

A forward-modelling method to infer the dark matter particle mass from strong gravitational lenses

Qiuhan He ¹, ¹★ Andrew Robertson ¹, James Nightingale ¹, Shaun Cole,¹ Carlos S. Frenk,¹ Richard Massey ¹, Aristeidis Amvrosiadis,¹ Ran Li,^{2,3} Xiaoyue Cao^{2,3} and Amy Etherington¹

¹*Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*

²*National Astronomical Observatories, Chinese Academy of Sciences, 20A Datun Road, Chaoyang District, Beijing 100012, China*

³*School of Astronomy and Space Science, University of Chinese Academy of Sciences, Beijing 100049, China*

Accepted 2022 January 16. Received 2022 January 16; in original form 2020 October 25

ABSTRACT

A fundamental prediction of the cold dark matter (CDM) model of structure formation is the existence of a vast population of dark matter haloes extending to subsolar masses. By contrast, other dark matter models, such as a warm thermal relic (WDM), predict a cutoff in the mass function at a mass which, for popular models, lies approximately between 10^7 and $10^{10} M_{\odot}$. We use mock observations to demonstrate the viability of a forward modelling approach to extract information about low-mass dark haloes lying along the line of sight to galaxy–galaxy strong lenses. This can be used to constrain the mass of a thermal relic dark matter particle, m_{DM} . With 50 strong lenses at *Hubble Space Telescope* resolution and a maximum pixel signal-to-noise ratio of ~ 50 , the expected median 2σ constraint for a CDM-like model (with a halo mass cutoff at $10^7 M_{\odot}$) is $m_{\text{DM}} > 4.10$ keV (50 per cent chance of constraining m_{DM} to be better than 4.10 keV). If, however, the dark matter is a warm particle of $m_{\text{DM}} = 2.2$ keV, our ‘approximate Bayesian computation’ method would result in a median estimate of m_{DM} between 1.43 and 3.21 keV. Our method can be extended to the large samples of strong lenses that will be observed by future telescopes and could potentially rule out the standard CDM model of cosmogony. To aid future survey design, we quantify how these constraints will depend on data quality (spatial resolution and integration time) as well as on the lensing geometry (source and lens redshifts).

Key words: gravitational lensing: strong – methods: statistical – dark matter.

1 INTRODUCTION

Constraining the identity and mass of the dark matter particle is one of the most challenging goals in physics today. Astrophysics can play a major role in this endeavour. Cosmological N -body simulations have shown how dark matter haloes form from initial density fluctuations in the Universe (Frenk et al. 1985). The halo mass function (i.e. the abundance as a function of mass) and internal density profiles are determined by the power spectrum of initial fluctuations, which itself depends on the nature of the dark matter particle.

In the case of cold dark matter (CDM), the halo mass function extends as a power law to very low masses (about Earth mass for a 100-GeV particle) (Wang et al. 2020). This is perhaps the most fundamental prediction of the standard model of cosmology, Λ CDM. By contrast, if the dark matter is ‘warm’ (WDM), the thermal velocities of the particles at early times damp small-mass fluctuations leading to a cutoff in the halo mass function at a scale that depends on the inverse of the particle mass; for particles of keV mass, this is a few times $10^8 M_{\odot}$. Thus, a measurement of the small-mass end of the halo mass function would set strong constraints on the properties of the dark matter particle, including its mass, and could, in principle, rule out some of the main current candidates, including CDM.

However, measuring the properties of dark matter haloes is not trivial. In sufficiently massive haloes, gas can flow to the centre

and give birth to a luminous galaxy, thus, in principle, enabling direct number counts. The observed clustering of luminous galaxies provided the first astrophysical constraints on the mass of dark matter particles, ruling out light particle candidates called ‘hot dark matter’, such as light neutrinos, as the predominant form of dark matter (Frenk, White & Davis 1983; White, Frenk & Davis 1983; see Frenk & White 2012 for a review). However, these observations are unable to constrain WDM particles, which would require measurements of haloes with masses less than $\sim 10^{10} M_{\odot}$, that are too small to have made a luminous galaxy (Efstathiou 1992; Thoul & Weinberg 1996; Benson et al. 2002; Sawala et al. 2016; Benitez-Llambay & Frenk 2020). Furthermore, baryonic physics can have a significant effect on dark haloes in this mass range (particularly for those living within a galaxy) adding further complexity to the task. Strong gravitational lensing is an observable that can, in principle, be used to detect dark haloes and constrain the dark matter particle mass. Other observables include the statistics of the Lyman- α forest (Viel et al. 2013; Iršič et al. 2017; Garzilli et al. 2019), features in stellar cluster tidal debris in the Milky Way (Bonaca et al. 2019) and populations of Milky Way satellites (Lovell et al. 2016; Schneider 2016; Lovell et al. 2017; Nadler et al. 2021; Newton et al. 2021).

Strong gravitational lensing is the phenomenon whereby light rays from a distant galaxy are deflected by a foreground object, resulting in the formation of large luminous arcs or multiple distinct galaxy images. In principle, any object that happens to lie near the light rays coming from the background source galaxy will leave an imprint on the image, such that small perturbations (or lack thereof) to the

* E-mail: qiuhan.he@durham.ac.uk

lensed source light can be used to infer the presence (or absence) of low-mass dark matter haloes. Depending on the source size and the kind of data available, two broad types of analysis are possible.

If the source is an extremely compact object, like a quasar, when lensed, one will observe multiple point sources. Low-mass ‘invisible’ dark matter haloes along the line of sight induce anomalies in the flux ratios among multiple images, and by analysing those anomalies one can constrain the dark matter particle mass (Mao & Schneider 1998; Metcalf & Madau 2001; Dalal & Kochanek 2002; Amara et al. 2006; Xu et al. 2012; Nierenberg et al. 2014; Xu et al. 2015; Nierenberg et al. 2017; Gilman et al. 2019; Hsueh et al. 2020; Gilman et al. 2020a). However, because of the limited amount of information contained in the fluxes and positions of multiple images of a point-like source, there are degeneracies between the smooth model assumed for the lens and the inferred number of low-mass perturbing haloes.

If the source is extended (e.g. an ordinary galaxy), it appears after lensing as multiple arcs or a complete Einstein ring. Sufficiently massive dark matter haloes along the line of sight perturb local regions of the arc in a way that is measurable using techniques such as the pixelized potential correction method (Vegetti et al. 2012, 2014, 2018; Ritondale et al. 2019). By combining results for individual haloes in many systems, one can draw conclusions regarding the halo mass function (e.g. Vegetti & Koopmans 2009; Li et al. 2016). However, this explicit subhalo modelling has demanding requirements on image quality and resolution and so far only a few halo detections have been reported with masses below $10^{10} M_{\odot}$ (Vegetti et al. 2010, 2012; Hezaveh et al. 2016).

The detection of individual haloes or subhaloes requires their mass and position to be such that they produce an observable perturbation to the lensed image of the source. The resolution and signal to noise of the data thus define the ‘sensitivity’ to halo detection. Haloes of mass $\sim 10^8 M_{\odot}$ are detectable with Keck AO data (Vegetti et al. 2012) while haloes as light as $10^6 M_{\odot}$ are detectable with radio data (McKean et al. 2015). Recent studies have proposed that lower masses may be accessible by instead extracting, statistically, the cumulative perturbation of all intervening haloes on strong lensing arcs, even though none of these haloes could be detected individually (Brewer, Huijser & Lewis 2016).

Most studies have focused on determining a theoretical relation between the power spectrum of the convergence field and that of the image residuals after fitting a smooth ‘macro model’ to the lens galaxy mass distribution (Chatterjee & Koopmans 2018; Díaz Rivero et al. 2018; Cyr-Racine, Keeton & Moustakas 2019). There is one example using real data, in which Bayer et al. (2018) analysed image residuals of the system SDSS J0252+0039, obtaining a detection of power much higher than the prediction of the CDM model. There remain a number of questions as to what assumptions are appropriate for this technique. For example, to simplify the theoretical calculations, a random Gaussian field is often assumed for the convergence, potentially omitting multiplane lensing effects (McCully et al. 2014; Schneider 2014). Indeed, recent studies have shown that the expected signal in a Λ CDM universe is dominated by line-of-sight haloes at different redshifts to the lens galaxy as opposed to subhaloes within the lens galaxy itself (Li et al. 2016, 2017; Despali et al. 2018; Çağan Şengül et al. 2020). Here, we will focus on how the line-of-sight low-mass dark matter haloes contribute to perturbations on lensing images.

To form a visual impression of how the line-of-sight low-mass dark matter haloes perturb the images, in Fig. 1 we show the effective convergence, as derived from the deflection angles (upper panels; Gilman et al. 2019; Çağan Şengül et al. 2020), and the corresponding

residuals obtained by fitting the smooth lensing model to the image (lower panels). As may be seen from the effective convergence map, even though all low-mass dark matter haloes are modelled as spherical profiles, some are heavily stretched into arc-like features due to multiplane lensing effects,¹ which are difficult to compute analytically. Comparing the left two columns, the ‘warm’ case in the middle and the ‘cold’ case in the left have very similar residuals, particularly the large patches, even though the warm case has far fewer low-mass dark matter haloes. This suggests that the residuals are dominated by massive dark matter haloes and to distinguish between different dark matter models, we need to identify the small patches. The system shown in the third column has a higher redshift source and thus many more low-mass dark matter haloes along the line of sight, which results in larger and more complicated residuals.

In this work, we investigate whether a forward modelling procedure built around the approximate Bayesian computation (ABC) statistical inference method can extract meaningful information about the properties of low-mass dark matter haloes perturbing strong lensing images. As opposed to determining a concise mathematical expression, our forward modelling method relies on robust modelling of the perturbations on lensing images induced by small dark haloes, which can directly build up a relation between the models of interest and the observations without any further assumptions. The approach was first applied to strong lensing by Birrer, Amara & Refregier (2017b), who analysed the strongly lensed quasar RXJ1131-1231. Subsequently, Gilman et al. (2019) used this approach to study flux ratio anomalies of point-like sources with a full realization of dark matter haloes including line-of-sight dark haloes.

We apply the ABC framework to a simulated sample of 50 *HST*-resolution strongly lensed extended sources, which is a comparable number of lenses to the high-quality strong lensing SLACS sample (Bolton et al. 2006, but the details of our sample are different from those of SLACS). Each simulation contains a full cosmological realization of small dark haloes whose redshift distribution requires analysis using multiplane ray tracing. To apply the ABC method, we simulate and refit each of these 50 lenses 20 000 times, producing a total of 1 million lens models. The scale of this analysis necessitates an automated framework for lens modelling for which we use the open source software PyAutoLens.² Our aim is to determine whether the cumulative distortions due to the many dark matter haloes perturbing the light of the lensing arcs can be extracted to determine the halo mass function and hence the dark matter particle mass.

Our main goal is to demonstrate that this signal is present in *HST* imaging of strong lenses and that it can, in principle, be extracted using modern lens modelling techniques, given realistic levels of noise. However, our study is based on idealized systems: we make a number of simplifying assumptions for the structure of the lens and source and neglect effects such as an imperfect point spread function (PSF) model, correlated noise, or inadequate lens light subtraction. These assumptions will need to be relaxed before our methodology can reliably be applied to real data, including, for example, a non-parametric source model (e.g. Warren & Dye 2003), additional complexity in the lens model (Vegetti & Koopmans 2009; Nightingale et al. 2019), and a proper treatment of the PSF and of correlated noise.

¹In strong lensing, as viewed from the image plane, not only is the distant light warped by foreground objects, like a luminous source galaxy deflected by a lens, but also the mass on distant planes is warped by near planes, and thus we see some arc-like low-mass dark matter haloes in these convergence maps even though they are all spherical.

²<https://github.com/Jammy2211/PyAutoLens>

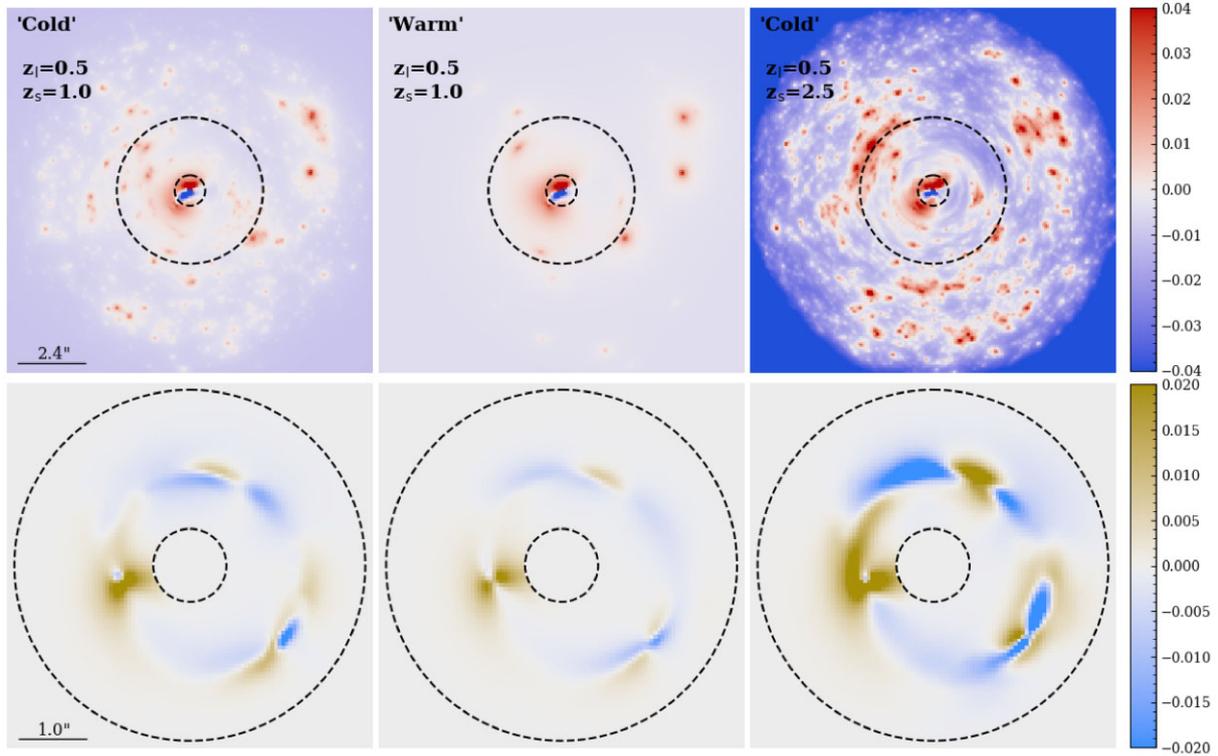


Figure 1. Upper panels: the effective convergence of low-mass perturburers. Lower panels: the corresponding best-fitting image residuals. The colour scale is in units of $e^- \text{pix}^{-1} \text{s}^{-1}$. The parameters of the lens and source galaxies shown in the three columns are the same as those used in Fig. 4, except for the source redshifts. The lensing systems in the left and middle columns have a source at $z = 1$, while that in the right column has a source at $z = 2.5$. The systems in both the left and right columns have a cutoff in the mass function at $10^7 M_\odot$, while the one in middle column has the cutoff at $10^9 M_\odot$. The inner and outer dashed circles in each panel have radii $0.5''$ and $2.4''$, respectively.

We first introduce the forward modelling procedure and lensing models in Section 2. In Section 3, we show tests of the accuracy of this method, the dependency on different lensing and observational settings, and compare our method to other methods, discussing its possible future applications and shortcomings. Finally, we conclude in Section 4. Throughout the paper, we adopt the cosmological parameters given by WMAP9 (Hinshaw et al. 2013).

2 PROCEDURE AND MODELS

In this section, we first provide an overview of our forward modelling procedure. We then describe the parametric models we use for the mass distributions of the main lens, dark matter haloes, and light distributions of the source galaxies. We then describe how we fit our simulated images with a combination of a smooth parametric lens and source model, and how the residuals of each are used within an ABC framework to place constraints on the mass function of dark matter haloes.

2.1 The forward modelling scheme

In Fig. 2, we provide an overview of the forward modelling procedure. Starting from an observed strong lensing image (which in this paper is simulated), we begin by fitting it with parametric lens and source models, omitting substructure from the lens model. This procedure gives us best-fitting smooth lens mass and source light models, as well as a map of the best-fitting image residuals (the observed image minus the best-fitting model image). In this work,

‘best-fitting’ refers to the maximum likelihood model determined by means of a non-linear search.

It is impossible to fit all the dark matter substructure in a similar way, because of its low mass and low signal to noise (Birrer et al. 2017b). However, we can use the best-fitting source and macroscopic lens model to simulate a set of images of this lens system, each including a random realization of dark matter substructure (we call this set of images the ‘forward models’). We refit each forward model in the same manner as we fitted the ‘observed’ image, providing best-fitting image residuals for each forward model. We compare the forward-modelled residuals with the observed residuals, as described in detail in Section 2.2, and apply ABC inference to obtain a constraint on the cutoff in the halo mass function.

2.2 Lensing simulations

The strong lens simulations used in this work represent the lensing system with three components: the lens galaxy, the source galaxy, and the line-of-sight dark matter haloes.

2.2.1 Lens and source

We simulate the smooth mass distribution of our lens galaxy as a two-dimensional singular isothermal ellipsoid (SIE) of the form,

$$\begin{aligned} \Sigma(x, y) &= \frac{\sigma^2}{2G} \frac{1}{\sqrt{x^2 q + y^2/q}} \\ &= \frac{c^2}{8\pi G} \frac{D_A(0, z_{\text{source}})}{D_A(z_{\text{lens}}, z_{\text{source}}) D_A(0, z_{\text{lens}})} \frac{R_E}{\sqrt{x^2 q + y^2/q}}, \quad (1) \end{aligned}$$

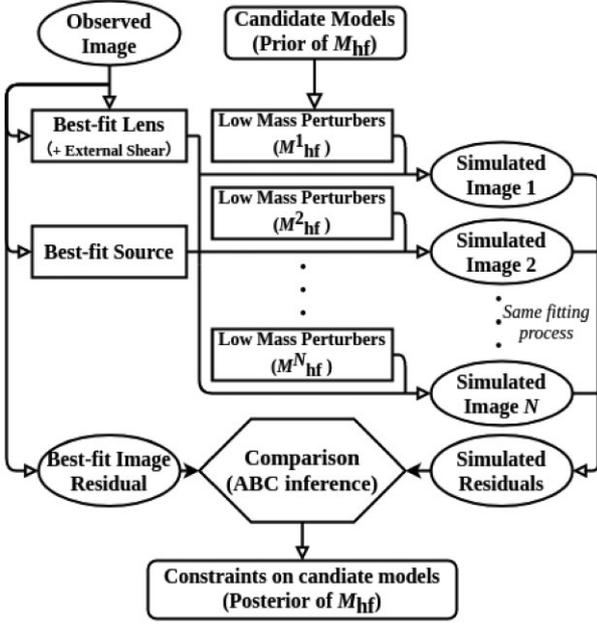


Figure 2. An overview of the forward modelling procedure for one lensing system. The observed image is fitted with a parametric lens mass model and a parametric source light model, producing a best-fitting lens model, source model, and image residuals. We use this best-fitting model to generate N -simulated images, including random realizations of low-mass perturbing haloes in the lens model. The number of perturbers as a function of mass depends on the properties of the dark matter as encoded in the ‘half-mode mass’, M_{hf} . The distribution of M_{hf} over the N ‘forward models’ follows our prior on M_{hf} . Each of the forward model images is fit in the same way as the observed image to produce N sets of image residuals. These forward model residuals are compared with the observed residuals, using ABC to constrain M_{hf} .

where R_E is the projected Einstein radius, which can be related to the velocity dispersion of the profile, σ , given the lens and source redshifts (z_{lens} and z_{source} , respectively) and a cosmology. The quantity, $D_A(z_1, z_2)$, is the angular diameter distance between z_1 and z_2 ; q is the axial ratio. For simplicity, we do not add any external shear to our mock observations. However, we include an external shear when modelling the lenses.

For the source, we adopt the Sérsic (1963) profile,

$$I(r) = I' \exp \left[-b_n \left(\frac{r}{r_c} \right)^{1/n} \right], \quad (2)$$

where I' is the scale intensity, n is the Sérsic index, and b_n is a function of n defined such that r_c is the half-light radius. However, the image residuals caused by perturbations to the lens model (in our case the perturbing haloes) are effectively a product of the deflection angles due to the perturbation and the gradient of the source surface brightness (Vegetti & Koopmans 2009). The cuspy centre of the Sérsic profile creates an infinite surface brightness gradient that numerically causes infinite surface brightness differences. This can be overcome by oversampling the source light but at the expense of increased computational run time. Instead, we use a Sérsic profile with a ‘core’ (Graham et al. 2003; Trujillo et al. 2004) to simulate the background source. This has the form,

$$I(r) = I' \exp \left[-b_n \left(\frac{r^\alpha + r_c^\alpha}{r_c^\alpha} \right)^{1/(n\alpha)} \right], \quad (3)$$

where r_c is the ‘core radius’, which marks the transition radius from the Sérsic profile to constant surface brightness. When $r_c \rightarrow 0$, the profile reduces to the Sérsic form. The parameter, α , quantifies how quickly the profile transitions from a regular Sérsic form to one with a constant surface brightness core. To simplify both our simulation and modelling processes, we fix $\alpha = 2$ and $r_c = 0.01''$. With these parameters fixed, the source model has the same three free parameters as the regular Sérsic profile: r_c , n , and I' .

2.2.2 Low-mass dark matter haloes

Low-mass dark matter haloes that can perturb lensing arcs can be of two types: subhaloes within the main lens, or central haloes at different redshifts that happen to lie close to the path of light from the source galaxy to the observer. We will refer to the latter as line-of-sight haloes.

The number and mass distribution of subhaloes within the lens are somewhat uncertain. Subhaloes are subject to tidal stripping and disruption; these effects are significantly enhanced in hydrodynamic simulations compared to dark matter-only simulations due to the presence of the dense stellar component at the centre of the main lens (Garrison-Kimmel et al. 2017; Sawala et al. 2017; Richings et al. 2021). The degree of disruption also depends on the details of the galaxy formation model (Richings et al. 2020). In contrast, line-of-sight haloes are not subject to these environmental effects and, in the mass range of interest here, they are entirely ‘dark’, having never formed stars.

Recent studies (Li et al. 2016, 2017; Despali & Vegetti 2017; Çağan Şengül et al. 2020) have shown that for the lens and source redshifts typical of SLACS lenses (Bolton et al. 2006), the lensing perturbations mainly arise from line-of-sight haloes. This is fortunate because the irrelevance of uncertain baryon effects makes line-of-sight haloes a particularly clean probe of a cutoff scale in the halo mass function. In this work, we focus on these low-mass, line-of-sight haloes whose mass function has the form given by Lovell et al. (2014):

$$\begin{aligned} F(M_{200}, z) &\equiv \frac{d^2 N}{dM_{200} dV}(M_{200}, z) \\ &= F_{\text{CDM}}(M_{200}, z) \left(1 + \frac{M_{\text{hf}}}{M_{200}} \right)^{-1.3}. \end{aligned} \quad (4)$$

This is composed of two parts, the CDM mass function, $F_{\text{CDM}}(M_{200}, z)$, for which we use the form derived by Sheth et al. (2001), and a cutoff parametrized by the ‘half-mode’ mass, M_{hf} , a characteristic mass corresponding to the mass scale at which the dark matter transfer function falls to half the CDM transfer function. The half-mode mass can be related to the mass of the dark matter particle, m_{DM} , as,

$$M_{\text{hf}} = 10^{10} \left(\frac{m_{\text{DM}}}{\text{keV}} \right)^{-3.33} M_\odot h^{-1}, \quad (5)$$

where h is the Hubble parameter in units of 100 km/s/Mpc (Lovell et al. 2014). From the perspective of strong lensing, M_{hf} may be regarded as an effective ‘cutoff’ mass in the dark matter mass function, below which there are very few haloes.

We draw low-mass dark haloes in a light-cone between the observer and the source galaxy within an angular radius, r_{lc} , chosen as,

$$r_{\text{lc}} = \begin{cases} 5.0'' & (z \leq z_{\text{lens}}) \\ 5.0'' - 2.0'' \times \frac{D_A(z_{\text{lens}}, z) D_A(0, z_{\text{source}})}{D_A(0, z) D_A(z_{\text{lens}}, z_{\text{source}})} & (z > z_{\text{lens}}) \end{cases}, \quad (6)$$

where D_A is the angular-diameter distance and z_{lens} and z_{source} are the redshifts of the lens and source galaxies, respectively. The angular radius of the light-cone is fixed at $5.0''$ (this is ~ 3 times the Einstein radius, and including haloes at larger radii hardly affects our lensing images) in front of the main lens and gradually decreases behind from $5.0''$ to $3.0''$ at the source redshift, which ensures that all perturbers that might influence the lensed image of the source are included. The light-cone is evenly divided into redshift bins, of width $\Delta z = 0.01$, which balances the accuracy and computational cost of the calculation (see Appendix A for details). Within each bin, we sample from the halo mass function, assign density profiles from the halo mass–concentration relation corresponding to that redshift, and place these perturbing haloes randomly on the central redshift plane. Neglecting the clustering of the line-of-sight haloes is a good approximation if, over the angular extent of interest, the typical line-of-sight separation of pairs of haloes is much larger than their correlation length. This is expected to be true because (a) the intrinsic clustering of low-mass haloes is weak (Kaiser 1984; Sheth & Tormen 1999), and (b) the angular extent of interest is very narrow since it is set by the width of the Einstein ring image.

We generate low-mass haloes with M_{200} between $10^6 M_\odot$ and $10^{10} M_\odot$ within the light-cone. Depending on the dark matter model, a significant fraction of the mass in the Universe can exist as collapsed objects in this mass range. To solve correctly the strong lensing equations (which have already assumed a background universe), we add negative mass sheets on to each plane to compensate for the otherwise ‘double-counted’ mass within perturbers in the light-cone (Petkova, Metcalf & Giocoli 2014; Birrer et al. 2017a; Gilman et al. 2019). For each plane, we include a mass sheet with constant negative surface density equal to the expected density in low-mass haloes on that plane. This leads to a mass sheet with a surface density

$$\Sigma_{\text{negative}} = -\frac{1}{A} \int F(M_{200}, z) M(M_{200}, z, s) dM_{200} ds dV, \quad (7)$$

where A is the physical area of the light-cone at the plane’s redshift, $F(M_{200}, z)$ is the mass function, and $M(M_{200}, z, s)$ is the total mass of a dark halo on the plane, which can differ from M_{200} due to halo truncation. The ‘ s ’ denotes some other parameters of the dark matter profile; in our case, as we discuss below, it represents the ‘concentration’ and ‘truncation radius’.

Instead of modelling haloes as Navarro–Frenk–White (NFW) profiles (Navarro, Frenk & White 1996, 1997), which have an infinite total mass (when integrating to infinite radius), we use a truncated NFW profile (Baltz, Marshall & Oguri 2009),

$$\rho(r) = \frac{M_0}{4\pi r (r_s + r)^2} \cdot \frac{r_t^2}{r_t^2 + r^2}, \quad (8)$$

which has an additional truncation term compared with the standard NFW form. Here, the scale mass, M_0 , and scale radius, r_s , are a pair of parameters that specify an NFW profile and r_t is the truncation radius. With the addition of the truncation term, $\rho(r) \propto r^{-5}$ at large radii, and the profile now has a finite mass given by

$$M_t = M_0 \frac{\tau^2}{(\tau^2 + 1)^2} [(\tau^2 - 1)\ln\tau + \tau\pi - (\tau^2 + 1)], \quad (9)$$

where $\tau = r_t/r_s$. We set the truncation radius to be the corresponding NFW r_{100} , the radius within which the average density of an NFW profile is 100 times the critical density of the universe. With τ fixed to r_{100}/r_s , the tNFW profiles are determined by the same two parameters as the NFW profile, M_{200} and the concentration, $c = r_{200}/r_s$, which quantifies how centrally concentrated a halo is. Haloes that form earlier typically have higher concentrations (Navarro et al.

1996, 1997). Since the reduction in small-scale power in a WDM model leads to later formation time for low-mass haloes, these are less concentrated than their CDM counterparts (Lovell et al. 2012; Schneider et al. 2012; Macciò et al. 2013; Bose et al. 2016; Ludlow et al. 2016).

When considering WDM models, we adopt the mass–concentration relation of Bose et al. (2016)

$$c(M_{200}, z) = c_{\text{CDM}}(M_{200}, z) \times \left((1+z)^{0.026z-0.04} \left(1 + 60 \frac{M_{\text{hf}}}{M_{200}} \right)^{-0.17} \right), \quad (10)$$

where $c_{\text{CDM}}(m, z)$ is the mass–concentration relation in the CDM case. For CDM, we use the mass–concentration relation of Ludlow et al. (2016), which agrees with recent simulation results from Wang et al. (2020), who measured halo concentrations in Λ CDM simulations over 20 orders of magnitude in halo mass. We use the public package COLOSSUS (Diemer 2018a) to compute the mass–concentration relation. The Ludlow et al. (2016) relation is for the median concentration at a given mass, but simulations predict scatter in concentration at fixed mass, which is well described by a lognormal distribution (Neto et al. 2007). As a result, we include lognormal scatter in the mass–concentration relation, with a standard deviation of 0.15 dex, as shown by Wang et al. (2020).

Multiplane ray tracing is necessary to calculate the effect of line-of-sight haloes on flux ratio anomalies with sufficient precision to match observations (Gilman et al. 2019). We use the publicly available software package PYAUTOLENS (Nightingale & Dye 2015; Nightingale, Dye & Massey 2018; Nightingale et al. 2019; Nightingale et al. 2021b) to simulate lenses using the light and mass profiles described above. To define our fiducial lens sample, we simulate 50 unique lensing systems with different values for the lens SIE mass model (e.g. Einstein radius, axial ratio, etc.) and source core-Sersic profile. The input M_{hf} is $10^7 M_\odot$. For all 50 systems, we assume lens redshifts of 0.5 and source redshifts of 1.

We produce observations representative of *Hubble Space Telescope* imaging, with a pixel size of 0.05 arcsec, a Gaussian PSF with a standard deviation of 0.05 arcsec, and background sky level of $1e^-/\text{pix/s}$. We set the exposure time to be 600 s. The source surface brightness of each simulation is chosen such that the maximum signal-to-noise ratio (e.g. in the brightest pixel) has a value of ~ 50 . Table 1 summarizes the models and parameters for our mock observations, where closed brackets indicate the range of values from which each parameter is randomly drawn. For convenience, we use $\{T\}_{\text{fiducial}}$ to denote the set of true parameter values used for our fiducial tests and a superscript i to represent the parameters of the i -th lensing system. Mock images of our 50 strong lensing systems are shown in Fig. 3.

2.3 Analysis process

Our analysis of each mock observation consists of two steps: smooth model fitting, followed by ABC inference. Below we describe the detailed steps in this process and illustrate them for one system.

2.3.1 Smooth model fitting

First, we fit a simple lens model to the mock lensing image. This consists of a smooth parametric model plus an external shear for the lens galaxy mass distribution (the *macro model*) and a smooth parametric model for the source. We adopt the same parametric

Table 1. Fiducial model parameters. Closed brackets indicate the range from which values are randomly drawn.

Parameter	Range or value
Lens	
Centre (x, y) ["]	Elliptical isothermal (0.0, 0.0)
R_E ["]	[1.2, 1.6]
Axial ratio	[0.6, 0.95]
Position angle [°]	30
Redshift	0.5
External shear	
Magnitude A	0
Position angle [°]	0
Source	
Centre (x, y) ["]	Elliptical core-Sersic ([−0.3, 0.3], [−0.3, 0.3])
r_e ["]	[0.1, 0.5]
n	[1.5, 2.5]
$\dot{\Gamma}$ [$e^- \text{pix}^{-1} \text{s}^{-1}$]	7.4
Axial ratio	[0.5, 0.95]
Position angle [°]	[10, 100]
r_c ["]	0.01
α	2.0
γ	0.0
Redshift	1.0
Dark haloes	
F_{CDM}	Truncated NFW profile Sheth, Mo & Tormen (2001)
c_{CDM}	Ludlow et al. (2016)
$\sigma_{\log_{10} c_{\text{CDM}}}$	0.15 dex
$t_{\text{NFW}} \tau$	r_{100}/r_s
$M_{\text{hf}} [M_{\odot}]$	10^7
Image	
Pixel size ["]	0.05
PSF σ ["]	0.05
t_{exp} [s]	600
Background sky [$e^- \text{pix}^{-1} \text{s}^{-1}$]	1.0
max S/N of pixels	~50

forms we used when generating the simulated mock image, an SIE lens mass model, and a core-Sersic source.

We use PYAUTOLENS to fit the mock images and adopt PYMULTINEST with constant sampling efficiency mode switched on (Feroz, Hobson & Bridges 2009a; Buchner et al. 2014a) for the non-linear optimization. Table 2 shows the inferred parameters for this fit, alongside the true input values for the first mock image. To distinguish between the true inputs and the best-fitting values, we label the latter as $\{F\}_{\text{fiducial}}$, where a superscript, i , denotes the best-fitting parameters for the i -th lensing system. This table shows that the inferred and input lens model is close but not in perfect agreement; this is due to the effects of the low-mass dark matter haloes. We will show later that this small mismatch does not affect our calculation.

In the upper and lower panels of Fig. 4, we show the ‘observed’ image and best-fitting residuals for one of our 50 mock observations. To develop intuition on how the low-mass perturbers affect the image residuals, we further show the ‘effective convergence’ (the divergence of deflection angles) of low-mass perturbers and the corresponding best-fitting residuals in Fig. 1. We assume a very small constant noise to show the residuals clearly. The three strong lensing systems shown have the same lens and source galaxy as in Fig. 4, except that the source for the system of the right-hand column is at $z = 2.5$. For a clear comparison, we make sure that

the three systems share the same realization of low-mass perturbers. To achieve this, instead of applying the mass function described by M_{hf} in equation (4), we draw haloes from the mass function of Sheth et al. (2001) but discard those with M_{200} below a certain value. Furthermore, we assume no correlation between M_{hf} and the mass–concentration relation. Please note that this change of low-mass perturber realization is implemented only for Fig. 1; in the rest of this paper, we use the mass function and mass–concentration relation described by equations (4) and (10), respectively.

In Fig. 1, the systems in the left and right columns have a cutoff in the mass function at $10^7 M_{\odot}$, while the system in the middle column has a cutoff at $10^9 M_{\odot}$. As seen from the left two columns, the system with a lower cutoff in the mass function has far more low-mass perturbers and thus more subtle structure in the residuals. When comparing the left and the right columns, we see that the number of low-mass perturbers increases significantly with the source redshift and the ‘profiles’ of individual perturbers are more heavily distorted because of the multiplane lensing effect. Also, for a higher redshift source galaxy, the residuals are larger.

2.3.2 ABC inference through forward modelling

Approximate Bayesian computation is a likelihood-free method suitable for problems where the likelihood is difficult to express analytically, but model predictions are relatively easy to simulate. It has been widely applied in astrophysics, for example in studies of large-scale structure, planet surveys, and reionization (Akeret et al. 2015; Hahn et al. 2017; Davies et al. 2018; Hsu, Ford & Terrien 2020). Birrer et al. (2017b), Gilman et al. (2019), Gilman et al. (2020a), Gilman et al. (2020b), and Enzi et al. (2020) have also used ABC to constrain dark matter substructure using strong lensing data. To begin, one defines a summary statistic that measures the similarity between the observational data and the simulations. In principle, any statistic (as long as it contains information of interest) can be used as a measure of similarity, but different statistics may have different requirements on data quality, sampling efficiency, etc. A good statistic captures as much of the characteristic information of different models as possible: in our case, the lensing perturbations from low-mass dark matter haloes. Finding a proper summary statistic is a challenging part of ABC and it is beyond the scope of this work to investigate the many possible candidate summary statistics. Instead, following Bayer et al. (2018), we use the power spectrum of the best-fitting image residuals to construct a summary statistic that we hope can extract information on the number of low-mass dark matter haloes perturbing the lens.

Based on the best-fitting source and macro model for the lens, $\{F\}_{\text{fiducial}}^1$, we simulate images of this best-fitting source and lens combination with the addition of low-mass haloes along the line of sight. We uniformly sample $\log_{10} M_{\text{hf}}/M_{\odot}$ between 7 and 10 (corresponding to assuming a flat prior on $\log_{10} M_{\text{hf}}$ between $10^7 M_{\odot}$ and $10^{10} M_{\odot}$) and, for each sampled M_{hf} value, we draw a random realization of low-mass dark matter haloes with the appropriate mass function and mass–concentration relation. We do this for 20 000 values of M_{hf} , producing a corresponding lensed image in each case and then fit each of these 20 000 images to obtain the best-fitting image residuals. We have confirmed that our results are converged with 20 000 samples. Note that the forward-modelled images are simulated from the known lens and source galaxies, so when performing the fitting, we add narrow priors centred on the parameter values of the known best-fitting lens and source models to accelerate the fitting processes. Since our only goal is to find

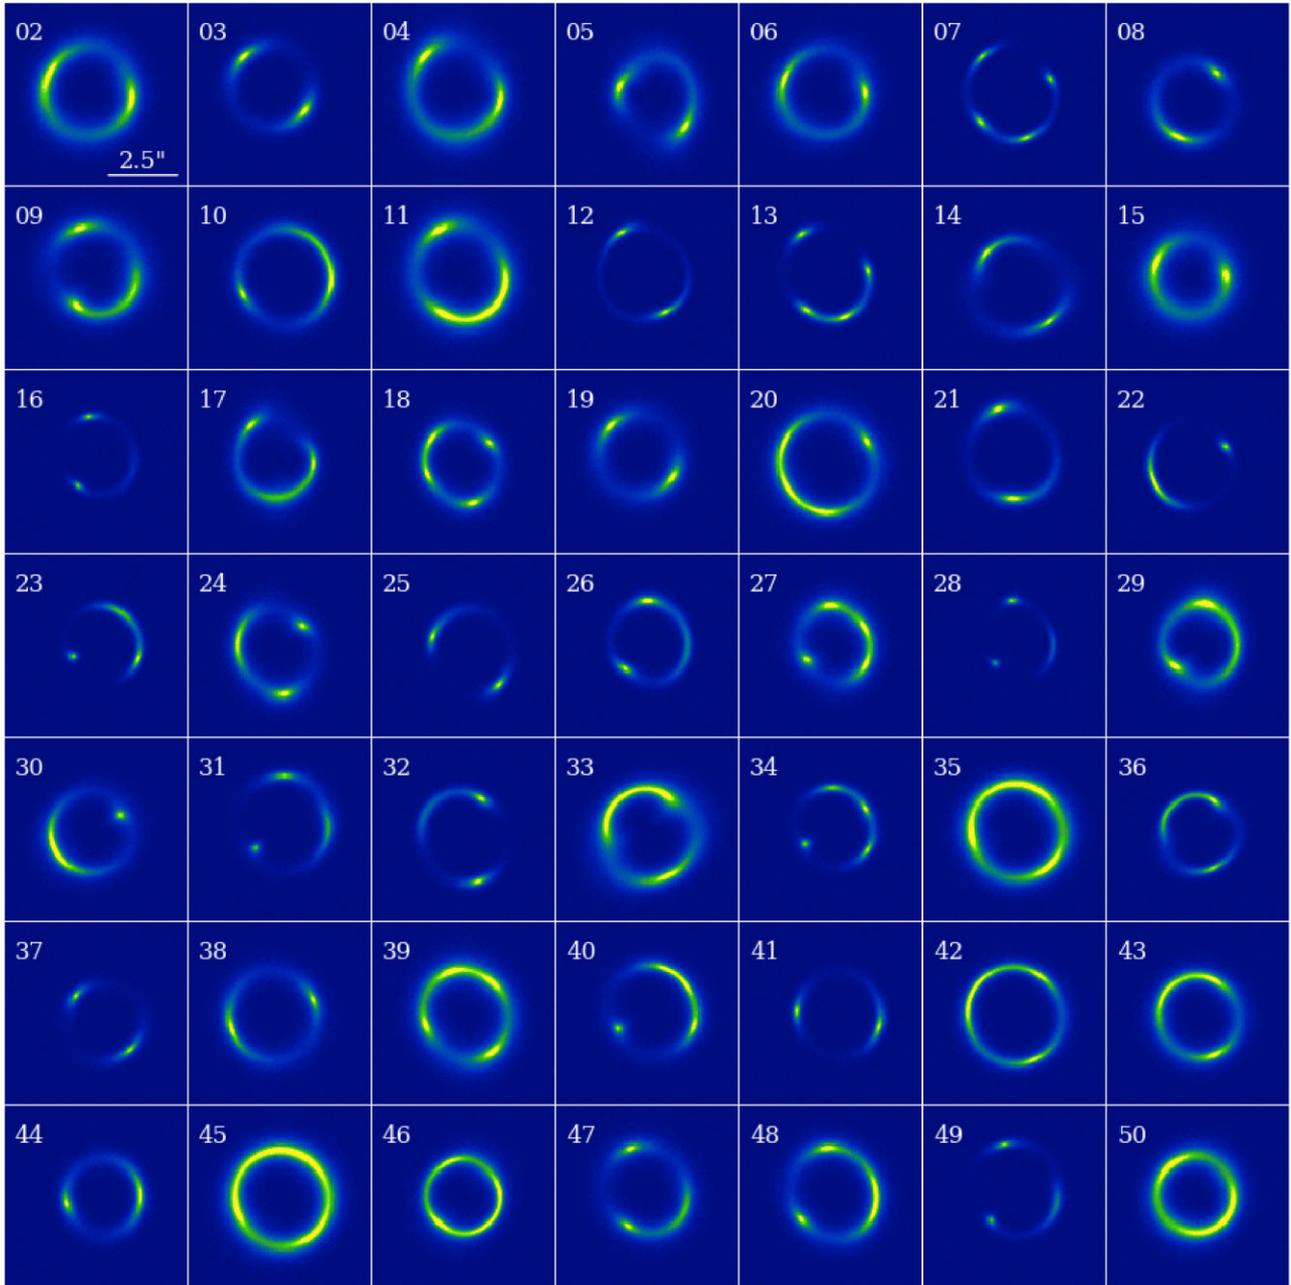


Figure 3. Fiducial mock strong lensing images. A total of 49 mock images are shown, with the last of our set of 50 shown in the upper panel of Fig. 4. The images are simulated with parameters randomly drawn according to Table 1.

the maximum likelihood model, once our priors contain the correct solution, the size of the priors does not affect our results. In Table 3, we summarize the type and size of the priors we used for the fitting processes. To make sure that the priors we take are broad enough to find the maximum likelihood model, we fit 4000 different realizations with both the listed priors and the five times larger priors and then compare the maximum likelihoods obtained. Here, ‘five times larger priors’ means that for the priors listed in Table 3, if it is a Gaussian prior, then the sigma is taken five times larger than the value listed. In Fig. 5, we show the histogram of the difference of the log maximum likelihood. As shown, the log maximum likelihood results of larger and smaller priors are very close with only ~ 1.0 difference, which

verifies that the priors we take are broad enough to contain the maximum likelihood model.

Next, for the images of the best-fitting residuals, we set pixels outside the annular region between $0.5''$ and $2.4''$ to be 0 and then Fourier transform the residual image and azimuthally average to obtain the one-dimensional (1D) power spectrum, $P(k)$. In the upper panel of Fig. 6, we show power spectra for the forward-modelled residuals as black lines (for clarity, we plot only 100 of them); the red curve marks the power spectrum of the image residuals of the original observation. Each line is composed of 15 $P(k_i)$ values with the k_i logarithmically spaced between $\sim 1.5 \text{ arcsec}^{-1}$ and 85 arcsec^{-1} [$\sim 2\pi / (4.2 \text{ arcsec})$ to $2\pi / (0.07 \text{ arcsec})$].

Table 2. The true input and best-fitting parameters for the first lensing system. The errors listed in this table are $3\text{-}\sigma$ limits. Notice that we fit a system with additional low-mass perturbers, and thus the inconsistency of the best-fitting parameters and the true values is expected. Parameters fixed during fitting are not shown here.

	$\{T\}_{\text{fiducial}}^1$	$\{F\}_{\text{fiducial}}^1$
Lens		
Elliptical isothermal		
Centre (x, y) ["]	(0.0, 0.0)	$(-0.001^{+0.002}_{-0.002}, -0.001^{+0.003}_{-0.003})$
R_E ["]	1.342	$1.342^{+0.002}_{-0.003}$
Axial ratio	0.618	$0.617^{+0.005}_{-0.005}$
Position angle [°]	30.0	$30.0^{+0.4}_{-0.4}$
External shear		
Magnitude	0	$0.002^{+0.002}_{-0.002}$
Position angle [°]	0	159^{+37}_{-41}
Source		
Elliptical core-Sersic		
Centre (x, y) ["]	(-0.072, 0.259)	$(0.071^{+0.001}_{-0.001}, 0.253^{+0.001}_{-0.001})$
r_e ["]	0.244	$0.240^{+0.005}_{-0.004}$
n	1.85	$1.81^{+0.05}_{-0.05}$
I' [$e^- \text{pix}^{-1} \text{s}^{-1}$]	7.4	$7.3^{+0.2}_{-0.2}$
Axial ratio	0.85	$0.85^{+0.01}_{-0.01}$
Position angle [°]	41	40^{+2}_{-2}

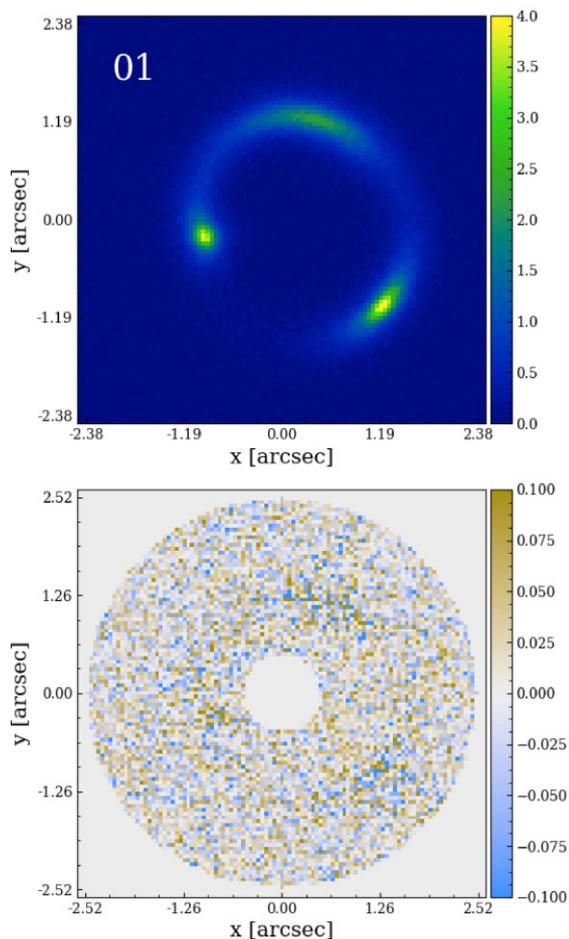


Figure 4. Upper panel: the mock image of the first system. Lower panel: the corresponding residuals after fitting a smooth macro model.

Table 3. Priors used in fitting processes. $G(a, b)$ represents a Gaussian prior centring on a with standard deviation b . $U(a, b)$ represents a uniform prior between a and b . The ‘input’ here refers to the corresponding input value used when simulating the mock image.

	Prior
Lens	
Centre (x, y) ["]	G (input, 0.01)
R_E ["]	G (input, 0.008)
Axial ratio	G (input, 0.03)
Position angle [°]	G (input, 2.0)
External shear	
Magnitude	G (0.0, 0.05)
Position angle [°]	U(0.0, 180.0)
Source	
Centre (x, y) ["]	G (input, 0.01)
r_e ["]	G (input, $0.015 \times \text{input}$)
n	G (input, 0.1)
I' [$e^- \text{pix}^{-1} \text{s}^{-1}$]	G (input, $0.16 \times \text{input}$)
Axial ratio	G (input, 0.02)
Position angle [°]	G (input, 2.0)

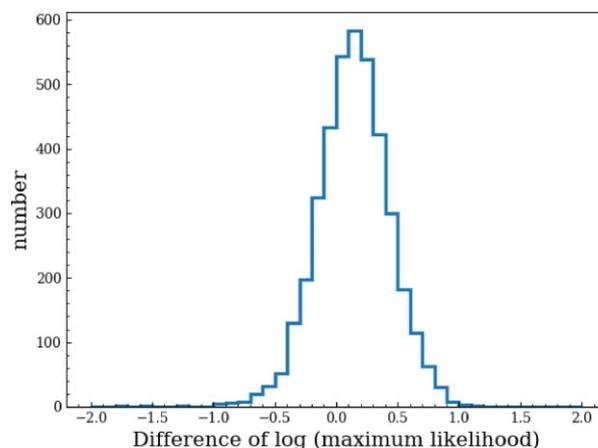


Figure 5. Histogram of the differences between the log maximum likelihood obtained with the priors listed in Table 3 and that obtained with five times larger priors.

Although we have converted the images into power spectra, we cannot simply use the 1D ‘power spectrum’ as our statistic. The reason is that each power spectrum is composed of 15 different values and for ABC to converge at a point in a 15-dimensional space, a very small acceptance rate must be set, which, in turn, requires far more forward models. To further reduce the ‘dimensionality’ of the statistic, following Fearnhead & Prangle (2012), we generate a summary statistic from a linear combination of the logarithm of the power spectrum values at different k ,

$$\log_{10} \hat{M}_{\text{hf}} = \sum_{i=1}^{15} \beta_i \log_{10} [P(k_i)/\text{arcsec}^2] + \beta_0. \quad (11)$$

We use $\log_{10} \hat{M}_{\text{hf}}$ as the summary statistic for the ABC inference. The coefficients $\{\beta_i\}$ (including β_0) are obtained by minimizing the differences between the input $\log_{10} M_{\text{hf}}$ values used in the forward models and their corresponding $\log_{10} \hat{M}_{\text{hf}}$. Specifically, we find the

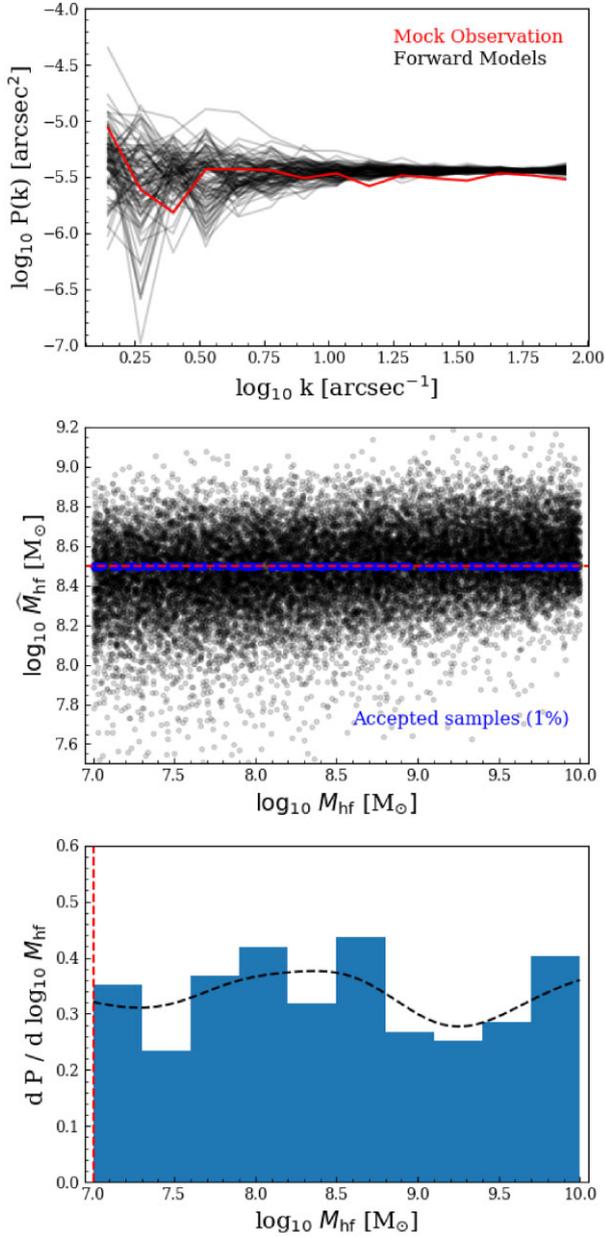


Figure 6. **Upper panel:** the power spectrum computed from the best-fitting image residuals. The red line is the power spectrum of the mock observation. Black lines are for 100 forward models. **Middle panel:** the summary statistic, $\log_{10} \hat{M}_{\text{hf}}$, computed from the power spectra. The black points are $\log_{10} M_{\text{hf}}$ for the forward models. The horizontal dashed red line marks the value for the mock observation and points 1 per cent closest to it are marked in blue. **Lower panel:** The posterior distribution of M_{hf} , formed by collecting together the M_{hf} of the blue points, is shown in the middle panel. The vertical dashed red line marks the true input of M_{hf} . The black dashed curve is a kernel density estimate corresponding to the histogram.

values of $\{\beta_i\}$ that minimize,

$$\sum_{j=1}^{20000} \left\| \log_{10} \hat{M}_{\text{hf}}^{(j)} - \log_{10} M_{\text{hf}}^{(j)} \right\|^2, \quad (12)$$

where $\log_{10} \hat{M}_{\text{hf}}^{(j)}$ is computed from the j -th forward-modelled power spectrum and $M_{\text{hf}}^{(j)}$ is the corresponding input half-mode mass. Since the forward-modelled power spectrum is systematically different for

different macro lensing configurations, we optimize $\{\beta_i\}$ separately for each of our 50 lensing systems. In the middle panel of Fig. 6, we plot the value of $\log_{10} \hat{M}_{\text{hf}}$ for each forward-modelled simulation as a black point. The horizontal dashed red line marks the value of $\log_{10} \hat{M}_{\text{hf}}$ given by the mock observation.

We accept the forward-modelled simulations with the 1 per cent of $\log_{10} \hat{M}_{\text{hf}}$ values closest to the observed $\log_{10} M_{\text{hf}}$ (equation 11), which are shown as blue points in the middle panel of Fig. 6. The set of M_{hf} values associated with those blue points is then a sample drawn from the posterior distribution of M_{hf} (following the ABC method); their density can be used to estimate the posterior density of M_{hf} .

In the lower panel of Fig. 6, we plot the posterior density for $\log_{10} M_{\text{hf}}$, where the vertical dashed red line shows the true input, $M_{\text{hf}} = 10^7 M_{\odot}$. It is clear that with just one system and the data quality of our fiducial setup, we are not able to derive a tight constraint on M_{hf} . While improved data quality might improve this somewhat, a single lens system will always be limited by the stochastic nature of the distribution of the low-mass dark matter haloes we are trying to detect. Just because a particular value of M_{hf} allows perturbers of a given mass, this does not mean that there will happen to be one in a location where it produces detectable image residuals. As such, tight constraints will rely on combining results from multiple systems.

3 RESULTS AND DISCUSSION

In this section, we will first show that the forward modelling procedure described above can correctly recover the input value of M_{hf} when combining results from multiple lenses. Then we will explore how the precision of the constraints depends upon the lensing configuration and image quality of the lenses in an observed sample. We will then compare the strength of our method with the method based on quasar flux ratio anomalies.

3.1 Tests of the accuracy of the method

By repeating the procedure described above for each observed system, we can obtain a constraint from our mock observations of 50 lensing systems. Just as for a real set of 50 observed systems, the posterior probability density for M_{hf} will not necessarily peak at the true value. To assess whether our method is biased, we create 500 sets of 50 observed lenses. Running the full procedure described above 500 times, each time for 50 different lenses, would be prohibitively expensive because of the very large number of forward models this would require. We therefore use the same set of 50 macro lens and source model parameters ($\{T\}_{\text{fiducial}}$) for each of our 500 sets, only changing the realizations of low-mass dark matter haloes between the sets of 50 observed images. This allows us to reuse the same forward models for each of our 500 sets.

In principle, we should generate a new set of forward models for each of the 500 sets (together 25 000 lenses) because the particular realization of the dark matter haloes affects the best-fitting macro model, which is the model used to generate the forward model images. We will see later that even though there is a slight mismatch in best-fitting models for different realizations, we can still correctly recover the true M_{hf} . Similarly, to examine how the method behaves with different input (true) values for M_{hf} , we further simulate 500 sets of 50 observations with $M_{\text{hf}} = 10^8$ and $10^9 M_{\odot}$ and apply ABC inference to them, all using the same forward models previously simulated.

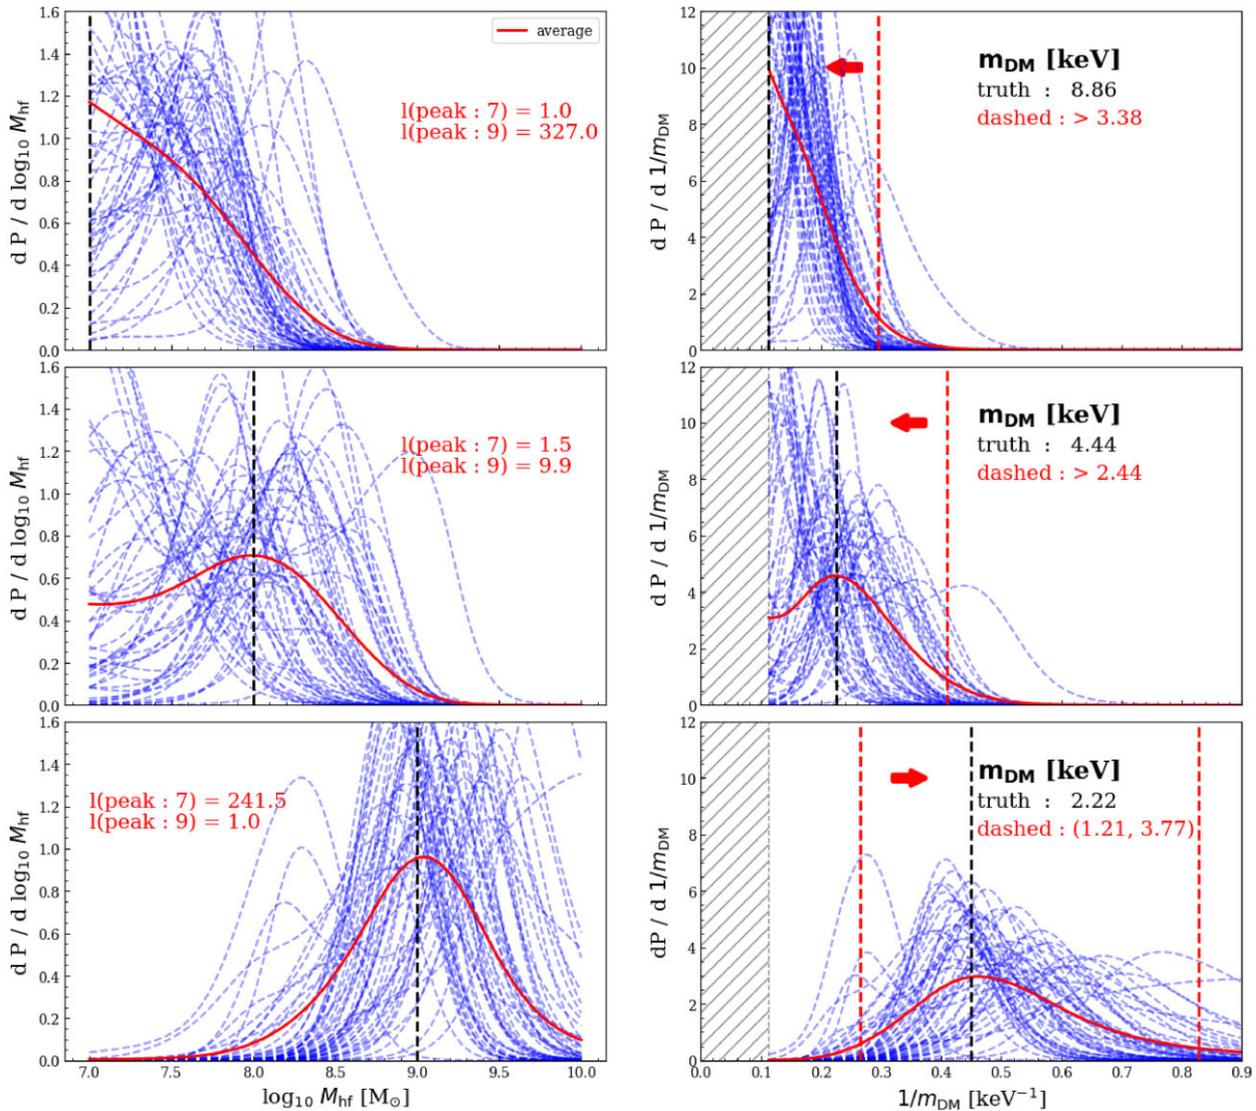


Figure 7. Tests on the fiducial setting. The left column shows the constraints on M_{hf} (adopting a flat prior on $\log M_{\text{hf}}$) and the right columns the constraints on $1/m_{\text{DM}}$ (adopting a flat prior on $1/m_{\text{DM}}$). From top to bottom, panels correspond to true M_{hf} values of 10^7 , 10^8 , and $10^9 M_{\odot}$, marked by the black vertical dashed lines. Blue dashed curves are individual constraints for each set of 50 lensing systems (for clarity, we plot only 50 out of 500 here). The red curve is the average constraint from the 500 sets. In the right column, we use the mean posteriors to place 2σ (95 per cent) credible interval limits (upper limits on $1/m_{\text{DM}}$ for the top two panels and both upper and lower limits for the bottom one). The hatched regions are where $M_{\text{hf}} < 10^7 M_{\odot}$, which is outside the range probed by our forward models.

Note that some care needs to be taken when combining results from multiple ABC calculations. This is because ABC produces an estimate of the posterior distribution, rather than the likelihood. Combining multiple measurements in a traditional Bayesian analysis (with a calculable likelihood function) is a case of multiplying the likelihoods of the different measurements together to get a total likelihood and then multiplying this by the prior in order to find (something proportional to) the posterior distribution. In this work, we use a flat prior on $\log_{10} M_{\text{hf}}$, and thus the posterior density (per unit $\log_{10} M_{\text{hf}}$) we obtain is proportional to the likelihood. This means that the posterior densities per unit $\log_{10} M_{\text{hf}}$ for individual systems (such as the blue histograms in the lower panel of Fig. 6) can be multiplied together to produce something proportional to the joint likelihood, which can then be multiplied by a prior (once) to get the posterior.

Before multiplying individual constraints, to reduce the noise and obtain a smooth likelihood, we apply a kernel density estimation

method to the distribution of $\log_{10} M_{\text{hf}}$ with a kernel width of 0.3 dex. We have tested smaller kernel sizes, but the constraints quickly become noisy without significant improvement. To correct for boundary effects, we choose the ‘renormalization’ correction for our kernel density estimations.³ As an example, the kernel density estimation for the fiducial result is shown as the black dashed curve in the lower panel of Fig. 6.

The left column of Fig. 7 shows our test results for three different cases with true input M_{hf} of $10^7 M_{\odot}$, $10^8 M_{\odot}$, and $10^9 M_{\odot}$ from top to bottom, respectively. Each blue line is a constraint from 50 systems (for clarity, we show only 50 blue lines in the plots). Any one set of 50 observations will not necessarily have a posterior that peaks at the true value, so in order to assess whether our results are systematically

³We use the public package PyQt-Fit to implement the kernel density estimation (see <https://pythonhosted.org/PyQt-Fit/>).

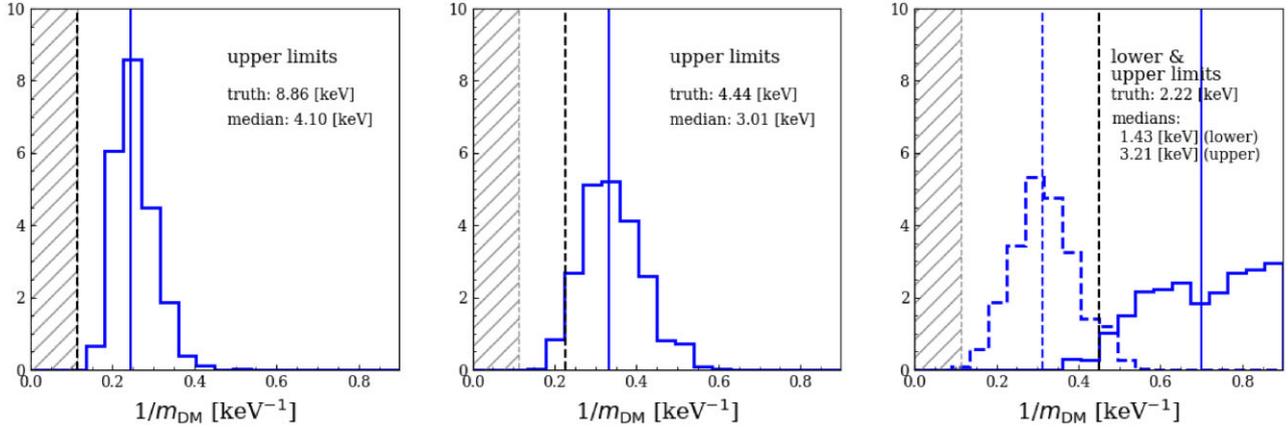


Figure 8. Histograms of 2σ (95 per cent) upper or both upper and lower limits on $1/m_{\text{DM}}$. From left to right, the input half-mode mass is $M_{\text{hf}} = 10^7$, 10^8 , and $10^9 M_{\odot}$, with the corresponding DM particle masses written in the relevant panel. The limits are upper limits on $1/m_{\text{DM}}$ (i.e. a lower limit on the DM particle mass), except for the case of $M_{\text{hf}} = 10^9 M_{\odot}$ where both histograms of upper (solid line) and lower (dashed line) limits on $1/m_{\text{DM}}$ are shown. The black dashed lines mark the true inputs. The vertical blue lines mark the medians of the histograms and corresponding median values are also listed.

biased – as opposed to just subject to random error – we plot a mean posterior distribution from the 500 sets of 50 observations. These are the red lines in Fig. 7; they peak close to the input values of M_{hf} , suggesting that our method is unbiased. Using these mean posteriors, we compute the ratio of the posterior between the peak value and the values at $M_{\text{hf}} = 10^7 M_{\odot}$ and $M_{\text{hf}} = 10^9 M_{\odot}$. With an input value of $M_{\text{hf}} = 10^7 M_{\odot}$, the recovered constraint shows that the model with $M_{\text{hf}} = 10^9 M_{\odot}$ is disfavoured with a posterior density of ~ 327 times smaller than that of the peak. In a WDM universe (e.g. with $M_{\text{hf}} = 10^9 M_{\odot}$), $M_{\text{hf}} = 10^7 M_{\odot}$ would be ruled out, with a mean posterior density of ~ 242 times smaller than at the peak.

In the right column of Fig. 7, we show our posterior distribution in terms of $1/m_{\text{DM}}$, which is the way in which constraints on the DM particle mass from the Lyman- α forest are typically expressed (Iršič et al. 2017). Note that the $1/m_{\text{DM}}$ posteriors are not simply the M_{hf} posteriors transformed to a new parametrization. Instead, we transform the likelihood as a function of M_{hf} to the likelihood as a function of $1/m_{\text{DM}}$ following equation (5) and then adopt a flat prior on $1/m_{\text{DM}}$ (as done in Lyman- α studies). This is different from a flat prior on $\log_{10} M_{\text{hf}}$ and so the posteriors are not actually the same in the two columns of Fig. 7.

Depending on the ‘behaviour’ of the red curves, we place either a 2σ (95 per cent) upper or lower limit on $1/m_{\text{DM}}$. For the cases with true $M_{\text{hf}} = 10^7$ and $10^8 M_{\odot}$, we place an upper limit, while for the case with $M_{\text{hf}} = 10^9 M_{\odot}$, we place both lower and upper limits. These limits are shown as vertical dashed red lines in the figure. We can see that with 50 fiducial-like lensing systems, at 2σ level, one can rule out particle candidates with m_{DM} less than 3.38 keV and 2.44 keV in universes with true $m_{\text{DM}} = 8.86$ and 4.44 keV, respectively, or rule out particles with $m_{\text{DM}} > 3.77$ keV and $m_{\text{DM}} < 1.21$ keV in a universe with true $m_{\text{DM}} = 2.22$ keV. We view the 2σ limits in the first two panels as conservative because we do not assign any posterior mass to the shaded regions in our computation. If we had simply assumed that the posterior density in the shaded region was the same as at $\sim 0.12 \text{ keV}^{-1}$, the upper limits would have been tighter.

Note that the constraints discussed above are average results. Observations of any 50 specific lensing systems would yield a constraint like one of the blue dashed lines in the figures, which might be tighter or looser than the average constraint. To demonstrate actual constraint, one would expect to get with 50 observations;

in Fig. 8, we show the histograms of 2σ limits from the 500 different realizations of a set of lensing observations (for true $m_{\text{DM}} = 8.86$, 4.44 keV, we compute lower 2σ limits, and for true $m_{\text{DM}} = 2.22$ keV, we compute lower and upper 2σ limits). As shown by the histograms, when true $m_{\text{DM}} = 8.86$ keV, the median 2σ constraint we would get is $1/m_{\text{DM}} < 0.24 \text{ keV}$, and thus $m_{\text{DM}} > 4.10$ keV. In other words, there is 50 per cent chance to constrain m_{DM} better than 4.10 keV with 50 lenses of similar settings. If true $m_{\text{DM}} = 4.44$ keV, the median 2σ constraint is $m_{\text{DM}} > 3.01$ keV. If true $m_{\text{DM}} = 2.22$ keV, the median constraint obtained is $1.43 \text{ keV} < m_{\text{DM}} < 3.21 \text{ keV}$.

3.2 Dependency on the lensing configuration

Having shown that with a large number of lensing systems the forward modelling procedure can correctly recover the true M_{hf} (albeit with fairly broad posteriors), we now explore how the constraints change when the properties of the lenses and sources are varied. In particular, we vary the lens and source redshifts, the image S/N (signal-to-noise) ratio, and the image resolution.

To show clearly how the constraints depend upon a specific parameter, we change only one parameter value at a time. For example, to investigate how the result changes for a higher redshift source, we change only the source redshift and leave the image resolution, Einstein radius, surface brightness, etc., unchanged from the $\{T\}_{\text{fiducial}}$ values. Notice that the parameters listed in Table 1 are observational rather than intrinsic physical quantities, so by fixing a parameter, we refer to fixing a particular observational quantity. For example, keeping the source surface brightness unchanged from a redshift of 1–2.5 means that the source galaxy at $z = 2.5$ is actually intrinsically brighter because of cosmological dimming. The reason for changing quantities in this way is that we want to focus on how the results change from an observational perspective and thus provide a basic idea of what type of lensing configuration has most constraining power, which may help inform future observational designs.

In total, we have carried out five additional tests: placing the lens at a lower redshift of 0.2, denoted as $\{T\}_{l_z=0.2}$; placing the source at a higher redshift of 2.5, denoted as $\{T\}_{s_z=2.5}$; doubling the exposure time, denoted as $\{T\}_{\text{exp}=1200}$; lowering the resolution, with pixel size of $0.1''$ and PSF σ of $0.08''$, similar to the expected resolution of images from the Euclid Space Telescope (Collett 2015), denoted

as $\{T\}_{\text{Euclid}}$; and assuming a similar resolution to that of the China Space Station Telescope (CSST), with pixel size of $0.075''$ and PSF σ of $0.08''$, denoted $\{T\}_{\text{CSST}}$. To first focus on the changes caused by different settings, without being affected by set-of-observations to set-of-observations noise, we compare the mean posteriors from 500 sets of 50 observed systems (i.e. the equivalents of the red lines in Fig. 7).

In the left column of Fig. 9, we plot the expected M_{hf} posteriors for different lensing configuration using different colours. Blue, green, red, cyan, orange, and purple correspond to the settings, $\{T\}_{\text{fiducial}}$, $\{T\}_{l_z=0.2}$, $\{T\}_{s_z=2.5}$, $\{T\}_{\text{exp}=1200}$, $\{T\}_{\text{Euclid}}$, and $\{T\}_{\text{CSST}}$, respectively. Different rows show results assuming different true inputs of M_{hf} , which are marked as vertical dashed black lines. As in Fig. 7, we plot $1/m_{\text{DM}}$ posteriors in the right column and place upper (in the top two panels) or both lower and upper (in the bottom panel) 2σ limits on $1/m_{\text{DM}}$, which are marked as vertical dashed lines. The red, green, and blue curves show how the results vary with the lens and source redshifts.

Increasing either the source or lens redshift improves the constraints because (at fixed Einstein radius) the volume in which low-mass dark matter haloes are projected close to the observed Einstein ring increases as either redshift is increased. The comparison between the cyan and blue curves shows that increasing the S/N ratio (exposure time) results in a better constraint. Also, better angular resolution gives better constraints when comparing the fiducial setting (blue curve) and the *Euclid*/CSST resolution setting (orange/purple curves), but the improvement is not significant. Although a worse constraint is obtained with lower image resolution, we see that even with lower angular resolution, we can distinguish between models with M_{hf} of $10^7 M_{\odot}$ and $10^8 M_{\odot}$. The advantage of *Euclid* and CSST will be sample sizes that are vastly larger than 50. Note that future *Euclid*/CSST observations will likely be different from the sample we simulate here: they will have, for example, different lens/source redshift distributions, data quality, and Einstein radii. Those differences can potentially make our predictions here different from future real observations.

Similarly, we use average constraints above to demonstrate the result dependency on lensing configurations, while for the constraints from individual sets of 50 strong lensing systems with different lensing configurations, we refer the 2σ limit histograms in Fig. 10. Colours of the histograms follow the meaning of Fig. 9. The vertical dashed line marks the true input and the left regions below our test range are plotted in shadow. For each setting, we list the medians of 2σ constraints in corresponding colours. As shown, the dependency reflected from the histograms is the same as that from the average constraints in Fig. 9. Among all the settings, the setting with a higher redshift source gives the best constraint, where when true $m_{\text{DM}} = 8.86$ keV, the median 2σ constraint one could get is $m_{\text{DM}} > 5.20$ keV.

To summarize, the tests shown here agree with the expectation that lensing systems with longer exposure times, higher resolution, and more low-mass dark matter haloes (higher source redshifts and a larger area around the Einstein arcs) give tighter constraints.

3.3 Model assumptions and limitations

Our method makes a number of simplifying assumptions that we now summarize and describe in more detail.

By simulating low-mass dark matter haloes with the mass function of equation (4), we neglect lensing perturbations from subhaloes within the lens galaxy. This is a good approximation since, as Li et al. (2017) have shown, for realistic lensing configurations

the signal is dominated by line-of-sight haloes, rather than by subhaloes in the lens. In any case, this assumption makes our results conservative, as including subhaloes would boost the lensing perturbations, increasing the signal we can extract from the lens model residuals. On the other hand, there are uncertainties regarding subhaloes that do not affect line-of-sight haloes, for example, the extent of tidal disruption (Despali et al. 2018; Richings et al. 2020); marginalizing over these uncertainties would weaken the constraint on M_{hf} .

We neglect any uncertainties on the amplitude and shape of the (CDM) halo mass function. The amplitude is fixed by the value of the cosmological parameter, σ_8 , which is known to better than 1 per cent from cosmic microwave background data (Planck Collaboration VI 2020). The shape is known very precisely from cosmological simulations (Wang et al. 2020). There is, however, some degeneracy between the amplitude of the halo mass function and the half-mode mass. Some lensing studies, such as that of Gilman et al. (2019) based on flux ratio anomalies, have included the amplitude of the mass function as a free parameter to be fit at the same time as the subhalo mass function.

In this work, we have assumed an SIE lens mass model, whereas in studies of real lenses a power-law mass model with external shear is widely used (e.g. Dye et al. 2014; Vegetti et al. 2014). This gives the mass model more freedom to fit the residuals, reducing the signal left over from substructures. Even then, the mass model may provide only an approximate fit to the lens and result in residuals not associated with subhaloes, at which point even more complex mass profiles (Nightingale et al. 2019) or a potential corrections-based approach may be required (Vegetti & Koopmans 2009). We assume a core-Sersic model for the source light, whereas studies using real data often use a non-parametric approach that reconstructs a source's irregular morphology (see Warren & Dye 2003; Suyu et al. 2006). Gilman, Birrer & Treu (2020c) show how such an approach may absorb part of the residual signal of the low-mass perturbers, reducing the available information on the halo mass function. PyAutoLens has all the necessary functionality to test these assumptions and this will be the topic of future work.

Our method implicitly attributes all image residuals to the presence of perturbing dark matter haloes when, in reality, there could be other sources of mismatch between a true lens and source model and the best-fitting macro model. In the idealized setup used here, we have shown that this allows us to obtain a correct measurement of the input half-mode mass. In reality, other sources of perturbations will contribute to the power spectrum of the image residuals, including a deficient model for the smooth component of the lens, an inaccurate description of the telescope PSF, and artefacts or correlated noises introduced by the data reduction process. Not taking these effects into account may bias our inference of the subhalo mass function. Currently, the impact of any individual effect is unclear and further investigation is required, noting that all effects that impact the signal can, in principle, be included in our forward modelling procedure and marginalized over.

On a positive note, our use of the power spectrum to define a summary statistic to extract the signal from the residuals is likely to be suboptimal. A more carefully crafted summary statistic or a machine learning-based approach (see Brehmer et al. 2019; Diaz Rivero & Dvorkin 2020) can potentially improve the signal that can be extracted from a lens system and thus provide better constraints on M_{hf} than shown here. It could also make the estimation less sensitive to systematic sources of residuals as shown by Birrer et al. (2017b) and so increase the constraining power of the forward modelling method.

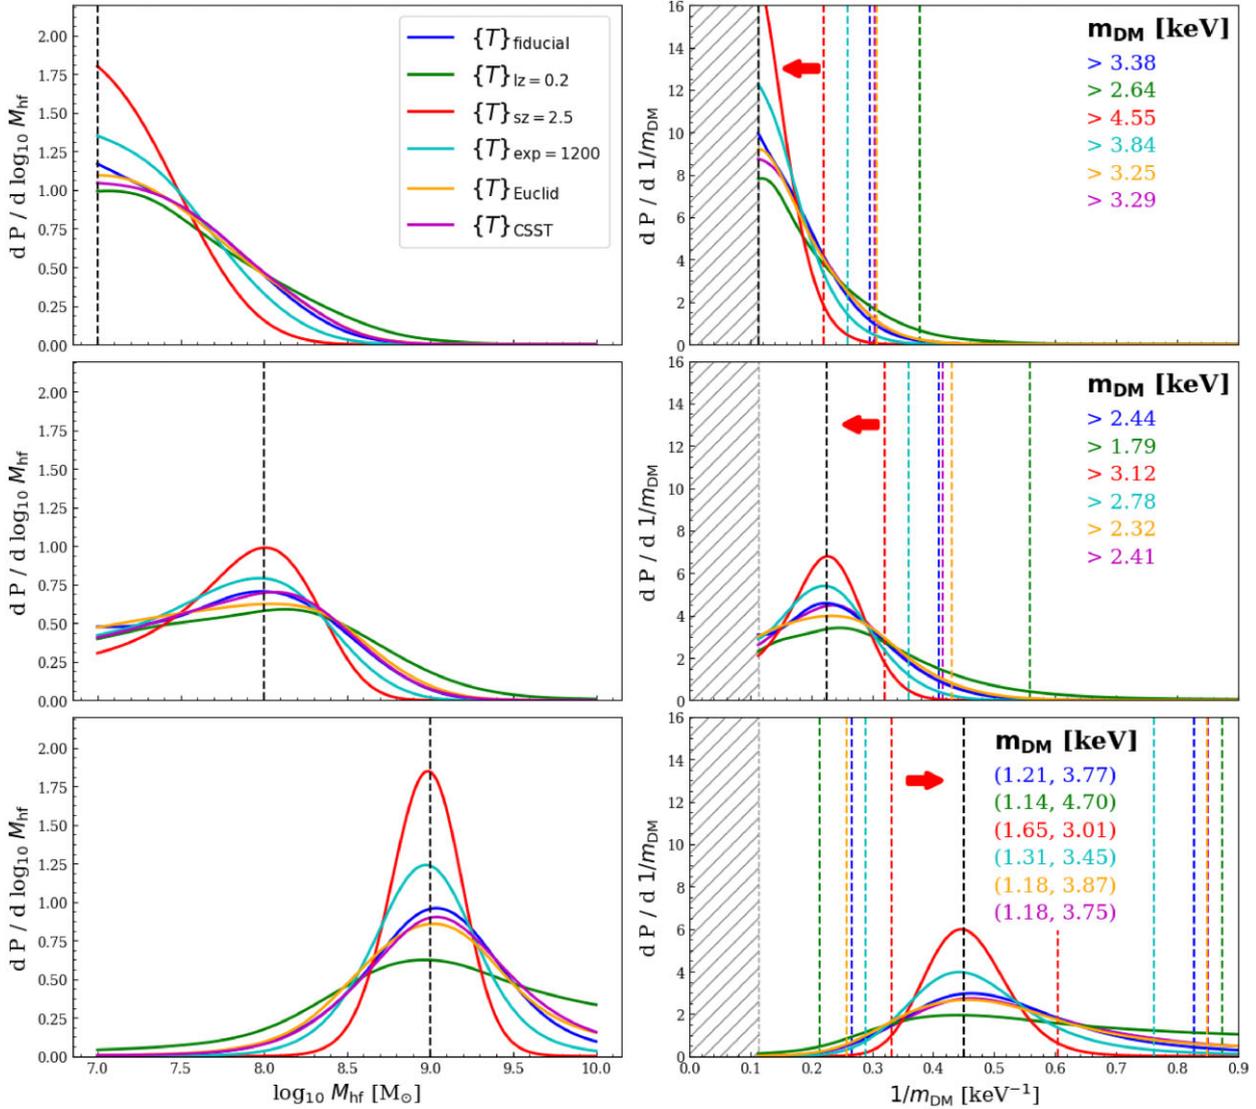


Figure 9. Average constraints from 50 lensing systems with different settings. As in Fig. 7, the left-hand panels show the constraints on M_{hf} and the right-hand panels the constraints on $1/m_{\text{DM}}$. Blue, green, red, cyan, orange, and purple lines correspond to the results of settings, $\{T\}_{\text{fiducial}}$, $\{T\}_{\text{lz}=0.2}$, $\{T\}_{\text{sz}=2.5}$, $\{T\}_{\text{exp}=1200}$, $\{T\}_{\text{Euclid}}$, and $\{T\}_{\text{CSST}}$, respectively. The vertical dashed black lines mark the true input values and the colour lines the corresponding 2σ upper limits (top 2 panels) or both upper and lower limits (bottom panel) on $1/m_{\text{DM}}$.

3.4 Comparison with flux ratio anomalies

A theoretical investigation of the constraining power of flux ratio anomalies was performed by Gilman et al. (2019), who subsequently applied their method to real observations, placing constraints on both the dark matter mass function and the mass–concentration relation (Gilman et al. 2020a,b). It is interesting to compare our approach with that of Gilman et al. (2019). Note, however, that, as we just discussed, these authors treat the amplitude and slope of the halo mass function as free parameters, whereas our tests assume a halo mass function with only one free parameter, M_{hf} . In our tests, for a universe with true $M_{\text{hf}} = 10^7 M_{\odot}$, our fiducial calculation gives an upper 2σ constraint (average constraint shown in Fig. 7), $M_{\text{hf}} = 10^{8.25} M_{\odot}$, which is slightly smaller than the upper constraints quoted by Gilman et al., $M_{\text{hf}} = 10^{8.34} M_{\odot}$, for a 2 per cent uncertainty on flux measurements (see fig. 8 of Gilman et al. 2019). Given the different assumptions in the two studies, this comparison serves to show that the constraints that they provide are roughly comparable.

When considering the relative performance of our method and those based on flux ratio anomalies, it is pertinent to consider the mass scales to which each method is sensitive. The ‘coldest’ case in our tests assumes $M_{\text{hf}} = 10^7 M_{\odot}$ and our signal is heavily influenced by the lensing effects of larger perturbers with $M_{200} \sim 10^{7-8} M_{\odot}$. The small detection area of flux ratio anomalies results in a lack of sensitivity to subhaloes with these relatively high masses (as these subhaloes are rare), but the ‘point-like’ sources (with highly variable surface brightness over a small region) provide sensitivity down to dark matter haloes with $M_{200} \sim 10^6 M_{\odot}$, a scale that is not accessible with our approach. However, another problem faced by flux ratio anomalies is that the sparsity of information means that the degeneracy between anomalies caused by substructure and by a complex smooth halo is not easily broken, which can bias the inference on the abundance of substructures (Xu et al. 2015; Hsueh et al. 2018). The two approaches are therefore complementary and, together, can provide constraints on the halo mass function over a broad range of masses.

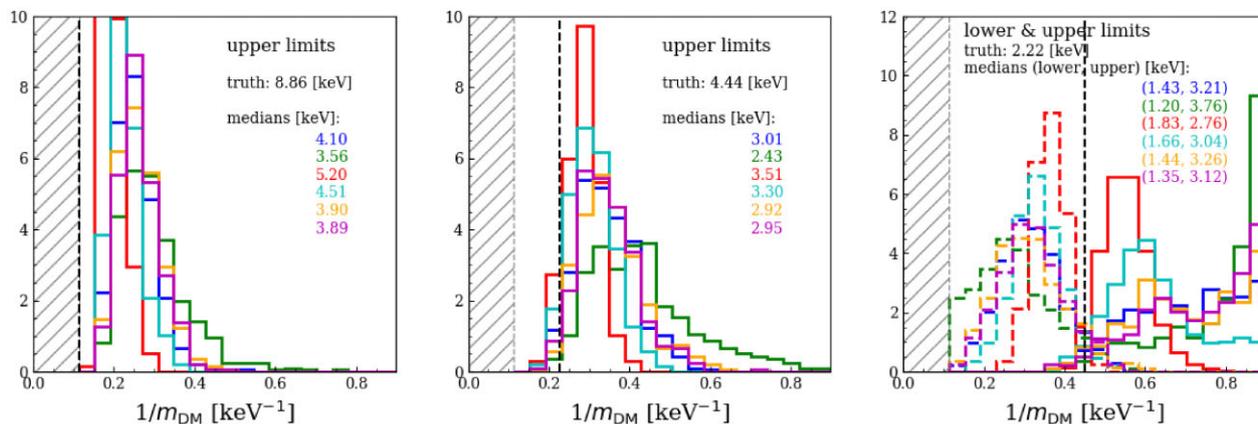


Figure 10. Histograms of 2σ (95 per cent upper) limits (solid lines) and lower limits (dashed lines, only in the rightmost panel) on $1/m_{\text{DM}}$. Blue, green, red, cyan, orange, and purple histograms correspond to the results of settings, $\{T\}_{\text{fiducial}}$, $\{T\}_{l_z=0.2}$, $\{T\}_{s_z=2.5}$, $\{T\}_{\text{exp}}=1200$, $\{T\}_{\text{Euclid}}$, and $\{T\}_{\text{CSST}}$, respectively. The vertical black dashed lines mark the true input m_{DM} in each panel. The medians of the histograms are also listed in corresponding colours in the panels.

4 CONCLUSIONS

The existence of a large population of dark matter haloes of all masses down to subsolar values is a fundamental prediction of the CDM model that distinguishes it from other currently popular models such as WDM in which the halo mass function is truncated below a mass $\sim 10^8 M_{\odot}$. Detecting the predicted population is a key test of the standard cosmological model, which would be ruled out if it could conclusively be shown that low-mass haloes are absent or that they are present in lower numbers than predicted.

In this work, we have presented a method based on forward modelling applied to the analysis of resolved strong lensing arc data that can constrain a cutoff in the dark matter halo mass function. The key idea is to simulate a large number of strong lensing images (forward models) based on the best-fitting macro model, with the addition of perturbations to the lens model from low-mass dark matter haloes (whose number depends on the assumed nature of the dark matter). These images can then be fit in the same manner as the original observed image, and the best-fitting image residuals for the observed system can be compared with the best-fitting image residuals for the forward models. This comparison is made by means of the power spectrum of the image residuals within the ABC framework and leads to a posterior distribution for the half-mode mass (M_{hf}), which describes the cutoff in the halo mass function. Our main results may be summarized as follows:

(i) For mock observed lenses constructed from parametric source light profiles and lens mass models, we confirm that information on low-mass dark matter haloes can be extracted from the power spectrum of the best-fitting image residuals. As demonstrated in Fig. 7, with a large number of observed lenses, a forward modelling procedure can correctly recover the true input half-mode mass, M_{hf} in average (the red curves). However, the scatter in the constraint is significant even though results for 50 systems have been combined together (blue curves), which may be attributed to the large intrinsic scatter in realizations of low-mass dark matter haloes. Taking the scatter into account, we compute the histograms of 2σ limits (see Fig. 8) and find that, with 50 lensing systems, for our fiducial settings (see Table 1), if the true (thermal) m_{DM} is 8.86 keV ($M_{\text{hf}} = 10^7 M_{\odot}$), the median 2σ constraint on m_{DM} is that $m_{\text{DM}} > 4.10$ keV (or there is 50 per cent chance that the m_{DM} can be constrained better than 4.10 keV at 2σ level). Conversely, in a WDM universe where the true

(thermal) m_{DM} is 2.22 keV ($M_{\text{hf}} = 10^9 M_{\odot}$), one could get a median measure of m_{DM} to be between 1.43 keV and 3.21 keV at 2σ level.

(ii) We have tested the dependency of the method on different lensing configurations and image quality settings. As shown in Figs 9 and 10, the dependency agrees with expectations: higher redshift sources and/or larger areas around the Einstein arcs, and better data quality (longer exposure time and/or higher resolution) all result in a tighter constraint. Among our tests, the one with sources placed at $z = 2.5$ produced the strongest constraint: in a universe with true $m_{\text{DM}} = 8.86$ keV ($M_{\text{hf}} = 10^7 M_{\odot}$), particles with mass less than 5.20 keV can be ruled out at the 2σ level (see medians of the histograms in Fig. 10). Although a slightly worse constraint is obtained from images with the resolution of *Euclid*/*CSST*, 50 systems can still provide a constraint on M_{hf} , and with *Euclid*/*CSST* data we will have many more than 50 strong lenses. We defer a thorough test using mock strong lensing samples similar to those expected from *Euclid*/*CSST* to future work.

Throughout this study, we have made several simplifying assumptions, particularly in using parametric models for the lens mass distributions and the source light profiles. The effects of removing these assumptions will need to be investigated before we can safely apply our technique to real observations. We have shown that, in principle, information on low-mass dark matter haloes can be statistically extracted from the power spectrum of the image residuals. In the near future, we hope that using a more powerful summary statistic together with more advanced lens modelling techniques (such as pixelized sources) can improve the power of this technique. With thousands of new strong lensing observations expected from future space telescopes, there is every prospect of pinning down the mass of the dark matter particles using this forward modelling technique.

SOFTWARE CITATIONS

This work uses the following software packages:

- (i) Astropy (Astropy Collaboration 2013; Price-Whelan et al. 2018)
- (ii) Colossus (Diemer 2018b)
- (iii) corner.py (Foreman-Mackey 2016)
- (iv) hmf (Murray, Power & Robotham 2013)
- (v) MATPLOTLIB (Hunter 2007)
- (vi) numba (Lam, Pitrou & Seibert 2015)
- (vii) NUMPY (van der Walt, Colbert & Varoquaux 2011)

- (viii) PyAutoFit (Nightingale, Hayes & Griffiths 2021a)
- (ix) PyAutoLens (Nightingale & Dye 2015; Nightingale et al. 2021b)
- (x) PyMulitNest (Feroz, Hobson & Bridges 2009b; Buchner et al. 2014b)
- (xi) PyQt-Fit (Reuille 2015)
- (xii) pyquad (Kelly 2020)
- (xiii) PYTHON (Van Rossum & Drake 2009)
- (xiv) scikit-image (Van der Walt et al. 2014)
- (xv) scikit-learn (Pedregosa et al. 2011)
- (xvi) SCIPY (Virtanen et al. 2020)

ACKNOWLEDGEMENTS

We thank the anonymous referee for useful suggestions that led to several improvements of the original draft. We also thank Nicola C. Amorisco for helpful discussions and suggestions, and Simon Birrer for insightful comments. QH, CSF, and SMC are supported by European Research Council (ERC) Advanced Investigator grant, DMIDAS [GA 786910], and also by the Science and Technology Facilities Council (STFC) Consolidated Grant for Astronomy at Durham (grant numbers ST/F001166/1, ST/I00162X/1, ST/P000541/1). This project has received funding from the European Union's Horizon2020 research and innovation programme under grant agreement no. 776247. RM acknowledges the support of a Royal Society University Research Fellowship. RL acknowledge the support of National Nature Science Foundation of China (nos. 11773032, 12022306), the science research grants from the China Manned Space Project (No CMS-CSST-2021-B01).

This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grants ST/H008519/1 and ST/K00087X/1, STFC DiRAC Operations grant ST/K003267/1, and Durham University. DiRAC is part of the National E-Infrastructure.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

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

Amara A., Metcalf R. B., Cox T. J., Ostriker J. P., 2006, *MNRAS*, 367, 1367

Astropy Collaboration, 2013, *A&A*, 558, A33

Baltz E. A., Marshall P., Oguri M., 2009, *J. Cosmology Astropart. Phys.*, 2009, 015

Bayer D., Chatterjee S., Koopmans L. V. E., Vegetti S., McKean J. P., Treu T., Fassnacht C. D., 2018, preprint ([arXiv:1803.05952](https://arxiv.org/abs/1803.05952))

Benitez-Llambay A., Frenk C., 2020, *MNRAS*, 498, 4887

Benson A. J., Frenk C. S., Lacey C. G., Baugh C. M., Cole S., 2002, *MNRAS*, 333, 177

Birrer S., Welschen C., Amara A., Refregier A., 2017a, *J. Cosmology Astropart. Phys.*, 2017, 049

Birrer S., Amara A., Refregier A., 2017b, *J. Cosmology Astropart. Phys.*, 2017, 037

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

Bonaca A., Hogg D. W., Price-Whelan A. M., Conroy C., 2019, *ApJ*, 880, 38

Bose S., Hellwing W. A., Frenk C. S., Jenkins A., Lovell M. R., Helly J. C., Li B., 2016, *MNRAS*, 455, 318

Brehmer J., Mishra-Sharma S., Hermans J., Louppe G., Cranmer K., 2019, *ApJ*, 886, 49

Brewer B. J., Huijser D., Lewis G. F., 2016, *MNRAS*, 455, 1819

Buchner J. et al., 2014a, *A&A*, 564, A125

Buchner J. et al., 2014b, *A&A*, 564, A125

Çağan Şengül A., Tsang A., Diaz Rivero A., Dvorkin C., Zhu H.-M., Seljak U., 2020, *Phys. Rev. D*, 102, 063502

Chatterjee S., Koopmans L. V. E., 2018, *MNRAS*, 474, 1762

Collett T. E., 2015, *ApJ*, 811, 20

Cyr-Racine F.-Y., Keeton C. R., Moustakas L. A., 2019, *Phys. Rev. D*, 100, 023013

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

Davies F. B., Hennawi J. F., Eilers A.-C., Lukić Z., 2018, *ApJ*, 855, 106

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

Despali G., Vegetti S., White S. D. M., Giocoli C., van den Bosch F. C., 2018, *MNRAS*, 475, 5424

Diaz Rivero A., Dvorkin C., 2020, *Phys. Rev. D*, 101, 023515

Diaz Rivero A., Dvorkin C., Cyr-Racine F.-Y., Zavala J., Vogelsberger M., 2018, *Phys. Rev. D*, 98, 103517

Diemer B., 2018a, *ApJS*, 239, 35

Diemer B., 2018b, *ApJS*, 239, 35

Dye S. et al., 2014, *MNRAS*, 440, 2013

Efstathiou G., 1992, *MNRAS*, 256, 43P

Enzi W., Vegetti S., Despali G., Hsueh J.-W., Metcalf R. B., 2020, *MNRAS*, 496, 1718

Fearnhead P., Prangle D., 2012, *J. R. Stat. Soc. B Stat. Methodol.*, 74, 419

Feroz F., Hobson M. P., Bridges M., 2009a, *MNRAS*, 398, 1601

Feroz F., Hobson M. P., Bridges M., 2009b, *MNRAS*, 398, 1601

Foreman-Mackey D., 2016, *J. Open Source Softw.*, 1, 24

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

Frenk C. S., White S. D. M., Davis M., 1983, *ApJ*, 271, 417

Frenk C. S., White S. D. M., Efstathiou G., Davis M., 1985, *Nature*, 317, 595

Garrison-Kimmel S. et al., 2017, *MNRAS*, 471, 1709

Garzilli A., Magalich A., Theuns T., Frenk C. S., Weniger C., Ruchayskiy O., Boyarsky A., 2019, *MNRAS*, 489, 3456

Gilman D., Birrer S., Treu T., Nierenberg A., Benson A., 2019, *MNRAS*, 487, 5721

Gilman D., Birrer S., Nierenberg A., Treu T., Du X., Benson A., 2020a, *MNRAS*, 491, 6077

Gilman D., Du X., Benson A., Birrer S., Nierenberg A., Treu T., 2020b, *MNRAS*, 492, L12

Gilman D., Birrer S., Treu T., 2020c, *A&A*, 642, A194

Graham A. W., Erwin P., Trujillo I., Asensio Ramos A., 2003, *AJ*, 125, 2951

Hahn C., Vakili M., Walsh K., Hearin A. P., Hogg D. W., Campbell D., 2017, *MNRAS*, 469, 2791

Hezaveh Y. D. et al., 2016, *ApJ*, 823, 37

Hinshaw G. et al., 2013, *ApJS*, 208, 19

Hsu D. C., Ford E. B., Terrien R., 2020, *MNRAS*, 498, 2249

Hsueh J.-W., Despali G., Vegetti S., Xu D., Fassnacht C. D., Metcalf R. B., 2018, *MNRAS*, 475, 2438

Hsueh J. W., Enzi W., Vegetti S., Auger M. W., Fassnacht C. D., Despali G., Koopmans L. V. E., McKean J. P., 2020, *MNRAS*, 492, 3047

Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90

Iršič V. et al., 2017, *Phys. Rev. D*, 96, 023522

Kaiser N., 1984, *ApJ*, 284, L9

Kelly A. J., 2020, *pyquad*

Lam S. K., Pitrou A., Seibert S., 2015, Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC - LLVM '15. p. 1

Li R., Frenk C. S., Cole S., Gao L., Bose S., Hellwing W. A., 2016, *MNRAS*, 460, 363

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

Lovell M. R. et al., 2012, *MNRAS*, 420, 2318

Lovell M. R. et al., 2016, *MNRAS*, 461, 60

Lovell M. R. et al., 2017, *MNRAS*, 468, 4285

Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, *MNRAS*, 439, 300

Ludlow A. D., Bose S., Angulo R. E., Wang L., Hellwing W. A., Navarro J. F., Cole S., Frenk C. S., 2016, *MNRAS*, 460, 1214

Macciò A. V., Ruchayskiy O., Boyarsky A., Muñoz-Cuartas J. C., 2013, *MNRAS*, 428, 882

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

McCully C., Keeton C. R., Wong K. C., Zabludoff A. I., 2014, *MNRAS*, 443, 3631

McKean J. et al., 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*. p. 84, preprint ([arXiv:1502.03362](https://arxiv.org/abs/1502.03362))

Metcalf R. B., Madau P., 2001, *ApJ*, 563, 9

Murray S. G., Power C., Robotham A. S. G., 2013, *Astron. Comput.*, 3, 23

Nadler E. O. et al., 2021, *Phys. Rev. Lett.*, 126, 091101

Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563

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

Neto A. F. et al., 2007, *MNRAS*, 381, 1450

Newton O. et al., 2021, *J. Cosmology Astropart. Phys.*, 2021, 062

Nierenberg A. M. et al., 2017, *MNRAS*, 471, 2224

Nierenberg A. M., Treu T., Wright S. A., Fassnacht C. D., Auger M. W., 2014, *MNRAS*, 442, 2434

Nightingale J. W. et al., 2021b, *J. Open Source Softw.*, 6, 2825

Nightingale J. W., Dye S., 2015, *MNRAS*, 452, 2940

Nightingale J. W., Dye S., Massey R. J., 2018, *MNRAS*, 478, 4738

Nightingale J. W., Massey R. J., Harvey D. R., Cooper A. P., Etherington A., Tam S.-I., Hayes R. G., 2019, *MNRAS*, 489, 2049

Nightingale J. W., Hayes R. G., Griffiths M., 2021a, *J. Open Source Softw.*, 6, 2550

Pedregosa F. et al., 2011, *J. Mach. Learn. Res.*, 12, 2825

Petkova M., Metcalf R. B., Giocoli C., 2014, *MNRAS*, 445, 1954

Planck Collaboration VI 2020, *A&A*, 641, A6

Price-Whelan A. M. et al., 2018, *AJ*, 156, 123

Reuille P. B., 2015, *PyQt-fit*

Richings J. et al., 2020, *MNRAS*, 492, 5780

Richings J., Frenk C., Jenkins A., Robertson A., Schaller M., 2021, *MNRAS*, 501, 4657

Ritondale E., Vegetti S., Despali G., Auger M. W., Koopmans L. V. E., McKean J. P., 2019, *MNRAS*, 485, 2179

Sawala T. et al., 2016, *MNRAS*, 456, 85

Sawala T., Pihajoki P., Johansson P. H., Frenk C. S., Navarro J. F., Oman K. A., White S. D. M., 2017, *MNRAS*, 467, 4383

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

Schneider A., 2016, *J. Cosmology Astropart. Phys.*, 2016, 059

Schneider A., Smith R. E., Macciò A. V., Moore B., 2012, *MNRAS*, 424, 684

Sérsic J. L., 1963, *Boletín de la Asociación Argentina de Astronomía La Plata Argentina*, 6, 41

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

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

Suyu S. H., Marshall P. J., Hobson M. P., Blandford R. D., 2006, *MNRAS*, 371, 983

Thoul A. A., Weinberg D. H., 1996, *ApJ*, 465, 608

Trujillo I., Erwin P., Asensio Ramos A., Graham A. W., 2004, *AJ*, 127, 1917

van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22

Van der Walt S., Schönberger J. L., Nunez-Iglesias J., Boulogne F., Warner J. D., Yager N., Gouillart E., Yu T., 2014, *PeerJ*, 2, e453

Van Rossum G., Drake F. L., 2009, *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA

Vegetti S., Koopmans L. V. E., 2009, *MNRAS*, 400, 1583

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

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, *MNRAS*, 442, 2017

Vegetti S., Despali G., Lovell M. R., Enzi W., 2018, *MNRAS*, 481, 3661

Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, *Phys. Rev. D*, 88, 043502

Virtanen P. et al., 2020, *Nature Methods*, 17, 261

Wang J., Bose S., Frenk C. S., Gao L., Jenkins A., Springel V., White S. D. M., 2020, *Nature*, 585, 39

Warren S. J., Dye S., 2003, *ApJ*, 590, 673

White S. D. M., Frenk C. S., Davis M., 1983, *ApJ*, 274, L1

Xu D., Sluse D., Gao L., Wang J., Frenk C., Mao S., Schneider P., Springel V., 2015, *MNRAS*, 447, 3189

Xu D. D., Mao S., Cooper A. P., Gao L., Frenk C. S., Angulo R. E., Helly J., 2012, *MNRAS*, 421, 2553

APPENDIX A: DETERMINING MULTIPLANE RESOLUTION

We simulate the effect of line-of-sight haloes by dividing the light-cone from the observer to the source galaxy into a number of intervals and approximating low-mass dark matter haloes within a given redshift interval as lying on a single plane at the central redshift. Reducing the redshift interval, Δz , between two neighbouring planes (increasing the number of multiplanes) will increase the accuracy of the approximation to the line-of-sight effects but will also increase the computational cost. The smallest allowable Δz should be determined by the accuracy one would like to achieve in the inference of M_{hf} . In practice, one needs to make sure that the ‘difference’ caused by approximating line-of-sight effects into multiplanes is smaller than the ‘difference’ caused by changing M_{hf} by the amount one would like to distinguish.

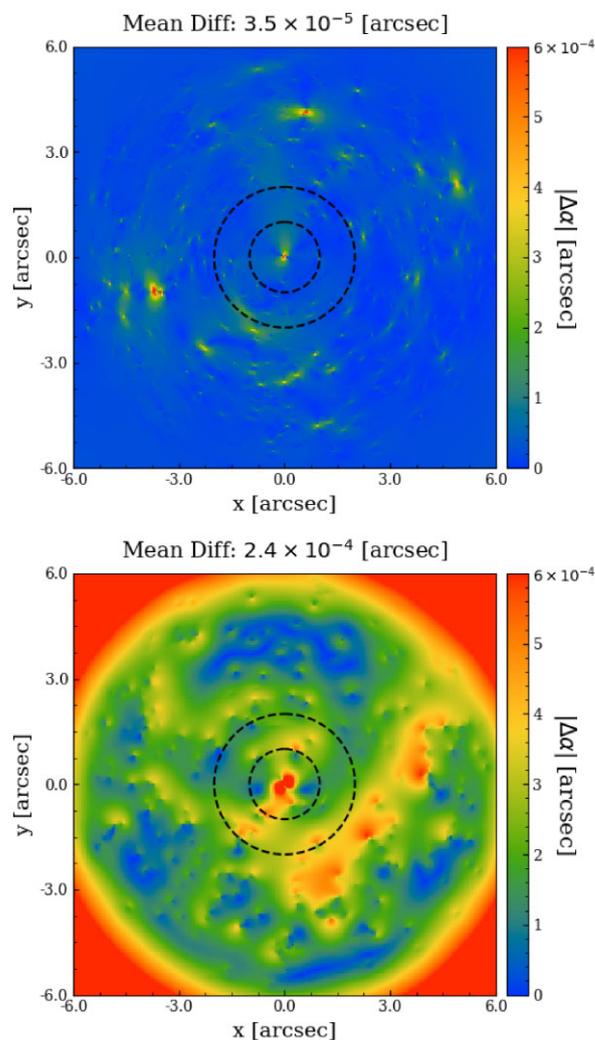


Figure A1. Upper panel: amplitude of the differences in deflection angles caused by changing Δz from 0.01 to 0.001. Lower panel: amplitude of the differences in deflection angles caused by changing M_{hf} from $10^7 M_{\odot}$ to $10^{7.1} M_{\odot}$. The region between two dashed circles is the annulus between $1.0''$ and $2.0''$. The unit of the colour bar is arcsecs.

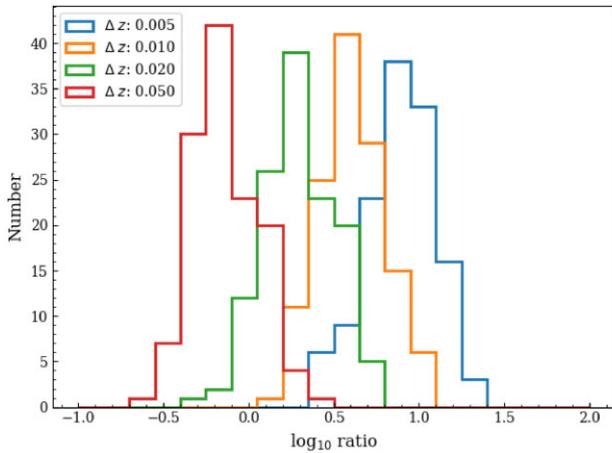


Figure A2. Histograms of the \log_{10} ratio of the mean amplitude of the differences in deflection angles caused by changing Δz and M_{hf} . Blue, orange, green, and red colours shows results for $\Delta z = 0.005, 0.01, 0.02,$ and 0.05 , respectively.

Fig. A1 shows an example of the amplitudes of the differences in deflection angles when changing the values of Δz and M_{hf} . The two panels are computed from the same macro settings (a spherical isothermal sphere model with R_E of $1.5''$ at $z = 0.5$ and a source at $z = 1.0$) and the same low-mass dark matter haloes (generated from a distribution with $M_{\text{hf}} = 10^7 M_\odot$). To simplify the test, instead of using the mass function described in the main body of the paper, we impose a ‘sharp cut’ on the mass function, such that no haloes are drawn with M_{200} smaller than M_{hf} . Also, we assume that there is no correlation between M_{hf} and the mass–concentration relation to make sure that massive haloes are unchanged when changing M_{hf} .

The upper panel of Fig. A1 shows the amplitudes of the differences in deflection angles when changing Δz from 0.01 to 0.001 (a very high-resolution setting which we take as the ‘accurate’ result). The lower panel shows the amplitudes of the differences in deflection angles when changing M_{hf} from $10^7 M_\odot$ to $10^{7.1} M_\odot$. As seen from the figure, the patterns of the differences caused by changing Δz and M_{hf} are different. We use the mean amplitude of the difference in deflection angles within the annular region $1.0'' - 2.0''$ (the region between the two dashed circles in the figure) to quantify the ‘differences’. The mean difference caused by changing Δz is 3.47×10^{-5} arcsec, which is smaller than the mean difference caused by changing M_{hf} by 0.1 dex, 2.42×10^{-4} arcsec, suggesting that Δz of 0.01 would not affect the ability to distinguish between models with $M_{\text{hf}} = 10^7 M_\odot$ and $10^{7.1} M_\odot$.

The results shown in Fig. A1 are for a special case, so we repeat the same procedure for 128 different realizations of low-mass dark matter haloes. The orange histogram in Fig. A2 shows the distribution of the \log_{10} ratio of the mean difference in amplitudes defined previously, that is the ratio of the mean difference in the deflection angles between models with $M_{\text{hf}} = 10^7$ and $M_\odot = 10^{7.1}$ divided by the mean difference in the deflection angles for the Δz given in the legend. As shown by the orange histogram, the ratios from all of the 128 realizations are larger than one, indicating that the difference caused by approximating line-of-sight effects using multiplanes separated by 0.01 in redshift is smaller than the difference cause by changing M_{hf} by 0.1 dex. Also, we plot the histogram of ratios for other values of Δz . As may be seen, increasing Δz increases the differences caused by changing Δz and, in some cases, the differences are larger than those caused by changing M_{hf} by 0.1 dex. To summarize, the value of Δz we adopt for our tests, $\Delta z = 0.01$, does not affect our inference of M_{hf} to an accuracy of 0.1 dex.

APPENDIX B: THE DISTRIBUTION OF CONSTRAINTS FROM SETS OF 50 LENSES

To obtain an indication of the constraints on the DM particle mass that one expects from 50 lensing systems, in Fig. 8 we plot histograms of the upper, or both upper and lower, 2σ limits on $1/m_{\text{DM}}$ for our fiducial lens and source properties. Each histogram shows the distribution of limits from 500 different realizations of a 50-lens sample, with upper limits shown when the true input, $M_{\text{hf}} = 10^{7.8} M_\odot$, while both upper (solid lines) and lower (dashed lines) limits are shown when the true input, $M_{\text{hf}} = 10^9 M_\odot$. For $M_{\text{hf}} = 10^8 M_\odot$, the probability that the true input will be above the 2σ upper limit is ~ 4 per cent, while for $M_{\text{hf}} = 10^9 M_\odot$, the true input is ruled out by the 2σ limit ~ 10 per cent of the time.⁴ Additionally, in Fig. 10 we show histograms of upper, or both upper and lower, 2σ limits from lens samples with different lensing configurations, corresponding to the different data quality and source and lens redshifts presented in Fig. 9.

⁴It may seem counterintuitive that the 2σ limit rules out the true input with a probability different from ~ 5 per cent. However, here we are performing two different calculations: the 2σ limits come from Bayesian inference with a prior of $1/m_{\text{DM}}$ assumed to be uniform over a certain range, while the multiple tests are carried out only for a particular value of $1/m_{\text{DM}}$. There is thus a mismatch between the prior for the inference and the tests. In other words, if one repeatedly computed 2σ limits for a true $1/m_{\text{DM}}$ uniformly distributed over the range, one would expect to find that the 2σ limits rule out the truth with a probability of ~ 5 per cent.

This paper has been typeset from a \LaTeX file prepared by the author.