cosmology and Astroparticle Physics*, 04, 049
 Birrer S., Refregier A., Amara A., 2018, *The Astrophysical Journal*, 852, L14
 Birrer S., et al., 2019, *Monthly Notices of the Royal Astronomical Society*, 484, 4726
 Birrer S., et al., 2021, *Journal of Open Source Software*, 6, 3283
 Blandford R., Narayan R., 1986, *The Astrophysical Journal*, 310, 568
 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

Dhanasingham B., Cyr-Racine F.-Y., Peter A. H. G., Benson A., Gilman D., 2022, Interlopers speak out: Studying the dark universe using small-scale lensing anisotropies ([arXiv:2203.13775](https://arxiv.org/abs/2203.13775))
 Diemer B., 2018, *The Astrophysical Journal Supplement*, 239, 35
 Diemer B., Joyce M., 2019, *The Astrophysical Journal*, 871, 168
 Etherington A., et al., 2023, arXiv e-prints, p. [arXiv:2301.05244](https://arxiv.org/abs/2301.05244)
 Falco E. E., Gorenstein M. V., Shapiro I. I., 1985, *The Astrophysical Journal Letters*, 289, L1
 Fleury P., Larena J., Uzan J.-P., 2021a, *Journal of Cosmology and Astroparticle Physics*, 08, 024
 Fleury P., Larena J., Uzan J.-P., 2021b, *Classical and Quantum Gravity*, 38, 085002
 Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306
 Gilman D., Birrer S., Treu T., 2020, *Astronomy & Astrophysics*, 642, A194
 Gomer M. R., Williams L. L. R., 2021, *Monthly Notices of the Royal Astronomical Society*, 504, 1340
 Hall A., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 4935
 He Q., et al., 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 5862
 Hinton S. R., 2016, *The Journal of Open Source Software*, 1, 00045
 Karamanis M., Beutler F., 2021, *Statistics and Computing*, 31, 61
 Karamanis M., Beutler F., Peacock J. A., 2021, *Monthly Notices of the Royal Astronomical Society*, 508, 3589
 Kormann R., Schneider P., Bartelmann M., 1994, *Astronomy & Astrophysics*, 284, 285
 Kovner I., 1987, *The Astrophysical Journal*, 316, 52
 Kuhlen M., Guedes J., Pillepich A., Madau P., Mayer L., 2013, *The Astrophysical Journal*, 765, 10
 Kuhn F. A., Birrer S., Bruderer C., Amara A., Refregier A., 2021, *Journal of Cosmology and Astroparticle Physics*, 04, 010
 Massey R., et al., 2015, *Monthly Notices of the Royal Astronomical Society*, 449, 3393
 McCully C., Keeton C. R., Wong K. C., Zabludoff A. I., 2014, *Monthly Notices of the Royal Astronomical Society*, 443, 3631
 Millon M., et al., 2020, *Astronomy & Astrophysics*, 639, A101
 Navarro J. F., Frenk C. S., White S. D. M., 1996, *The Astrophysical Journal*, 462, 563
 Nightingale J. W., Dye S., 2015, *Monthly Notices of the Royal Astronomical Society*, 452, 2940
 Nightingale J., Dye S., Massey R., 2018, *Monthly Notices of the Royal Astronomical Society*, 478, 4738
 Nightingale J. W., Massey R. J., Harvey D. R., Cooper A. P., Etherington A., Tam S.-I., Hayes R. G., 2019, *Monthly Notices of the Royal Astronomical Society*, 489, 2049
 Nightingale J. W., et al., 2022, Scanning For Dark Matter Subhalos in Hubble Space Telescope Imaging of 54 Strong Lenses ([arXiv:2209.10566](https://arxiv.org/abs/2209.10566))
 Ostdiek B., Diaz Rivero A., Dvorkin C., 2022, *The Astrophysical Journal*, 927, 83
 Schneider P., 1997, *Monthly Notices of the Royal Astronomical Society*, 292, 673
 Schneider P., Sluse D., 2014, *Astronomy & Astrophysics*, 564, A103
 Seitz S., Schneider P., 1994, *Astronomy & Astrophysics*, 287, 349
 Sengül A. Ç., Dvorkin C., Ostdiek B., Tsang A., 2022, *Monthly Notices of the Royal Astronomical Society*, 515, 4391
 Shajib A. J., et al., 2020, *Monthly Notices of the Royal Astronomical Society*, 494, 6072
 Suess K. A., Kriek M., Price S. H., Barro G., 2019, *The Astrophysical Journal*, 877, 103
 Tessore N., Metcalf R. B., 2015, *Astronomy & Astrophysics*, 580, A79
 Tinker J. L., Kravtsov A. V., Klypin A., Abazajian K., Warren M. S., Yepes G., Gottlober S., Holz D. E., 2008, *The Astrophysical Journal*, 688, 709
 Van de Vyvere L., Sluse D., Mukherjee S., Xu D., Birrer S., 2020, *Astronomy & Astrophysics*, 644, A108
 Van de Vyvere L., Gomer M. R., Sluse D., Xu D., Birrer S., Galan A., Vernardos G., 2022, *A&A*, 659, A127
 Vegetti S., Koopmans L. V. E., 2009, *Monthly Notices of the Royal Astronomical Society*, 400, 1583

⁸ <https://github.com/1313e/CMasher>.

- Vegetti S., Lagattuta D. J., McKean J. P., Auger M. W., Fassnacht C. D., Koopmans L. V. E., 2012, *Nature*, 481, 341
- Vegetti S., Koopmans L. V. E., Auger M. W., Treu T., Bolton A. S., 2014, *Monthly Notices of the Royal Astronomical Society*, 442, 2017
- Van der Velden E., 2020, *The Journal of Open Source Software*, 5, 2004
- Williams L. L. R., Zegeye D., 2020, Two-component mass models of the lensing galaxy in the quadruply imaged supernova iPTF16geu ([arXiv:2006.09391](https://arxiv.org/abs/2006.09391)), doi:10.21105/astro.2006.09391
- Windhorst R. A., et al., 2011, *The Astrophysical Journal Supplement*, 193, 27
- Wong K. C., et al., 2020, *Monthly Notices of the Royal Astronomical Society*, 498
- Zhang G., Mishra-Sharma S., Dvorkin C., 2022, Inferring subhalo effective density slopes from strong lensing observations with neural likelihood-ratio estimation ([arXiv:2208.13796](https://arxiv.org/abs/2208.13796))

APPENDIX A: MODIFYING `lenstronomy`

The `lenstronomy` software is a modular code designed primarily for modelling strong gravitational lenses, with numerous cosmological applications. The software contains a number of packages and subpackages which deal with the separate aspects of this task. For example, the `LensModel` package contains subpackages such as `Profiles`, which in turn contains a class for each lens mass profile which is available in `lenstronomy`.

The modifications to `lenstronomy` made for this work follow this modular structure. We introduced a new subpackage under the `LensModel` package called `LineOfSight`, which contains a new class, `SinglePlaneLOS`. This new class contains functions to compute the displacement angle, Hessian and Fermat potential in the LOS formalism. The class also inherits the ray shooting function (i.e. the computation of the source position β), the lensing potential function and the functions for computing the lens mass enclosed in a given sphere or projected radius and lens mass density from the native `lenstronomy` class `SinglePlane`, as these quantities are unchanged by LOS effects.

Furthermore, our new `LineOfSight` package contains a subpackage called `LOSModels`, inside which are the `LOS` and `LOS_MINIMAL` classes. The functions encoding the LOS effects are defined in the `LOS` class, and these functions are then called in the `SinglePlaneLOS` class. The additional parameters associated with the LOS effects are also defined in the `LOS` class, and the parameters associated with the minimal model are defined in `LOS_MINIMAL`. The additional parameters are:

- os: $\gamma_1^{\text{os}}, \gamma_2^{\text{os}}, \kappa^{\text{os}}, \omega^{\text{os}}$;
- od: $\gamma_1^{\text{od}}, \gamma_2^{\text{od}}, \kappa^{\text{od}}, \omega^{\text{od}}$;
- ds: $\gamma_1^{\text{ds}}, \gamma_2^{\text{ds}}, \kappa^{\text{ds}}, \omega^{\text{ds}}$;

in the `LOS` model and

- od: $\gamma_1^{\text{od}}, \gamma_2^{\text{od}}, \kappa^{\text{od}}, \omega^{\text{od}}$;
- LOS: $\gamma_1^{\text{LOS}}, \gamma_2^{\text{LOS}}, \kappa^{\text{LOS}}, \omega^{\text{LOS}}$;

in the `LOS_MINIMAL` model.

The final modifications we made to `lenstronomy` were to introduce additional unit tests which ensure the validity of the functions and the quantities which they return each time the code is built. We created two types of unit test. The first checks that when all the additional LOS parameters are fixed to zero, the `SinglePlaneLOS` functions return the same values as the `SinglePlane` functions for a given source and lens. The second exploits the correspondence that exists between the LOS formalism and multi-plane lensing with three shear planes. This allows us to construct unit tests which check

that the LOS modifications return the same values as those returned by the equivalent multi-plane set-up using the `MultiPlane` class of `lenstronomy`.

For example, in Table A1, we show the output of the altered single-plane functions that we implemented in the `SinglePlaneLOS` class with the equivalent `MultiPlane` settings for a simple lens model consisting of an elliptical power law lens model with three shear planes. The lens and shear parameters, image positions, and the lens and source redshifts were chosen at random to produce these results.

The displacement angle computed by the `SinglePlaneLOS` class is identical to that computed in the multi-plane case to machine precision, $\Delta\alpha \sim 10^{-16}$ arcsec. The elements of the Hessian matrix in the LOS case are the same as in the multi-plane case at the level of at least 10^{-8} . Finally, the time delay resulting from the modified Fermat potential in the `SinglePlaneLOS` class matches that of the multi-plane case at the level of $\mathcal{O}(10^{-11})$ days.

APPENDIX B: MOCK IMAGE CATALOGUE MODEL PARAMETERS

In this Appendix, we describe the full set of model parameters used to create the images in our mock catalogue. The numbers used to describe galaxy masses, half-light radii and Sérsic indices are freely inspired from [Suess et al. \(2019\)](#).

B1 Baryonic core: Sérsic ellipse

This component is modelled using the `SERSIC_ELLIPSE_POTENTIAL` in `lenstronomy`, which requires the angular half-light radius $R_{\text{Sérsic}}$, the Sérsic index $n_{\text{Sérsic}}$, the ellipticity components e_1, e_2 , the effective convergence κ_{eff} and the centre position x, y .

We begin by drawing a redshift for each lens from a uniform distribution, $z \sim \mathcal{U}(0.4, 0.6)$ and from this compute the angular diameter distance to each lens, D_{od} . Next, we draw the mass of the baryonic component randomly from a log-normal distribution, $\mathcal{M} \sim \ln \mathcal{N}[\ln(\mathcal{M}_{\text{mean}}), \ln(2)/2]$, where the mean mass $\mathcal{M}_{\text{mean}} = 2 \times 10^{11} M_{\odot}$. This distribution ensures that 95% of the lens galaxies we simulate have a mass which is at most a factor of two smaller or larger than the mean mass. From this we compute the half-light radius of the baryonic core in arcseconds,

$$\frac{R_{\text{Sérsic}}}{\text{arcsec}} = \frac{\mathcal{M}}{\mathcal{M}_{\text{mean}}} \frac{2 \text{ kpc}}{D_{\text{od}}} \frac{180 \times 3600}{\pi}. \quad (\text{B1})$$

We also draw the values for the Sérsic index from a log-normal distribution, $n_{\text{Sérsic}} \sim \ln \mathcal{N}[\ln(n_{\text{mean}}), \ln(1.5)/2]$, where $n_{\text{mean}} = 4$.

We compute the ellipticity components using the following expressions,

$$e_1 = \frac{(1-q)}{(1+q) \cos(2\phi)}, \quad (\text{B2})$$

$$e_2 = \frac{(1-q)}{(1+q) \sin(2\phi)}, \quad (\text{B3})$$

where q is the aspect ratio and ϕ the orientation angle of the ellipse. We draw the aspect ratio and orientations for each lens from uniform distributions, $q \sim \mathcal{U}(0.7, 1.0)$ and $\phi \sim \mathcal{U}(0, 2\pi)$ respectively.

We compute the effective convergence at the half-light radius,

$$\kappa_{\text{eff}} = \frac{\mathcal{M}}{\Sigma_{\text{crit}} I}, \quad (\text{B4})$$

where \mathcal{M} is the lens mass in units of M_{\odot} , Σ_{crit} is the critical surface

Table A1. Example output of the `SinglePlaneLOS` and `MultiPlane` classes for comparison.

	SinglePlaneLOS	MultiPlane	Difference
Displacement angle	(0.92460116, 0.77805976)''	(0.92460116, 0.77805976)''	$(4.44089210 \times 10^{-16}, 3.33066907 \times 10^{-16})''$
Hessian	$\begin{bmatrix} 0.11183298 & -0.0151521 \\ -0.01004559 & 0.15331121 \end{bmatrix}$	$\begin{bmatrix} 0.11183299 & -0.0151521 \\ -0.01004559 & 0.15331119 \end{bmatrix}$	$\begin{bmatrix} 6.50405486 \times 10^{-9} & -3.40569642 \times 10^{-10} \\ -1.94965391 \times 10^{-10} & -2.69573809 \times 10^{-8} \end{bmatrix}$
Time delay	6.19293093 days	6.19293093 days	$-1.18127730 \times 10^{-11}$ days

density in the lens plane in units of $M_{\odot}/\text{arcsec}^2$ and the normalisation I is the integrated convergence for the given lens when $\kappa_{\text{eff}} = 1$ in units of arcsec^2 .

Finally, the centre of each lens is fixed to zero, $x = y = 0$, i.e. exactly aligned with the optical axis.

B2 Lens light: Sérsic ellipse

We model the lens light with an identical Sérsic ellipse profile so that the lens light traces the lens mass exactly. We use the `SERSIC_ELLIPSE` light profile in `lenstronomy`, which takes the same parameters as the `SERSIC_ELLIPSE_POTENTIAL` mass profile, less the effective convergence and with the addition of an apparent magnitude.

We start by computing the absolute magnitude of each lens,

$$M = M_{\odot} - 2.5 \log_{10} \left(\frac{M}{Y} \right), \quad (\text{B5})$$

where $M_{\odot} = 4.74$ is the absolute magnitude of the Sun and $Y = 2$ is the mass-to-light ratio for the lens galaxy. The apparent magnitude is then

$$m = M + 5 \log_{10} \left(\frac{D}{1 \text{ Mpc}} \right) + 25, \quad (\text{B6})$$

where D is the luminosity distance to the lens, computed from the redshift and angular diameter distance,

$$D = (1 + z)^2 D_{\text{od}}. \quad (\text{B7})$$

The remainder of the lens light parameters for each lens are the same as those used for the baryonic core.

B3 Dark matter halo: NFW ellipse

We model the dark matter halo using the `NFW_ELLIPSE` mass profile in `lenstronomy`, which requires the scale radius, R_s , the deflection angle at the scale radius α , the ellipticity components e_1, e_2 and the centre position, x, y .

We start by drawing the ratio of dark matter to baryons in the lens galaxy from a uniform distribution, $\lambda \sim \mathcal{U}(0.3, 1.0)$, and thus compute the mass of each halo via

$$M_{\text{halo}} = M/\lambda. \quad (\text{B8})$$

We draw the concentration of each halo randomly from a log-normal distribution, $c \sim \ln \mathcal{N}(c_{\text{mean}}, 0.2)$, where $c_{\text{mean}} = 10^{1.1}$. We then use `lenstronomy` to compute the scale radius and deflection angle for each halo from the mass and concentration.

We follow the same procedure as for the ellipticity of the baryonic component to compute the ellipticity of each NFW halo. Lastly, we displace the centre of each halo from the baryonic core by drawing an offset (in parsecs) from a normal distribution, $o \sim \mathcal{N}(0, 300)$. This distribution is motivated by the findings of [Kuhlen et al. \(2013\)](#), who

found a 300–400 pc offset in a simulation of a Milky Way-like galaxy. An even larger offset of 1.62 kpc has been observed in a central galaxy in the Abell 3827 cluster ([Massey et al. 2015](#)).

B4 Source light: Sérsic ellipse

As with the lens light, we model the source component using the `SERSIC_ELLIPSE` light profile in `lenstronomy`, and our methodology for computing the model parameters for each source is the same as followed for the lens light, except for two differences: the half-light radius and apparent magnitude are computed using the angular diameter distance from observer to source, D_{os} , rather than observer to lens, D_{od} ; and the centre of the source, x, y , is allowed to be offset from the optical axis. We draw the source centre randomly from a uniform distribution on a disk, with the maximum allowed displacement from the optical axis being given by whichever is smallest: the half-light radius or the effective Einstein radius of the lens.

B5 Einstein radius

We estimate the Einstein radius of our composite lenses by neglecting the ellipticity and offset of their components, i.e. as if the lenses were axially symmetric; θ_E is, thus, the solution of

$$\theta_E - \alpha_{\text{Sérsic}}(\theta_E) - \alpha_{\text{NFW}}(\theta_E) = 0, \quad (\text{B9})$$

where $\alpha_{\text{Sérsic}}$ and α_{NFW} are the displacement angles individually produced by the baryon and dark-matter component of the lens. After drawing randomly their parameters, we compute θ_E by solving Equation (B9) numerically, and we only keep the lenses such that $\theta_E > 0.5''$, to ensure that the images we produce have arcs or rings which are reasonably well separated from the lens light.

B6 LOS shear

We add the LOS shear to the images using the LOS model that we implemented in `lenstronomy`. The two components of the od, os and ds shears, γ_1 and γ_2 , are drawn from uniform distributions on a disk. This is done by drawing the square of each shear component randomly from a uniform distribution, $\gamma^2 \sim \mathcal{U}(0, \gamma_{\text{max}}^2)$, where we set $\gamma_{\text{max}} = 0.025$ to ensure that the LOS shear, γ_{LOS} , is never larger than 5% – as we discussed in Section 2, large shears can induce rotation in the image with respect to the source.

The two components of each shear are then computed via

$$\gamma_1 = \gamma \cos(2\phi), \quad (\text{B10})$$

$$\gamma_2 = \gamma \sin(2\phi), \quad (\text{B11})$$

where ϕ is drawn randomly from a uniform distribution, $\phi \sim \mathcal{U}(0, \pi)$.