Halo concentration strengthens dark matter constraints in galaxy–galaxy strong lensing analyses

Nicola C. Amorisco,[1]⋆† James Nightingale [1], Qiuhan He [1], Aristeidis Amvrosiadis,[1] Xiaoyue Cao,[2,3] Shaun Cole,[1] Amy Etherington,[1] Carlos S. Frenk,[1] Ran Li,[2,3] Richard Massey [1] and Andrew Robertson [1]

[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*

ABSTRACT

A defining prediction of the cold dark matter cosmological model is the existence of a very large population of low-mass haloes. This population is absent in models in which the dark matter particle is warm (WDM). These alternatives can, in principle, be distinguished observationally because haloes along the line of sight can perturb galaxy–galaxy strong gravitational lenses. Furthermore, the WDM particle mass could be deduced because the cut-off in their halo mass function depends on the mass of the particle. We systematically explore the detectability of low-mass haloes in WDM models by simulating and fitting mock lensed images. Contrary to previous studies, we find that haloes are harder to detect when they are either behind or in front of the lens. Furthermore, we find that the perturbing effect of haloes increases with their concentration: Detectable haloes are systematically high-concentration haloes, and accounting for the scatter in the mass–concentration relation boosts the expected number of detections by as much as an order of magnitude. Haloes have lower concentration for lower particle masses and this further suppresses the number of detectable haloes beyond the reduction arising from the lower halo abundances alone. Taking these effects into account can make lensing constraints on the value of the mass function cut-off at least an order of magnitude more stringent than previously appreciated.

Key words: gravitational lensing: strong – dark matter.

1 INTRODUCTION

The nature and identity of dark matter (DM) remain fundamental open questions in contemporary astrophysics; enormous effort is currently being directed at finding the answer. Numerical simulations of the cosmological process of structure formation (e.g. Davis et al. 1985; Springel et al. 2005; Frenk & White 2012) have shown that a model based on the assumption that the DM consists of cold dark matter (CDM) particles can very successfully reproduce a number of large-scale astrophysical measurements (e.g. Wang et al. 2016; Alam et al. 2017; Planck Collaboration VI 2020). Several plausible DM candidates behave like CDM on large scales, but luckily, their different physical properties can make them distinguishable on subgalactic scales. The defining property of standard CDM is the nearly scale-invariant primordial power spectrum of density fluctuations, which results in an equally distinctive halo mass function, characterized by a large population of haloes down to masses comparable to the Earth's mass (Jenkins et al. 2001; Green, Hofmann & Schwarz 2005; Diemand et al. 2008; Angulo et al. 2012;

Wang et al. 2020). Most alternative DM models predict a suppression of the primordial power spectrum on small scales and an associated truncation of the halo mass function at a mass M_{cut}. For example, in the popular warm dark matter (WDM) model, free streaming arising from the thermal velocities of the particles at early times is the cause of the suppression that occurs at a mass scale that is roughly inversely proportional to the WDM particle mass (e.g. Avila-Reese et al. 2003; Lovell et al. 2012; Schneider et al. 2012; Bose et al. 2016).

Current constraints on the WDM model stem primarily from a combination of the abundance of satellite galaxies in the Milky Way (Lovell et al. 2012, 2016; Kennedy et al. 2014; Newton et al. 2020) and the properties of the Lyman α forest inferred from high-redshift QSO spectra (Viel et al. 2013; Baur et al. 2016; Iršič et al. 2017; Boyarsky et al. 2019). A joint analysis of these, together with constraints from gravitational lensing (see below), places M_{cut} at $\approx 4.3 \times 10^7$ M$_\odot$ for a thermal WDM relic. These bounds, however, are subject to possible systematics such as uncertainties in the galaxy formation physics in the case of satellites (Newton et al. 2020) and assumptions on the thermal history of the intergalactic medium at high redshift in the case of the Lyman α forest (e.g. Garzilli, Boyarsky & Ruchayskiy 2017; Garzilli et al. 2019). Constraints from independent probes, such as we discuss here, are therefore a priority.

⋆ E-mail: nicola.c.amorisco@durham.ac.uk
† Ernest Rutherford Fellow.

Strong gravitational lensing has emerged as an independent way to quantify the abundance of low-mass DM haloes and thus constrain the WDM particle mass. This technique uses galaxy–galaxy strong gravitational lenses (e.g. Bolton 2005; Shu et al. 2016) to detect low-mass haloes through the perturbations they cause to the lens image (see also the alternative approach based on flux ratio anomalies of lensed quasars; e.g. Xu et al. 2015; Gilman et al. 2018, 2019, 2020; Harvey et al. 2020). These perturbations make it possible to detect both satellite haloes in the main lens (subhaloes) and low-mass 'central' haloes along the line of sight (LOS; Li et al. 2017; Despali et al. 2018), even if they contain negligible baryonic mass. In fact, in the mass range of interest, $\lesssim 10^8 \, \mathrm{M}_\odot$, haloes are too small to have made a galaxy and so are completely dark (Benitez-Llambay & Frenk 2020). This is a great advantage as the abundance and structure of isolated DM haloes are unaffected by complications associated with baryonic processes and are very robustly determined by cosmological simulations. Distortions of strong lenses are therefore an especially clean way to probe the DM particle mass.

A small number of detections have already been claimed, albeit for subhaloes more massive than those that can test WDM models (Vegetti et al. 2010, 2012; Hezaveh et al. 2016). The most challenging aspect of lensing lies in using these detections – and non-detections as well – to extract quantitative inferences about the DM model (e.g. Vegetti et al. 2014, 2018; Ritondale et al. 2019). To do so, it is necessary to formulate robust predictions. For example, how many detections are to be expected in a given lensing system assuming CDM, or as a function of the WDM particle mass?

Quantifying the number of detectable haloes, N_d, in a specific lens means identifying which DM haloes, out of the cosmological population of haloes, can cause 'observable' perturbations to that system. More specifically, for a WDM particle with cut-off mass, M_cut,

$$N_\mathrm{d}(M_\mathrm{cut}) = \int n(x, y, z, \zeta_\mathrm{h} | M_\mathrm{cut}) \times p(x, y, z, \zeta_\mathrm{h} | \theta, \mathbf{n}) \, \mathrm{d}V \mathrm{d}\zeta_\mathrm{h}, \quad (1)$$

where n is the cosmological number density of DM haloes[1] at sky-projected location, (x, y), redshift, z, and with properties, ζ_h (such as mass, concentration, etc.), while p is the probability of actually detecting such haloes were they to be truly present in the observed system, given the properties of the lens itself, θ, and those of the data – for instance, the noise properties, $\mathbf{n}$. In other words, were a halo to be truly present:

(i) $p = 0$ if its perturbations are too small to be observable, implying that a perturbing halo mass component would not be required in the modelling to describe the data;

(ii) $p = 1$ if its perturbations make a model including a perturbing mass component statistically preferable to one that does not.

The increase in Bayesian evidence between the two models (or the increase in log-likelihood) is often used as a deciding metric, and most studies (e.g. Vegetti et al. 2014, 2018) have indeed reduced p to a binary classification: Were the halo to be truly present, $p = 1$, if including a perturbing mass component causes the Bayesian evidence or log-likelihood to increase beyond some given threshold. This is usually referred to as the *sensitivity function* and, simply put, it identifies the region of parameter space (comprising both physical cosmological volume and halo properties) that can be actually probed by lensing. In contrast to the cosmological number density of DM

haloes, n, the sensitivity function itself does not directly depend on the DM model,[2] but it does shape expectations for the number of detectable haloes, N_d – as well as expectations for any other observable obtained through structure lensing studies.

While advanced tools to model optical strong lensing data have been developed (e.g. Vegetti & Koopmans 2009; Nightingale et al. 2021), it remains computationally expensive to calculate the sensitivity function and formulate these predictions. Systematic exploration is required to establish the range of properties that make a perturber detectable. A minimum list of the independent variables includes the halo mass and halo concentration, the projected location of the halo with respect to the lensing system, and its redshift. In addition, the sensitivity function is unique to each individual lensing system because degeneracies in the lensing effects are such that different lensing configurations can 'reabsorb' the perturbations of identical DM haloes with different efficiencies. In practice, mapping the entire parameter space for each lens is often computationally prohibitive, and a number of simplifying assumptions have been used to obtain estimates of the integral in equation (1). Here, we explore the effect of these simplifications. Among the independent variables mentioned above, halo concentration and halo redshift are the most important.

For instance, Minor et al. (2021) have recently shown that the halo concentration must be included as a free parameter when modelling a perturber: If the concentration is fixed, the inferred perturber's mass may be biased by a factor of up to 6. They also show that higher halo concentrations make perturbers more easily detectable, as the lensing effect of any mass distribution is driven by its surface density. However, the intrinsic scatter in the concentration of DM haloes (e.g. Neto et al. 2007; Ludlow et al. 2016; Wang et al. 2020) has so far been ignored in sensitivity mapping studies; instead, the mass–concentration relation has been collapsed on to the concentration axis entirely, forcing all haloes on to the mean value for their mass. Additionally, the dependence of the mean mass–concentration relation on the DM model itself has also been neglected (e.g. Vegetti et al. 2018; Ritondale et al. 2019). This latter assumption leaves cosmological halo abundances as the single measure to differentiate between WDM models of different particle masses, despite the fact that warmer DM models produce haloes that are increasingly less dense than their equal-mass CDM counterparts (e.g. Lovell et al. 2012; Bose et al. 2016).

As regards the perturber's redshift, this axis has often been collapsed by adopting a one-to-one scaling relation that recasts a halo's redshift in terms of its effective mass (Li et al. 2017; Despali et al. 2018). This is obtained by requiring that the lensing convergence – i.e. the strength of the lensing effect – should remain nearly constant. This is not equivalent to performing a full non-linear search, as done on real data, and therefore does not fully take into account the modelling degeneracies that can occur in the real case. Lastly, we also briefly reflect on the influence of noise properties and the role of a specific noise realization. We should stress that the above issues pertain specifically to galaxy–galaxy lensing analyses focused on individual perturber detections (e.g. Vegetti et al. 2014; Hezaveh et al. 2016). Some recent analyses based on forward-modelling approaches and studies of flux ratio anomalies take full account of the effects above, including multiplane lensing effects and the scatter in the mass–concentration relation (see e.g. Gilman et al. 2019; He et al. 2020).

---

[1] Here, we focus on LOS haloes. The same formalism applies to subhaloes in the main lens, albeit with a different density, n.

[2] Although it does depend on the density profile of the perturbing haloes, which, in turn, may itself depend on the DM model.

To explore the parameter space of a sensitivity function fully, we use mock data and models that are somewhat simpler than those employed in state-of-the-art lens modelling studies (e.g. Vegetti & Koopmans 2009; Nightingale et al. 2019; Powell et al. 2021). Specifically, we use models featuring parametric sources rather than non-parametric pixelized sources (e.g. Warren & Dye 2003; Dye & Warren 2005; Nightingale & Dye 2015; Birrer, Amara & Refregier 2016). This allows us to develop and apply a fitting procedure based on a gradient descent algorithm, which is efficient enough to enable the exploration of all of the independent dimensions of the parameter space relevant to this problem. This work compares how previous assumptions regarding the sensitivity function affect the power of strong lensing to discriminate between different DM models, which we assume is not strongly dependent on the specific approach to source modelling.

In this work, we concentrate on LOS isolated perturbers, which we model as pure Navarro–Frenk–White (NFW) profiles (Navarro, Frenk & White 1997). Satellite subhaloes have different density profiles and their number density in the main lens is affected by a variety of physical processes. He et al. (in preparation) use the high-resolution cosmological hydrodynamical simulations of Richings et al. (2021) to facilitate a comparison between the lensing perturbations caused by satellite and LOS haloes, thereby allowing an estimate of their relative importance. In their study, they also employ different sets of lens configurations and modelling and fitting techniques. While they do not focus on the effect of halo concentration, their independent analysis finds quantitatively similar results on the redshift dependence of a perturber's detectability, which further reinforces the need to move away from the approximations used so far.

We stress that this work is not intended as a substitute for analyses aimed at quantifying the sensitivity function of actual sets of observed lenses – which should be tailored to the lens configurations featured in the real data and should be performed using the same modelling techniques as applied to the real data.

This paper is organized as follows: Section 2 provides a quick overview of the standard procedure used to estimate the sensitivity function; Section 3 describes our modelling framework and fitting procedures; Section 4 describes our results, focusing on the dependence of halo detectability on redshift and concentration; Section 5 uses our sensitivity maps to estimate the number of expected detections for different WDM models; and Section 6 examines the consequences of our results for future substructure lensing studies.

2 SENSITIVITY MAPPING OVERVIEW

In the interest of clarity, we start by outlining the procedure usually adopted to measure the sensitivity function. Let us assume, for example, that we wish to predict the number of detectable haloes (equation 1) for a specific strong lens. One would start by modelling the lens image in order to infer (i) a mass model for the lens galaxy and (ii) a light model for the source galaxy. Then,

(1) using these and the noise properties of the data themselves, one simulates a strong lens image that includes a perturbing DM halo with a set of properties (such as mass and concentration), located at projected location, (x, y), and redshift, z;

(2) one fits these mock data in two full but distinct non-linear searches, the first with a model that includes a perturbing halo mass component and the second without.

(3) one compares the two model fits, by means of the Bayesian evidence or the maximum log-likelihood. If a model including a halo

mass component provides a significantly better fit, the original data are sensitive to a halo with those specific properties, thereby mapping the probability of detection, p.

This procedure is repeated multiple times so as to sample the entire parameter space of perturbers' locations and properties.

We delve into the different steps of this procedure in the next section. In particular, we shall show that, in practice, we do not need to perform one of the fits at all.

3 MODELLING FRAMEWORK

We assume we have optical (mock) data, $\mathbf{d}$, for a lensing system characterized by the presence of some perturbing LOS halo with properties, ζ_h, and we wish to assess its detectability.[3] We do so by quantifying the log-likelihood[4] difference

$$\Delta \mathcal{L}(\zeta_h) = \mathcal{L}_{m,h}\left(\hat{\theta}_m, \hat{\zeta}_h\right) - \mathcal{L}_m\left(\bar{\theta}_m\right), \tag{2}$$

where

(i) $\mathcal{L}_m\left(\bar{\theta}_m\right)$ is the log-likelihood value corresponding to the best-fitting model that does not include a perturbing halo mass component. This model is optimized over the parameters of the so-called macromodel alone, θ_m, which include the parameters describing the source, θ_s, and those describing the lens, as well as any external shear, θ_l;

(ii) $\mathcal{L}_{m,h}\left(\hat{\theta}_m, \hat{\zeta}_h\right)$ is the log-likelihood value corresponding to the best-fitting model that does include a perturbing halo mass component, which is optimized over both macromodel and halo parameters, with best-fitting values, $\hat{\theta}_m$ and $\hat{\zeta}_h$, respectively.

We take the detection probability, $p(\zeta_h)$, in equation (1) to be a function of the log-likelihood gain, $\Delta \mathcal{L}$:

$$p(\zeta_h) = p(\Delta \mathcal{L}(\zeta_h)); \tag{3}$$

i.e. perturbing haloes that result in higher values of the log-likelihood gain, $\Delta \mathcal{L}$, are more easily detectable; for the moment, we defer fixing the functional link between detection probability and log-likelihood increase.

All likelihood values are obtained by comparing the pixelized flux values, $\mathbf{d}$, with the model flux distributions, $\mathbf{f}$. We ignore the effect of noise covariance [due to effects like point spread function (PSF) convolution] so that we simply have

$$\mathcal{L} = -\frac{1}{2} \sum_{\text{pixels}} \left| \frac{\mathbf{d} - \mathbf{f}}{\mathbf{n}} \right|^2, \tag{4}$$

where $\mathbf{n}$ represents the noise map associated with the data and we can neglect the normalization term in $\ln \mathbf{n}$ because we are only interested in log-likelihood differences. We assume that the data only include flux from the source; i.e. that both sky background and lens fluxes have been subtracted before performing the fits.

A number of authors (e.g. Vegetti et al. 2018; Ritondale et al. 2019) have argued for adopting the gain in Bayesian evidence, rather than in the log-likelihood, as a basis for quantifying halo detectability. We recall that the evidence is defined as the integral of the posterior over the entire parameter space. We agree that the gain in log evidence is a sounder statistical metric for model comparison.

---

[3] For compactness of notation, we include the perturber's sky coordinates, (x, y), and redshift, z, in the halo properties, ζ_h.

[4] By 'log-likelihood' we mean the natural logarithm of the likelihood. All other instances of 'log' in this work represent the base 10 logarithm.

However, calculating the evidence is orders of magnitude more computationally expensive than identifying the best-fitting model, as it requires sampling the likelihood surface over the entire parameter space in order to integrate it. This has become one of the main reasons behind the need for making simplifications when calculating the sensitivity function.

It is important to stress that, as long as the same criterion is consistently employed to both (i) detect perturbers on real data and (ii) measure the sensitivity function and make predictions for the expected number of detections, an evidence-based strategy and a likelihood-based one are both perfectly acceptable and will both yield correct results for the DM particle mass. It is indeed possible that an evidence-based criterion may help to weed out false detections better in studies of real data. On the other hand, this is not strictly necessary if the models used in sensitivity mapping share the same complexity of the real data, and therefore also share any spurious detections. In fact, a computationally more efficient criterion may facilitate a robust characterization of the properties and frequency of false positives, and therefore help to take them into account when making predictions. For the present purposes, a likelihood-based criterion is beneficial in that it allows us to explore the entire parameter space systematically. Furthermore, for data of sufficiently high quality, the gain in evidence and in log-likelihood become equivalent (as embodied by the BIC; see e.g. Hastie, Tibshirani & Friedman 2008; Konishi & Kitagawa 2007), and numerical experiments seem to indicate that the quality of the *Hubble Space Telescope* (*HST*) data is, in fact, high enough for the two approaches to often provide very similar results (He et al., in preparation).

3.1 Model families and mock data

Our lensing systems feature the following:

(i) a power-law mass profile to represent the main lens (Tessore et al. 2016), characterized by the following free parameters: (x_l, y_l), the centre of the mass distribution; ϵ_l, the Einstein radius; β_l, the slope of the mass profile; $(e_{1,l}, e_{2,l})$, the two independent components of the profile's ellipticity; and (γ_1, γ_2), the two independent components of the external shear.

(ii) a parametric source with a Sersic profile: projected centre, (x_s, y_s); effective radius, r_{eff}; ellipticity, $(e_{1,s}, e_{2,s})$; Sersic index, n_s; and total flux, I_s.

These two components define our macromodel: θ_m is therefore a 15-dimensional vector.

The perturbing haloes are modelled with spherically symmetric NFW mass profiles (Navarro et al. 1997), introducing the following additional five parameters: projected centre, (x_h, y_h); redshift, z_h; mass, M_h; and concentration c_h. Throughout the paper, we take halo masses, M_h, to be the virial mass, M_{200}, i.e. the mass contained within a sphere of density 200 times the critical density. We use the open-source software PYAUTOLENS[5] (Nightingale & Dye 2015; Nightingale, Dye & Massey 2018; Nightingale et al. 2021) to generate all of our mock data and for all our lensing modelling.

Within this framework, we choose two different lensing configurations for our exploration. Both have their source at $z_s = 1$ and lens at $z_l = 0.5$, but one is in a quad configuration while the other features two asymmetric arcs. Values for the ground-truth macromodel parameters, θ_m, are recorded in Table A1; Fig. 1

---

[5]PYAUTOLENS is an open-source software and available at https://github.com/Jammy2211/PyAutoLens.
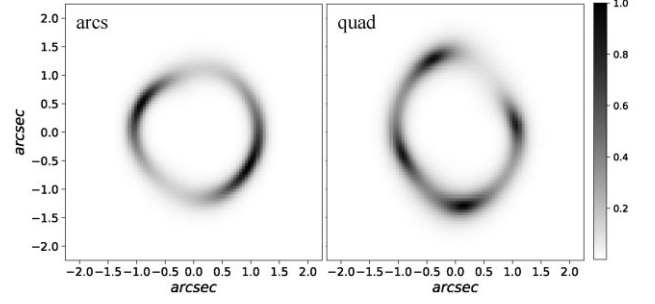
Figure 1. Model fluxes for the two lensing configurations used in this study, giant arcs (left) and quad (right) (linear scale with arbitrary units).

illustrates their geometry, displaying the corresponding model fluxes, $\mathbf{f}(\theta_m)$. The exact values are of no importance here: We simply choose sets of parameters that are qualitatively in line with what is found in lensing studies of real systems, and we use a pair of different configurations to ensure that the trends we identify are not a peculiarity of the specific system we happen to adopt. Throughout this work, we use pixel size and PSF width typical for the *HST* data, fixing both quantities at 0.05 arcsec.

3.2 Signal-to-noise and noise realization

The sensitivity function of a lens scales with the quality of the available data, quantified here by the noise map, $\mathbf{n}$, and the associated maximum value of the signal-to-noise, $\mathbf{n} \propto 1/\mathrm{SN}_{max}$. Notice, however, that while the noise map itself is known, the actual noise realization of the observed data is not. We are therefore interested in assessing how different noise realizations affect the value of the log-likelihood change.

Let us assume that the data are characterized by a noise realization, $\mathbf{r}$, so that $\mathbf{d} = \langle \mathbf{d} \rangle + \mathbf{r}$, where $\langle \cdot \rangle$ denotes an average over noise realizations. In the case of mock data, $\langle \mathbf{d} \rangle$ is the input model flux corresponding to the ground truth model parameters, which we indicate with (θ_m, ζ_h): $\langle \mathbf{d} \rangle = \mathbf{f}(\theta_m, \zeta_h)$. The log-likelihood gain is as in equation (2), with terms in the same order:

$$\Delta\mathcal{L}(\zeta_h, \mathbf{r}) = -\frac{1}{2}\left|\frac{\mathbf{d} - \mathbf{f}(\hat{\theta}_m, \hat{\zeta}_h)}{\mathbf{n}}\right|^2 + \frac{1}{2}\left|\frac{\mathbf{d} - \mathbf{f}(\bar{\theta}_m)}{\mathbf{n}}\right|^2, \quad (5)$$

where we have implicitly assumed a sum over image pixels. The fluxes, $\mathbf{f}(\hat{\theta}_m, \hat{\zeta}_h)$, and $\mathbf{f}(\bar{\theta}_m)$, correspond to the model that best fits the noise-corrupted data, $\mathbf{d}$, with and without an extra halo. It is convenient to consider instead the model fluxes that provide the best fit to the noise-free data, $\langle \mathbf{d} \rangle$, which we refer to as $\mathbf{f}_h$ and $\mathbf{f}_m$. These models do not achieve the maximum log-likelihood values we require in equation (5). For example, for the model including a halo mass component,

$$\frac{1}{2}\left|\frac{\mathbf{d} - \mathbf{f}(\hat{\theta}_m, \hat{\zeta}_h)}{\mathbf{n}}\right|^2 = \frac{1}{2}\left|\frac{\mathbf{d} - \mathbf{f}_h}{\mathbf{n}}\right|^2 - l_h^{bf}(\mathbf{r}), \quad (6)$$

where the difference, l_h^{bf}, is a function of the noise realization and has a positive value. Similarly,

$$\frac{1}{2}\left|\frac{\mathbf{d} - \mathbf{f}(\bar{\theta}_m)}{\mathbf{n}}\right|^2 = \frac{1}{2}\left|\frac{\mathbf{d} - \mathbf{f}_m}{\mathbf{n}}\right|^2 - l_m^{bf}(\mathbf{r}). \quad (7)$$

Furthermore, in order to highlight the dependence on the noise realization, $\mathbf{r}$, we can recast the model fluxes, $\mathbf{f}_h$ and $\mathbf{f}_m$, in terms

of their associated residuals, $\mathbf{f}_h = \langle \mathbf{d} \rangle + \delta_h$ and $\mathbf{f}_m = \langle \mathbf{d} \rangle + \delta_m$, which gives

$$\frac{1}{2} \left| \frac{\mathbf{d} - \mathbf{f}\left(\hat{\theta}_m, \hat{\zeta}_h\right)}{\mathbf{n}} \right|^2 = \frac{1}{2} \left| \frac{\mathbf{r} - \delta_h}{\mathbf{n}} \right|^2 - l_h^{bf}(\mathbf{r}), \tag{8}$$

and

$$\frac{1}{2} \left| \frac{\mathbf{d} - \mathbf{f}\left(\bar{\theta}_m\right)}{\mathbf{n}} \right|^2 = \frac{1}{2} \left| \frac{\mathbf{r} - \delta_m}{\mathbf{n}} \right|^2 - l_m^{bf}(\mathbf{r}). \tag{9}$$

Equation (5) is therefore the difference between the right-hand sides of the two equations above. We are interested in the mean and the standard deviation of this difference across noise realizations.

Let us first consider the two shifts, l_h^{bf}, and l_m^{bf}. These are non-zero, but they are not the leading terms of equation (5) that we seek. We can show this by estimating their average magnitude across varying noise realizations. This can be calculated analytically under the assumption that the likelihood surface is Gaussian. If so, we can see that

$$\langle l_{bf} \rangle = \frac{k}{2} \ln 2, \tag{10}$$

which is valid for both l_h^{bf} and l_m^{bf}, and in which k is the number of independent parameters in the likelihood. In our case, the model that includes a halo mass component features five additional free parameters, so that

$$\langle l_h^{bf} - l_m^{bf} \rangle = \frac{5}{2} \ln 2 \approx 1.73. \tag{11}$$

This is considerably smaller than the log-likelihood differences we are after and we therefore ignore these terms from now on.

By expanding the chi-square terms in equations (8) and (9), we finally obtain the leading terms we are interested in:

$$\Delta \mathcal{L}(\zeta_h, \mathbf{r}) \approx \frac{1}{2} \left(\left| \frac{\delta_h}{\mathbf{n}} \right|^2 - \left| \frac{\delta_m}{\mathbf{n}} \right|^2 \right) + \frac{\mathbf{r}}{\mathbf{n}} \cdot \frac{\delta_h - \delta_m}{\mathbf{n}}. \tag{12}$$

Here, the first term quantifies the inability of a model that does not include a perturbing halo mass component to describe the perturbed data. This is what sensitivity mapping is after, and is independent of the noise realization. The second term introduces scatter in the measurement of the log-likelihood gain as a consequence of varying noise realizations, $\mathbf{r}$. The case we are interested in is the one in which a model featuring a halo mass component provides a satisfactory fit to the data, $\delta_h/\mathbf{n} \ll 1$. This is also the case of mock data in which the model used to generate data is the same used to fit it: $\langle \mathbf{d} \rangle = \mathbf{f}_h$. In this case,

$$\Delta \mathcal{L}(\zeta_h, \mathbf{r}) \approx \frac{1}{2} \left| \frac{\delta}{\mathbf{n}} \right|^2 + \frac{\delta}{\mathbf{n}} \cdot \frac{\mathbf{r}}{\mathbf{n}}, \tag{13}$$

where we have used $\delta \equiv \delta_m$ for compactness. By definition, the noise realization, $\mathbf{r}$, is a random variable with zero mean; furthermore, by construction, the residuals, δ, are not correlated with $\mathbf{r}$. As a result, the second term in equation (13) averages to zero:

$$\langle \Delta \mathcal{L}(\zeta_h) \rangle \approx \frac{1}{2} \left| \frac{\delta}{\mathbf{n}} \right|^2. \tag{14}$$

We can estimate the magnitude of the scatter introduced by the same term by taking the ratio, $\mathbf{r}/\mathbf{n}$, to be a set of independent normal random variables with unit variance, which results in a standard deviation of

$$\text{std}(\Delta \mathcal{L}(\zeta_h)) \approx \sqrt{\left| \frac{\delta}{\mathbf{n}} \right|^2} \sim \sqrt{2 \langle \Delta \mathcal{L}(\zeta_h) \rangle}. \tag{15}$$

We test this scaling in Appendix B, where Fig. B1 shows experiments that highlight the scaling of equation (15).

In conclusion, from equation (14) we deduce that the mean log-likelihood increase scales with the square of the maximum signal-to-noise ratio, SN_{max}, and we note that the Bayesian evidence will also feature in the same scaling. From equations (12) and (13), we see that, in real data, a scatter of the order of $\sqrt{2\Delta \mathcal{L}}$ should be expected. In fact, the same scatter should be expected when mapping the sensitivity function using mock data that include a random noise realization. This implies that multiple noise realizations should be used and results averaged. However, the above analysis also shows that this can be avoided by using noise-free mock data (i.e. $\mathbf{r} = 0$), while at the same time using the appropriate noise map, $\mathbf{n}$, featuring the same maximum signal-to-noise as in the real data. This is the strategy we adopt in this work.

3.3 Fitting procedure

Having chosen our macromodels, θ_m, we can introduce intervening LOS haloes with input parameters, ζ_h, and simulate the resulting model fluxes, $\mathbf{f}(\theta_m, \zeta_h) = \mathbf{f}$. As outlined in Section 2, each determination of the likelihood gain, $\Delta \mathcal{L}(\zeta_h, \mathbf{r})$, requires two non-linear searches. However, in our case, $\mathbf{r} = 0$, so that we have $(\hat{\theta}_m, \hat{\zeta}_h) = (\theta_m, \zeta_h)$, or equivalently, $\mathcal{L}(\hat{\theta}_m, \hat{\zeta}_h) = 0$, by construction, and therefore,

$$\Delta \mathcal{L}(\zeta_h) = \mathcal{L}\left(\bar{\theta}_m\right). \tag{16}$$

Thus, for each set of halo parameters, ζ_h, we only require one non-linear search in order to determine the best-fitting parameters, $\bar{\theta}_m$, of the model that does not include a halo mass component. This is also the fit with fewer free parameters – and therefore both the fastest to run, and the least likely to get stuck in local minima during a likelihood optimization.

3.3.1 Gradient descent approach

We perform these optimizations using an iterative gradient descent algorithm. In essence, at each step, i, provisional estimates of the best-fitting parameters, θ_m^i, and corresponding model fluxes, $\mathbf{f}_i$, are used to calculate the increments, $\delta\theta_m^i$, which provide the best linear improvement of the model fluxes themselves. That is, the increment, $\delta\theta_m^i$, minimizes the log-likelihood,

$$\Delta \mathcal{L}_{i+1} = \frac{1}{2} \sum_{\text{pixels}} \left| \frac{1}{\mathbf{n}} \left[\mathbf{d} - \left(\mathbf{f}_i + \left. \frac{\partial \mathbf{f}}{\partial \theta_m} \right|_{\theta_m^i} \cdot \delta\theta_m^i \right) \right] \right|^2, \tag{17}$$

where $\partial \mathbf{f}/\partial \theta_m|_{\theta_m^i}$ is the gradient of the model fluxes calculated at θ_m^i. This minimization is easily solved by the corresponding least-squares problem. The parameters at the subsequent step are therefore $\theta_m^{i+1} = \theta_m^i + \eta\delta\theta_m^i$, where η is the so-called learning rate. Iterations are stopped when the corresponding likelihood value converges. In order to avoid convergence at possible local maxima, we repeat the procedure over a set of different initialization parameters, close to the input parameters, θ_m. In practice, we find this to be rarely necessary, possibly because for most perturbing haloes the best-fitting parameters, $\bar{\theta}_m$, are sufficiently close to the input values themselves, and the log-likelihood surface is smooth in our noise-free setting.

Despite allowing for this redundancy, we find gradient descent to be efficient and inexpensive for models featuring parametric sources. This is because the (noise-free) gradient maps, $\partial \mathbf{f}/\partial \theta_m$, are well behaved and easily estimated. In contrast, this is not so when non-parametric pixelized sources are used. We have tried using gradient descent to optimize the parameters of the lens while, at each iterative

step, a linear source inversion (e.g. Warren & Dye 2003; Dye & Warren 2005) determines the source model. However, we find this approach to be unsuitable, despite the fact that in this case, gradient descent is used on a significantly smaller parameter space (featuring 8 dimensions instead of 15). Due to the nature of the semilinear source inversion, the residuals, $\mathbf{d} - \mathbf{f}_i$, contain little information on the mass model itself, unless unrealistically high regularization values (e.g. Suyu et al. 2006) are used.

4 MAPPING THE SENSITIVITY FUNCTION

For each of the two macromodels we investigate, we map the log-likelihood gain, $\Delta\mathcal{L}$, over the space of halo parameters, $\zeta_h = (x_h, y_h, z_h, M_h, c_h)$, using a rectangular grid as follows:

(i) The halo mass is varied between $8.0 \leq \log M_h/M_\odot \leq 10.0$, at intervals of 0.5 dex;

(ii) Halo redshift is varied between $0.05 \leq z_h \leq 0.95$, at intervals of $\delta z = 0.15$.

(iii) Halo concentrations deviate from the mass–concentration relation between $\log c - \log c(M, z) \equiv \delta \log c = 4\sigma_{\log c}$ and $\delta \log c = -2\sigma_{\log c}$, in intervals of $\sigma_{\log c}$. Here, $\sigma_{\log c}$ is the lognormal scatter of the mass–concentration relation, which we take to be independent of halo mass and redshift, and fix at $\sigma_{\log c} = 0.15$ dex (Wang et al. 2020). For this exploration, we assume that $\delta \log c = 0$ means the mass–concentration relation, $c(M, z)$, of CDM haloes, as measured by Ludlow et al. (2016).

(iv) Projected locations, x_h, y_h, are mapped over 50 intervals in both coordinates. We scale the total extent of our maps with redshift so as to achieve better spatial resolution in the (x, y) plane when $z_h > z_l$.

Figs 2 and 3 illustrate some of our maps as sky-projections of the log-likelihood increase, $\Delta\mathcal{L}$, for our 'arcs' and 'quad' configurations, respectively. The first figure shows results for a perturbing halo of mass $M_h = 10^{10}\,\mathrm{M}_\odot$, the second for $M_h = 10^{9.5}\,\mathrm{M}_\odot$. Columns correspond to different values of the perturber redshift, z_h. Rows are for different halo concentrations.

4.1 Dependence on redshift

A common assumption in sensitivity mapping is that the perturber's redshift can be recast in terms of its effective mass. Li et al. (2017) noticed that the mass and redshift of a perturber are highly degenerate, and used this to introduce the idea of rescaling a perturber's mass as a function of redshift. Despali et al. (2018) investigated this equivalence further and provided a universal scaling between redshift and effective mass, which is obtained by requiring that the map of deflection angles be minimally changed. In practice, at any fixed projected location, the perturbing characteristics of a halo of mass M_h at a redshift z_h have been equated to those of a halo located at the redshift of the lens and having an 'effective' mass of $\log M_{\mathrm{sh}} = \log M_h + \delta M(z_h)$. The mass shift, $\delta M(z_h)$, is clearly zero at $z_h = z_l$, and is found to be monotonically increasing with redshift, so that detecting a halo of fixed mass becomes more challenging with increasing redshift, and is easiest for close-by perturbers.

Figs 2 and 3 show that our calculations do not support this working hypothesis in previous work. When fitting image fluxes (rather than deflection angles), we find that haloes are harder to detect when they are behind or in front of the main lens. The details depend on the precise lensing configuration, mass and concentration of the perturbing halo, and on the precise projected location. However, a decrease in $\Delta\mathcal{L}$ when the halo redshift deviates from the lens redshift

is a universal qualitative feature in our maps displayed by both our adopted lensing configurations and across halo masses. It is worth stressing that, as found in previous work, perturbers in the foreground appear to be more easily detected than perturbers in the background. However, we do not find the sensitivity to increase with decreasing redshift at $z_h < z_l$. This means that, with respect to a perturber at the redshift of the lens, $z_h = z_l$, suitable changes to the macromodel can more easily reabsorb the effect of a foreground halo of equal mass when this is at low redshifts. In other words, degeneracies in the lens modelling make the detection of perturbers of the same mass increasingly more difficult with decreasing redshift.

Fig. 4 summarizes the redshift dependence of the log-likelihood increase, showing the ratio between the log-likelihood increase for a perturber at some redshift, $\Delta\mathcal{L}(z_h, M_h, c_h)$, divided by that at the redshift of the lens, $\Delta\mathcal{L}(z_h = z_l, M_h, c_h)$. This ratio is then averaged over the perturbers' projected locations, (x, y). The top row refers to our 'quad' configuration, and the bottom one to our 'arcs'. The right column is for a perturber of mass $M_h = 10^{10}\,\mathrm{M}_\odot$, the left column to one of mass $M_h = 10^{9.5}\,\mathrm{M}_\odot$. Profiles of different colours refer to different values of the perturbers' concentrations, corresponding to different rows in Figs 2 and 3, for $\delta_{\log c} \in \{-2, -1, 0, 1, 2, 3, 4\} \times \sigma_{\log c}$. As described, for most haloes, the log-likelihood increase *decreases* for redshifts that are both higher and lower than the lens' redshift, z_l. The size of this decrease depends systematically and monotonically on halo concentration, with the less-concentrated haloes displaying sharper fall-offs. This suggests that the decrease in sensitivity at lower redshifts is driven by the increasing angular size of the perturbers. Haloes at lower redshifts and/or with lower concentrations cause perturbations to the lensing signal over wider areas. In turn, these are more readily reabsorbed by 'global' changes of the macromodel and/or source properties.

Depending on the lensing configuration, for haloes on the mass–concentration relation, $\Delta\mathcal{L}(z_h)$ decreases by a factor between 1.15 and 1.6 between $z_h = z_l$ and 0.2, and then drops more sharply towards lower redshift. For the highest halo concentrations, we see, instead, a mild apparent increase in detectability at low redshift. However, analysis of the top rows of Figs 2 and 3 shows that this increase is due to the fact that Fig. 4 displays averages over projected locations. At the highest concentrations, the projected area in which perturbers result in 'intermediate' log-likelihood values, $50 \lesssim \Delta\mathcal{L} \lesssim 100$ (corresponding to orange hues in Figs 2 and 3), increases at the lowest redshifts. In the same regions, log-likelihood values are lower at $z_h = z_l$, which drives the mild increase apparent in the average quantities shown in Fig. 4. On the other hand, even at the highest concentrations, it remains true that the peak values of the log-likelihood gain, $\Delta\mathcal{L}$, *decrease* with decreasing redshift. In any case, this mild increase is limited to extremely concentrated haloes, and, therefore, is not a representative behaviour.

The qualitative contradiction between the predictions using deflection angle maps and our results implies that previous estimates of the number of detectable haloes obtained using the relation between mass and redshift proposed in Despali et al. (2018) are likely to overestimate the number of low-redshift haloes. We will return to this point in Section 5.2.

4.2 Dependence on concentration

Most previous studies have fixed the halo concentration to the mean for their mass for the adopted mass–concentration relation. However, Figs 2 and 3 clearly show that concentration makes a significant difference to halo detectability. From top to bottom, the values of the log-likelihood gain decrease monotonically: The perturbations
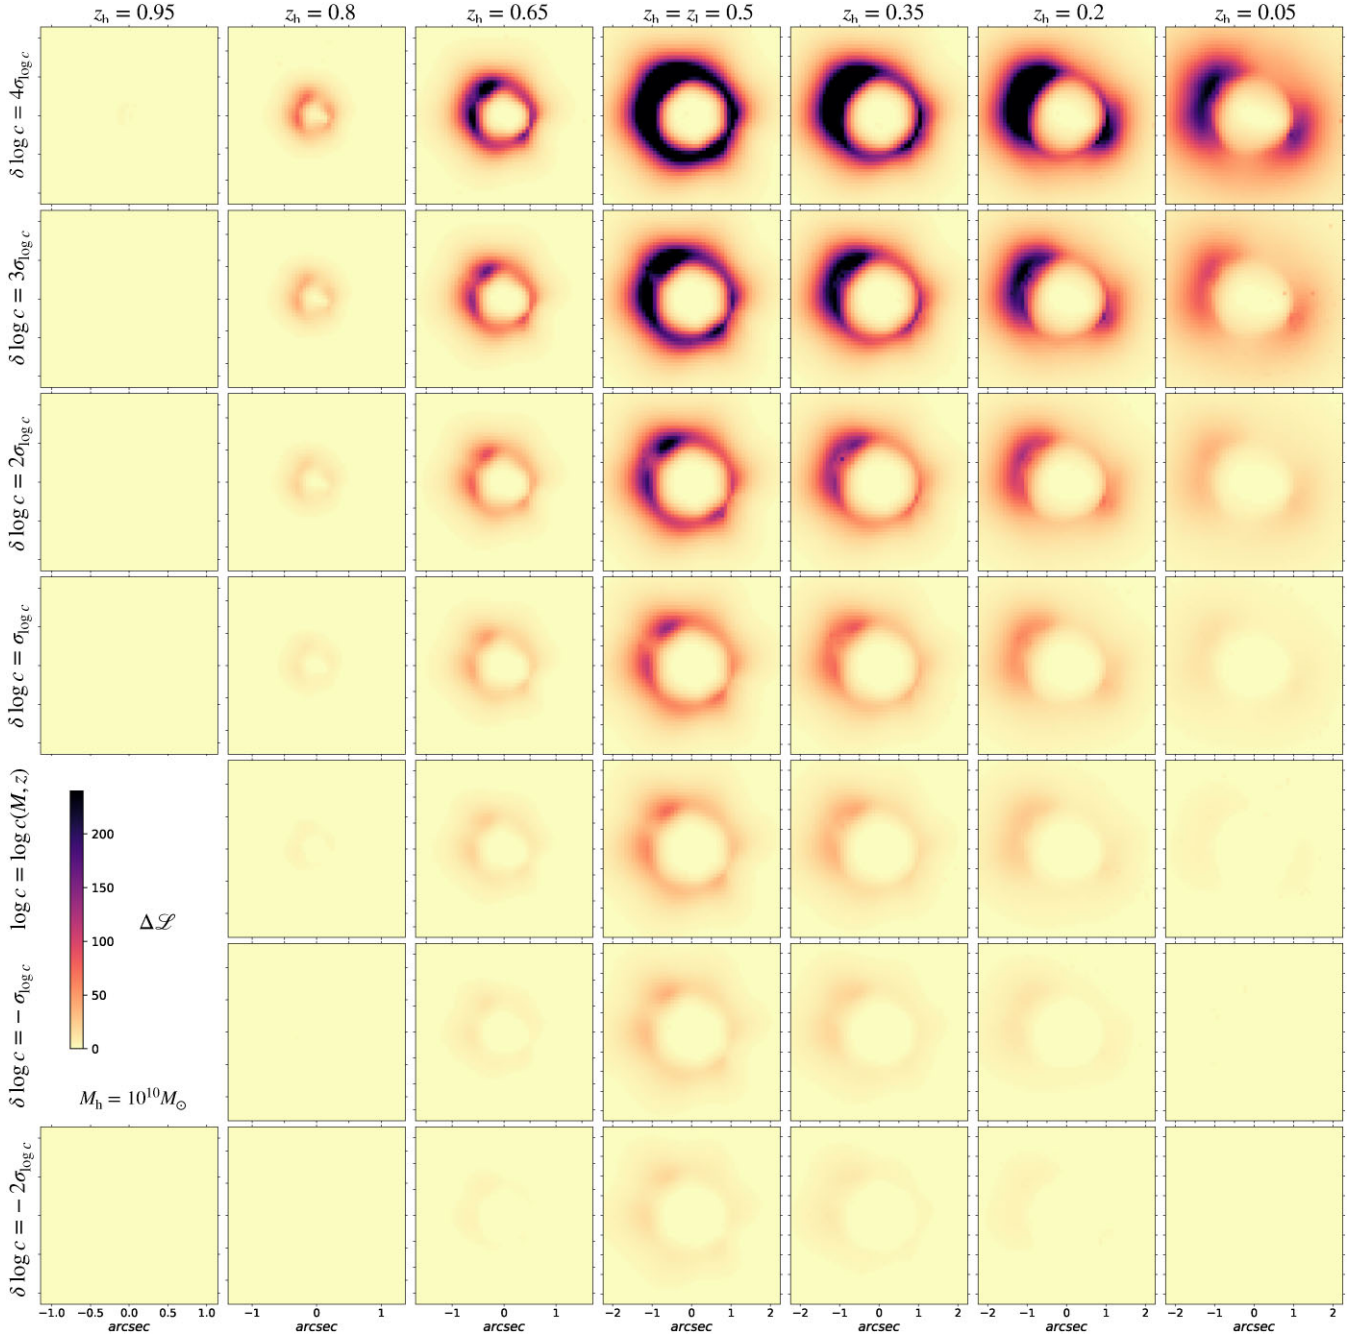
Figure 2. An illustration of the sky-projections of our maps of the log-likelihood increase, $\Delta\mathscr{L}$, for our 'arcs' lensing configuration. Columns show a grid of different perturber redshifts. Rows are for different perturber concentrations. The perturber mass is fixed at $M_{\rm h} = 10^{10}$ M$_\odot$ in all panels. Individual panels share the same colour scale. It is apparent that more concentrated haloes result in larger $\Delta\mathscr{L}$ values. Also, the $\Delta\mathscr{L}$ values decrease away from the redshift of the main lens, $z_{\rm l} = 0.5$, for both higher and lower perturber redshifts.

of less concentrated haloes are more easily reabsorbed by changes of the macromodel parameters. In turn, at fixed mass, more concentrated haloes are more easily detected.

Fig. 5 provides a summary of the dependence of the log-likelihood increase, $\Delta\mathscr{L}$, on halo concentration. This shows the ratio between the log-likelihood increase, $\Delta\mathscr{L}(z_{\rm h}, M_{\rm h}, c_{\rm h})$, for a perturber of any concentration, c, divided by that for one on the mass–concentration relation, $\Delta\mathscr{L}(z_{\rm h}, M_{\rm h}, c_{\rm h} = c(M_{\rm h}, z_{\rm h}))$. We reiterate that $\delta\log c = 0$ in our mapping of the sensitivity function corresponds to the mass–concentration relation measured by Ludlow et al. (2016). Similarly

to Fig. 4, these ratios are then averaged over projected locations, (x, y). The top row refers to our 'quad' configuration, the bottom to one to our 'arcs'. The right column is for a perturber of mass $M_{\rm h} = 10^{10}$ M$_\odot$, and the left column for one of mass $M_{\rm h} = 10^{9.5}$ M$_\odot$. Profiles of different colours refer to different values of the perturbers' redshift, corresponding to different columns in Figs 2 and 3. It is clear that $\Delta\mathscr{L}$ increases monotonically with concentration in all cases. The scalings appear qualitatively similar in all four panels, although, as for the dependence on redshift, quantitative details are still dependent on the lensing configuration and other halo parameters. In
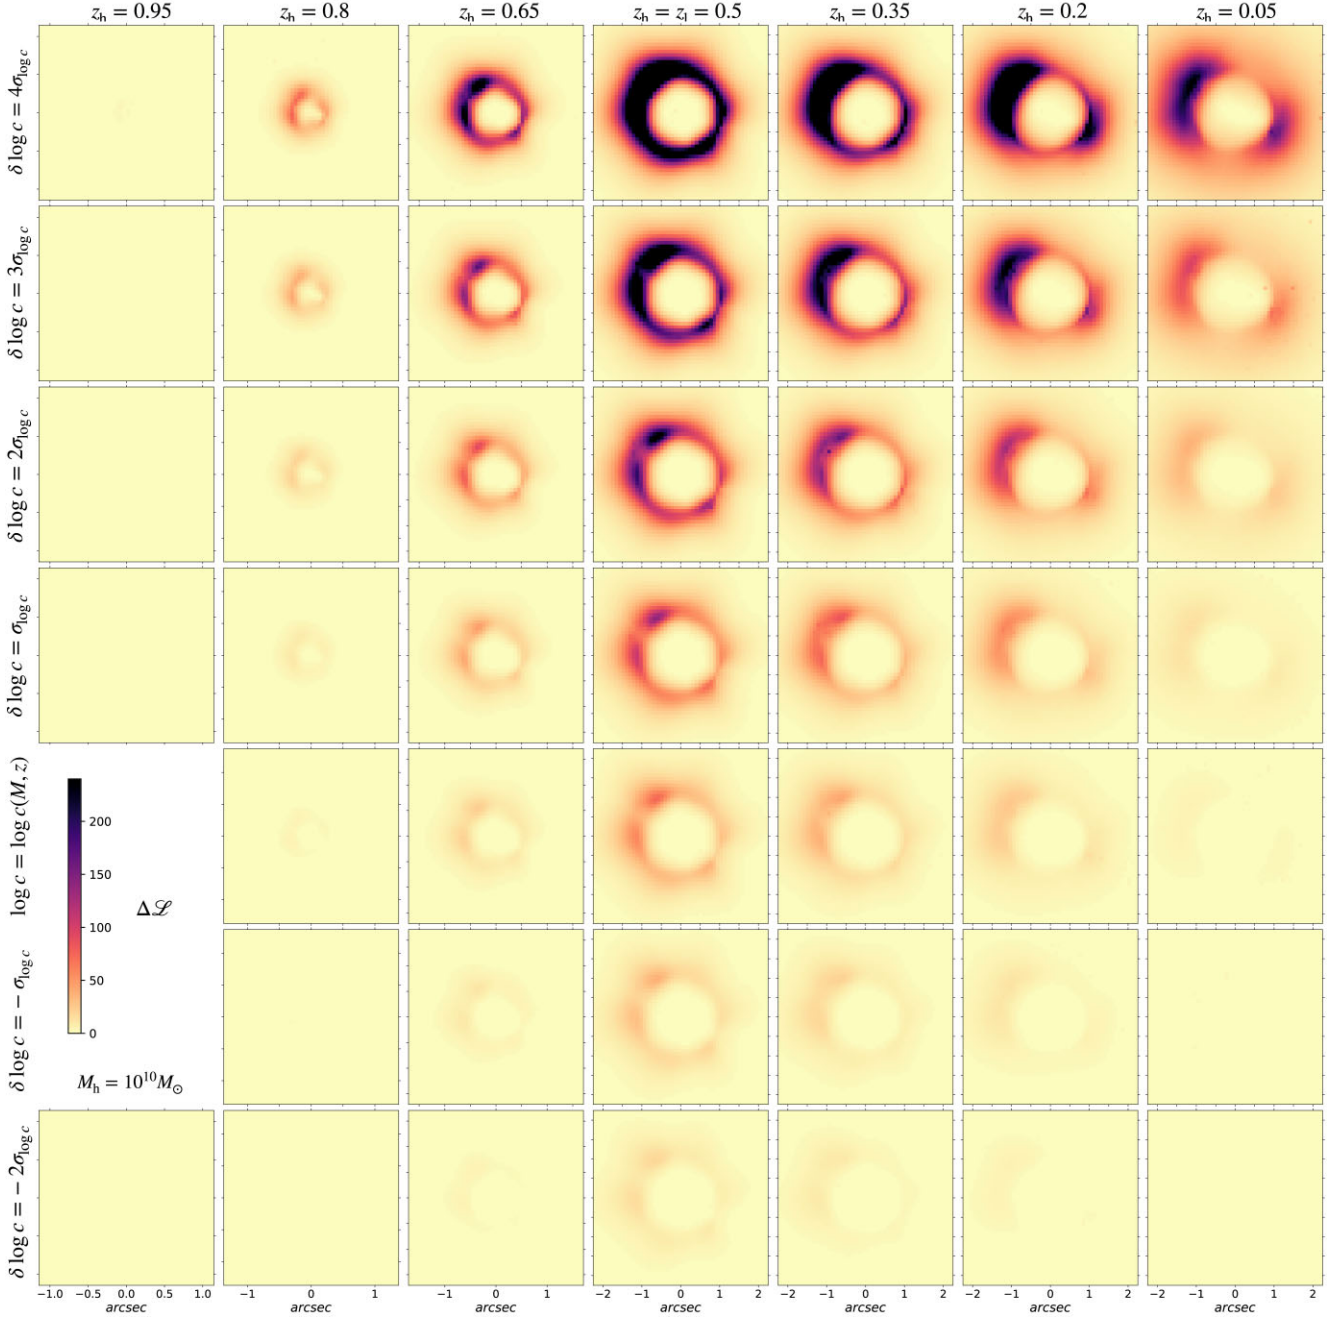
Figure 3. Same as Fig. 2, for our 'quad' lensing configuration. The perturber halo has a mass of $M_h = 10^{9.5}$ M$_\odot$.

particular, we record a significant secondary dependence on redshift: The detectability of haloes at the lowest redshifts is most strongly boosted by concentration. The magnitude of this boost then decreases for redshifts approaching the redshift of the lens, where it has a minimum, to then increase again towards higher redshifts.

Notably, we find the dependence on concentration to be essentially exponential, at least when averaged over projected locations:

$$\langle \Delta \mathscr{L}(\delta_{\log c}) \rangle_{(x,y)} \sim 10^{\alpha \cdot \delta_{\log c}} \langle \Delta \mathscr{L}(\delta_{\log c} = 0) \rangle_{(x,y)} . \qquad (18)$$

The top-right panel of Fig. 5 displays guiding lines for the exponent $\alpha = \{0.18, 0.28, 0.4\}$. The log-likelihood increase grows roughly by a factor between 1.5 and 2.5 for each additional $+1\sigma$ deviation from

the mass–concentration relation. We will analyse the consequences of this on the expected number of haloes in Section 5.3.

5 THE POPULATION OF DETECTABLE HALOES

Using our maps of the log-likelihood increase, $\Delta \mathscr{L}(\zeta_h)$, we are ready to perform the integral in equation (1). As done in previous work, we use a sharp threshold, $\Delta \mathscr{L}_{th}$, in the log-likelihood increase to separate detectable haloes from non-detectable haloes:

$$p(\zeta_h) = \begin{cases} 1 & \text{if } \Delta \mathscr{L}(\zeta_h) \geq \Delta \mathscr{L}_{th} \\ 0 & \text{if } \Delta \mathscr{L}(\zeta_h) < \Delta \mathscr{L}_{th}, \end{cases} \qquad (19)$$
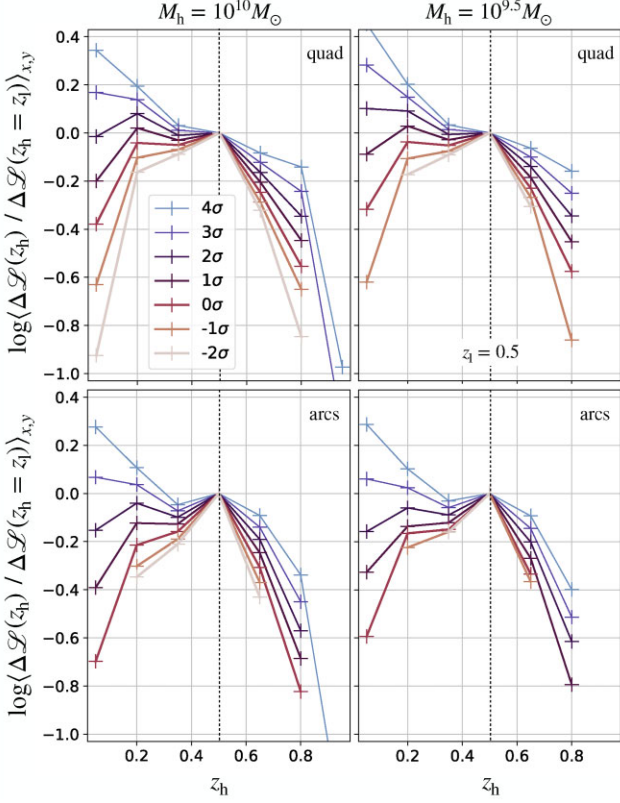
Figure 4. The average dependence of the log-likelihood increase, $\Delta\mathscr{L}$, on redshift (see the text). The top row refers to our 'quad' configuration, and the bottom one to our 'arcs'. The right column is for a perturber of mass $M_h = 10^{10}$ $M_\odot$, and the left column for one of mass $M_h = 10^{9.5}$ $M_\odot$. Profiles of different colours refer to different values of the perturber concentrations (in terms of the shift relative to the median concentration at the relevant redshift; see the text).

Figure 5. The average dependence of the log-likelihood increase, $\Delta\mathscr{L}$, on halo concentration (see the text). The top row refers to our 'quad' configuration, and the bottom one to our 'arcs' configuration. The right column is for a perturber of mass $M_h = 10^{10}$ $M_\odot$, and the left column to one of mass $M_h = 10^{9.5}$ $M_\odot$. Lines of different colours refer to different values of the perturber redshift.

although we notice that the scatter characterized in equation (15) would provide a natural scaling for a smooth transition in the detection probability. This will be useful when preparing detailed predictions for real data, but would not affect our conclusions here, so in this study we retain a sharp transition for simplicity.

We parametrize the cosmological number density of DM haloes as suggested by Lovell et al. (2014):

$$n(M_h, z|M_{cut}) = n_{CDM}(M_h, z)\left(1 + \frac{M_{cut}}{M_h}\right)^{-1.3}, \quad (20)$$

where $n_{CDM}(M_h, z)$ is the CDM halo number density, for which we adopt the form derived by Sheth, Mo & Tormen (2001). We take this distribution to be uniform over projected sky coordinates, and assume that the distribution in concentration is lognormal, with a spread of $\sigma_{\log c} = 0.15$ dex (independent of mass, redshift, and DM model). We take the median concentration to be dependent on the DM model, and adopt the parametrization proposed by Bose et al. (2016):

$$c(M_h, z|M_{cut}) =$$
$$c_{CDM}(M_h, z)\left[(1 + z)^{0.026z - 0.04}\left(1 + 60\frac{M_{cut}}{M_h}\right)^{-0.17}\right], \quad (21)$$

in which the median concentration of CDM haloes is as recorded by Ludlow et al. (2016).

These prescriptions allow us to calculate the integral of equation (1) using a Monte Carlo strategy. We randomly sample the

candidate haloes' ζ_h according to their cosmological number density, and then check whether they would be detectable using our maps of log-likelihood increase, which we linearly interpolate between our rectangular grid points.

5.1 The effect of data quality

Although our objective is not to provide absolute figures for the number of detectable haloes, Fig. 6 shows the cumulative distributions of detectable LOS haloes, N_d, we obtain for our two lens configurations. Both panels are for a CDM universe and are calculated from our full maps of the log-likelihood increase, that is including both the full redshift dependence and the scatter in the mass concentration relation. We have used $SN_{max} = 50$. We stress that these figures cannot be directly applied to real data analysed with different techniques and featuring different lensing configurations. For definitiveness, we include all haloes with projected coordinates in a 4.5 arcsec $\times$ 4.5 arcsec area at $z_h \leq z_l$, decreasing to 2.1 arcsec $\times$ 2.1 arcsec at $z_h = z_s = 1$, as displayed in Figs 2 and 3.

Our two lensing configurations provide substantially different numbers of detectable haloes: A quad configuration appears less prone to modelling degeneracies and therefore more promising for the detection of perturbers. The number of expected detections is also a strong function of the imposed detection threshold. In our quad lens configuration, thresholds of $\Delta\mathscr{L}_{th} = \{10, 20, 35, 50\}$ yield
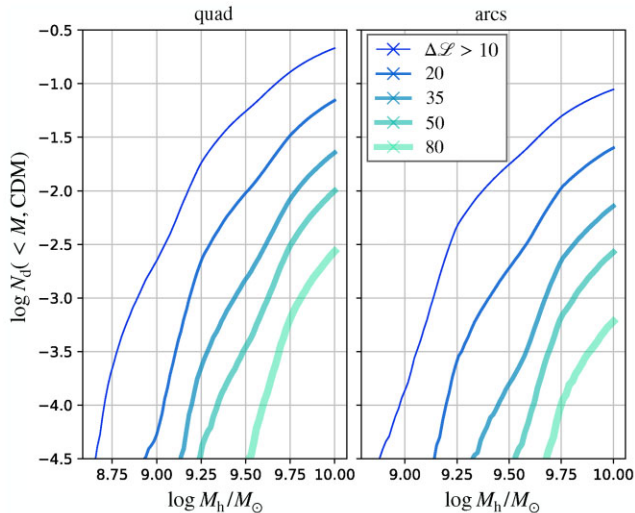
Figure 6. The cumulative number of detectable haloes, $N_d(< M)$, in a CDM universe. Lines of different colours refer to different values of the log-likelihood threshold required for detectability. The left-hand panel refers to our quad configuration, and the right-hand panel to our configuration featuring asymmetric arcs. The values displayed are for $SN_{max} = 50$.

$N_d = \{0.2, 0.07, 0.02, 0.01\}$ detections with $M_h < 10^{10}\, M_\odot$ per lens, respectively.

The analysis of Section 3.2 also allows us to address systematically the dependence of the total number of detected haloes with the quality of the data, which in this work we have characterized by the value of SN_{max}. Equation (14) allows us to equate changes in the value of SN_{max} with changes in the value of the log-likelihood threshold required for detection, $\Delta\mathcal{L}_{th}$:

$$N_d(SN_{max}, \Delta\mathcal{L}_{th}) = N_d(\alpha SN_{max}, \alpha^{-2}\Delta\mathcal{L}_{th}), \qquad (22)$$

for any factor α. We use this equivalence to focus on how the number of expected detections for a CDM universe would change for higher or lower values of the signal-to-noise, i.e. for longer or shorter exposure times.

Fig. 7 shows the number of expected detections of haloes of mass $M_h < 10^{10}\, M_\odot$ (left-hand panel) and $M_h < 10^{9.5}\, M_\odot$ (right-hand panel) for varying values of SN_{max}, normalized by the number of detections predicted for $SN_{max} = 50$. The figure displays the case of our 'quad' configuration; results for our 'arcs' configuration are similar, albeit with a more marked dependence on the SN itself. Lines of different colours refer to different values of the log-likelihood threshold required for detectability. As expected, the number of detectable haloes is a rapidly increasing function of the signal-to-noise ratio. We find that, for a likelihood threshold of $\Delta\mathcal{L}_{th} = 20$, an increase in signal-to-noise ratio from 50 to 60 corresponds to a doubling in the number of expected detections $N_d(M_h < 10^{10}\, M_\odot)$. The same increase in data quality makes the expected detections $N_d(M_h < 10^{9.5}\, M_\odot)$ grow by a factor of 3. As shown in the same figure, these factors are even larger if higher values of the log-likelihood ratio are required for detectability. For a value of $\Delta\mathcal{L}_{th} = 50$, analogous to what has been used in most previous studies, we find the corresponding figures to be 2.5 and 5 for haloes of $M_h < 10^{10}\, M_\odot$ and $M_h < 10^{9.5}\, M_\odot$, respectively. It should be noted that an increase in the maximum SN ratio from 50 to 60 corresponds to an increase in the exposure time of a factor of ≈ 1.44, which is therefore smaller than the corresponding gain in the number of detectable haloes.

5.2 The effect of redshift dependence

We now examine the consequences of dropping the simplifying assumption of a tight relationship between halo mass and redshift. We isolate this effect by considering a population of CDM haloes assumed to lie on the Ludlow et al. (2016) mass–concentration relation; i.e. with reference to the previous section, here we ignore the scatter in halo concentration. (We will consider this shortly.) Fig. 8 shows the redshift distribution of the population of detected haloes of mass $M_h < 10^{10}\, M_\odot$ we obtain when: for our two lensing configurations

(i) using the mass shift proposed by Despali et al. (2018) and described in Section. 4.1, shown by a dashed line[6];

(ii) using the full redshift dependence of our log-likelihood maps, shown by a solid line.

As expected, the two curves match at $z_h = z_l$, but we predict significantly fewer detections for foreground haloes, a reflection of the dependence on redshift of the log-likelihood increase described in Section 4.1. We also find that collapsing the redshift axis leads to an overestimate of detectable haloes also at $z_h > z_l$, though by a smaller factor.

The magnitude of the global overestimate varies with the lensing configuration. For definitiveness, we use a threshold of $\Delta\mathcal{L}_{th} = 20$ in Fig. 8. For the 'arcs' configuration, the overestimate is a factor of 1.95; for the quad configuration, it is a factor of 1.63. We stress that these figures should not be used to 'correct' previous measurements of the number of detectable haloes and are meant only as an estimate of the magnitude of the effect.

5.3 The effect of the scatter in concentration

We now consider the effect of accounting for scatter in the mass–concentration relation. We again focus on a population of CDM haloes, and compare the case in which all haloes are assumed to lie exactly on the median mass–concentration relation to the case in which a lognormal scatter is included. For definitiveness, in both cases we use the full redshift dependence of the log-likelihood increase, $\Delta\mathcal{L}$, and, for simplicity, we restrict attention to our quad configuration. Results are analogous for our 'arcs' lensing morphology.

Fig. 9 shows the distribution of detectable haloes in the space of halo mass and shift relative to the median halo concentration, for the case in which the scatter in concentration is accounted for. The two-dimensional histogram in the top panel is for an assumed threshold, $\Delta\mathcal{L}_{th} = 35$. The vertical dashed line shows the median mass–concentration relation, $\delta \log c = 0$. It is clear that, thanks to the dependence on concentration of the log-likelihood increase described in Section 4.2, most detectable haloes are high-concentration haloes. This is quantified in the bottom panel of the figure, which collapses the mass axis to show the distribution of detected haloes over concentration shifts. For reference, the dashed line shows the Gaussian distribution of all cosmological haloes. Coloured lines show the distribution of the population of detectable haloes for different values of the log-likelihood threshold, $\Delta\mathcal{L}_{th}$.

Haloes with high concentration achieve $\Delta\mathcal{L} > \Delta\mathcal{L}_{th}$ more easily, so that higher thresholds for detectability correspond to increasingly concentrated populations of detectable haloes. For $\Delta\mathcal{L}_{th} = 50$ and our quad configuration, we find $\langle \delta \log c / \sigma_{\log c} \rangle = 1.25$ when including

---

[6]We use the same mass shift, $\delta M(z)$, for all projected coordinates, (x, y).
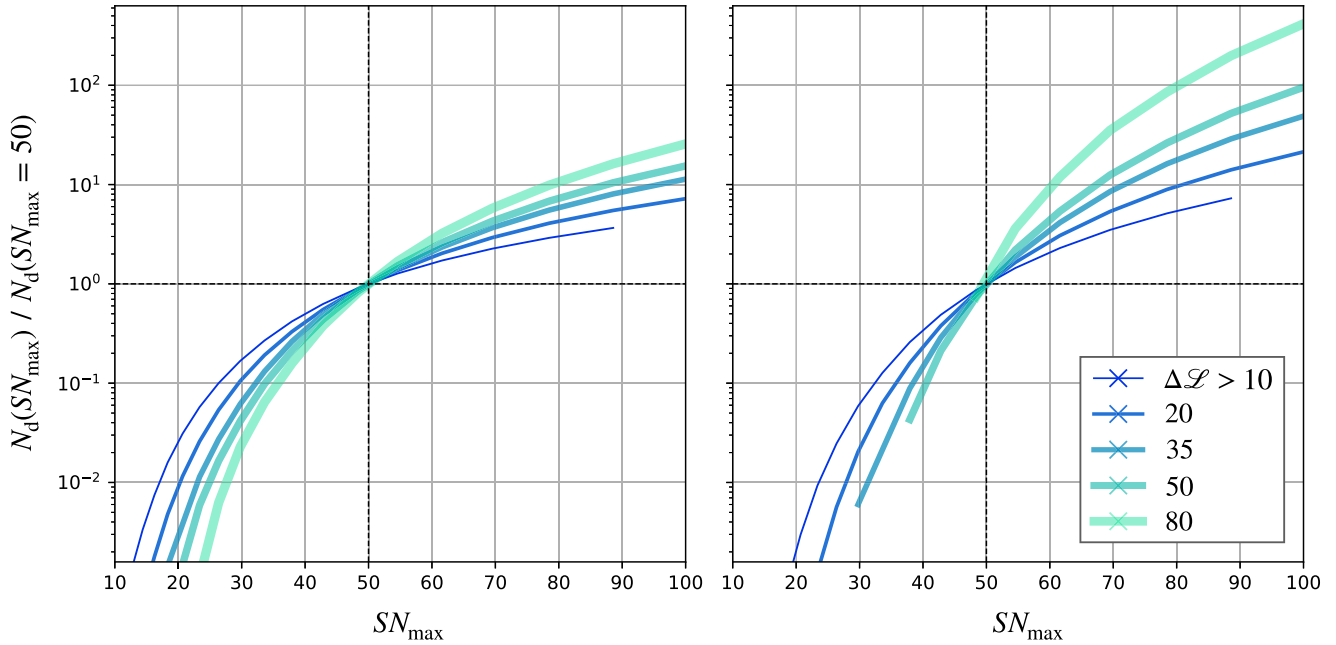
Figure 7. The number of detectable haloes of mass $M_h < 10^{10}$ M$_\odot$ (left-hand panel), and $M_h < 10^{9.5}$ M$_\odot$ (right-hand panel) in a CDM universe as a function of the maximum SN of the data. Values are normalized to the number of detections predicted for $SN_{max} = 50$, which we use as a fiducial value in this work. Lines of different colours refer to different values of the log-likelihood threshold required for detectability.
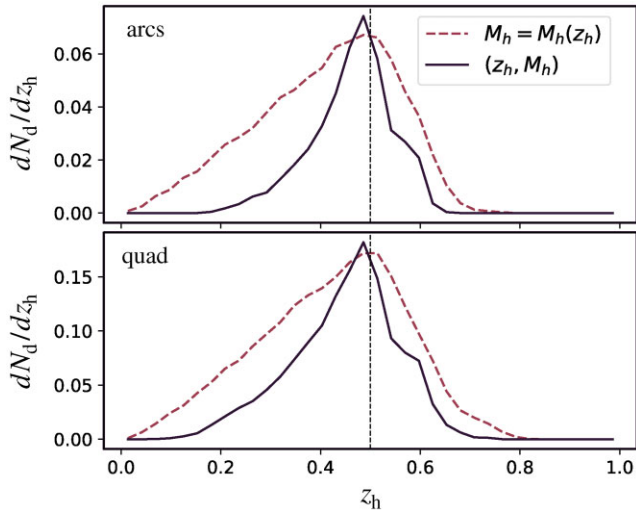
Figure 8. Predictions for the redshift distribution of detected CDM haloes obtained when: (i) using the relation between halo mass and redshift proposed by Despali et al. (2018) based on deflection angles (dashed line) and (ii) using the full redshift dependence of the log-likelihood increase (solid line). The top panel is for the 'arcs' configuration, and the bottom panel for the quad configuration. Concentration effects are not included in this comparison.

all haloes of $M_h < 10^{10}$ M$_\odot$. The average concentration shift increases further for decreasing perturber masses, as shown by the coloured lines in the top panel. These represent the 'mass–concentration relation of detectable haloes'. Lower mass haloes require stronger concentration boosts to achieve detectability, so that for $\Delta\mathcal{L}_{th} = 50$, $\langle \delta\log c/\sigma\log c \rangle = 2.2$ for haloes of $M_h \approx 10^{9.2}$ M$_\odot$.

The result is that the scatter in the mass–concentration relation boosts the number of expected detections. At low halo masses, for any fixed halo mass, there is a fraction of haloes with high enough

concentration that become detectable if the scatter in the mass–concentration relation is accounted for, and which is lost when $c = c(M, z)$. This is quantified in Fig. 10, which shows the ratio between the cumulative number of detectable haloes, $N_c(< M_h)$, predicted when accounting for the scatter in the mass–concentration relation, to the corresponding number obtained when $c = c(M, z)$ for all haloes. Lines of different colours refer to different values of the detection threshold, $\Delta\mathcal{L}_{th}$. For all thresholds, the ratios are a strong function of halo mass. The onset of the sharp rise identifies the halo mass that is not detectable at that threshold if $c = c(M, z)$, but for which detections are possible because of concentration effects. For $\Delta\mathcal{L}_{th} = 50$, even including all haloes with $M_h < 10^{10}$ M$_\odot$, concentration effects boost the number of detectable haloes by a factor of 2.75 for our 'quad' configuration. Since our 'arcs' configuration leads to lower values of the log-likelihood increase and fewer detections, the corresponding boost is a factor ≈ 26, exemplifying, on the one hand, the importance of concentration effects, and, on the other, the need for estimates tailored to the specific lensing configuration.

Fig. 11 displays the same boost factor as a function of the number of expected detections, N_c (which include the effect of concentration). Different values of the expected number of detections correspond to different values of the data SN (or equivalently, different values of the log-likelihood threshold for detection). Once again, we see that the exact figures depend on the lens configuration; for example, here the 'arcs' configuration appears more sensitive to concentration effects than the quad.

5.4 The effect of concentration on distinguishing DM models

In most previous estimates of the dependence of the number of detectable haloes on the properties of the DM model, particularly the mass of a WDM particle (or, equivalently, the cut-off mass in the mass function), it was assumed that all haloes lie exactly on the median mass–concentration relation of CDM haloes. Not only
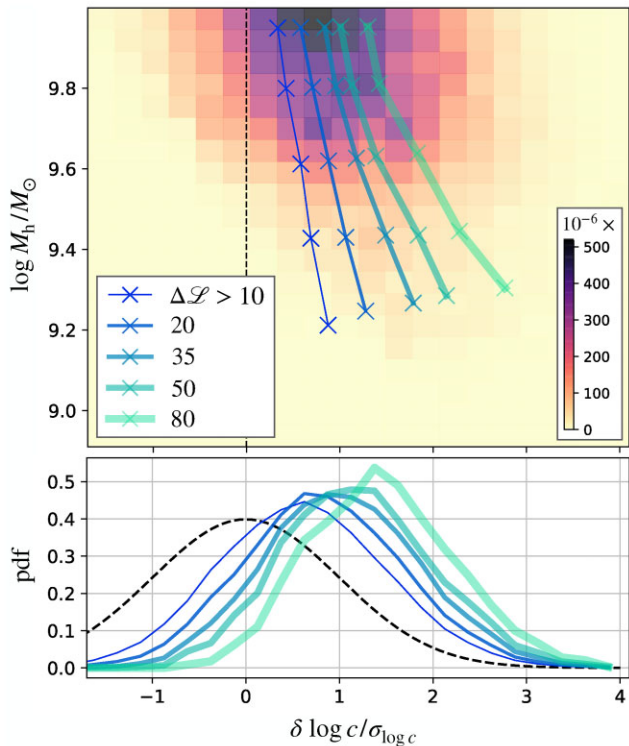
Figure 9. The distribution of halo concentration for the population of detectable haloes in a CDM universe. The top panel shows a two-dimensional histogram of all detectable haloes in the plane of halo mass and concentration shift from the median relation, in units of the lognormal spread in the mass–concentration relation. The colour scale shows the number of perturbers in each M-c pixel for a detection threshold of $\Delta\mathscr{L}_{\rm th} = 35$. Lines of different colours display the mean concentration shift, $\log c/\sigma_{\log c}$, as a function of mass for different thresholds for detectability (see the text). The bottom panel shows the distribution of concentration for all detectable haloes of $M_{\rm h} < 10^{10}$ M$_\odot$, as a function of the thresholds for detectability. For reference, the dashed line shows the parent distribution of all cosmological haloes.

is this a poor approximation but it left the differences in the halo abundances alone as the statistic to differentiate among different models. Our results show that concentration is a crucial ingredient for halo detectability. Warmer WDM models make low-mass haloes that are progressively less concentrated, as quantified by equation 21. This makes the concentration parameter potentially helpful in boosting the spread among the expected numbers of detections in different DM models.

We test this in Fig. 12, in which we show, side by side, the cumulative number of expected detections, $N_{\rm d}(< M_{\rm h})$, for a range of WDM models, parametrized by their cut-off mass, $M_{\rm cut}$. The panel on the right displays results for the case in which concentration effects are neglected, whereas the panel on the left includes both: (i) the scatter in the mass–concentration relation and (ii) the dependence of the median concentration on the DM model. For definitiveness, we show results pertaining to our 'quad' configuration, using a threshold for detection of $\Delta\mathscr{L}_{\rm th} = 30$.

It is clear that concentration effects significantly enhance the dependence of the expected number of detections on the DM model. Reduced concentration values act together with reduced cosmological abundances to determine the expected number of detections. Once again, while the qualitative trend is clear and the magnitude of the effect significant, precise values are dependent on the lensing
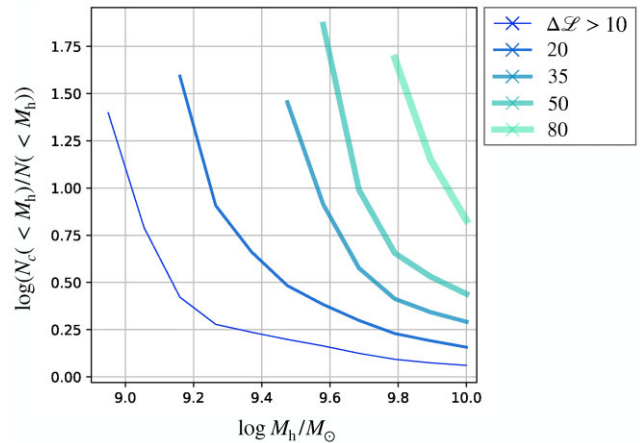
Figure 10. The boost to the cumulative number of expected detections, $N_c(< M_{\rm h})$, resulting from including the scatter in the mass–concentration relation for our 'quad' configuration. Lines of different colours correspond to different thresholds for detectability.
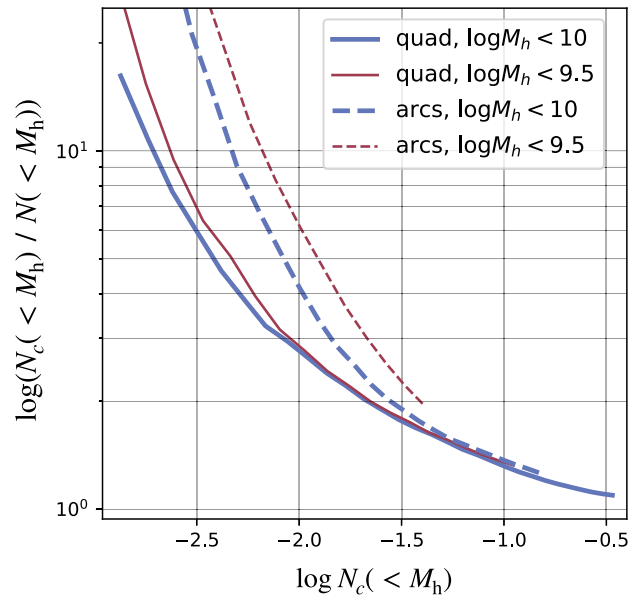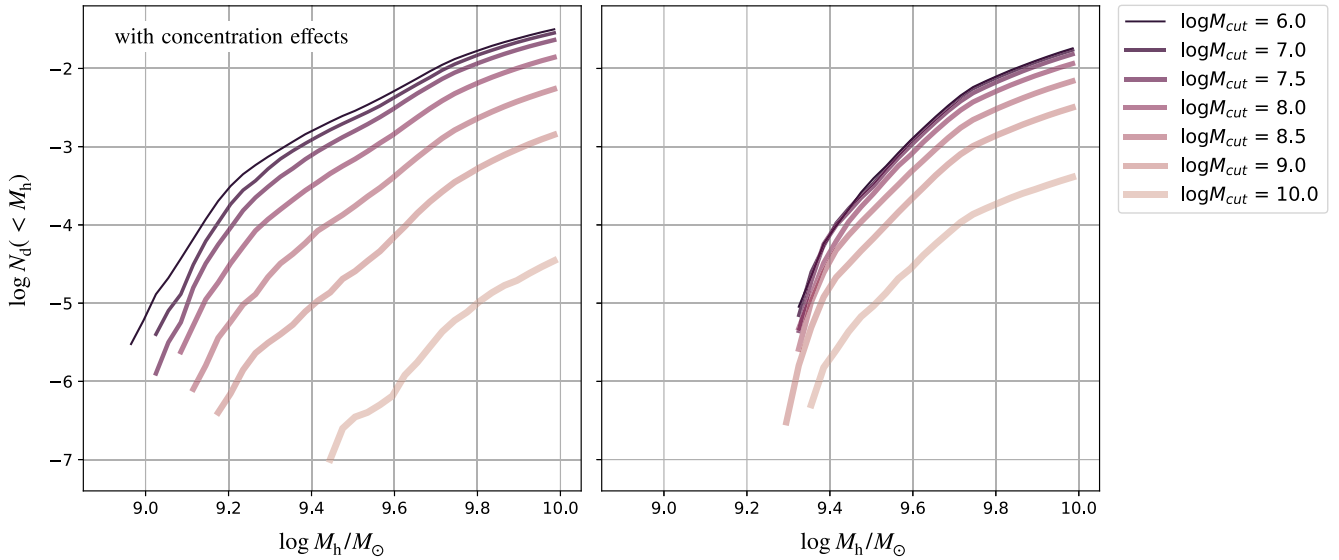

Figure 11. The boost to the cumulative number of expected detections, $N_c(< M_{\rm h})$, resulting from including the scatter in the mass–concentration relation, as a function of the expected number of detections.

configuration and on the value of the threshold chosen for detection. We attempt to quantify how much concentration effects can actually sharpen substructure lensing constraints in Section 6.1.

6 DISCUSSION AND CONCLUSIONS

We have quantified the ability of low-mass DM haloes along the LOS to perturb strong gravitational lenses, and explored how this depends on halo properties. This is a fundamental ingredient of sensitivity mapping, i.e. the process of assessing which perturbers, out of the cosmological population of haloes, would actually be detectable when modelling strong lensing data. It is impossible to quantify the number of expected detections in different DM models in a given observational data set without building the sensitivity function. Therefore, sensitivity mapping is a key aspect of placing

Figure 12. A comparison between the cumulative number of expected detections for our 'quad' configuration (threshold for detection, $\Delta \mathcal{L}_{th} = 30$) when including both (i) the scatter in the mass–concentration relation and (ii) the dependence of the median concentration on the DM model (left), and when concentration effects are neglected (right). Concentration substantially enhances the spread between the expected detections in WDM models with different cut-off masses, M_{cut}.

constraints on the identity and properties of DM from the number of perturbing haloes detected in strong lensing studies.

We have adopted a likelihood-based approach; i.e. we differentiate between detectable and non-detectable haloes according to the likelihood gain, $\Delta \mathcal{L}$, associated with including a halo mass component in the lens modelling. Some previous studies have proposed using instead the Bayesian evidence for this comparison. It should be stressed that both approaches are equally valid. As long as the same criterion is consistently applied to both measure the sensitivity function – and therefore make predictions for the different DM models – and detect perturbers in the data, both approaches will return the correct inference on the DM properties if the models used in sensitivity mapping share the same complexities of the real data.

At this stage, it cannot be excluded that the sensitivity functions derived using likelihood or Bayesian evidence may exhibit some differences when compared side by side, possibly reflecting that the two criteria may lead to different sensitivity to perturbers in different regions of the parameter space. However, it seems quite unlikely that this would happen systematically at the high levels of significance that we are considering here and that have become the norm in structure lensing. At present, an evidence-based sensitivity function has not been derived as a function of either redshift or halo concentration, where our major findings are. Therefore, we are unable to make a direct comparison and have limited our analysis to the differences with the approximate strategies that have been adopted so far.

Rather than attempting to quantify the number of expected detections for specific observed strong lenses, we have focused on building an understanding of the sensitivity function itself, and how this scales with some of the crucial parameters at play. We have concentrated on the importance of image depth, and shown that, as the log-likelihood difference scales with SN^2, high-signal-to-noise data are extremely beneficial in that they allow the detection of a larger number of perturbers, particularly of low-mass haloes (see Fig. 7). We have also shown that the specific noise realization introduces scatter in the log-likelihood difference, of magnitude $\approx \sqrt{2\Delta \mathcal{L}}$ (see Section 3.2),

which suggests that a smooth link between the probability of detection, p, and the log-likelihood gain, $\Delta \mathcal{L}$, rather than a sharp threshold, may be a better choice when comparing with real data.

We find that our two different lens configurations yield significantly different numbers of expected halo detections, which indicates that some lensing morphologies (a quad configuration in our case) are more valuable for strong lensing analyses (see Fig. 6). This will be useful in selecting lenses to target with deeper observations. We also note that our estimates for the total number of detectable haloes are somewhat lower than what has been suggested by similar studies for the same values of lens and source redshifts, and for similar data quality. This may be due to the increased flexibility of our lens modelling, including the possibility of shifts in the lens centre (see Vegetti et al. 2014), as well as the assumed power-law profile slope (see Li et al. 2017; Despali et al. 2018). With particular reference to Li et al. (2017) and Despali et al. (2018), the fact that we perform fully non-linear searches when optimizing our macromodels certainly enables them to reproduce better the perturbed data without the need of including a halo mass component, hence lowering the log-likelihood gain. If anything, this highlights the importance of using exactly the same techniques to both (i) model real data and make perturber detections and (ii) produce estimates of the expected detections, as any mismatch would inevitably introduce systematic biases.

We then concentrated on the role of halo redshift and halo concentration. In previous work, simplifications had been made to collapse these axes, in order to make the calculation of the sensitivity function computationally feasible.

Concerning the redshift of the perturber, we have shown that, contrary to previous understanding, it becomes increasingly challenging to detect perturbing haloes in front of the main lens when they get closer to the observer (see Figs 2–4). This implies that previous studies of the number of detectable haloes have likely overestimated the number of foreground detections, at redshifts $z_h < z_l$. The exact magnitude of this overestimation appears to depend on the specific lensing configuration. As a reference, our experiments show this factor to be between 1.5 and 2 (see Fig. 8). These previous

estimates are not based on a calculation of the Bayesian evidence at $z_h \neq z_l$. Therefore, it is not currently known whether an evidence-based criterion for detection would indeed yield a dependence of the perturbers' detectability on redshift that is analogous to the one we measure. Certainly, we find that the strategy adopted so far (of using deflection angles as proxies) underestimates the degree of degeneracy in the lens modelling and therefore artificially makes the detection of foreground perturbers easier than actually is in reality.

Concerning concentration, we find that detectability is a strong function of halo concentration, such that the population of detectable haloes is, in fact, a population of systematically high-concentration haloes (see Fig. 9). The shift in the average concentration relative to the mass–concentration relation becomes increasingly large for haloes of lower masses, and increases when a higher threshold for detectability is adopted. For a threshold of $\Delta \mathcal{L}_{th} = 50$, the average shift in the concentration of haloes with masses below 10^{10} M$_\odot$ is about $1.25 \sigma_{\log c}$, where $\sigma_{\log c}$ is the lognormal scatter in the mass–concentration relation.

Crucially, accounting for the scatter in the mass–concentration relation results in a boost to the number of detectable haloes. This boost is a strong function of the lensing configuration and of the threshold for detectability (or, equivalently, of the data quality as quantified by the maximum signal-to-noise; see Figs 10 and 11). As reference, for a combination of a lens configuration and detection threshold that results in a total of 0.03 detections with $M_h < 10^{10}$ M$_\odot$ per lens in a CDM universe – which is roughly comparable to what was previously predicted for lenses with the *HST* data – this boost amounts to a factor of ≈ 2.5, and quickly grows to $\gtrsim 10$ for the detections expected at $M_h \lesssim 10^{9.5}$ M$_\odot$.

Unfortunately, without a tailored study, it is impossible to provide a precise quantification of how the above two effects would combine to affect previous estimates of the expected number of detections in real observed strong lenses, especially since the two effects have opposite signs. The overestimate related to the redshift dependence is sensitive to the lensing configuration and certainly to the redshifts of both lens and source, which here we have kept fixed. The underestimate due to the concentration dependence is a strong function of lensing configuration and data quality. It would appear that the correction due to concentration is larger than that due to the redshift dependence, but further study is required to ascertain in which regime that is the case, and by how much.

What we can already establish in this study is how concentration effects can facilitate differentiating WDM models with different cut-off halo masses. We have shown that taking into account the dependence of the median halo concentration on the DM model increases the spread among the number of expected detections (see Fig. 12). For warmer models, lower halo concentrations conspire with lower cosmological halo abundances increasingly to suppress the number of detectable haloes. The effect of halo concentration had not been included in previous studies, leaving only halo abundances to differentiate among DM models, therefore making it harder to distinguish them in strong lensing studies.

6.1 Sharper DM constraints from substructure lensing

In order to quantify the extent to which concentration effects can sharpen future substructure lensing constraints, we assume that we have a set of strong lenses such that the total expected number of detections in CDM is

$$N_{d,tot}\left(M_h < 10^{9.5} \, \mathrm{M}_\odot, \mathrm{CDM}\right) = 1. \tag{23}$$
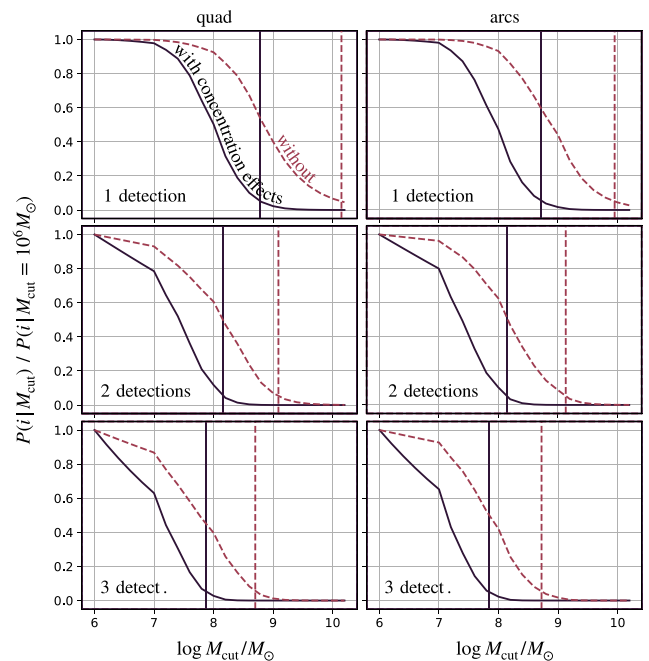
Figure 13. The change on limits to the WDM cut-off mass, M_{cut}, from including concentration effects, at a fixed number of expected detections for a CDM universe: $N_{d,CDM} = 1$. Lines show likelihood ratios (see the text) resulting from the detection of (1, 2, 3) perturbers, respectively, from the top to the bottom row. Dashed lines display the inference based on predictions that ignore concentration effects. These are included in the solid lines.

We ignore the contribution of satellite haloes, and assume that the above figure only includes haloes along the LOS, on which we have focused in this work. While we are not able to tailor our analysis quantitatively to any specific set of observed strong lenses, this figure is representative of what is achievable with current *HST* data (Vegetti et al. 2018; Ritondale et al. 2019), and therefore provides a useful reference point. We use our maps of log-likelihood increase, $\Delta \mathcal{L}$, to calculate the number of expected detections in the same set of lenses for WDM models with different cut-off masses, M_{cut}. We do so separately for our 'quad' and 'arcs' lensing configurations, requiring that the number of lenses in the two separate sets be such as to satisfy equation (23) separately.[7] Furthermore, we set up predictions for both, the case in which concentration effects are accounted for and the case in which they are ignored. Then, we compare the inferences on the DM model that would result from actually detecting $i = (1, 2, 3)$ individual haloes, in the two different configurations. These are displayed in Fig. 13, which shows the likelihood ratio:

$$R = \frac{P(i|N_{d,tot}(M_{cut}))}{P(i|N_{d,tot}(M_{cut} = 10^6 \, \mathrm{M}_\odot))}, \tag{24}$$

where $P(\cdot \,|m)$ is the Poisson probability distribution with mean, m, and i is the number of actual detections. Inference resulting from predictions that ignore concentration effects is shown with dashed lines, while the likelihood ratio obtained when accounting for concentration effects is shown with solid lines. The vertical lines indicate the limits on the WDM cut-off mass corresponding to a likelihood ratio (R) of 0.05. The right and left columns

---

[7]We require a detection threshold, $\Delta \mathcal{L}_{th} = 30$, for our quad configuration and $\Delta \mathcal{L}_{th} = 20$ for our 'arcs' configuration, which, as we have shown, leads to systematically fewer detections per lens.

correspond, respectively, to the two lensing configurations. Details are only marginally different and the magnitude of the effect is very similar in the two cases: At fixed number of expected haloes in a CDM universe, concentration effects make constraints on the WDM cut-off mass significantly more stringent. The suppression in the concentration of WDM haloes enhances the effect of lower cosmological halo abundances, allowing constraints that are about one order of magnitude more stringent in M_{cut}.

6.2 Outlook

Our results bring renewed confidence to the field of halo detection with strong lensing data, and boost confidence that meaningful constraints can be obtained from analysis of current optical data. A number of previous works have contributed to the realization that it is extremely challenging to use current optical *HST* data to obtain constraints on the cut-off of the DM halo mass function that are competitive with those obtained from the satellites of the Milky Way or measurements of the Lyman α forest (see Enzi et al. 2021, and references therein). This is because, if halo abundances alone are used to differentiate between CDM and WDM, in order to be able to probe a WDM model with a cut-off mass, M_{cut} (the mass below which the abundance of haloes declines sharply), it is necessary to be sensitive to perturbers of that halo mass and below. However, evidence is mounting that detecting haloes of mass $M_h \approx 10^{8.5}$ $M_\odot$ is extremely challenging with current lensing data, and therefore that it would be very difficult to place competitive constraints.

Concentration effects change this picture completely. For example, the limits displayed in Fig. 13 stem from detections of haloes of mass $M_h > 10^{9.5}$ $M_\odot$ – which is realistic with current *HST* data – but they can rule out values of the cut-off mass scale, $M_{cut} \gtrsim 10^8$ $M_\odot$. This is a direct reflection of the effects of halo concentration, which, in contrast to halo abundances, first affects haloes of masses significantly *above* the cut-off mass itself. For this reason, concentration effects allow substructure lensing studies to probe WDM models with cut-off masses at least one order of magnitude *below* the lowest sensitivity mass scale. This implies that substructure lensing is, in fact, a much more sensitive probe of the identity of the DM than had been previously recognized.

This work used the following software packages:

(i) ASTROPY (Astropy Collaboration 2013, 2018);
(ii) COLOSSUS (Benedikt et al. 2018);
(iii) MATPLOTLIB (Hunter et al. 2007);
(iv) NUMPY (van der Walt et al. 2011);
(v) PYAUTOLENS (Nightingale & Dye 2015; Nightingale et al. 2018, 2021);
(vi) PYTHON (Van Rossum et al. 2009);
(vii) SCIPY (Virtanen et al. 2020).

DATA AVAILABILITY

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

REFERENCES

Alam S. et al., 2017, MNRAS, 470, 2617
Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, MNRAS, 426, 2046
Astropy Collaboration, 2013, A&A, 558, A33
Astropy Collaboration, 2018, AJ, 156, 123
Avila-Reese V., Colín P., Piccinelli G., Firmani C., 2003, ApJ, 598, 36
Baur J., Palanque-Delabrouille N., Yèche C., Magneville C., Viel M., 2016, J. Cosmol. Astropart. Phys., 2016, 012
Benedikt D., 2018, ApJS, 2, 35
Benitez-Llambay A., Frenk C., 2020, MNRAS, 498, 4887
Birrer S., Amara A., Refregier A., 2016, J. Cosmol. Astropart. Phys., 2016, 020
Bolton A. S., 2005, Ph.D dissertation, 2005. Massachusetts Institute of Technology, Massachusetts, USA
Bose S., Hellwing W. A., Frenk C. S., Jenkins A., Lovell M. R., Helly J. C., Li B., 2016, MNRAS, 455, 318
Boyarsky A., Drewes M., Lasserre T., Mertens S., Ruchayskiy O., 2019, Prog. Part. Nucl. Phys., 104, 1
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
Despali G., Vegetti S., White S. D. M., Giocoli C., van den Bosch F. C., 2018, MNRAS, 475, 5424
Diemand J., Kuhlen M., Madau P., Zemp M., Moore B., Potter D., Stadel J., 2008, Nature, 454, 735
Dye S., Warren S. J., 2005, ApJ, 623, 31
Enzi W. et al., 2021, MNRAS, 506, 5848
Frenk C. S., White S. D. M., 2012, Ann. Phys., Lpz., 524, 507
Garzilli A., Boyarsky A., Ruchayskiy O., 2017, Phys. Lett. B, 773, 258
Garzilli A., Ruchayskiy O., Magalich A., Boyarsky A., 2019, preprint (arXiv:1912.09397)
Gilman D., Birrer S., Treu T., Keeton C. R., Nierenberg A., 2018, MNRAS, 481, 819
Gilman D., Birrer S., Treu T., Nierenberg A., Benson A., 2019, MNRAS, 487, 5721
Gilman D., Du X., Benson A., Birrer S., Nierenberg A., Treu T., 2020, MNRAS, 492, L12
Green A. M., Hofmann S., Schwarz D. J., 2005, J. Cosmol. Astropart. Phys., 2005, 003
Harvey D., Valkenburg W., Tamone A., Boyarsky A., Courbin F., Lovell M., 2020, MNRAS, 491, 4247
Hastie T., Tibshirani R., Friedman J., 2008, The Elements of Statistical Learning. Springer, New York
He Q. et al., 2020, preprint (arXiv:2010.13221)
Hezaveh Y. D. et al., 2016, ApJ, 823, 37
Hunter J. D., 2007, AJ, 9, 3
Iršič V. et al., 2017, Phys. Rev. D, 96, 023522
Jenkins A., Frenk C. S., White S. D. M., Colberg J. M., Cole S., Evrard A. E., Couchman H. M. P., Yoshida N., 2001, MNRAS, 321, 372
Kennedy R., Frenk C., Cole S., Benson A., 2014, MNRAS, 442, 2487

Konishi S., Kitagawa G., 2007, Information Criteria and Statistical Modeling. Springer, New York

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

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

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

Minor Q., Kaplinghat M., Chan T. H., Simon E., 2021, MNRAS, 507, 1202

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., 2020, J. Cosmol. Astropart. Phys., 2021, 062

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. et al., 2021, J.Open Source Softw., 6, 2825

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

Powell D., Vegetti S., McKean J. P., Spingola C., Rizzo F., Stacey H. R., 2021, MNRAS, 501, 515

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

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

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

Shu Y. et al., 2016, ApJ, 824, 86

Springel V. et al., 2005, Nature, 435, 629

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

Tessore N., Bellagamba F., Metcalf B., 2016, MNRAS, 463, 3115

van der Walt S., Colbert S. C., Varoquaux G., 2011, Comput. Sci. Eng., 13, 2

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

Vegetti S., Koopmans L. V. E., 2009, MNRAS, 392, 945

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., 2011, Nat. Methods, 17, 261

Wang W., White S. D. M., Mandelbaum R., Henriques B., Anderson M. E., Han J., 2016, MNRAS, 456, 2301

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

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

APPENDIX A: INPUT MACROMODELS

Table A1 contains values of the adopted model parameters for the two different lensing configurations used in this study.

Table A1. Input values for the two sets of macromodel parameters used in this study.

Model parameter	Quad	Arcs
(x_1, y_1)	(0.0051 arcsec, 0.0765 arcsec)	(0.0396 arcsec, 0.08 arcsec)
ϵ_1	1.165 arcsec	1.095 arcsec
β_1	1.93	1.99
$(e_{1,1}, e_{2,1})$	(0.022, 0.011)	(−0.013, 0.007)
(γ_1, γ_2)	(−0.037, −0.099)	(−0.008, 0.001)
z_1	0.5	0.5
(x_s, y_s)	(0.024 arcsec, 0.032 arcsec)	(−0.024, 0.036)
r_{eff}	0.15 arcsec	0.12 arcsec
$(e_{1,s}, e_{2,s})$	(0.147, −0.135)	(0.05, −0.25)
n_s	1.1	1.2
I_s	1	1
z_1	1.0	1.0

APPENDIX B: DEPENDENCE ON THE NOISE REALIZATION

Fig. B1 shows the link between the mean and the scatter (standard deviation) of the log-likelihood gain $\Delta\mathscr{L}$. Each point corresponds to a different set of halo properties ζ_h. For each, 10 different random noise realizations have been considered and modelled. The corresponding set of values for the log-likelihood increase has been used to estimate both mean value and standard deviation. The red dashed line illustrates the prediction of equation (15).
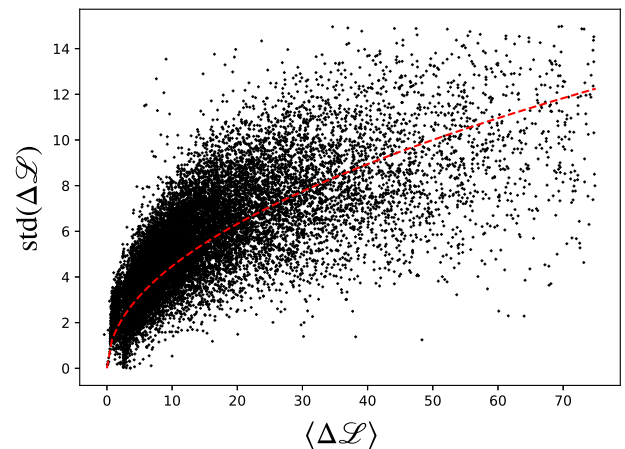
Figure B1. The scaling of the scatter in the log-likelihood increase std($\Delta\mathscr{L}$) resulting from different noise realizations (at fixed signal-to-noise) and the mean log-likelihood increase $\langle\Delta\mathscr{L}\rangle$. The red line shows the scaling of equation (15).

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