Gravitational lens recovery with GLASS: measuring the mass profile and shape of a lens

Jonathan P. Coles,[1,2]⋆ Justin I. Read[3] and Prasenjit Saha[2]

[1]*Exascale Research Computing Lab, Campus Teratec, 2 Rue de la Piquetterie, F-91680 Bruyeres-le-Chatel, France*
[2]*Physik-Institut, Universität Zürich, 190 Winterthurerstrasse, CH-8057 Zürich, Switzerland*
[3]*Department of Physics, University of Surrey, Guildford, Surrey GU2 7XH, UK*

ABSTRACT

We use a new non-parametric gravitational modelling tool – GLASS – to determine what quality of data (strong lensing, stellar kinematics, and/or stellar masses) are required to measure the circularly averaged mass profile of a lens and its shape. GLASS uses an underconstrained adaptive grid of mass pixels to model the lens, searching through thousands of models to marginalize over model uncertainties. Our key findings are as follows: (i) for pure lens data, multiple sources with wide redshift separation give the strongest constraints as this breaks the well-known mass-sheet or steepness degeneracy; (ii) a single quad with time delays also performs well, giving a good recovery of both the mass profile and its shape; (iii) stellar masses – for lenses where the stars dominate the central potential – can also break the steepness degeneracy, giving a recovery for doubles almost as good as having a quad with time-delay data, or multiple source redshifts; (iv) stellar kinematics provide a robust measure of the mass at the half-light radius of the stars $r_{1/2}$ that can also break the steepness degeneracy if the Einstein radius $r_E \neq r_{1/2}$; and (v) if $r_E \sim r_{1/2}$, then stellar kinematic data can be used to probe the stellar velocity anisotropy β – an interesting quantity in its own right. Where information on the mass distribution from lensing and/or other probes becomes redundant, this opens up the possibility of using strong lensing to constrain cosmological models.

Key words: gravitational lensing: strong – methods: numerical – methods: statistical.

1 INTRODUCTION

Strong gravitational lenses are rare. Since the discovery of the first lens Q0957+561 (Walsh, Carswell & Weymann 1979), just ∼400 have been discovered to date.[1] However, this number is expected to increase to several thousand over the next 10 years as new surveys, both ground based[2,3] and space based[4] – together with a community of citizen-science volunteers examining the image data for candidates[5] – come online.

Since lensing depends only on gravity, strong lenses offer a unique window on to dark matter and cosmology (Bartelmann 2010; Amendola et al. 2013). However, extracting dark matter properties or cosmological constraints from these lensing data will require sophisticated modelling. In particular, with an unprecedented data set imminent, it is prudent to look again at systematic errors in the lens models to determine what quality of data (in particular complementary data from stellar/gas kinematics, lens time delays and/or stellar-mass constraints) are required to address problems of interest. It is towards that goal that this work is directed.

To see why lens modelling details are of crucial importance, let us recall the essential quantities that appear in lensing (see also Section 2 for a more detailed exposition). First, we have the distances. Let D_L, D_S, D_{LS} be the angular-diameter distances to the lens, source, and from lens to source; these are all proportional to c/H_0 but have factors that depend on the particular choice of cosmology.[6] Typically

$$D_L \approx z_L \frac{c}{H_0} \text{ and } \frac{D_S}{D_{LS}} \sim 1, \qquad (1)$$

where z_L is the redshift of the lens. For multiple images, the sky-projected density must exceed the critical lensing density in some

[1]See e.g. http://masterlens.astro.utah.edu for a catalogue.
[2]http://pan-starrs.ifa.hawaii.edu
[3]http://www.darkenergysurvey.org
[4]http://www.euclid-ec.org
[5]http://spacewarps.org

[6]Here, c is the speed of light in vacuo and H_0 is the Hubble parameter.

region:

$$\Sigma_{\text{crit}} = \frac{c^2}{4\pi G D_{\text{L}}} \sim \frac{1\,\text{kg}\,\text{m}^{-2}}{z_{\text{L}}}, \tag{2}$$

where G is Newton's gravitational constant. The angular separation between the lensed images is of the order of the Einstein radius θ_{E}, which is related to the mass by

$$\theta_{\text{E}} \sim \frac{R_{\text{G}}}{D_{\text{L}}} \frac{D_{\text{LS}}}{D_{\text{S}}}, \tag{3}$$

where $R_{\text{G}} = GM/c^2$ (with M the projected mass enclosed within θ_{E}) is the gravitational radius. If the source is a quasar or otherwise rapidly variable, a time-delay Δt in the variability will be present, where

$$\Delta t \sim R_{\text{G}}/c. \tag{4}$$

So, in principle, one can not only measure the mass of the lens, one can use the dependence on the cosmology-dependent D factors to extract the cosmological model and all its parameters. Zwicky (1937) drew attention to the former, and Refsdal (1964, 1966) pointed out the latter, all long before lenses were discovered. The difficulty with actually doing this, however, became apparent soon after the discovery of the first lens by Walsh et al. (1979). In the first ever paper on lens modelling, Young et al. (1981) found that many plausible mass distributions could reproduce the data.

Young et al. (1981) were remarkably prescient about the subsequent development of lens modelling. First, they introduced the technique of choosing a parametric form for the lensing mass and then fitting for the parameters, which is still the most common strategy (see for example Keeton 2010; Kneib & Natarajan 2011). Secondly, they pointed out the non-uniqueness of lens models – lensing degeneracies. Thirdly, they suggested combining lensing data with stellar kinematics and X-rays, to reduce the effect of the degeneracies. Later work, as well as following up these suggestions, has introduced some further new ideas. Five of these are important for this work.

(i) *Free-form modelling*: In 'free-form' or non-parametric modelling, there is no specified parametric form for the mass distribution. There are still assumptions (or priors) on the mass distribution, such as smoothness or being centrally concentrated (Saha & Williams 1997; Diego et al. 2005; Merten et al. 2009; Coe et al. 2010), but these are much less restrictive than parametric forms. A particularly elegant prior is implemented by Liesenborgs, De Rijcke & Dejonghe (2006), requiring that the mass distribution be non-negative and no extra images allowed. To be concrete, we define from here on: 'non-parametric, or 'free-form' ≡ more parameters than data constraints (i.e. deliberately underconstrained)'.

Being underconstrained, it is then *necessary* to explore model degeneracies rather than finding a single 'best-fitting' solution. Free-form models are more commonly used with cluster lenses (Saha, Read & Williams 2006; Merten et al. 2009; Saha & Read 2009; Sendra et al. 2014), but can be used with galaxy lenses as well, where their less restrictive assumptions can be important. For example, in time-delay galaxy lenses, parametric model measures of the Hubble parameter H_0 have historically been at tension with independent measures (e.g. Kochanek 2002a,b); these are resolved once the less restrictive assumptions of free-form models are permitted (Read, Saha & Macciò 2007). Hybrid methods, using a mass grid on top a parametric model, have also been explored (e.g. Vegetti et al. 2010).

(ii) *Model ensembles*: Model ensembles, exploring a diverse range of possible mass distributions that none the less all fit the data, are a way of combating the non-uniqueness of models. Such ensembles are possible in parametric models (e.g. Bernstein & Fischer 1999; Jullo et al. 2010; Coe, Bradley & Zitrin 2014; Johnson et al. 2014; Richard et al. 2014), but are more common in free-form models, where – since such models are deliberately underconstrained – they become vital (Williams & Saha 2000; Saha & Read 2009; Lubini & Coles 2012).

(iii) *Stellar kinematic constraints*: This was first suggested by Treu & Koopmans (2002) as a means to break lensing degeneracies. The idea is that stellar kinematics can provide an independent estimate of the Einstein radius, via the virial theorem

$$\frac{\langle v_{\text{los}}^2 \rangle}{c^2} \approx \frac{\theta_{\text{E}}}{6\pi} \frac{D_{\text{S}}}{D_{\text{LS}}}, \tag{5}$$

where $\langle v_{\text{los}}^2 \rangle$ is the line of sight stellar velocity dispersion, and the above relation becomes exact for isothermal lenses. This can then be used to probe cosmological parameters if lenses are known to be isothermal (e.g. Collett et al. 2012); or to break the steepness degeneracy in the more general situation (see Section 2.3). The technique has since been applied to many lenses (e.g. Koopmans et al. 2006; Bolton et al. 2008). Going further, the use of two-dimensional kinematics (Barnabè et al. 2011) is especially interesting.

(iv) *Stellar-mass constraints*: The stellar mass in a lens can be inferred from photometry and compared with the total mass (e.g. Keeton, Kochanek & Falco 1998; Kochanek et al. 2000; Rusin et al. 2003; Ferreras, Saha & Williams 2005; Ferreras, Saha & Burles 2008; Leier et al. 2011). Since the inferred stellar mass depends on the assumed initial mass function (IMF), lenses in which stellar mass dominates can be used to derive upper bounds on the stellar M/L (Ferreras et al. 2010). Lower bounds on stellar M/L have also recently been claimed by fitting ΛCDM (Λ cold dark matter) semi-analytic models to the tilt of the Fundamental Plane (Dutton et al. 2013).

(v) *Testing modelling strategies*: Using mock data to see how well a given model can recover simulated lenses is increasingly being recognized as essential. Simple blind tests have appeared in earlier work (for example, fig. 2 in Williams & Saha 2000), but more recently, tests against dynamically simulated galaxies or clusters are favoured (Liesenborgs et al. 2007; Read et al. 2007; Barnabè et al. 2009; Merten et al. 2009; Coe et al. 2010).

There are three further key modelling ideas in the literature that we will not touch upon in this work: to use X-ray intensity and temperature profiles as a mass constraint (e.g. Newman et al. 2013); and to model multiple lenses simultaneously, with one or more cosmological parameters variable but shared between the lenses. This latter strategy has been used to constrain H_0 from time-delay lenses (Saha et al. 2006; Coles 2008; Paraficz & Hjorth 2010) and recently the cosmological parameters Ω as well (Sereno & Paraficz 2014). Third, it is in principle possible to estimate the Ω parameters even from a single lens, if there are lensed sources at multiple redshifts (Lubini et al. 2014) or by using additional priors (Jullo et al. 2010; Suyu et al. 2014).

In this paper, we introduce a new non-parametric lens modelling framework – GLASS (Gravitational Lensing AnalysiS Software). This shares some aspects with an earlier code PIXELENS (Saha & Williams 2004; Coles 2008). However, GLASS – which contains all new code written from the ground up – significantly improves upon PIXELENS in several key ways.

(i) At the heart of GLASS is a new uniform sampling algorithm for high-dimensional spaces (Lubini & Coles 2012). This

allows for large ensembles of $>10\,000$ models to be efficiently generated.

(ii) GLASS provides a modular framework that allows new priors to be added and modified easily.

(iii) The basis functions approximating a model can be easily changed (in this paper, we assume pixels as in PIXELENS).

(iv) With so many models in the final ensemble, we can afford to apply non-linear constraints (for example stellar kinematic data; or the removal of models with spurious extra images) to accept/reject models in a post-processing step.

(v) The central region of the mass map can have a higher resolution to more efficiently capture steep models.

(vi) Stellar density can be used as an additional constraint on the models.

(vii) Point or extended mass objects can be placed in the field.

As a first application, we use GLASS on mock data to determine which combination of lensing, stellar mass and/or stellar kinematic constraints best constrain the projected mass profile and shape of a gravitational lens. We will apply GLASS to real lens data in a series of forthcoming papers.

This paper is organized as follows. In Section 2, we review the key elements of lensing theory, stellar population synthesis, and stellar dynamics we will need. In Section 3, we describe the GLASS code. In Section 4, we describe our mock data. In Section 5, we present our results from applying GLASS to these mock data. Finally, in Section 6 we present our conclusions.

2 THEORY

2.1 Lensing essentials

In the following summary, we follow Blandford & Narayan (1986) with some differences in notation, in particular putting back the speed of light c and the gravitational constant G.

The lens equation

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \frac{D_{\text{LS}}}{D_{\text{S}}}\boldsymbol{\alpha}(\boldsymbol{\theta})$$

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \frac{4G}{c^2 D_{\text{L}}} \int \Sigma(\boldsymbol{\theta}') \frac{(\boldsymbol{\theta} - \boldsymbol{\theta}')}{|\boldsymbol{\theta} - \boldsymbol{\theta}'|^2}\, \mathrm{d}^2\boldsymbol{\theta}' \tag{6}$$

maps an observed image position $\boldsymbol{\theta}$ to a source position $\boldsymbol{\beta}$. Using the thin lens approximation, the lens can be thought of as a projected surface density Σ which diverts the path of a photon instantaneously through the bending angle $\boldsymbol{\alpha}$. The D factors, as in the previous section, are angular-diameter distances, which depend on the cosmological density-parameters Ω, the redshifts z_{L}, z_{S} of the lens and the source, and the Hubble parameter H_0; thus,

$$D_{\text{LS}} = \frac{c}{H_0} \frac{1 + z_{\text{S}}}{1 + z_{\text{L}}} \int_{z_{\text{L}}}^{z_{\text{S}}} \frac{\mathrm{d}z}{\sqrt{\Omega_{\text{m}}(1 + z)^3 + \Omega_\Lambda}} \tag{7}$$

and $D_{\text{L}} \equiv D_{0,\,\text{L}}, D_{\text{S}} \equiv D_{0,\,\text{S}}$. One way to understand the lens equation is via Fermat's principle. We can think of light as travelling only along extremum paths where lensed images occur. Such paths occur at the extrema of the photon *arrival time* $t(\boldsymbol{\theta})$ that depends on the geometric path the photon takes and the general relativistic gravitational time dilation due to a thin lens at redshift z_{L}:

$$\frac{ct(\boldsymbol{\theta})}{(1 + z_{\text{L}})D_{\text{L}}} = \tfrac{1}{2}|\boldsymbol{\theta} - \boldsymbol{\beta}|^2 \cdot \frac{D_{\text{S}}}{D_{\text{LS}}}$$

$$- \frac{4G D_{\text{L}}}{c^2} \int \Sigma(\boldsymbol{\theta}') \ln|\boldsymbol{\theta} - \boldsymbol{\theta}'|\, \mathrm{d}^2\boldsymbol{\theta}'. \tag{8}$$

We can simplify the above equation by introducing a dimensionless time τ and density κ:

$$\tau(\boldsymbol{\theta}) = \frac{ct(\boldsymbol{\theta})}{(1 + z_{\text{L}})D_{\text{L}}} \quad ; \quad \kappa(\boldsymbol{\theta}) \equiv \frac{\Sigma(\boldsymbol{\theta})}{\Sigma_{\text{crit}}} \tag{9}$$

and hence rewrite equation (8) as

$$\tau(\boldsymbol{\theta}) = \tfrac{1}{2}|\boldsymbol{\theta} - \boldsymbol{\beta}|^2 \cdot \frac{D_{\text{S}}}{D_{\text{LS}}} - \frac{1}{\pi} \int \kappa(\boldsymbol{\theta}') \ln|\boldsymbol{\theta} - \boldsymbol{\theta}'| \mathrm{d}^2\boldsymbol{\theta}'. \tag{10}$$

The scaled arrival time τ is like a solid angle. It is of the order of the area (in steradians) of the full lensing system. The expression $|\boldsymbol{\theta} - \boldsymbol{\beta}|^2$ is of the order of the image-separation squared, and the other terms are of similar size. For this reason, it is convenient to measure τ in arcsec^2.

Lensing observations provide information only at $\boldsymbol{\theta}$ where there are images. Hence, the arrival-time surface $\tau(\boldsymbol{\theta})$ is not itself observable. Its usefulness lies in that observables can be derived from it. An image observed at $\boldsymbol{\theta}_1$ implies that $\nabla\tau(\boldsymbol{\theta}_1) = 0$. A measurement of time delays between images at θ_1 and θ_2 implies that $t(\boldsymbol{\theta}_1) - t(\boldsymbol{\theta}_2)$ is known. Interestingly, both these types of observations give constraints that are linear in κ and $\boldsymbol{\beta}$.

The rather complicated dependence of lensing observables on the mass distribution $\kappa(\boldsymbol{\theta})$ has an important consequence: very different mass distributions can result in similar observables. This is the phenomenon of lensing degeneracies. While the non-uniqueness of lens models noted by Young et al. (1981) already hinted at degeneracies, their existence was first derived by Falco, Gorenstein & Shapiro (1985). The most important is the so-called mass-sheet degeneracy, which is that image positions remain invariant if $\tau(\boldsymbol{\theta})$ is multiplied by an arbitrary constant. This corresponds to rescaling the surface density at the images $\kappa(\boldsymbol{\theta})$. In fact there are infinitely many degeneracies (Saha 2000) because any transformation of the arrival-time surface away from the images has no effect on the lensing observables. In particular, there are degeneracies that involve the shape of the mass distribution (Saha & Williams 2006; Schneider & Sluse 2014). Degeneracies tend to be suppressed if there are sources at very different redshifts or 'redshift contrast' (AbdelSalam, Saha & Williams 1998; Saha & Read 2009), because the presence of different factors of $D_{\text{S}}/D_{\text{LS}}$ in the image plane makes it more difficult to change the mass distribution and the arrival-time surface without affecting the lensing observables. But degeneracies are still present with multiple source redshifts (Liesenborgs et al. 2008; Schneider 2014).

2.2 Stellar populations

For many galaxy lenses, the gravitational potential in the inner region is dominated by the stellar mass. Stellar mass can be estimated by combining photometry and colours with models of the stellar populations. Such estimates are reasonably robust, even if the star formation history is very uncertain: given a stellar-population model (such as Bruzual & Charlot 2003) and an IMF, the stellar mass can be inferred to 0.1–0.2 dex using just two photometric bands (see e.g. fig. 1 in Ferreras et al. 2008). By comparing the lensing-mass and stellar-mass profiles in elliptical galaxies, it is possible to extract the radial dependence of the baryonic versus dark matter fraction (Ferreras et al. 2005, 2008; Leier et al. 2011).

The major uncertainty at present in the stellar mass is probably the IMF. In the lensing galaxy of the Einstein Cross, the IMF cannot be much more bottom-heavy than Chabrier (2003), because otherwise the stellar mass would exceed the lensing mass (Ferreras et al. 2010). More massive galaxies, however, do appear to have more of

their stellar mass in low-mass stars. This is indicated by molecular spectral features characteristic of low-mass stars (Cenarro et al. 2004; Conroy & van Dokkum 2012; Ferreras et al. 2013). The Chabrier (2003) IMF would, however, still provide a robust lower limit on the stellar mass and hence, also a limit on the total mass. Accordingly, GLASS allows a constraint of the form

$$M(\boldsymbol{\theta}) \geq M_{\text{stel}}(\boldsymbol{\theta}) \tag{11}$$

on the total mass.

2.3 Stellar kinematics

Another useful constraint follows from the velocity of stars within the lensing galaxy. Assuming spherical symmetry, stars obey the projected Jeans equations (e.g. Binney & Tremaine 2008):

$$\sigma_{\text{p}}^2(R) = \frac{2}{I(R)} \int_R^\infty \text{d}r \left(1 - \beta \frac{R^2}{r^2}\right) \frac{\nu \sigma_{\text{r}}^2 r}{\sqrt{r^2 - R^2}}; \tag{12}$$

$$\sigma_{\text{r}}^2(r) = \frac{r^{-2\beta}}{\nu} \int_r^\infty r'^{2\beta} \nu \frac{G M(r')}{r'^2} \, \text{d}r', \tag{13}$$

where σ_{p} is the projected velocity dispersion of the stars as a function of projected radius R; $I(R)$ is the surface density of the stars; $\nu(r)$ is the three-dimensional stellar density; $\sigma_{\text{r, t}}(r)$ are the radial and tangential velocity dispersions, respectively; $\beta(r) = 1 - \sigma_{\text{t}}^2/2\sigma_{\text{r}}^2 = $ const. is the velocity anisotropy (here assumed to be constant, and not to be confused with $\boldsymbol{\beta}(\boldsymbol{\theta})$ from lensing); G is Newton's gravitational constant; and $M(r)$ is the mass profile that we would like to measure. By convention, we always write R for a projected radius, and r for a 3D radius.

It is immediately clear from equation (12) that, even assuming spherical symmetry, we have a degeneracy between the enclosed mass profile $M(r)$ and the velocity anisotropy $\beta(r)$. This can be understood intuitively since $\beta(r)$ measures the relative importance of radial versus circular orbits and is intrinsically difficult to constrain given only one component of the velocity vector for each star. None the less, $\beta(r)$ can be constrained given sufficiently many stars, since radial Doppler velocities sample eccentric orbits as $r \to 0$ and tangential orbits as $r \to \infty$ (e.g. Wilkinson et al. 2002). It can also be estimated if an independent measure of $M(r)$ is available – for example coming from strong lensing.

While $M(r)$ is difficult to measure from stellar kinematics alone, the mass within the half-light radius is robustly recovered (e.g. Walker et al. 2009; Wolf et al. 2010; Agnello & Evans 2012) since stellar systems in dynamic quasi-equilibrium obey the virial theorem (equation 5). This means that stellar kinematics can break the steepness degeneracy if $r_{1/2} \neq r_{\text{E}}$, where $r_{\text{E}} = D_{\text{L}} \theta_{\text{E}}$ is the physical Einstein radius. We test this expectation in Section 5.

We describe our numerical solution of equation (12) in Section 3.7 and present tests applied to mock data in Section 5.

3 NUMERICAL METHODS

3.1 A new lens modelling framework: GLASS

GLASS is the Gravitational Lensing AnalysiS Software. It extends and develops some of the concepts from the free-form modelling tool PIXELENS (Saha & Williams 2004; Coles 2008), but with all new code. The most compute intensive portion was written in C but PYTHON was chosen because of its flexibility as a language and for its large scientific library support. The flexibility allows GLASS to have

quite sophisticated behaviour while at the same time simplifying the user experience and reducing the overall development time. One of the striking features is that the input file to GLASS is itself a PYTHON program. Understanding PYTHON is not necessary for the most basic use, but this allows a user to build complex analysis of a model directly into the input file. GLASS may furthermore be used as an external library to other PYTHON programs. The software is freely available for download or from the first author.[7]

The key scientific and technical improvements are:

(i) A new uniform sampling algorithm for high-dimensional spaces.

At the heart of GLASS lies a new algorithm for sampling the high-dimensional linear space that represents the modelling solution space. This algorithm was described and tested in Lubini & Coles (2012); it is multithreaded allowing it to run efficiently on many-cored machines.

(ii) A modular framework that allows new priors to be added and modified easily.

Each prior is a simple function that adds linear constraints that operate on either a single lens object or the entire ensemble of objects. GLASS comes with a number of useful priors (the default ones will be described in Section 3.3), but a user can write their own directly in the input file, or by modifying the source code.

(iii) The basis functions approximating a model can be changed.

GLASS currently describes the lens mass as a collection of pixels, but the code has been designed to support alternative methods. In particular, there are future plans to develop a module using Bessel functions. This will require a new set of priors that operate on these functions.

(iv) Non-linear constraints can be imposed in an automated post-processing step.

Once GLASS has generated an ensemble of models given the linear constraints, any number of post-processing functions can be applied. Not only can these functions be used to derive new quantities from the mass models, they can also be used as a filter to accept or reject a model based on some non-linear constraint. For example, we can reject models that have spurious extra images (Section 3.6), or models that do not match stellar kinematic constraints (Section 3.7). The plotting functions within GLASS will correctly display models that have been accepted or rejected.

(v) The central region can have a higher resolution to capture steep models.

With the default basis set of pixels, the mass distribution of the lens is described by a uniform grid. However, in the central region of a lensing galaxy where the mass profile may rise steeply, the centre pixel uses a higher resolution. This allows the density to increase smoothly but still allow for a large degree of freedom within the inner region without allowing the density to be arbitrarily high.

(vi) Stellar density can be used as an additional constraint.

The mass in inner regions of galaxies is often dominated by the stellar component which one can estimate using standard mass-to-light models. These data can be added to the potential as described

---

[7]http://www.jpcoles.com

later in Section 5.3. By using the stellar mass one can place a lower bound on the mass and help constrain the inner most mass profile.

(vii) Point or extended mass objects can be placed in the field.

A shear term can be added to the potential, as shown later in equation (15), to account for mass external to the modelled region. This is useful to capture the gross effects of a distant neighbour, since there is a degeneracy between the ellipticity of a lens and its shear field (the greater the allowed shear, the more circular the lens may be). GLASS also allows further analytic potential components to be included. These can be used to model substructure or multiple neighbours close to the main lens. The substructure may have only a small effect if the lens is a single galaxy, but if the lens is a group or cluster then a potential can be added for each of the known member galaxies. A few standard functions are already included in GLASS including those for a point mass, a power law distribution, or an isothermal (a particular case of the power law).

3.2 Analysis tools

GLASS is not only a modelling tool but also an analysis engine. GLASS provides many functions for viewing and manipulating the computed models. These functions can either be called from a program written by the user or by using the program VIEWSTATE.PY included with GLASS. There is also a tool, LENSPICK.PY for creating a lens, either analytically or from an *N*-body simulation file. To load the simulation data, GLASS relies on the PYNBODY library (Pontzen et al. 2013) and can thus load any file supported by that package.

3.3 Pixelated models

For this paper, we will restrict ourselves to using a pixelated basis set as used by PIXELENS (Saha & Williams 2004; Coles 2008), but note that it is straightforward to add other basis function expansions to GLASS. The algorithm for generating models in GLASS samples a convex polytope in a high-dimensional space whose interior points satisfy both the lens equation and other physically motivated *linear* priors (Lubini & Coles 2012). A limitation of our sampling strategy is that only linear constraints may be applied when building the model ensemble; however, non-linear constraints can be applied in post-processing (see Sections 3.6 and 3.7). We therefore formulate all of our equations as equations linear in the unknowns. We describe the density distribution κ as a set of discrete grid cells or pixels κ_i and rewrite the potential (equation A4) as

$$\psi(\boldsymbol{\theta}) = \sum_n \kappa_n Q_n(\boldsymbol{\theta}) \tag{14}$$

where the sum runs over all the pixels and Q_n is the integral of the logarithm over pixel *n*. The exact form for Q is described in Appendix B. We can find the discretized lens equation by simply taking the gradient of the above equations.

The pixels only cover a finite circular area with physical radius R_{map} and pixel radius R_{pix} with the central cell centred on the lensing galaxy. To account for any global shearing outside this region from, e.g. a neighbouring galaxy, we also add to equation (14) two shearing terms:

$$\gamma_1(\theta_x^2 - \theta_y^2) + 2\gamma_2\theta_x\theta_y. \tag{15}$$

We can continue adding terms to account for other potentials. For instance, we may want to impose a base potential over the field, or add potentials from the presence of other galaxies in the field. GLASS already includes potentials for a point mass or an exponential form,

but custom potentials are straightforward to add and can be included directly in the input file. If the stellar density κ_{s} has been estimated, we can use this as a lower bound where the stellar potential is a known constant of the form equation (14), e.g. $\kappa_n = \kappa_{\mathrm{dm},n} + \kappa_{\mathrm{s},n}$ for a two-component model.

3.3.1 Priors

The lens equation and the arrival times alone are typically not enough to form a closed volume in the solution space. We therefore require additional linear constraints – *priors*. Some of these are 'physical' in the sense that they are unarguable – for example demanding that the mass density is everywhere positive; others are more subjective, for example demanding that the mass map is smooth over some region. Such 'regularization' priors may be switched off for all or some of the mass map if the data are sufficiently constraining.

The priors built in to GLASS are similar to those used in PIXELENS (Coles 2008). The physical priors are always used by default; the regularization priors are used sparingly – i.e. only if the data are not sufficiently constraining to obtain sensible solutions without them:

Physical priors

(i) The density must be non-negative everywhere.
(ii) Image parity is enforced.

Regularization priors

(i) The local gradient everywhere must point within $45°$ of the centre.
(ii) The azimuthally averaged density profile must have a slope everywhere ≤ 0.
(iii) The density is inversion symmetric.

For typical lens data, the regularization priors are very important for creating physically sensible solutions. Prior (i) demands that the peak in the mass density is at the centre of the mass map. Secondary 'plateaus' in the mass map are possible, but not secondary peaks. Note that this prior still successfully allows merging galaxy systems to be correctly captured, provided that the two galaxies are not equally dense in projection (see for example the PIXELENS model of the merger system B1608 in Read et al. 2007), and for the successful detection of 'meso-structure' in strong lensing galaxy clusters (Saha, Williams & Ferreras 2007). Prior (ii) is arguably a physical prior since a positive slope in the azimuthally averaged density profile would be unstable (e.g. Binney & Tremaine 2008). Note that this prior does not preclude successful modelling of mergers or substructure unless the total projected mass in substructure is comparable to the projected mass of the host in an azimuthal annulus (Read et al. 2007; Saha et al. 2007). Prior (iii) is only used for doubles that ought to be inversion symmetric and quads where inversion symmetry is clear from the image configuration.

Finally, we remind the reader that all of the regularization priors can be switched off or changed/improved depending on the data quality available. For clusters, substructure can be explicitly modelled by adding analytic potentials at the known locations of galaxies; furthermore, the above priors can be relaxed in regions of the mass map where the data are particularly constraining (for example near the images). We will apply GLASS to a host of strong lensing clusters in forthcoming work, where we will explicitly test the prior on mock data that has significant substructure.

3.4 Building the model ensemble

In the simplest form, a single model for a lens is a tuple $\mathcal{M} = (\kappa, \boldsymbol{\beta}, \gamma_1, \gamma_2)$. A single model represents a single point in the solution space polytope. Using the MCMC sampling strategy described in Lubini & Coles (2012), we uniformly sample this space. Collectively, the sampled models are referred to as an ensemble $\mathcal{E} = \{\mathcal{M}_i\}$, where we usually generate $|\mathcal{E}| \sim 1000$ models. One can choose to further process these models to impose priors that may be difficult to enforce during the modelling process. For instance, non-linear constraints, or simply filtering of models that do not meet some criteria can be excluded, or weighted against as discussed previously. In this paper, we do not exclude any models and treat all models as equally likely.

The time to generate the model ensemble is mostly a function of the size of the parameter space. The MCMC algorithm has a 'warm-up' phase where it estimates the size and shape of each dimension in the solution space. Once this has been completed, the models are sampled very quickly. In fact, there is little difference between generating 1000 models and generating 10 000 models, although we find little statistical difference after 1000 models. For the mock lenses, the typical 'warm-up' time was about 4 s, and the modelling time was 20 s using a parallel shared-memory machine with 40 cores. The ability to rapidly generate so many models is what allows us to then accept/reject models to apply non-linear constraints (see Sections 3.6 and 3.7). This is a key advantage over our earlier pixelated strong lens tool PIXELENS.

3.5 Ray-tracing

GLASS can also determine the position of images and time delays from particle-based simulation output given a source position $\boldsymbol{\beta}$. This is used to generate the lens configurations used in the parameter study. The particles are first projected on to a very high resolution grid representing the lens plane. The centres $\boldsymbol{\theta}_i$ of each of the grid cells are mapped back on to the source plane using equation (6). If the location on the source plane $\boldsymbol{\beta}_i$ is within a user specified $\varepsilon_{\mathrm{accept}}$ of $\boldsymbol{\beta}$ then $\boldsymbol{\theta}_i$ is accepted and further refined using a root finding algorithm until the distance to $\boldsymbol{\beta}$ is nearly zero. If multiple points converge to an $\varepsilon_{\mathrm{root}}$ of each other then only one point is taken. Care must be taken that the grid resolution is high enough that the resulting image position error is below the equivalent observational error. Time delays are then calculated in order of the arrival time at each image (equation A2).

3.6 Removing models with extra images

While linear constraints are applied in GLASS by the nature of the sampling algorithm, non-linear constraints must be applied in post-processing. Models that are inconsistent with such constraints must then be statistically discarded via a likelihood analysis. An example of such a non-linear constraint is the spurious presence of unobserved images. This 'null-space' prior was first proposed and explored by Liesenborgs et al. (2006) and found to be extremely powerful. We find that our gradient prior in GLASS (see Section 3.3) performs much of the same function as Liesenborgs et al.'s null-space prior, but some models can still rarely turn up spurious images. We reject these in a post-processing step, where we sweep through the model ensemble applying the ray-tracing algorithm described in Section 3.5.

3.7 A post-processing module for stellar kinematics

Similarly to the null-space constraint (Section 3.6), stellar kinematic constraints constitute a non-linear prior on the mass map and must be applied in post-processing. We sweep through the model ensemble performing an Abel de-projection to determine $M(r)$ from the projected surface density $\Sigma(R)$ assuming spherical symmetry (e.g. Binney & Tremaine 2008; Broadhurst & Barkana 2008):

$$M(r) = M_{\mathrm{p}}(<r)$$
$$-4r^2 \int_0^{\pi/2} \Sigma(x) \left[\frac{1}{\cos^2\theta} - \frac{\sin\theta}{\cos^3\theta} \arctan\left(\frac{\cos\theta}{\sin\theta}\right) \right] \mathrm{d}\theta, \tag{16}$$

where

$$M_{\mathrm{p}}(<r) = 2\pi \int_0^r R\Sigma(R)\mathrm{d}R \tag{17}$$

is the projected enclosed mass evaluated at 3D radius r; and $x = r/\cos\theta$.

This de-projection algorithm was tested on triaxial figures in Saha et al. (2006). They found that for triaxialities typical of our current cosmology, the method works extremely well unless the triaxial figure is projected directly along the line of sight such that we see the galaxy or galaxy cluster 'down the barrel'. Such a situation is unlikely, but in any case avoidable since the resultant figure appears spherical in projection. This leads to the seemingly counter-intuitive result that the kinematic constraints – that rely on the above de-projection – are most secure for systems that do not appear spherical in projection (unless independent data can confirm the three-dimensional shape is indeed very round).

We use the de-projected mass to numerically solve equation (12) for constant $\beta(r)$, assuming either $\beta(r) = 1$ or $\beta(r) = 0$ at all radii to bracket the two extremum situations. Where the data are good enough, these two may be distinguished giving dynamical information about $\beta(r)$. In more typical situations, however, we seek to simply marginalize over the effect of $\beta(r)$, using the stellar kinematics as a robust measure of $M(r_{1/2})$ (see Section 2.3).

4 THE MOCK DATA

We now present a study of four mock galaxies with known analytic forms. These are used to verify that GLASS is able to correctly recover the mass profile, and – more importantly – to determine what type and quality of data best constrain the mass profile and shape of a lens.

4.1 The triaxial *N*-body mock galaxies

We generate four two-component mock galaxies, where the dark matter and stellar profiles are allowed to be both steep and shallow. The enclosed mass of the stars and dark matter are both fixed to be $M_{*,\mathrm{DM}} = 1.8 \times 10^{10}$ M$_\odot$ at the stellar scale radius $a_* = 2$ kpc, such that the stars and dark matter contribute equally to the total mass at a_*. The dark matter scalelength is fixed for all models at $a_{\mathrm{DM}} = 20$ kpc. These values were chosen to closely resemble the lensing galaxy PG1115+080 (Weymann et al. 1980). We place the galaxy at a redshift of $z_{\mathrm{L}} = 0.31$ for lensing. Throughout, we assume a cosmology where $H_0^{-1} = 13.7$ Gyr, $\Omega_{\mathrm{M}} = 0.28$, and $\Omega_\Lambda = 0.72$. The critical lensing density is $\kappa_{\mathrm{crit}} \sim 1.8 \times 10^9$ M$_\odot$ kpc^{-2}.

Table 1. Profile parameters for the four mock galaxies. The name indicates whether the galaxy is centrally dark matter or stellar dominated with a shallow or cuspy dark matter density profile. Masses are in units of $1.8 \times 10^{10}\,\mathrm{M_{\odot}}$. The scalelengths for all lenses are $(a_{\star}, a_{\mathrm{DM}}) = (2, 20)$ kpc. R_{map} is the 2D projected radius used to generate the lens configurations. In the case of STAR1.5-DMCUSP, the profile is sufficiently steep that the profile could be truncated at $R_{\mathrm{map}} = 10$ kpc.

Galaxy	$\gamma_{\star}$	$M_{\star}$	γ_{DM}	M_{DM}	R_{map}
STAR1.0-DMCORE	1	4	0.05	$11^{2.95}$	50 kpc
STAR1.0-DMCUSP	1	4	1	11^{2}	50 kpc
STAR1.5-DMCORE	1.5	$2^{1.5}$	0.16	$11^{2.84}$	50 kpc
STAR1.5-DMCUSP	1.5	$2^{1.5}$	1	11^{2}	10 kpc

The galaxies were generated as three-dimensional particle distributions as in Dehnen (2009). Each component follows the profile

$$\rho(\tilde{r}) = \frac{M}{4\pi a^3}(3 - \gamma)(\tilde{r}/a)^{-\gamma}(1 + \tilde{r}/a)^{\gamma - 4}, \qquad (18)$$

where a is the component scale radius mentioned in Table 1; $\tilde{r}^2 = (x/\lambda_1)^2 + (y/\lambda_2)^2 + (z/\lambda_3)^2$ is the ellipsoidal radius; and the axis ratios are $\lambda_1 : \lambda_2 : \lambda_3 = 6 : 4 : 3$. In the case where the central density profile index γ is unity (and in the limit of spherical symmetry), this is the Hernquist profile (Hernquist 1990). The four combinations of profile indices are shown in Table 1.

In Fig. 1, we show the 3D radial density, the 2D projected density, and the 2D enclosed mass for each galaxy.

4.2 Lens configurations

For each of the four galaxies, we used the ray-tracing feature of GLASS described in Section 3.5 to construct six basic lensing morphologies:

(i) one double and one extended double;
(ii) one quad and one extended quad;
(iii) two 2-source quads with varying redshift contrast.

The 'extended' configurations use multiple point sources at the same redshift to simulate an extended source that will produce an arc-like image. Fig. 2 shows the lens configurations for the STAR1.5-DMCUSP galaxy. The configurations for the other galaxies are similar. The labels Z1, Z2, Z3 within the names refer to the redshift of the sources. We have chosen Z1 = 1.72, Z2 = 0.72, and Z3 = 0.51 so that the radial distribution of the images is roughly equally spaced. For all mocks, we do not apply any external shear field. Only the central image of the Z1 source is used to avoid overconstraining the models, otherwise all the central images would fall within the central pixel and no solution exists that satisfies all locations simultaneously for one pixel value.

Each of these configurations were modelled with and without time delays, with and without a central image, and with and without the stellar mass as a lower bound, for a total of 48 test cases. The central image is typically highly demagnified. For galaxy lenses it is very difficult to find since it lies along the sight line to the bright lensing galaxy; in clusters, however, such images have been seen – e.g. Inada et al. (2005). We assumed for all our tests that the lensing mass was radially symmetric (Prior vi). For our mock data, this is known to be true; it is most often the case with real galaxies, unless there is an obvious observed asymmetry. We explore the effect of switching off the symmetry prior in Appendix C. For the quads,

the difference is small; for the doubles – as expected – the results are significantly degraded without this prior. We use, by default, 8 pixels from the centre to the edge of the mass map; the central pixel was further refined into 5×5 pixels to capture any steep rise in the profile (two of the four mock galaxies have a steeply rising inner profile). We demonstrate that our results are robust to changing the grid resolution in Appendix D.

In all cases – despite applying no external shear to the mock lenses – we allow a broad range of external shear in our lens model reconstructions. GLASS correctly returns a small or zero shear in all cases. It is possible that more complex shear fields present in real lensing galaxies could introduce further degeneracies beyond those discussed here. However, any such shear field can, at least in principle, be constrained by data (e.g. combining weak lensing constraints, or assuming that the shear field correlates with visible galaxies – e.g. Merten et al. 2009; Wong et al. 2011).
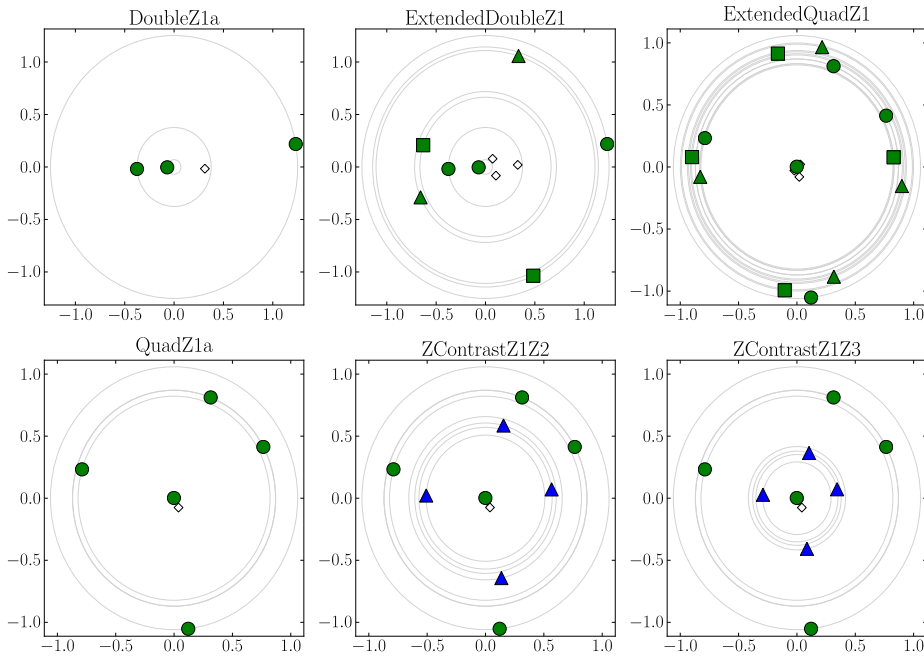
5 RESULTS

5.1 Radial profile recovery

Figs 3 and 4 show shows some example reconstructions of the radial profile of our mock lenses. The left-hand column shows the ensemble average arrival-time surface with images marked as circles and the inferred source positions as diamonds. The centre column shows the radial density profile. The error bars cover a 1σ range around the median; the grey bands show the full ensemble range. The true density profile from the mock data is also plotted for comparison. The vertical lines mark the radial position of the images. The right-hand column shows the enclosed mass. From top to bottom, the rows in Fig. 3 correspond to an extended double for STAR1.5-DMCUSP, and an extended double with stellar-mass constraints for STAR1.5-DMCUSP. The rows in Fig. 4 correspond to a quad with time-delay data for STAR1.5-DMCUSP, and a quad with time delays for STAR1.0-DMCORE. Fig. 5 shows an example 2D reconstruction for STAR1.5-DMCUSP for a quad; we discuss shape recovery further in Section 5.2.

As expected, the accuracies and precisions are best in the range of radii with lensed images where the most information about the lens is present. Even in the weakly constrained case of the extended double where the radial profile is poor, the true enclosed mass $M(<R)$ is well-recovered at the image radii and our ensemble always encompasses it. We have verified this is the case in all of our tests, although for brevity we have not included the plots here. In all cases, there is a dip in the profile at large R due to the cut-off in mass in the lensing map. This is of little importance, though, as there is no lensing information there.

Note that the extended double (top row) gives the poorest constraints, as expected. Adding stellar mass (second row) significantly improves the constraints, for this example where the stars contribute significantly to the potential. Moving to a quad with time delays gives constraints almost as strong as the double with stellar mass, but note that *focusing only on the goodness of the fit can be misleading*. In the third row of Fig. 3, we obtain a better recovery than in the bottom row for *precisely the same data quality*. This occurs because the GLASS prior favours steeper models consistent with STAR1.5-DMCUSP, but not STAR1.0-DMCORE. It is the GLASS prior, rather than the data that is driving the good recovery for STAR1.5-DMCUSP in this example. This emphasizes the importance of using a wide range of mock data tests to determine the role of data versus prior in strong lensing.
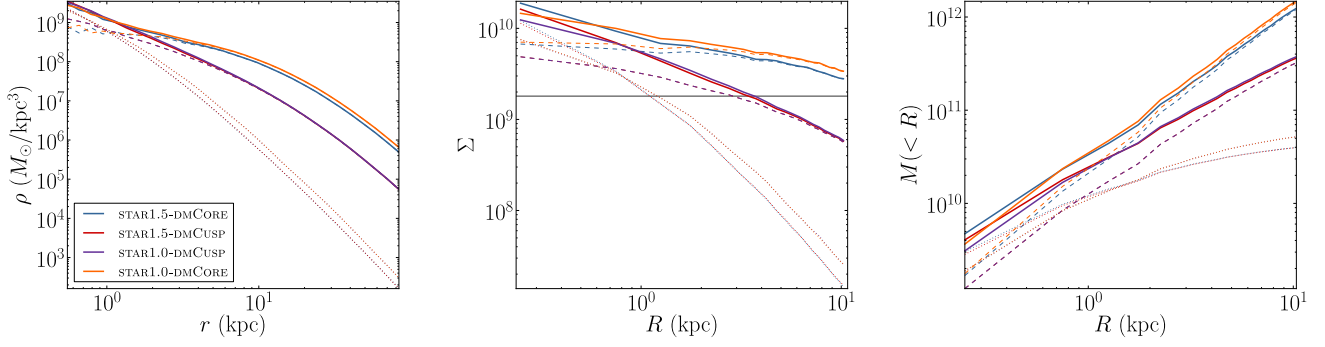
Figure 1. Profiles of the four mock galaxies showing the stellar (dotted) and dark matter (dashed) components and the total (solid). Left: the spherically averaged density. The stars in models STAR1.5-DMCORE and STAR1.5-DMCUSP contribute significantly to the central potential. Middle: the radially averaged two-dimensional projected density. The critical lensing density at $z_L = 0.31$, $\kappa_{crit} \sim 1.8 \times 10^9 \, M_\odot \, \text{kpc}^{-2}$, is marked by the horizontal line. Right: the enclosed projected mass.

Figure 2. The lens configurations for the six test cases using the STAR1.5-DMCUSP mock galaxy. The other mock galaxies produce similar results. Here, the central image is shown, although not all tests include it. The naming convention indicates the redshift of the sources with Z1 = 1.72, Z2 = 0.72, and Z3 = 0.51. The central image only belongs to the Z1 source to avoid overconstraining the models (see Section 4.2 for further details). Small diamonds identify the location of the source(s) and images of the same shape share a common source. The extended source examples have been constructed so that the images will form arclets. The maximum separation of the sources in the source plane is 2.23 kpc in the extended double and 0.92 kpc in the extended quad. Grey circles are a visual aid to help determine radial separation between images. The axes are in arcseconds.

Figs 6 and 7 show the results for our full mock data ensemble. Each subplot corresponds to a different mock galaxy, as marked. We show the fractional error of the mass distribution for each of the test configurations with (red) and without (black) stellar mass. In Fig. 6, we define the error

$$f_R = \frac{\sum_i \left| M(i) - \widehat{M}(i) \right|}{\sum \widehat{M}(i)}, \tag{19}$$

based on the mass $M(i)$ of each pixel ring i and the mass $\widehat{M}$ from the mock galaxy. In Fig. 7, the error is defined over all the pixels $\boldsymbol{\theta}$:

$$f_\theta = \frac{\sum_\theta \left| M(\boldsymbol{\theta}) - \widehat{M}(\boldsymbol{\theta}) \right|}{\sum_\theta \widehat{M}(\boldsymbol{\theta})}. \tag{20}$$

Since both error measurements consider the mass of each pixel, we are implicitly weighting the recovered density by the varying size of the pixels. The value f_R emphasises the error one would see from radial profiles, while f_θ is useful as a measure of how well each individual pixel is recovered. For both f_R and f_θ, we only consider mass up to one pixel length past the outermost image, since there is no longer any lensing information beyond that point. This means we typically use eight bins, linearly spaced, ignoring the outermost three bins. The spacing changes, however, within the high-resolution region in the middle.

The abundance of strong lensing data increases from left to right within each plot. As a result, there is a general trend for the reconstruction quality to increase (and therefore for f to decrease). When both time delays and a central image are present (TD+central), the quality is highest. A double is known to provide very little
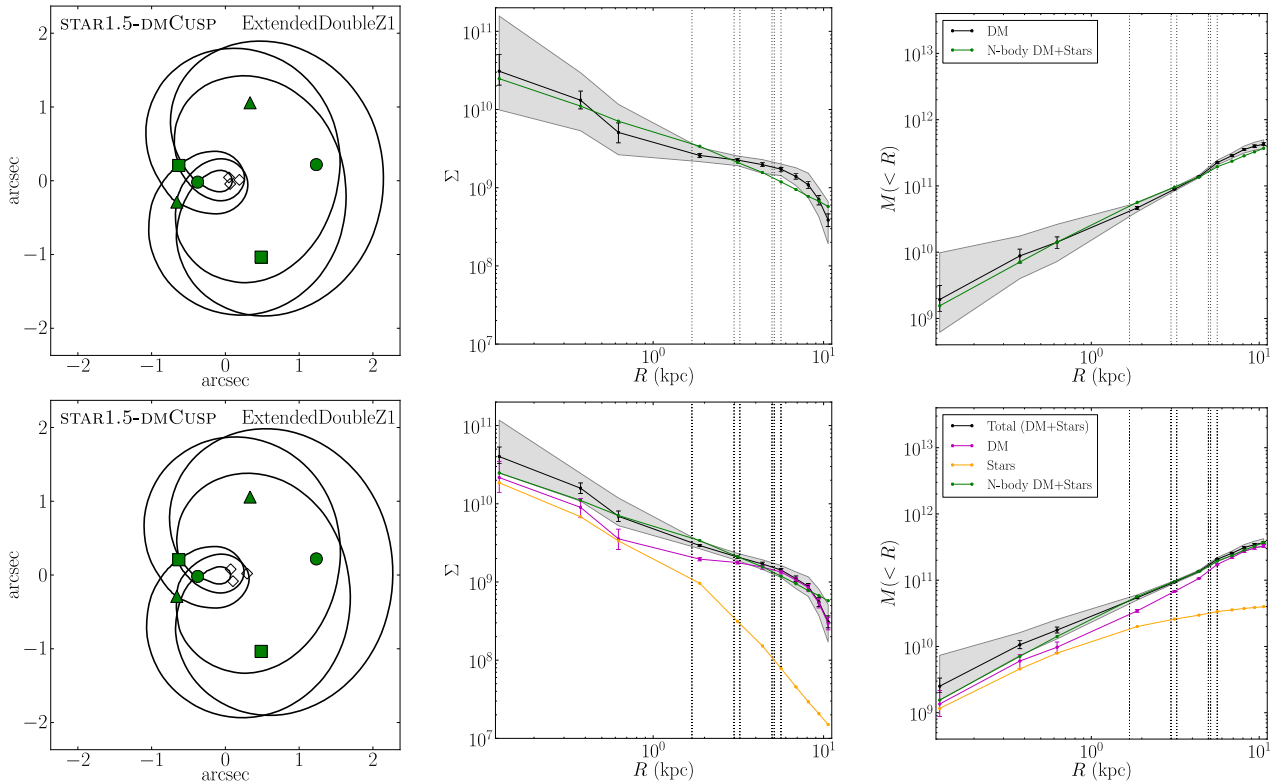
Figure 3. Two reconstructions of the mock galaxy STAR1.5-DMCUSP for an extended double without stellar mass (top) and with stellar mass (bottom). No time delays were assumed. The improved constraints on the mass distribution when a lower bound is given by the stellar mass is evident in the reduced range of allowable models. Left: the ensemble average arrival-time surface with just the iso-contours for the saddle points drawn. The central diamonds show the reconstructed source positions. Middle: the surface density of the dark matter (DM; magenta), the stars (yellow), and the total (black). The original *N*-body mass model (with stars) used to create the lens is shown in green. The vertical lines mark the radial positions of the images. The higher resolution feature of GLASS has been used on the central pixel allowing the steep profile to be captured. Right: the cumulative mass. The error bars on all plots are 1σ; the grey bands show the full range of models.

constraint on the mass distribution. This is particularly evident in galaxies STAR1.0-DMCUSP and STAR1.5-DMCUSP where the mass profile is steepest and the reconstruction of the double is poorest. However, the addition of an arc from the extended source is sufficient to correct this. Notice that, as in Fig. 3, the recovery for STAR1.5-DMCUSP quickly saturates; there is little improvement as the data improves beyond a single quad. This occurs because the GLASS sample prior in the absence of data favours steep models like STAR1.5-DMCUSP over shallower models like STAR1.0-DMCORE (see also Fig. 3).

5.2 Shape recovery

Fig. 7 already gives us important information about how well we can recover the *shape* of a lens. The trends are very similar to the radial profile recovery in Fig. 6, suggesting that if the radial profile is well-recovered then, typically, the shape is too. A notable exception is for the STAR1.5-DMCUSP models where adding stellar-mass constraints aids the shape recovery, but improves the radial mass profile very little. A visual example of the shape recovery is given in Fig. 5.

We can also more directly probe the recovery of the shape of the mass distribution by considering the ratio of the major and minor axes λ_1, λ_2 of the inertia ellipse. If they are equal, the mass is distributed uniformly on the projected disc. The more dissimilar

they are, the more elliptical the mass distribution. We define the global measure of lens shape as

$$s \equiv \lambda_1/\lambda_2, \tag{21}$$

where λ_1 and λ_2 are the eigenvalues of the 2D inertia tensor:

$$\begin{pmatrix} \sum_{\boldsymbol{\theta}} M(\boldsymbol{\theta})\theta_y^2 & -\sum_{\boldsymbol{\theta}} M(\boldsymbol{\theta})\theta_x\theta_y \\ -\sum_{\boldsymbol{\theta}} M(\boldsymbol{\theta})\theta_x\theta_y & \sum_{\boldsymbol{\theta}} M(\boldsymbol{\theta})\theta_x^2 \end{pmatrix}. \tag{22}$$

We always take λ_1 to be the largest value. As with f_R and f_θ, we only consider mass up to one radial position past the outermost image and compute the fractional error as

$$f_{\text{shape}} = \left| s - \widehat{s} \right| / \widehat{s}, \tag{23}$$

where $\widehat{s}$ is the shape of the mock galaxy. The distribution of f_{shape} for each mock galaxy and each test case is shown in Fig. 8. Interestingly, for this global shape parameter recovery it appears more important to have time-delay data and/or a central image (TD,TD+central) than to have a quad or multiple sources with wide redshift separation. In all cases, the stellar mass does little to aid the recovery, reflecting the fact that s is heavily weighted towards the shape e.g. of the mass map, rather than at the centre where the stars may dominate the potential (see equation 22).
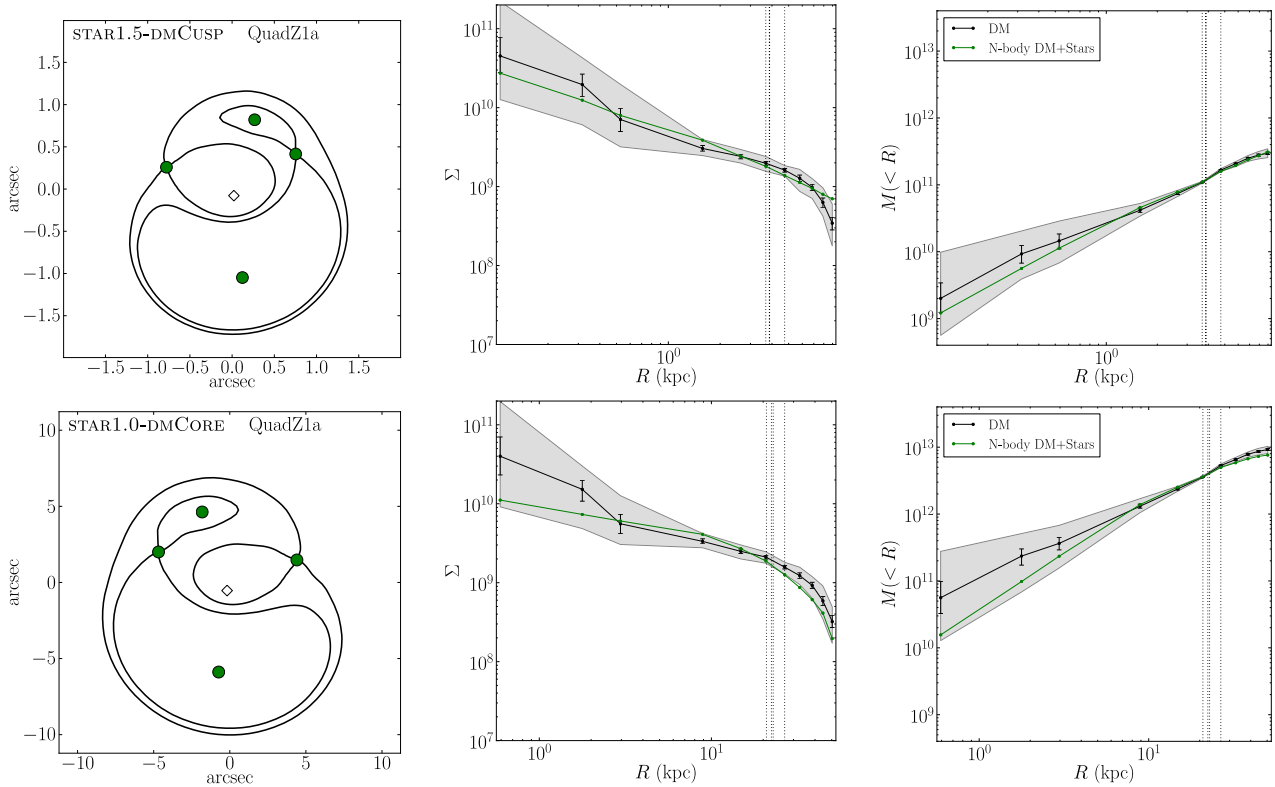
Figure 4. Two further reconstructions similar to Fig. 3. Top: the mock galaxy STAR1.5-DMCUSP but including time delays for a single quad and no stellar mass. With the added information from the quad, the density in the outer regions and at the image radii, is better constrained. Bottom: a quad with time delays, but using the STAR1.0-DMCORE mock galaxy. This galaxy has a shallower stellar density index, and a core in the dark matter. Due to the priors used in GLASS, the modelling favours steep solutions without additional information.
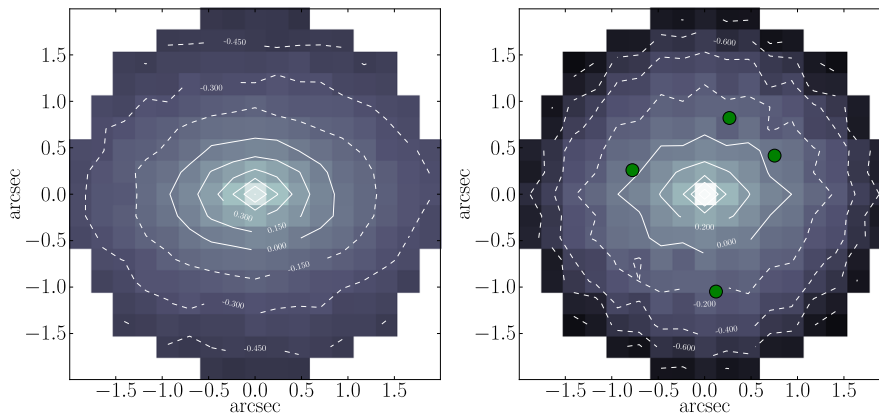
Figure 5. Left: the mock data distribution for STAR1.5-DMCUSP projected on to a coarse grid. Right: the recovered ensemble average κ distribution for the single quad with time delays. The contours are logarithmic base 10 values, where level 0 corresponds to the critical lensing density. Contours below the critical lensing density are drawn with dashed lines.

5.3 Stellar mass

The stellar-mass distribution gives a lower bound on the total mass. Where the stars dominate the central potential, it can provide a powerful constraint extra to the strong lensing data. We took the stellar mass directly from the generated galaxies and projected the particles on to the pixels. GLASS also offers an option to interpolate

any map of stellar mass (e.g. from an observation) on to the pixels. The linear constraint is added to GLASS by writing $\kappa_n = \kappa_{dm,n} + \kappa_{s,n}$ as the sum of the dark matter and stellar-mass components in the potential (equation 14). Since each $\kappa_{s,n}$ is just a constant, we do not add new, separate equations for each pixel. Although we assume a perfect recovery of the stellar mass with no error on the lower
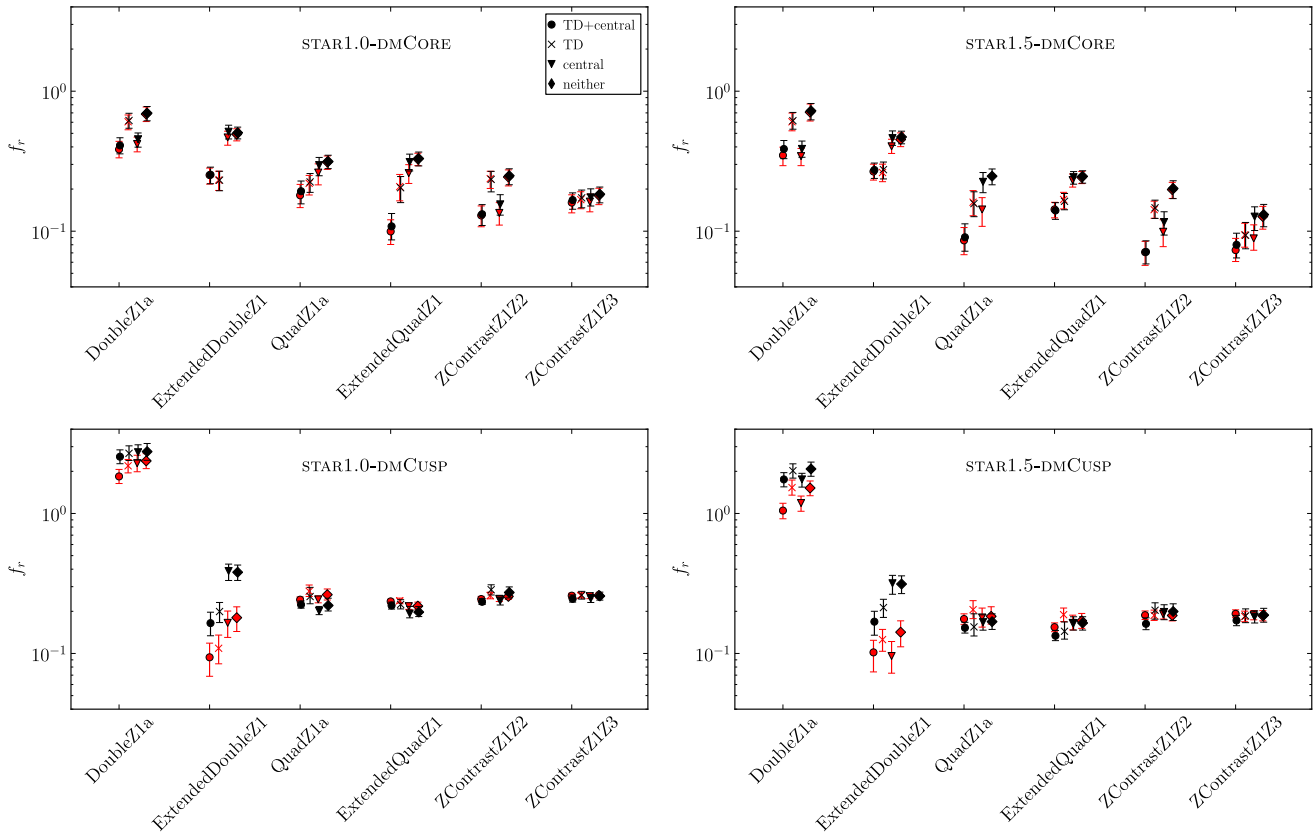
Figure 6. Our main results showing the quality of the radially averaged model recovery (equation 19) for all our test cases. Within each panel are six groups of results for each of six lens morphologies. Each morphology considered the presence of time delays (TD) and a central image (central). The black markers are for tests that did not include the stellar mass as a lower bound constraint, while the red markers indicate where the stellar mass has been included. Error bars show the 1σ range of the model ensemble.

mass bound, it is straightforward to add errors as the stellar-mass constraint remains linear: $\kappa_n = \kappa_{\mathrm{dm}, n} + \epsilon \kappa_{\mathrm{s}, n}$, where $\epsilon \sim 1$ is an additional error parameter.

With the stellar mass lower bound, there is a significant improvement of the reconstruction quality shown in Figs 6 and 7 for the doubles in the steepest mock galaxies (STAR1.0-DMCUSP and STAR1.5-DMCUSP). This is because these models are dominated by stars in the inner region. By contrast, the other two galaxies – where the stars contribute negligibly to the potential – are largely unaffected.

5.4 Stellar kinematics

As outlined in Section 3, GLASS can also run post-processing routines on the model ensemble which can be used to apply non-linear constraints. As an example, we consider here constraints from stellar kinematics. The models in the GLASS ensemble are processed as described in Section 3.7. To illustrate the power of stellar kinematic constraints, in Fig. 9, we plot the projected velocity dispersion calculated for one model (extracted from the full ensemble) of the STAR1.5-DMCUSP quad with time delays and no stellar mass (left), and the same but with stellar mass (middle). In both cases, we calculate curves for two extrema velocity anisotropies: $\beta = 0$ (green) and $\beta = 1$ (red). Overplotted is the correct answer for the STAR1.5-DMCUSP model (black). The stellar half-mass radius (yellow) and Einstein radius (black) are marked by vertical lines. For this configuration, these two radii are well-separated.

Without even sweeping through the model ensemble and formally accepting/rejecting models, Fig. 9 already illustrates what we can obtain from stellar kinematics. The left-hand plot shows the radially averaged projected velocity dispersion $\sigma_{\mathrm{p}}(R)$ (equation 12) for a single quad from the STAR1.5-DMCUSP galaxy without the stellar-mass constraint. The blue data points show the 1σ distribution from the ensemble assuming $\beta = 0$ (solid) and $\beta = 1$ (dashed); the grey bands show the full distributions. Also marked are the $\sigma_{\mathrm{p}}(R)$ calculated from the mock data assuming $\beta = 0$ (solid purple) and $\beta = 1$ (dashed purple), and the true $\sigma_{\mathrm{p}}(R)$ measured directly from the stars (black). This latter has a non-constant $\beta(r)$ (right-hand panel) and differs also from the purple and blue curves in that these all assume spherical symmetry, whereas the stellar distribution is really triaxial. Such triaxiality and varying $\beta(r)$ explains why the purple curves do not match the black one. However, they do largely bracket the correct solution. More interestingly, the curves approximately cross for $\beta = 0$ at the stellar half-light radius (yellow vertical line). This demonstrates, as has previously been reported in the literature, that $\sigma_{\mathrm{p}}(R)$ gives a good estimate of the mass enclosed within $\sim$ the half-light radius $M_{1/2}$ (e.g. Walker et al. 2009; Wolf et al. 2010). The mass *profile*, however, depends on β which is poorly constrained by these data. If we add stellar-mass constraints (middle panel), the situation is improves very little. The true answer already lay close to the bottom of the ensemble distribution; it now is forced to lie right at the edge.

From Fig. 9, it is clear that $\sigma_{\mathrm{p}}(R)$ provides two useful pieces of information. First, it is a powerful probe of $M_{1/2}$. Given a
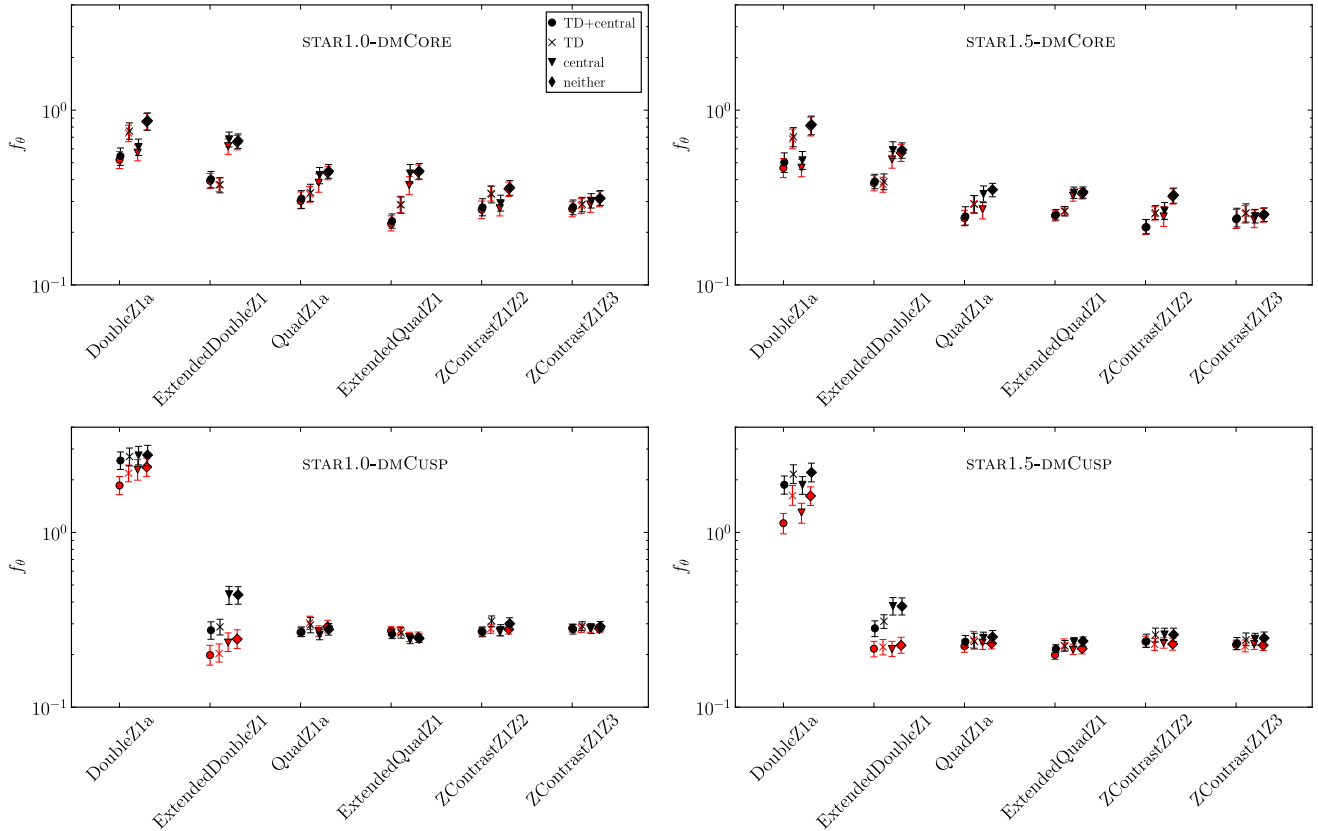
Figure 7. Similar to Fig. 6 but for the fractional error in the pixel-wise recovery (equation 20) of all models. The colours and labels are the same as previously. Error bars show the 1σ range of the model ensemble.

measurement of $\sigma_{\rm p}(r_{1/2}) \sim 150\,{\rm km\,s}^{-1}$, we could usefully reject many models in the ensemble as being overly steep in the centre. We would not, however, obtain a strong constraint on $\beta(r_{1/2})$. We could rule out $\beta(r_{1/2}) = 1$ (blue dashed line), but since our $\beta = 0$ model crosses the true $\beta \sim 0.5$ line at $r_{1/2}$ it is clear that many $\beta(r)$ profiles will be consistent with the data. On the other hand, if we have a situation where $r_{1/2} \sim r_{\rm E}$ (i.e. the vertical yellow and black lines in Fig. 9 overlap), then we will obtain tight constraints on β since we then have two strong constraints on $M(r_{1/2})$ that become redundant. This latter situation of redundancy is also exactly what we would like to constrain cosmological parameters. In this case, we require a third piece of redundant information – in this case in the form of strong lensing time delays. We will discuss such cosmological constraints in a forthcoming paper.

The results for stellar kinematics match our expectations from Section 2.3. Where the lens data already constrain the mass distribution at $r \sim r_{1/2}$, stellar kinematics provide valuable information about the velocity anisotropy of the stars, β (see Fig. 9). Where the lens data poorly constrain the mass distribution at $r_{1/2}$, we may 'integrate out' the effect of unknown β to obtain a robust measure of $M(r_{1/2})$ from the stellar kinematics. This latter is robust to both uncertainties in $\beta(r)$ and to our assumption of spherical symmetry in the kinematic models (Agnello & Evans 2012).

6 CONCLUSIONS

We have introduced a new gravitational lens modelling tool – GLASS – and used it to test the recovery of the mass profile and shape of mock strong lensing galaxies. Our key findings are as follows.

(i) For pure lens data, multiple sources with wide redshift separation give the strongest constraints as this breaks the well-known mass-sheet or steepness degeneracy.

(ii) A single quad with time delays also performs well, giving a good recovery of both the mass profile and its shape.

(iii) Stellar masses – for lenses where the stars dominate the central potential – can also break the steepness degeneracy, giving a recovery for doubles almost as good as having a quad with time-delay data, or multiple source redshifts.

(iv) If the radial density profile is well-recovered, so too is the shape of a lens.

(v) Stellar kinematics provide a robust measure of the mass at the half-light radius of the stars $M(r_{1/2})$ that can also break the steepness degeneracy if $r_{1/2} \neq r_{\rm E}$ – the Einstein radius.

(vi) If $r_{\rm E} \sim r_{1/2}$, then stellar kinematic data can be used to probe the stellar velocity anisotropy β – an interesting quantity in its own right.

Where information on the mass distribution from lensing and/or other probes becomes redundant, this opens up the possibility of using strong lensing to constrain cosmological models. We will study this, and present the first results from GLASS applied to real data, in forthcoming papers.
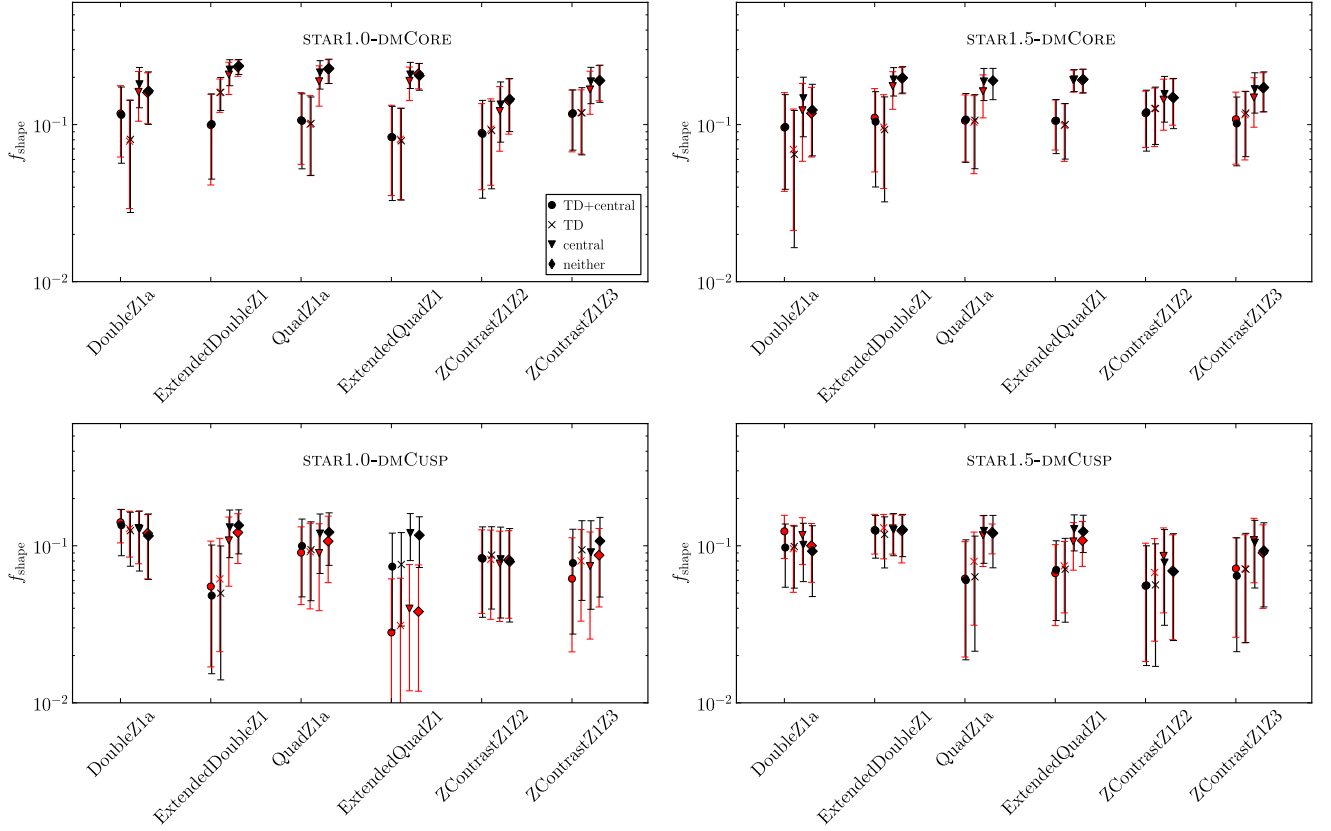
Figure 8. Here, we demonstrate our ability to recover the shape of the lensing mass. The shape ratio λ_1/λ_2 is measured from the principal components λ_1, λ_2 of the mass up to the outermost image. We plot the distribution of fractional error compared with the shape of the mock galaxies (equation 23).
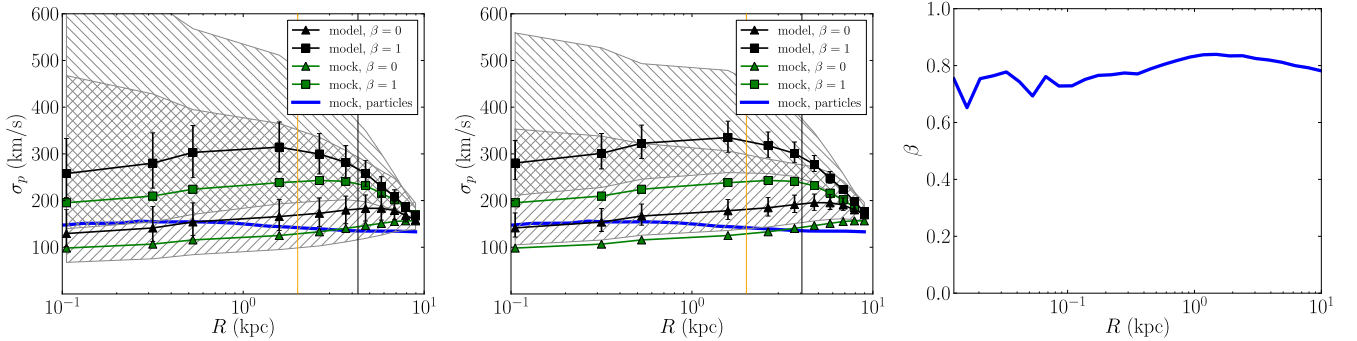
Figure 9. Estimated projected radially averaged velocity dispersion σ_p (equation 12) for a single quad from the STAR1.5-DMCUSP mock galaxy without stellar mass (left) and with stellar mass (middle) assuming an anisotropy $\beta = 0$ (black triangles) and $\beta = 1$ (black squares). Error bars are 1σ and two overlapping hatched areas indicate the full range of models. The equivalent curves are also shown for the projected mock data after using the same analysis routines (green). The solid blue line is the actual cylindrically averaged velocity dispersion of the original mock particle data. The stellar half-mass radius (orange) and Einstein radius (black) are marked by vertical lines. For this configuration, these two radii are well-separated. The actual variation in $\beta(r)$ is also shown (right).

REFERENCES

AbdelSalam H. M., Saha P., Williams L. L. R., 1998, AJ, 116, 1541
Agnello A., Evans N. W., 2012, ApJ, 754, L39
Amendola L. et al., 2013, Living Rev. Relativ., 16, 6
Barnabè M., Nipoti C., Koopmans L. V. E., Vegetti S., Ciotti L., 2009, MNRAS, 393, 1114
Barnabè M., Czoske O., Koopmans L. V. E., Treu T., Bolton A. S., 2011, MNRAS, 415, 2215

Bartelmann M., 2010, Class. Quantum Gravity, 27, 233001
Bernstein G., Fischer P., 1999, AJ, 118, 14
Binney J., Tremaine S., 2008, Galactic Dynamics: Second Edition. Princeton Univ. Press, Princeton, NJ
Blandford R., Narayan R., 1986, ApJ, 310, 568
Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Gavazzi R., Moustakas L. A., Wayth R., Schlegel D. J., 2008, ApJ, 682, 964
Broadhurst T. J., Barkana R., 2008, MNRAS, 390, 1647
Bruzual G., Charlot S., 2003, MNRAS, 344, 1000
Cenarro A. J., Sánchez-Blázquez P., Cardiel N., Gorgas J., 2004, ApJ, 614, L101
Chabrier G., 2003, PASP, 115, 763

Coe D., Benítez N., Broadhurst T., Moustakas L. A., 2010, ApJ, 723, 1678

Coe D., Bradley L., Zitrin A., 2014, preprint (arXiv:e-prints)

Coles J., 2008, ApJ, 679, 17

Collett T. E., Auger M. W., Belokurov V., Marshall P. J., Hall A. C., 2012, MNRAS, 424, 2864

Conroy C., van Dokkum P., 2012, ApJ, 747, 69

Dehnen W., 2009, MNRAS, 395, 1079

Diego J. M., Protopapas P., Sandvik H. B., Tegmark M., 2005, MNRAS, 360, 477

Dutton A. A. et al., 2013, MNRAS, 428, 3183

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

Ferreras I., Saha P., Williams L. L. R., 2005, ApJ, 623, L5

Ferreras I., Saha P., Burles S., 2008, MNRAS, 383, 857

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

Ferreras I., La Barbera F., de la Rosa I. G., Vazdekis A., de Carvalho R. R., Falcón-Barroso J., Ricciardelli E., 2013, MNRAS, 429, L15

Hernquist L., 1990, ApJ, 356, 359

Inada N. et al., 2005, PASJ, 57, L7

Johnson T. L., Sharon K., Bayliss M. B., Gladders M. D., Coe D., Ebeling H., 2014, preprint (arXiv:e-prints)

Jullo E., Natarajan P., Kneib J.-P., D'Aloisio A., Limousin M., Richard J., Schimd C., 2010, Science, 329, 924

Keeton C. R., 2010, Gen. Relativ. Gravit., 42, 2151

Keeton C. R., Kochanek C. S., Falco E. E., 1998, ApJ, 509, 561

Kneib J.-P., Natarajan P., 2011, A&AR, 19, 47

Kochanek C. S., 2002a, Astrophysics, preprint (arXiv:e-prints)

Kochanek C. S., 2002b, ApJ, 578, 25

Kochanek C. S. et al., 2000, ApJ, 543, 131

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

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

Liesenborgs J., De Rijcke S., Dejonghe H., 2006, MNRAS, 367, 1209

Liesenborgs J., de Rijcke S., Dejonghe H., Bekaert P., 2007, MNRAS, 380, 1729

Liesenborgs J., de Rijcke S., Dejonghe H., Bekaert P., 2008, MNRAS, 386, 307

Lubini M., Coles J., 2012, MNRAS, 425, 3077

Lubini M., Sereno M., Coles J., Jetzer P., Saha P., 2014, MNRAS, 437, 2461

Merten J., Cacciato M., Meneghetti M., Mignone C., Bartelmann M., 2009, A&A, 500, 681

Newman A. B., Treu T., Ellis R. S., Sand D. J., 2013, ApJ, 765, 25

Paraficz D., Hjorth J., 2010, ApJ, 712, 1378

Pontzen A., Roškar R., Stinson G. S., Woods R., Reed D. M., Coles J., Quinn T. R., 2013, Astrophysics Source Code Library, record ascl:1305.002

Read J. I., Saha P., Macciò A. V., 2007, ApJ, 667, 645

Refsdal S., 1964, MNRAS, 128, 307

Refsdal S., 1966, MNRAS, 132, 101

Richard J. et al., 2014, MNRAS, 444, 268

Rusin D. et al., 2003, ApJ, 587, 143

Saha P., 2000, AJ, 120, 1654

Saha P., Read J. I., 2009, ApJ, 690, 154

Saha P., Williams L. L. R., 1997, MNRAS, 292, 148

Saha P., Williams L. L. R., 2004, AJ, 127, 2604

Saha P., Williams L. L. R., 2006, ApJ, 653, 936

Saha P., Read J. I., Williams L. L. R., 2006, ApJ, 652, L5

Saha P., Williams L. L. R., Ferreras I., 2007, ApJ, 663, 29

Schneider P., 2014, A&A, 568, L2

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

Sendra I., Diego J. M., Broadhurst T., Lazkoz R., 2014, MNRAS, 437, 2642

Sereno M., Paraficz D., 2014, MNRAS, 437, 600

Suyu S. H. et al., 2014, ApJ, 788, L35

Treu T., Koopmans L. V. E., 2002, MNRAS, 337, L6

Vegetti S., Koopmans L. V. E., Bolton A., Treu T., Gavazzi R., 2010, MNRAS, 408, 1969

Walker M. G., Mateo M., Olszewski E. W., Peñarrubia J., Wyn Evans N., Gilmore G., 2009, ApJ, 704, 1274

Walsh D., Carswell R. F., Weymann R. J., 1979, Nature, 279, 381

Weymann R. J. et al., 1980, Nature, 285, 641

Wilkinson M. I., Kleyna J., Evans N. W., Gilmore G., 2002, MNRAS, 330, 778

Williams L. L. R., Saha P., 2000, AJ, 119, 439

Wolf J., Martinez G. D., Bullock J. S., Kaplinghat M., Geha M., Muñoz R. R., Simon J. D., Avedo F. F., 2010, MNRAS, 406, 1220

Wong K. C., Keeton C. R., Williams K. A., Momcheva I. G., Zabludoff A. I., 2011, ApJ, 726, 84

Young P., Gunn J. E., Oke J. B., Westphal J. A., Kristian J., 1981, ApJ, 244, 736

Zwicky F., 1937, ApJ, 86, 217

APPENDIX A: IMPLEMENTATION DETAILS

Since we want to model the density distribution with a computer it is convenient to choose units that make the relevant quantities of order unity. We therefore measure lengths in light years, time in years, positions in arcseconds, and choose $c = 1$ and $4\pi G = N^2$, where $N^2 \equiv 206\,265$ arcsec rad^{-1}. The mass unit is then $11.988\,\mathrm{M_\odot}$. It will also be useful to define a proxy to the Hubble constant $\zeta \equiv N^2 H_0$. We now express the equations from Section 2 in terms of these new units and introduce some other useful quantities.

The lens equation in its complete form becomes

$$N^2 ct(\boldsymbol{\theta}) = (1 + z_\mathrm{L})\frac{D_\mathrm{L} D_\mathrm{S}}{D_\mathrm{LS}}\frac{1}{2}|\boldsymbol{\theta} - \boldsymbol{\beta}|^2$$
$$- (1 + z_\mathrm{L})\frac{4GD_\mathrm{L}^2}{c^2}\int \Sigma(\boldsymbol{\theta}')\ln|\boldsymbol{\theta} - \boldsymbol{\theta}'|d^2\boldsymbol{\theta}', \quad \text{(A1)}$$

where the factor of D_L^2 in the second term comes from the fact that Σ has units of $\mathrm{M_\odot\,lyr^{-2}}$. We can clean this up by first writing down a dimensionless time delay

$$\tau = [(1 + z_\mathrm{L})d_\mathrm{L}]^{-1}\,\zeta t \quad \text{(A2)}$$

in terms of our previous definitions and defining $D_\mathrm{L} \equiv (c/H_0)d_\mathrm{L}$. We further define a dimensionless density

$$\kappa_\infty = \frac{4\pi G}{c^2}\frac{c}{H_0}d_\mathrm{L}\Sigma = \frac{d_\mathrm{L}}{\zeta}\Sigma \quad \text{(A3)}$$

and a lensing potential

$$\psi(\boldsymbol{\theta}) = \frac{1}{\pi}\int \kappa_\infty(\boldsymbol{\theta}')\ln|\boldsymbol{\theta} - \boldsymbol{\theta}'|d^2\boldsymbol{\theta}'. \quad \text{(A4)}$$

Now we can express equation (A1) very compactly as

$$\tau(\boldsymbol{\theta}) = \frac{1}{2}\xi|\boldsymbol{\theta} - \boldsymbol{\beta}|^2 - \psi(\boldsymbol{\theta}), \quad \text{(A5)}$$

where $\xi = d_\mathrm{S}/d_\mathrm{LS}$. We explicitly write κ_∞ to remind ourselves that there is no source distance factor involved. This is useful when we consider multiple sources.

APPENDIX B: DERIVATION OF PIXELATED DENSITY COEFFICIENTS

When the lens plane is pixelized we need a discrete form of the integral

$$\int \kappa(\boldsymbol{\theta}')\ln|\boldsymbol{\theta} - \boldsymbol{\theta}'|\mathrm{d}^2\boldsymbol{\theta}'.$$

In particular we want

$$\sum_n \kappa_n Q_n(\boldsymbol{\theta}),$$

where Q_n is the logarithm evaluated over the nth pixel at position $\boldsymbol{\theta}_n = (x_n, y_n)$. Let the pixel side length be a. Instead of working with a position vector $\boldsymbol{\theta}$ we work in Cartesian coordinates where $|\boldsymbol{\theta}| = r = \sqrt{x^2 + y^2}$. The integral now becomes

$$Q_n(x, y) = \frac{1}{2} \int_{y_-}^{y_+} \int_{x_-}^{x_+} \ln(x'^2 + y'^2) \mathrm{d}x' \mathrm{d}y',$$

where $x_\pm = x + x_n \pm (a/2)$ and similarly for $y_\pm$. Using the identity

$$\int \ln(x^2 + y^2) \mathrm{d}x = x \ln(x^2 + y^2) - 2x + 2y \arctan(x/a),$$

we can express Q_n as the sum of four parts

$$\begin{aligned} Q_n(x, y) = \frac{1}{2}[&\tilde{Q}_n(x_+, y_+) + \tilde{Q}_n(x_-, y_-) \\ &- \tilde{Q}_n(x_-, y_+) - \tilde{Q}_n(x_+, y_-)], \end{aligned}$$

where

$$\tilde{Q}_n(x, y) = xy(\ln r^2 - 3) + x^2 \arctan(y/x) + y^2 \arctan(x/y).$$

APPENDIX C: NO RADIAL SYMMETRY PRIOR

In this appendix, we explore the effect of the radial symmetry prior. Fig. C1 shows results for a single quad (top two rows) and an extended double (bottom two rows) without the radial prior; in both cases, we do not use the stellar-mass constraints. The bottom row of each group uses time-delay data. Without time-delay data or the radial symmetry prior, the results for the quad are poor – particularly the shape recovery. Including time delays, the results are similar to the case with the radial prior (Figs 3 and 5). Similarly, for the extended double the results without time delays are poor. Even with time delays, the shape is not well-recovered without the radial prior, as expected.

APPENDIX D: PIXEL RESOLUTION CONVERGENCE TEST

In this appendix, we present a convergence test of our results with the grid resolution – `pixrad`. By default, this is set to 8 pixels from the centre of the mass map to the edge. As can be seen from Fig. D1, our results are typically well-converged for `pixrad` > 5. The results for `pixrad` = 5 become systematically biased away from the central regions (where we have the higher resolution adaptive mesh), because our regularization prior combined with a low `pixrad` biases us towards shallow models. This effect diminishes with increasing resolution and is already negligible by `pixrad` = 7. Notice that the mass increases in size with decreasing resolution. This is because we always demand that there are four pixels beyond the outermost image.
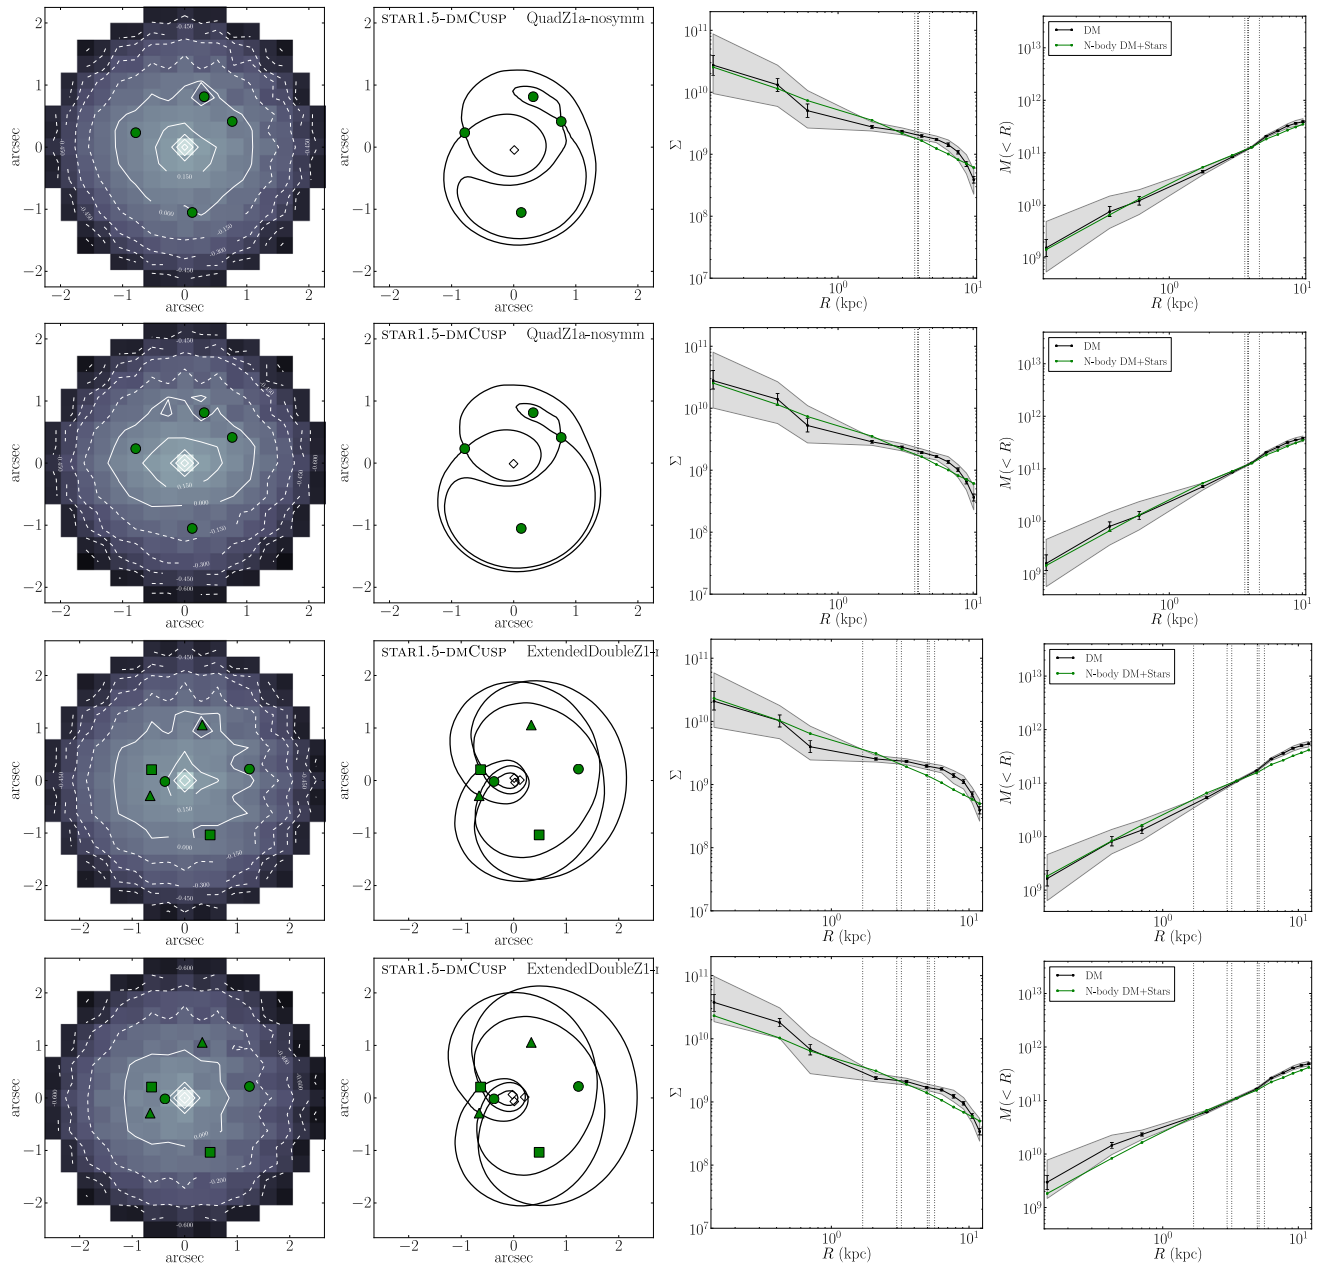
Figure C1. Results for a single quad (upper two rows) and an extended double (lower two rows) without the radial symmetry prior; the bottom row of each group uses time-delay data. The stellar-mass constraint has not been used in any of these examples. Figures and symbols are as in Figs 3 and 5.
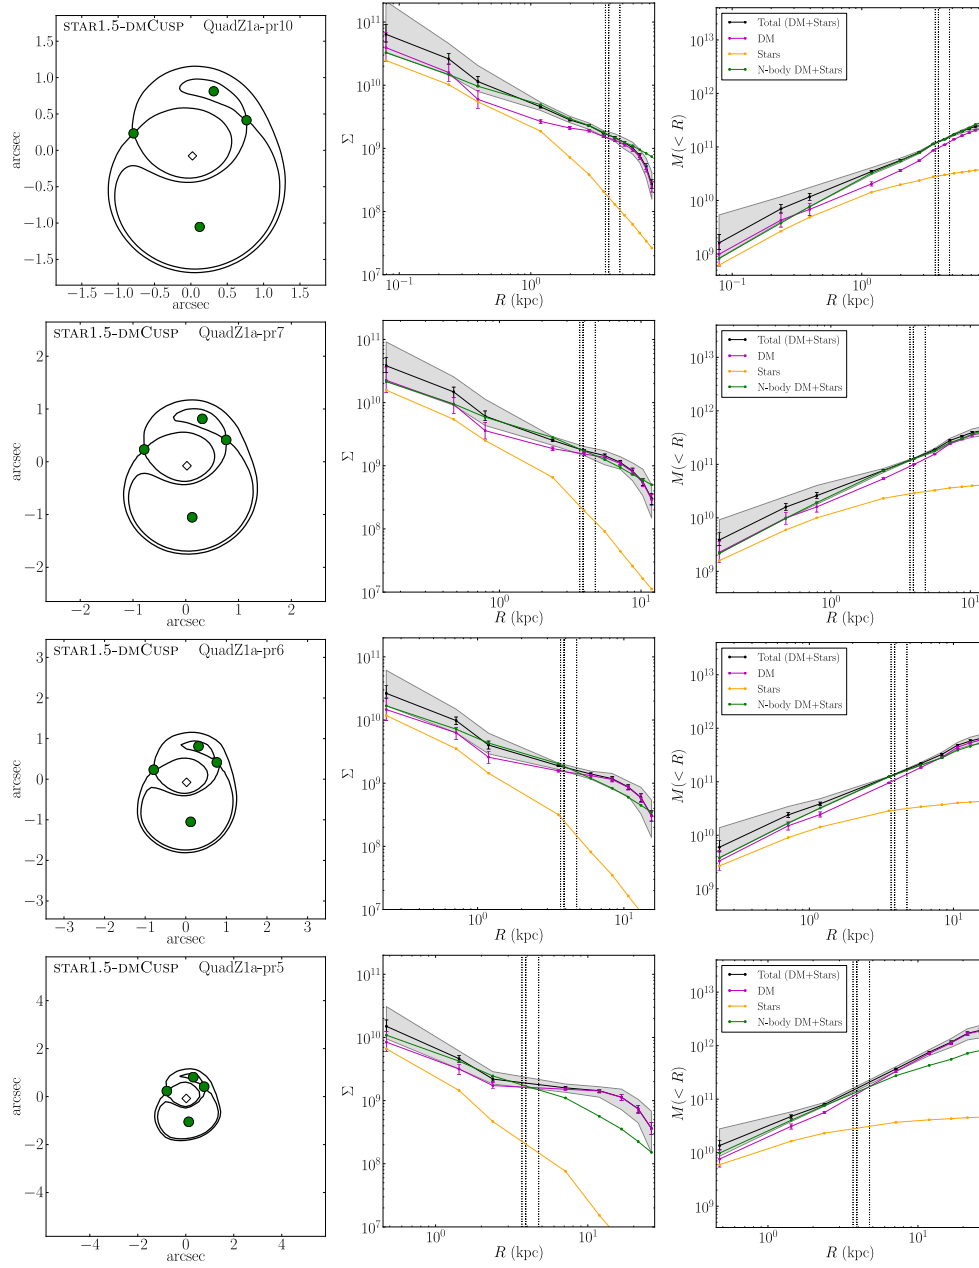
Figure D1. The effect of changing the grid resolution parameter `pixrad`. From top to bottom, the panels show results for a single quad with time delays and with stellar-mass constraints using `pixrad` = 10, 7, 6, 5, respectively. We always demand that there are four radial bins outside the outermost image, which causes the total mass to increase with decreasing `pixrad`, and the plot to shrink with increase in pixel size. In this paper we have used `pixrad` = 8 in all the tests.

This paper has been typeset from a TEX/LATEX file prepared by the author.