








What does strong gravitational lensing? The mass and redshift distribution of high-magnification lenses

Andrew Robertson ¹★, Graham P. Smith ², Richard Massey ¹, Vincent Eke ¹,
Mathilde Jauzac ^{1,3,4}, Matteo Bianconi ² and Dan Ryczanowski ²

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

²*School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK*

³*Centre for Extragalactic Astronomy, Department of Physics, Durham University, Durham DH1 3LE, UK*

⁴*Astrophysics and Cosmology Research Unit, School of Mathematical Sciences, University of KwaZulu-Natal, Durban 4041, South Africa*

Accepted 2020 May 19. Received 2020 April 6; in original form 2020 February 4

ABSTRACT

Many distant objects can only be detected, or become more scientifically valuable, if they have been highly magnified by strong gravitational lensing. We use EAGLE and BAHAMAS, two recent cosmological hydrodynamical simulations, to predict the probability distribution for both the lens mass and lens redshift when point sources are highly magnified by gravitational lensing. For sources at a redshift of 2, we find the distribution of lens redshifts to be broad, peaking at $z \approx 0.6$. The contribution of different lens masses is also fairly broad, with most high-magnification lensing due to lenses with halo masses between 10^{12} and $10^{14} M_{\odot}$. Lower mass haloes are inefficient lenses, while more massive haloes are rare. We find that a simple model in which all haloes have singular isothermal sphere density profiles can approximately reproduce the simulation predictions, although such a model overpredicts the importance of haloes with mass $< 10^{12} M_{\odot}$ for lensing. We also calculate the probability that point sources at different redshifts are strongly lensed. At low redshift, high magnifications are extremely unlikely. Each $z = 0.5$ source produces, on average, 5×10^{-7} images with magnification greater than 10; for $z = 2$, this increases to about 2×10^{-5} . Our results imply that searches for strongly lensed optical transients, including the optical counterparts to strongly lensed gravitational waves, can be optimized by monitoring massive galaxies, groups, and clusters rather than concentrating on an individual population of lenses.

Key words: gravitational lensing: strong – gravitational waves – cosmology: theory.

1 INTRODUCTION

Gravitational lensing and gravitational waves are two phenomena predicted by Einstein’s general theory of relativity (GR), both of which have now been observed. Evidence for strong gravitational lensing of electromagnetic radiation was first reported by Walsh, Carswell & Weymann (1979), who discovered a multiply imaged quasar. Since then, strong gravitational lensing has become a key tool in astrophysics, allowing the mass distribution in galaxies and galaxy clusters to be mapped (e.g. Bolton et al. 2008; Richard et al. 2010), and for distant objects to be studied that would otherwise be too faint and/or small (see Kelly et al. 2018, for a spectacular example). Gravitational waves were first detected only recently (Abbott et al. 2016) but also hold promise as a new tool for investigating our Universe (Miller & Yunes 2019).

The prediction of GR is that gravitational waves have their trajectories bent by gravity in the same way as light, that is to say that gravitational waves can also be gravitationally lensed. Ignoring the lensing of gravitational waves could lead to incorrect conclusions about the population of merging compact objects. To leading order, the frequency evolution, $\nu(t)$, of a gravitational wave signal is determined by the chirp mass, $\mathcal{M} = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$, where m_1 and m_2 are the masses of the two merging objects (Abbott et al. 2017c). The intrinsic amplitude of the emitted signal can be predicted from the chirp mass, so the measured amplitude of the signal can be used to determine the luminosity distance, d_L , to the coalescing objects.

For objects at redshift z , the received frequency will be lower than the intrinsic frequency by a factor of $1 + z$. With an assumed cosmology and ignoring peculiar velocities, there is a one-to-one relationship between d_L and z . This means that the chirp mass and luminosity distance (and therefore also redshift) can be inferred from the amplitude and frequency evolution of a gravitational

* E-mail: andrew.robertson@durham.ac.uk

wave signal. Gravitational lensing introduces a new quantity, the magnification μ , which alters the amplitude of the detected gravitational wave strain by a factor of $\sqrt{\mu}$ compared with the no lensing ($\mu = 1$) case. A gravitational wave signal that – ignoring lensing – would correspond to a chirp mass $\mathcal{M}_{\mu=1}$ at a redshift $z_{\mu=1}$ could actually be at a true redshift z , with true chirp mass \mathcal{M} where $\mathcal{M} = \mathcal{M}_{\mu=1}(1 + z_{\mu=1})/(1 + z)$. The magnification in this case would be $\mu = [d_L(z)/d_L(z_{\mu=1})]^2$. As such, low-mass, distant, and highly magnified gravitational wave sources will masquerade as more massive and closer gravitational wave sources if gravitational magnification is not considered in interpreting the gravitational wave signal.

At LIGO/Virgo’s current sensitivity, the detection rate of strongly lensed GWs is expected to be small (Li et al. 2018; Smith et al. 2018). This can be understood in terms of the lens magnification that is required to render a lensed GW detectable. At present, unlensed GWs can only be detected out to a luminosity distance of ~ 1 Gpc (Abbott et al. 2018), corresponding to $z \simeq 0.2$. Given that the typical redshift of a strong lensing galaxy or cluster is $z \simeq 0.2$ – 0.5 , it is inevitable that the typical redshift of a lensed GW is $z \gtrsim 1$, which corresponds to a luminosity distance of $\gtrsim 7$ Gpc. Given the dependence of lens magnification on luminosity distance discussed above, these distances imply $\mu \gtrsim 50$, and thus a low probability of detection thus far, bearing in mind that $P(> \mu) \propto \mu^{-2}$. Despite some early proof-of-concept follow-up observations (Smith et al. 2019b) and discussion of LIGO’s early detections (Broadhurst, Diego & Smoot 2018, 2019), there is scant evidence that any of LIGO/Virgo’s detections to date are strongly lensed (Hannuksela et al. 2019; Singer, Goldstein & Bloom 2019).

Calculating the expected number of detectable strongly lensed GWs and the probability that a given GW detection is strongly lensed requires a good understanding of the population of gravitational lenses. Knowledge of this population can then also shape the observing strategy for electromagnetic follow-up of candidate lensed GWs. There has been disagreement in the literature as to what population of lenses is most important for the high magnifications required to reinterpret the LIGO/Virgo events as being less massive but at high redshift. For example, Smith et al. (2018) used the results of ray tracing through a large N -body simulation (Hilbert et al. 2008, hereafter H08) to conclude that galaxy clusters were the most important population of lenses. Meanwhile, Haris et al. (2018) forward modelled a population of lensed gravitational wave signals, assuming that lensed images dominantly arise due to galaxy lenses, modelling their population of lenses on SDSS early-type galaxies (Choi, Park & Vogeley 2007). Claims that some of the gravitational wave signals already detected are likely strongly lensed (Broadhurst et al. 2018, 2019) have also assumed that galaxies dominate the optical depth for strong lensing, which they justify from the population of lenses responsible for lensed high-redshift star-forming galaxies detected by the *Herschel* satellite (Negrello et al. 2010; Bussmann et al. 2013; Wardlow et al. 2013). However, other lensed objects paint a different picture, with the first detected strongly lensed supernova (Kelly et al. 2015) and highly magnified individual star (Kelly et al. 2018), as well as the most magnified lensed quasars (Sharon et al. 2005; Oguri et al. 2013; Sharon et al. 2017), being lensed by galaxy clusters. Inferring the importance of different lenses for strong lensing using the distribution of lenses responsible for observed highly magnified objects is complicated by selection effects, which motivates answering this question from a theoretical perspective.

In this paper, we seek to answer the question of which objects are responsible for producing strong gravitational lensing. Given

the small size of the GW emission region in compact binary coalescence, and the high magnifications required to reinterpret gravitational wave events as having come from objects significantly less massive than originally inferred, we will concentrate on high magnifications ($|\mu| > 10$) for point sources. We do this by combining lensing calculations performed on two state-of-the-art hydrodynamical simulations, and show that the results can be approximately reproduced by a model in which the total mass distribution of each gravitationally collapsed object is modelled as a singular isothermal sphere.

This paper is organized as follows. In Section 2, we describe a simple model for the strong lensing optical depth, based upon combining the halo mass function with a simple density profile for each halo. Then, in Section 3 we describe our simulations and how we calculated the cross-sections for strong gravitational lensing from individual simulated haloes. In Section 4, we present our results, combining the strong lensing properties of all our simulated haloes to find the contribution of different lens masses and lens redshifts to the probability for strong lensing. We then present our conclusions in Section 5. We assume a Planck Collaboration II (2014) cosmology throughout this paper,¹ unless stated otherwise.

2 A SIMPLE MODEL FOR THE STRONG LENSING OPTICAL DEPTH

Answering the question of ‘*what does strong lensing?*’ requires two ingredients. First, the *lensing cross-section* of individual lenses as a function of their mass and redshift, and secondly, the *halo mass function* – the number density of haloes of different masses at different cosmic times. In this section, we discuss both of these in the context of a simple model where all lenses are assumed to be spherically symmetric *singular isothermal spheres* (SISs), which have 3D density profiles with $\rho \propto r^{-2}$. The advantage of using this density profile is that it provides a simple analytic method to relate the Einstein radius of a lens to the mass of its dark matter halo, across a broad range of halo mass, as detailed in Sections 2.1 and 2.2. The SIS model has been shown to be a good description of strong lensing galaxies (Gavazzi et al. 2007). Whilst this model is a poor description of individual strong-lensing clusters due to the complexity of such systems (e.g. Richard et al. 2010), the slope of the SIS profile is representative of cluster density profile slopes on the relevant scales and is therefore appropriate for population studies such as this. The goal of this section is to introduce notation and gain intuition that will help to understand the results from hydrodynamical simulations in Section 3.

2.1 Lensing cross-sections of individual lenses

For a mass distribution at redshift z^L , we can define the gravitational lensing cross-section for some criterion as the solid angle that satisfies that criterion. An example criterion could be that the magnification, μ , for a source at redshift z^S is greater than some magnification threshold, μ_0 . Gravitational lensing maps an infinitesimal solid angle in the source plane, into a solid angle in the image plane that differs by a factor of the magnification.² As such, we can

¹With $\Omega_m = 0.307$, $\Omega_b = 0.04825$, $\Omega_\Lambda = 0.693$, $\sigma_8 = 0.8288$, $n_s = 0.9611$, and $h = 0.6777$.

²The image plane solid angle does not need to be larger, as $|\mu|$ can be less than unity, but throughout this paper we are concerned with high magnifications.

define cross-sections both in the image plane and the source plane, and in general these can be different. To introduce our notation we take a concrete example of a cross-section for magnification greater than 10 for a redshift 2 source. A particular lens would then have an image plane cross-section, $\sigma_{\text{lens}}^{\text{I}}(|\mu| > 10, z^{\text{S}} = 2)$, and source plane cross-section, $\sigma_{\text{lens}}^{\text{S}}(|\mu| > 10, z^{\text{S}} = 2)$, with the former corresponding to the lens-plane solid angle over which lines of sight from the observer to the lens have $|\mu| > 10$ and the latter corresponding to the (un-lensed) solid angle in the source plane that will be magnified by $|\mu| > 10$. The source plane cross-section is therefore the relevant cross-section for calculating the probability that randomly located sources at $z^{\text{S}} = 2$ will be magnified by this lens by $|\mu| > 10$.

An SIS with Einstein radius θ_{E} , has an image-plane magnification profile

$$\mu_{\text{SIS}} = \frac{\theta}{\theta - \theta_{\text{E}}}, \quad (1)$$

where θ is the angular distance from the centre of the SIS profile. Armed with the magnification as a function of radius, we can then ask what the lens (or source) plane solid angle above a particular magnification threshold is. Note that the magnification can be both positive and negative, with negative values corresponding to images with inverted parity. In this study, we are interested in the brightness of lensed point sources for which the absolute value of μ is the relevant quantity. The image plane solid angle with $|\mu| > \mu_0$ is

$$\sigma_{\text{lens}}^{\text{I}}(> \mu_0) = \int_{\theta_0}^{\theta_1} 2\pi\theta \, d\theta, \quad (2)$$

where from equation (1) we have that $\theta_0 = \theta_{\text{E}} \mu_0 / (\mu_0 + 1)$ and $\theta_1 = \theta_{\text{E}} \mu_0 / (\mu_0 - 1)$, such that

$$\sigma_{\text{lens}}^{\text{I}}(> \mu_0) = \theta_{\text{E}}^2 \frac{4\pi\mu_0^3}{(\mu_0^2 - 1)^2}. \quad (3)$$

An infinitesimal solid angle of the image plane will map back into an infinitesimal solid angle of the source plane, which is smaller by a factor $|\mu|$ (i.e. $d\sigma_{\text{lens}}^{\text{I}} = |\mu| d\sigma_{\text{lens}}^{\text{S}}$), so the source plane solid angle with $|\mu| > \mu_0$ is

$$\sigma_{\text{lens}}^{\text{S}}(> \mu_0) = \int_{\theta_0}^{\theta_1} \frac{2\pi\theta}{|\mu(\theta)|} \, d\theta. \quad (4)$$

Note that there is a subtlety when discussing the source plane solid angle with some property (such as $|\mu| > \mu_0$), because regions of the source plane can map to multiple regions in the image plane. As equation (4) is defined as an integral over the image plane, points on the source plane that map to multiple points on the image plane will be counted multiple times, once for each image that meets the respective property. This means that for a number density of sources per unit solid angle in the source plane, n_{s} , we expect to see $\sigma_{\text{lens}}^{\text{S}}(> \mu_0) \times n_{\text{s}}$ images of those sources magnified by greater than μ_0 by a particular lens. This will typically be greater than the number of sources that have at least one image magnified by greater than μ_0 , because highly magnified lines of sight are typically multiply imaged.

Evaluating the integral in equation (4) we find

$$\sigma_{\text{lens}}^{\text{S}}(> \mu_0) = \theta_{\text{E}}^2 \frac{2\pi(\mu_0^2 + 1)}{(\mu_0^2 - 1)^2}. \quad (5)$$

In equations (3) and (5), $\sigma_{\text{lens}}(> \mu_0)$ is separable into the product of θ_{E}^2 and a function of μ_0 . As such, the relative contribution of different haloes to the optical depth is independent of the exact

definition of the optical depth, whether it is the solid angle within the Einstein radius, or the source/image plane solid angle with a magnification greater than some μ_0 . While different definitions of what constitutes ‘strong lensing’ will change the total optical depth to strong lensing, it will change the optical depth of each system in the same manner, keeping their relative contributions fixed. In particular, for large magnification thresholds ($\mu_0 \gg 1$), $\sigma_{\text{lens}}^{\text{S}}(> \mu_0) \propto 1/\mu_0^2$.

While we have shown here that $P(|\mu| > \mu_0) \propto 1/\mu_0^2$ for sources behind an SIS lens, this behaviour can be shown to be more general. Critical curves are curves in the image plane along which the magnification is formally infinite in the geometric optics limit, while caustics are curves in the source plane found from mapping the critical curve to the source plane using the gravitational deflection angles. By Taylor expanding the gravitational potential about a point on a critical curve, it can be shown that the magnification varies inversely as the square root of the perpendicular distance of the source from the caustic (e.g. Gaudi & Petters 2002, equation 35). The source plane solid angle above some magnification, μ_0 , is the length of the caustic multiplied by the distance from the caustic at which the magnification drops to μ_0 , l_0 . As $l_0 \propto 1/\mu_0^2$, the high-magnification cross-section is $\sigma_{\text{lens}}^{\text{S}}(> \mu_0) \propto 1/\mu_0^2$.

2.2 The halo mass function

The halo mass function has been well studied, both in the context of analytical predictions (Press & Schechter 1974; Sheth, Mo & Tormen 2001) and measurements from N -body simulations (e.g. Jenkins et al. 2001; Tinker et al. 2008). Here, we use the Tinker et al. (2008) mass function with a Planck Collaboration II (2014) cosmology, as implemented in the Python library HMF (Murray, Power & Robotham 2013). Throughout this paper we use M_{200} to define halo masses; where r_{200} is the radius at which the mean enclosed density is 200 times the critical density, and M_{200} is the mass within r_{200} .

We define $n(M, z)$ as the comoving number density of haloes with $M_{200} < M$ at redshift z . $\partial n(M, z) / \partial \log_{10} M$ is then the number density of haloes per decade in halo mass, which is plotted for different redshifts in the top panel of Fig. 1. For brevity, we will often drop the ‘200’ from halo masses, and will not explicitly write the base (always 10) of logarithms.

2.2.1 Tying M_{200} to an SIS density profile

In order to relate an SIS with a given Einstein radius to a halo described by its M_{200} , it helps to be concrete about the normalization of the SIS density profile. Starting with a density profile with a normalization described by the velocity dispersion, σ_{v} ,

$$\rho_{\text{SIS}}(r) = \frac{\sigma_{\text{v}}^2}{2\pi G r^2}, \quad (6)$$

we can integrate along a line of sight a projected radius R from the centre of the halo to get the projected surface density

$$\Sigma_{\text{SIS}}(R) = \frac{\sigma_{\text{v}}^2}{2GR}. \quad (7)$$

The mass enclosed within a 2D radius is then

$$M_{\text{SIS}}(< R) = \frac{\pi\sigma_{\text{v}}^2}{G} R. \quad (8)$$




Figure 1. Top: the halo mass function at different redshifts. At all redshifts, the low-mass end of the mass function is approximately a power law with a slope slightly shallower than -1 as indicated by the dashed line, which is a power law with a slope of -0.94 . Bottom: the optical depth per decade in halo mass and per unit redshift, at the same redshifts as the mass function is shown in the top panel, assuming all haloes to be singular isothermal spheres. The power-law dependence of the lensing cross-section on halo mass (Section 2.2.1) multiplied by the mass function leads to a gently increasing power law at low masses, which is exponentially suppressed at high masses due to the suppression of the mass function.

For a given cosmology and lensing geometry we can define the critical surface density for lensing, Σ_{crit} , as

$$\Sigma_{\text{crit}} = \frac{c^2}{4\pi G} \frac{d_A(z^S)}{d_A(z^L) d_A(z^L, z^S)}. \quad (9)$$

Here, $d_A(z^S)$, $d_A(z^L)$, and $d_A(z^L, z^S)$ are the angular diameter distances between the observer and the source, the observer and the lens, and the lens and the source, respectively. For an axisymmetric lens, the average surface density within the Einstein radius is Σ_{crit} (i.e. $M_{\text{SIS}}(< R_E) = \Sigma_{\text{crit}} \pi R_E^2$). Combining this with the fact that the physical Einstein radius is related to an angular Einstein radius by the angular diameter distance to the lens ($R_E = \theta_E d_A(z^L)$) we find that

$$\theta_E = \frac{\sigma_v^2}{G \Sigma_{\text{crit}} d_A(z^L)}. \quad (10)$$

Returning to equation (6), we can integrate the density profile to find that the mass within a 3D radius is

$$M_{\text{SIS}}(< r) = \frac{2\sigma_v^2}{G} r. \quad (11)$$

Equating the mass within r_{200} with M_{200} allows us to write σ_v , and hence θ_E , in terms of the halo mass

$$\theta_E(M, z_{\text{lens}}) = \frac{M}{2r_{200} \Sigma_{\text{crit}} d_A(z^L)}. \quad (12)$$

At fixed z^L , Σ_{crit} and $d_A(z^L)$ are constant, and $r_{200} \propto M^{1/3}$, such that $\theta_E \propto M^{2/3}$. Then, using equation (5), $\sigma_{\text{lens}}^S \propto M^{4/3}$.

2.3 Strong-lensing optical depths

If $\sigma_{\text{tot}}(M, z)$ is the sum of σ_{lens} over all haloes with $M_{200} < M$, within some volume V , at redshift z , then

$$\frac{\partial^2 \sigma_{\text{tot}}(M, z)}{\partial V \partial \log M} = \frac{\partial n(M, z)}{\partial \log M} \sigma_{\text{lens}}(M, z). \quad (13)$$

For a patch of sky with solid angle, Ω , and with a total source-plane cross-section, σ_{tot}^S , we can define a source-plane *optical depth* $\tau^S \equiv \sigma_{\text{tot}}^S / \Omega$. Note that this definition of τ^S differs somewhat from the source plane optical depth used by (for example) Schneider, Ehlers & Falco (1992), because of how it treats multiply imaged regions of the source plane (see Section 2.1 for a discussion of this with an isolated lens). We illustrate our definition with an example, consisting of N^S point-source objects randomly distributed in the source plane. If we suppose that these objects must be magnified by $|\mu| > \mu_0$ in order to be detectable, then the expected number of detectable images will be $\tau_{|\mu| > \mu_0}^S N^S$, which can be larger than the number of different sources that are detected, due to some sources being detected multiple times. This is in contrast to the Schneider et al. (1992) definition, which we call $\tilde{\tau}_{|\mu| > \mu_0}^S$ following a similar discussion in Hilbert et al. (2007). $\tilde{\tau}_{|\mu| > \mu_0}^S N^S$ is the expected number of different sources that we would detect, some of which we may detect multiple times. Using τ^S over $\tilde{\tau}^S$ has two key advantages: first, that it is much easier to compute, because the calculation can be done in the lens plane without having to find which regions of the lens plane map on to a common region of the source plane; and secondly, because the number of images (rather than sources) is easier to count observationally.

Note that by performing our calculations exclusively in the lens plane we cannot consider complications that arise when multiple images are blended into one. There could, for example, be blended images with a total magnification $\mu_{\text{tot}} > \mu_0$ even though neither individual image has $|\mu| > \mu_0$. This is highly unlikely for lensed gravitational waves, where the length of a ‘chirp’ is typically less than a second, and so multiple images will only rarely overlap in time, however, it could be important for understanding other populations of lensed objects. For an SIS lens, the time delay between multiple images of the same source is $\Delta t \propto \theta_E^2 / \mu_{\text{tot}}$ (e.g. Oguri et al. 2002), and would be around an hour for an SIS with $M_{200} = 10^{11.5} M_\odot$, $z^L = 0.5$, $z^S = 2$, and $\mu_{\text{tot}} = 100$. This halo mass is the lowest relevant for high-magnification strong lensing (see Fig. 4) and so extremely large magnifications would be required for time delays of a second or less.

Defining $V(z)$ as the comoving volume out to redshift z , and $\Omega = 4\pi$ as the solid angle of the whole sky, we find that

$$\frac{\partial^2 \tau}{\partial \log M \partial z} = \frac{1}{4\pi} \frac{\partial^2 \sigma_{\text{tot}}(M, z) dV}{\partial V \partial \log M dz}, \quad (14)$$

which using equation (13) can be calculated from the halo mass function and $\sigma_{\text{lens}}(M, z)$.

2.3.1 Strong-lensing optical depths in an SIS universe

Using $\sigma_{\text{lens}}(M, z)$ for SIS density profiles in equation (14), we can calculate the relative contribution of different lens masses at different lens redshifts to the total optical depth. This is plotted (for an optical depth corresponding to $|\mu| > 10$ for $z^S = 2$ sources) in the bottom panel of Fig. 1, where it can be seen that the relative contribution of different lens masses shifts towards higher masses at lower redshifts, peaking at $M_{200} \approx 10^{13} M_\odot$ at $z^L = 1.4$ and $M_{200} \approx 10^{14} M_\odot$ at $z^L = 0.2$. This is driven by the increase in

the halo mass above which the mass function is exponentially suppressed (often known as M_*), through cosmic time. At all times there is a peak in $\partial^2 \tau / \partial \log M \partial z$. Lensing by low-mass haloes is relatively unimportant because the power-law slope of the mass-function ($\propto M^{-0.94}$) is not steep enough to compensate for the decreasing cross-section of low-mass lenses ($\propto M^{4/3}$), while the high-mass end is suppressed by the high-mass end cut-off in the mass function.

Aside from the relative contribution to the total lensing optical depth of different lens masses, Fig. 1 also demonstrates the relative contribution from different lens redshifts. Of the redshifts we show, $z^L = 0.5$ and 0.8 have the largest optical depths per unit redshift (the areas under the curves in the bottom panel of Fig. 1). As well as the evolution of the mass function, there are a number of different factors that contribute to this. One is the comoving volume per unit redshift, $\frac{dV}{dz}$, which increases with increasing redshift out to around $z = 2.5$. Another is the *lensing efficiency*, the inverse of the critical surface density, which at fixed source redshift is proportional to $d_A(z^L)d_A(z^L, z^S)$. In a flat and non-expanding universe, $d_A(z^L) + d_A(z^L, z^S) = d_A(z^S)$. At fixed $d_A(z^S)$, the product of $d_A(z^L)$ and $d_A(z^L, z^S)$ is then maximized when they are equal, i.e. when the lens is halfway between the observer and source. This is complicated somewhat by a universe that is expanding, but it remains true that lensing is most efficient when the lens is neither close to the observer nor close to the source, but roughly mid-way between. In the case of a $z^S = 2$ source, and a Planck Collaboration II (2014) cosmology, the lensing efficiency peaks for $z^L = 0.52$.

3 LENSING IN HYDRODYNAMICAL SIMULATIONS

Cosmological hydrodynamical simulations have increased dramatically in their realism over the past decade and can now reproduce many of the key properties of observed galaxies (e.g. Vogelsberger et al. 2014; Schaye et al. 2015; Davé, Thompson & Hopkins 2016; Kaviraj et al. 2017; Pillepich et al. 2018a; Davé et al. 2019) as well as galaxy groups and clusters (e.g. Sembolini et al. 2013; Le Brun et al. 2014; Bahé et al. 2017; Barnes et al. 2017a,b; McCarthy et al. 2017; Cui et al. 2018; Pillepich et al. 2018b). Importantly for lensing, simulations can produce populations of galaxies with the correct distribution of stellar mass, and the correct stellar mass–size relation (Schaye et al. 2015). As such, these simulations can now be used to answer the question of how important different objects are as gravitational lenses, from a theoretical perspective.

In order to determine the contribution of different lens masses to the strong lensing optical depth as predicted from a state-of-the-art hydrodynamical simulation, we require the simulation to resolve the strong lensing region of the lowest mass haloes important for lensing ($\sim 3 \times 10^{11} M_\odot$, see Fig. 4), while simultaneously covering a large enough volume to accurately sample the high-mass end of the halo mass function. A single hydrodynamical simulation like this does not currently exist, so instead we combine two simulations, using the high-resolution publicly available³ (100 Mpc)³ EAGLE simulation (Crain et al. 2015; Schaye et al. 2015) to resolve low-mass haloes, and the large volume of the (400 h^{-1} Mpc)³ BAHAMAS simulation

(McCarthy et al. 2017) to have an adequate number of high-mass haloes.

3.1 Simulation descriptions

EAGLE and BAHAMAS have similar sub-grid physics models for galaxy formation physics on scales below the resolution limits of the simulations. These include models for gas cooling, star formation, and feedback both from stars and active galactic nuclei (AGNs). EAGLE was run using a Planck Collaboration II (2014) cosmology,⁴ and for BAHAMAS (which has been run with many different cosmologies, with and without massive neutrinos) we used the simulation based on the same Planck Collaboration II (2014) results as EAGLE’s cosmology, with zero neutrino mass, which is described in McCarthy et al. (2018).⁵

For EAGLE, the DM and initial baryon particle masses are 9.7×10^6 and $1.8 \times 10^6 M_\odot$ respectively, while for BAHAMAS they are $6.6 \times 10^9 M_\odot$ and 1.2×10^9 . The corresponding Plummer-equivalent gravitational softening lengths are 0.7 kpc for EAGLE and 6.0 kpc for BAHAMAS, with these being fixed physical lengths (i.e. not comoving) at all redshifts we consider.

3.2 Lensing description

Our lensing procedure treats each halo as an isolated lens, ignoring the effect of other structures along the line of sight. This is done because we want to assign light-rays that have high magnifications to a single lensing object, in order to answer the question of what was responsible for the lensing. Ignoring the effects of multiple lens planes is justified by the fact that only a very small fraction of light-rays meet our criterion for being strongly lensed, such that the probability of an object being sufficiently aligned to be strongly lensed by two separate haloes is negligible. This was verified explicitly by Hilbert et al. (2007), who did full multiplane ray tracing through the Millennium simulation (Springel et al. 2005), and found that strong lensing events can almost always be traced to a single dominant lensing object. We stress that this does not mean that line-of-sight structures can be ignored in detailed lens modelling of individual systems, where including line-of-sight structures in the model can improve the match between the predicted and observed positions of multiply imaged sources (for example Chirivì et al. 2018). It means that the lensing cross-sections of lenses are not significantly altered on average by objects along the line of sight.

Our procedure for generating lensing maps from simulated haloes follows Robertson et al. (2019), who studied the Einstein radii of galaxy clusters from BAHAMAS simulations run with different DM models. For each halo at each snapshot redshift, we first find all mass within $5 r_{200}$ of the particle with the lowest gravitational potential energy. We then calculate the projected surface density, Σ , on a regular grid using an SPH-like smoothing scheme based on the distance to the 16th nearest neighbour. We make 3 maps of each halo – projecting the mass along 3 orthogonal lines of sight, for which we use the simulation x , y , and z axes. These maps are square, with a side length of $2 r_{200}$, and with 1024×1024 pixels. The resolution

⁴With $\Omega_m = 0.307$, $\Omega_b = 0.04825$, $\Omega_\Lambda = 0.693$, $\sigma_8 = 0.8288$, $n_s = 0.9611$, and $h = 0.6777$.

⁵This differs slightly from the EAGLE cosmology, with $\Omega_m = 0.3175$, $\Omega_b = 0.049$, $\Omega_\Lambda = 0.6825$, $\sigma_8 = 0.8341$, $n_s = 0.9624$, and $h = 0.6711$. We assume the EAGLE cosmology for all of our lensing calculations, including those with BAHAMAS.

³The galaxy and halo catalogues of the simulation suite, as well as the particle data, are publicly available at <http://www.eaglesim.org/database.php> (McAlpine et al. 2016; The EAGLE team 2017).




Figure 2. A schematic illustration of our volume of lenses, using the snapshot redshifts from BAHAMAS. The redshift range shown is covered by the three lowest redshift snapshots from BAHAMAS, with the red, green, and blue filled regions covering redshifts closest to $z^{\text{snap}} = 0, 0.125$, and 0.25 , respectively. The ticks along the left edge of the wedge show the redshifts at which we evaluate the contribution to the strong-lensing optical depth, which is done by taking the comoving volume corresponding to a redshift interval δz . When this redshift interval includes regions closer in redshift to two different snapshots (i.e. for z_2^L), the optical depth calculation uses only the snapshot closest in redshift to the centre of this redshift interval.

therefore increases in lower mass haloes, where the critical curve is on a smaller physical scale. At $z = 0$ this corresponds to a pixel scale of 4.5 kpc for a $10^{15} M_{\odot}$ halo and 0.2 kpc for a $10^{11} M_{\odot}$ halo. Note that the numerical values quoted here were motivated by the convergence tests discussed in Section 3.4.

Our calculation of $\partial^2 \tau / \partial \log M \partial z$ is done at more lens redshifts than we have simulation snapshots, as illustrated in Fig. 2. At each redshift that we use as a lens plane, z^L , we find the snapshot closest to it in redshift. We then use the Σ maps of the haloes from this snapshot to calculate the contribution to the strong lensing optical depth from this particular z^L . This is done by first dividing Σ by Σ_{crit} , to get the dimensionless convergence, κ . As both κ , and the gravitational shear, γ , are second derivatives of the projected Newtonian potential, they can be readily calculated from one another using discrete Fourier transforms (see e.g. Robertson et al. 2019). We can then make a map of the magnification μ using

$$\mu = \frac{1}{(1 - \kappa)^2 - |\gamma|^2}. \quad (15)$$

For lens planes at a different redshift from the snapshot used, we keep Σ as a function of physical coordinates fixed. For different lens planes that use the same snapshot, the differences between them are that Σ_{crit} changes with z^L (leading to a difference in the relationship between Σ and the convergence, κ) as does the relationship between physical distances in the lens plane and angles on the sky.

3.3 Cross-sections and optical depths

With μ on a regular grid in the image plane, we calculate the source plane solid angle with $|\mu| > \mu_0$ behind some lens as

$$\sigma_{\text{lens}}^S = \sum_{|\mu| > \mu_0} \frac{\sigma_{\text{pixel}}}{|\mu|}, \quad (16)$$

where σ_{pixel} is the solid angle of each image-plane pixel, and the sum is over all pixels with $|\mu| > \mu_0$.⁶

At each lens plane redshift ($z_i^L = 0.04i - 0.02$ for $i = 1, 2, \dots, 50$), we calculate σ_{lens}^S of each halo with $M_{200} > 10^{11} M_{\odot}$ in EAGLE and $M_{200} > 10^{13.5} M_{\odot}$ in BAHAMAS. We bin haloes by $\log_{10}(M_{200}/M_{\odot})$ and sum up σ_{lens}^S within each bin, where the bin width is $\Delta \log_{10} M_{200} = 0.18$. Dividing this sum by three times,⁷ the simulation volume and the log-mass bin width we get the left side of equation (13), which we can convert to $\partial^2 \tau / \partial \log M \partial z$ using equation (14).

3.4 Numerical convergence of strong lensing cross-sections

Before we did our full lensing analysis of the EAGLE and BAHAMAS simulations, we first performed a number of tests to determine that our lensing procedure produced numerically converged results. We did these tests on a single snapshot from each of EAGLE and BAHAMAS, so that we could test a large number of possible numerical parameters. These snapshots were chosen to be at a redshift that contributes significantly to the lensing of high-redshift sources, and to be at a redshift where EAGLE and BAHAMAS have a similarly timed output. We chose to use the $z = 0.366$ snapshot from EAGLE, and the $z = 0.375$ snapshot from BAHAMAS.

Accurately calculating the distribution of magnifications due to some mass distribution requires that the critical curves be adequately captured. This in turn requires that the mass distribution within the critical curves is sampled with a reasonable number of pixels. Smaller haloes have smaller critical curves, and hence require higher resolutions, while larger haloes need large fields of view to include the full halo's mass distribution. In order to achieve both of these, we used a pixel size and field of view that both increase with increasing halo mass, specifically in proportion to the virial radius ($\propto M_{200}^{1/3}$).

As both our pixel size and field-of-view scale with the virial radius, the number of pixels used for the mass maps is independent of halo mass. We experimented with 256, 512, 1024, and 2048 pixels on a side, making square maps with a side length of $2r_{200}$. The pixel scale acts as a scale on which the mass distribution is smoothed, and with the largest pixels ($r_{200}/128$), our lensing cross-sections were substantially reduced compared to smaller pixel cases. The two smallest pixel scales ($r_{200}/512$ and $r_{200}/1024$) produced converged results, suggesting that the larger of those two pixel scales (corresponding to 1024 pixels on a side) is more than adequate.

The reason that our lensing procedure becomes insensitive to the pixel scale is because it already smooths out the mass of individual particles on a scale that depends on the local number density of particles. This means that so long as the pixels are not too large, smaller pixels do not lead to smaller structures being resolved, but rather just a finer sampling of a smooth mass distribution. Because the resolution of our lensing maps is set by the density of simulation particles, we needed to check that a different simulation resolution,

⁶We stress again that this is not strictly the source plane solid angle that is magnified by greater than μ_0 , as it counts multiply imaged regions of the source plane multiple times.

⁷To reflect the three projection axes.

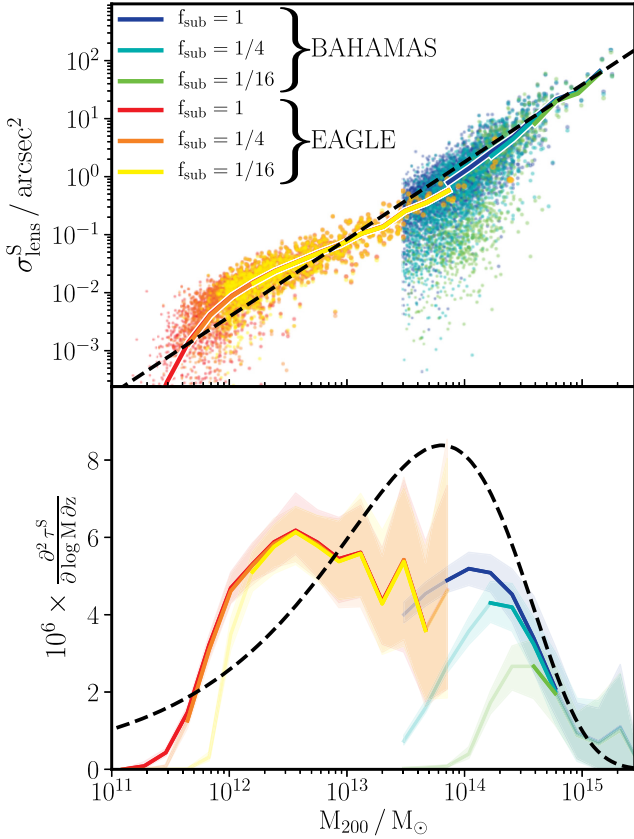


Figure 3. Top: the $z^S = 2$, $|\mu| > 10$, source-plane cross-sections for individual lenses from EAGLE and BAHAMAS, from the respective snapshots with $z \approx 0.37$. The different colours correspond to different f_{sub} values as described in Section 3.4. Running medians as a function of halo mass are shown as solid lines, which are drawn only in the mass range where the relevant simulation with the particular f_{sub} value is ‘converged’. The black dashed line shows the prediction for haloes modelled as singular isothermal spheres all the way out to their virial radii. Bottom: the optical depth per decade in halo mass and per unit redshift, calculated from the cross-sections in the top panel. Lines are shown as solid in the halo mass ranges where they are converged, and are faded at lower masses. We stop showing EAGLE lines at the high-mass end due to the poor sampling of the mass function there. The shaded regions around the lines show the 2.5th to 97.5th percentiles when the optical depth is calculated from a bootstrap resampling of all haloes in the top panel. For the black dashed line a Tinker et al. (2008) mass function was assumed.

with a different number density of simulation particles, would have provided converged results.

To investigate this, we made lensing maps of all haloes in our test snapshots, both with the full simulation data, and when only using a fraction, f_{sub} , of the simulation particles. When making these *subsamped* maps, the mass of each particle was increased by $1/f_{\text{sub}}$ to create a lower resolution version of the same simulated mass distribution. With these subsampled versions, we could then find the halo mass at which the lensing properties of the subsampled haloes disagreed with those of the full haloes, which indicates the halo mass down to which we can trust the lensing maps being generated from full simulation data.

In the top panel of Fig. 3, we show the $z^S = 2$ source plane cross-sections for $|\mu| > 10$, for individual haloes from both EAGLE and BAHAMAS, using subsampling factors of 1, 1/4, and 1/16. The bottom panel then shows $\partial^2 \tau / \partial \log M \partial z$ with these different subsampling

factors. For the case of EAGLE we see that other than a slight suppression at masses below $4 \times 10^{11} M_{\odot}$, the $f_{\text{sub}} = 1/4$ optical depth is indistinguishable from the full simulation case. We can also see that EAGLE has just enough resolution to resolve the lowest mass haloes important for lensing, as had the mass resolution been 16 times worse (corresponding to the yellow line), the low- M_{200} cut off in $\partial^2 \tau / \partial \log M \partial z$ would have been numerical rather than properly resolved.

The mass-scale at which convergence is achieved in BAHAMAS is, unsurprisingly, different, given the much poorer resolution of the BAHAMAS simulations. Defining *convergence* as an agreement on $\partial^2 \tau / \partial \log M \partial z$ better than 20 per cent between subsequent f_{sub} levels, $f_{\text{sub}} = 1/16$ is converged down to $M_{200} \approx 3.2 \times 10^{14} M_{\odot}$ and $f_{\text{sub}} = 1/4$ down to $M_{200} \approx 1.6 \times 10^{14} M_{\odot}$. Assuming a similar fractional improvement in the mass down to which we are converged going from $f_{\text{sub}} = 1/4$ to $f_{\text{sub}} = 1$ as we had when going from $f_{\text{sub}} = 1/16$ to $f_{\text{sub}} = 1/4$, we expect that the full BAHAMAS results are converged down to a halo mass of around $10^{14} M_{\odot}$. For this reason, we use BAHAMAS to make predictions for the contribution of $M_{200} > 10^{14} M_{\odot}$ lenses to the strong lensing optical depth, using EAGLE for halo masses below this. Owing to its relatively small box size, EAGLE has few haloes at masses $10^{13.5} - 10^{14} M_{\odot}$. As we will see in the next section, this leads to our lensing calculation being most uncertain at these intermediate masses, where BAHAMAS is not well resolved, but EAGLE suffers from a small volume.

4 RESULTS

The differential optical depth for strong lensing, $\partial^2 \tau / \partial \log M \partial z$, is plotted as a function of lens redshift and halo mass in Fig. 4, where we calculate this quantity from BAHAMAS for $M_{200} > 10^{14} M_{\odot}$ and from EAGLE for $M_{200} < 10^{14} M_{\odot}$. We remind the reader that this is a source plane optical depth for magnifications greater than 10. We also plot $\partial \tau / \partial z$ (from integrating over lens mass) and $\partial \tau / \partial \log M$ (from integrating over lens redshift).

The first comment to make is that both $\partial \tau / \partial z$ and $\partial \tau / \partial \log M$ are quite broad, so there is no one population of lenses that dominates the optical depth for high-magnification lensing. Integrating over both lens mass and lens redshift, the optical depth for $|\mu| > 10$ with $z^S = 2$ is approximately 2×10^{-5} . As explained earlier, for high-magnification thresholds, the optical depth is inversely proportional to the square of the threshold, such that the high magnifications required to significantly alter the true masses of compact binaries that have been detected with gravitational waves ($\mu \gtrsim 50$) will happen to only a very small fraction of all $z \sim 2$ compact binary coalescences (fewer than one in a million).

Results from the SIS model are also included in Fig. 4, where it can be seen that it does a reasonable job of reproducing the simulation-derived optical depth, including the relative contribution from different lens redshifts ($\partial \tau / \partial z$). This suggests that it will capture the dependence of τ on source redshift, which is the quantity required for calculating the expected number of lensed gravitational wave events that we should detect. Where this SIS model works less well is in the relative importance of different lens masses, with it predicting that low-halo masses ($< 10^{12} M_{\odot}$) make a reasonable contribution to the optical depth, while the hydrodynamical simulations have a steep drop in $\partial \tau / \partial \log M$ below $10^{12} M_{\odot}$. The mass scale at which the simulated systems become inefficient lenses is related to the stellar-to-halo mass relation for galaxies, whose behaviour changes at halo masses of around $10^{12} M_{\odot}$. We discuss this further in Section 4.2.

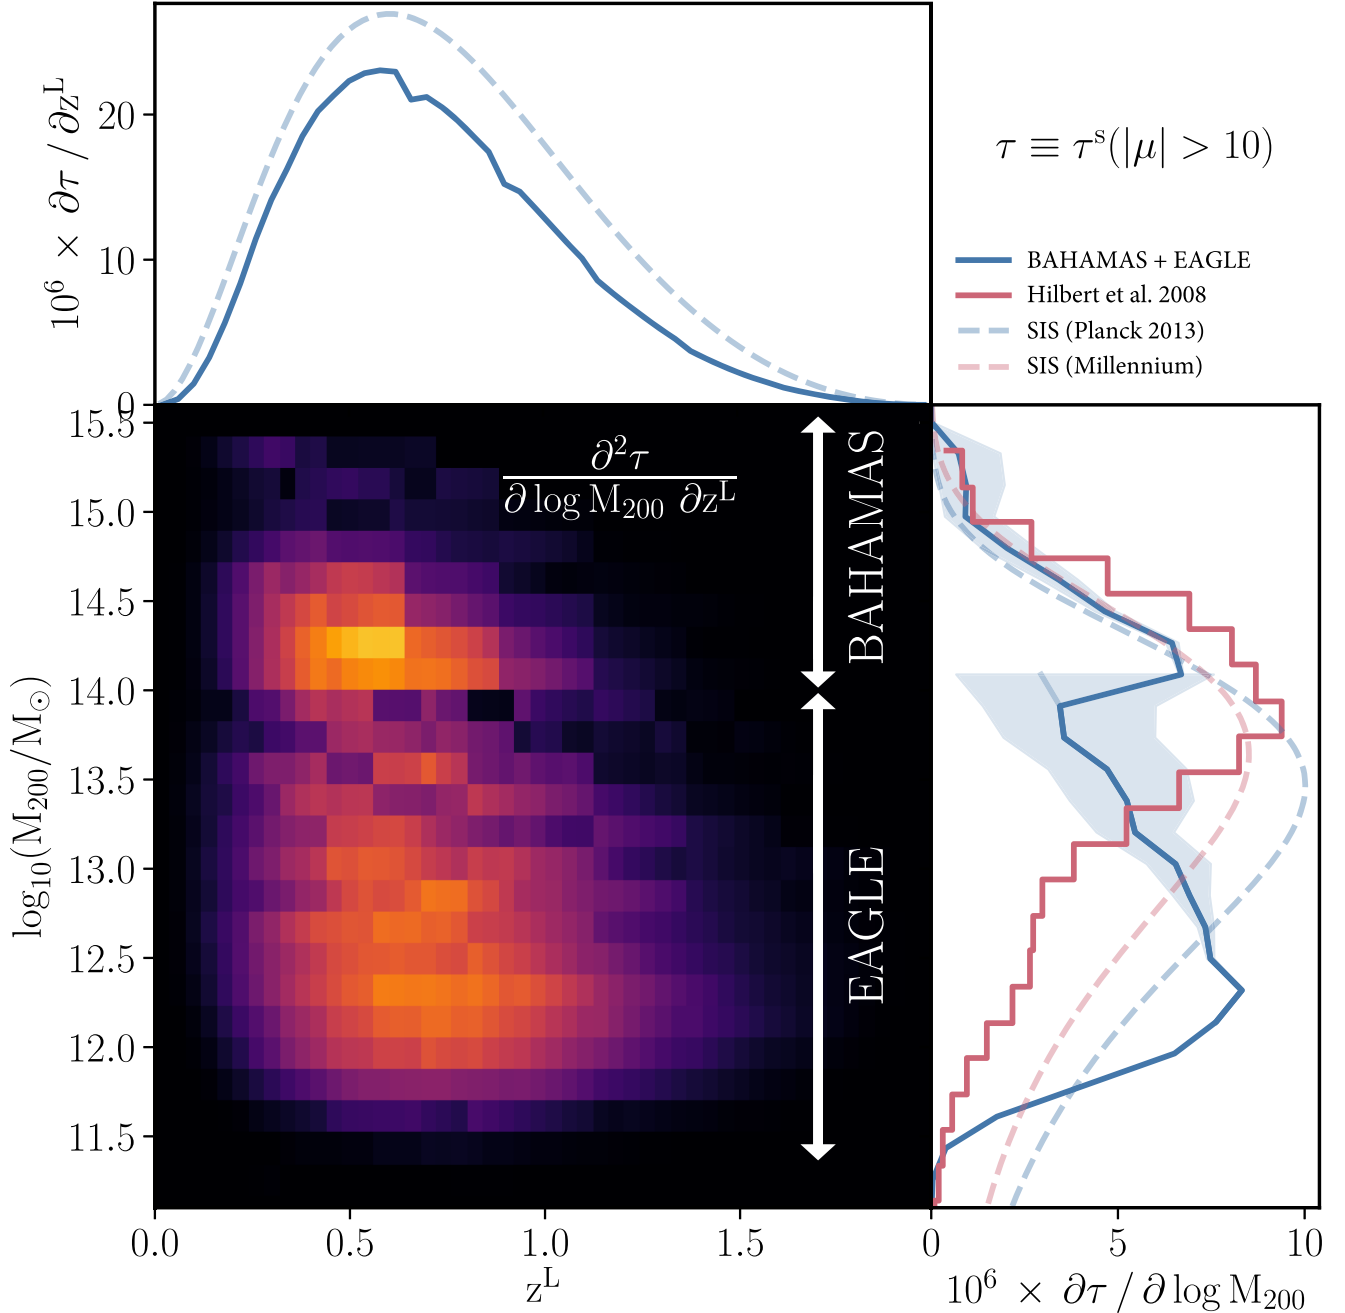


Figure 4. The main panel shows the source-plane optical depth for lensing by a magnification greater than 10, at a source redshift of 2, per unit lens redshift and per decade in lens halo mass. At halo masses above $10^{14} M_\odot$, we use the BAHAMAS simulation while EAGLE is used at lower masses. Above and to the right of the main panel we show this quantity marginalized over halo mass and lens redshift, respectively, as solid blue lines. For comparison, we show these same marginalized quantities for our SIS + Tinker et al. (2008) mass function model as blue dashed lines. We also show $\partial\tau/\partial\log M$ from H08 as the red stepped line. H08 used a different cosmology and mass definition from us, and to show the contribution of these different definitions to the differences between our results and those from H08, we show an SIS model with the cosmology and mass definition used by H08 as the red dashed line. The dip in the BAHAMAS + EAGLE results just below $10^{14} M_\odot$ is where BAHAMAS would not have sufficient resolution (so is not being used), but EAGLE’s volume leads to a poorly sampled mass function.

As another comparison, in Fig. 4 we plot $\partial\tau/\partial\log M$ from H08, who also used $z^S = 2$ and $|\mu| > 10$ (private communication). In H08, the optical depth was calculated from ray tracing through a DM-only simulation, with the addition of analytic gravitational potentials associated with galaxies, where the mass distribution within the galaxies was taken from the results of a semi-analytic galaxy formation model. The agreement between our $\partial\tau/\partial\log M$

and that from H08 is fairly good in general, although there are differences that we discuss further in Section 4.3.

4.1 Random error on $\partial\tau/\partial\log M$

The discontinuity in $\partial\tau/\partial\log M$ at $10^{14} M_\odot$ is because we change the simulation used at this mass. The discrepancy between the two

simulations at this mass scale, with EAGLE predicting a factor of 2 less lensing from $M_{200} \approx 10^{14} M_{\odot}$ haloes than BAHAMAS, could be for a number of reasons. Given that poor resolution leads to a decrease in the strong lensing cross-section of haloes (Fig. 3), it is unlikely to be resolution related. However, it could be that the different baryonic physics prescriptions lead to different predictions for the distribution of mass within haloes at this mass scale. Alternatively, it could just be random error associated with the EAGLE prediction, because the low volume of EAGLE means that the mass function at the high-mass end is poorly sampled. This can lead to noise both on the number of haloes and their mean lensing cross-section. The top panel of Fig. 3 shows that the median $\sigma_{\text{lens}}^{\text{S}}$ of EAGLE galaxies at this mass scale is a factor of almost two lower than in BAHAMAS, which explains the discrepancy in $\partial\tau/\partial\log M$. However, without more EAGLE haloes at this mass scale, we cannot say whether this is a systematic difference, or just a random quirk of the particular sample of massive haloes in EAGLE.

Using either the measured mass function from our simulations, or a mass function such as that from Tinker et al. (2008) combined with the volume of our simulations, it is possible to work out the expected number of haloes at each redshift in a given mass bin. However, using this with an assumption of Poisson statistics for the number of haloes in a given mass and redshift bin would underestimate the random error on $\partial\tau/\partial\log M$ for two reasons. First, the different snapshots are not independent, because it is the same haloes (evolving through time) that appear in the different snapshots. Secondly, there is fairly large scatter in $\sigma_{\text{lens}}^{\text{S}}$ at fixed lens redshift and halo mass. For BAHAMAS, the distribution of $\sigma_{\text{lens}}^{\text{S}}$ at $M_{200} \approx 10^{14} M_{\odot}$, $z^{\text{L}} = 0.375$ is well fit by a lognormal distribution with a standard deviation of 0.35 dex (this scatter can be seen in the top panel of Fig. 3). Such a distribution has 50 per cent of the signal coming from only 20 per cent of the objects, so that the scatter in the total lensing signal from a mass bin is significantly larger than the Poisson expectation.

In order to get an estimate of the random error on $\partial\tau/\partial\log M$ that reflects the points above, we use a bootstrap technique. Separately for both EAGLE and BAHAMAS, we extract the 100 most massive haloes from the $z = 0$ snapshot and find their primary progenitors in the preceding snapshots. We then draw 100 haloes with replacement from these 100 $z = 0$ haloes, and calculate $\partial\tau/\partial\log M$ with the 100 most massive haloes and their progenitors replaced by the bootstrapped sample of 100 haloes and their progenitors. We do this 1000 times and find the 2.5 per cent and 97.5 per cent percentiles for $\partial\tau/\partial\log M$, which are shown by the blue shaded regions in Fig. 4. Note that the choice of the 100 most massive haloes is fairly arbitrary. In principle we could bootstrap sample all haloes, but we avoided this because our simple method for finding progenitors (using the most massive halo with a comoving centre within $4r_{200}$ of the halo comoving centre at the preceding snapshot) works only for the most massive haloes in the box. Owing to the sharply rising mass function towards low-halo masses, the random error associated with $\partial\tau/\partial\log M$ at low masses should be negligible. Taking the noisy sampling of the high-mass end into account, the BAHAMAS and EAGLE $\partial\tau/\partial\log M$ at $M_{200} \sim 10^{14} M_{\odot}$ are consistent at the 2σ level.

4.2 The lowest halo masses that are efficient lenses

An obvious feature in Fig. 4 is that $\partial\tau/\partial\log M$ drops rapidly at halo masses below $10^{12} M_{\odot}$. We attribute this to the rapid fall off in stellar mass below this halo mass. Using abundance matching, the observed stellar mass function can be combined with N -body simulations to predict the stellar mass that resides in different

halo masses at different redshifts (Behroozi, Wechsler & Conroy 2013; Moster, Naab & White 2013). The resulting stellar-to-halo mass relation is matched reasonably well by the EAGLE simulations (Schaye et al. 2015). At a given redshift, the relationship between stellar mass and halo mass is approximately a double power law, which breaks at a characteristic halo mass of $M_1 = 10^{11.5-12} M_{\odot}$. At halo masses below this, the stellar mass (M^*) falls off rapidly with decreasing halo mass. In the local Universe, this fall off is approximately $M^* \propto M_{200}^{5/2}$, while at $z = 2$ it is closer to $M^* \propto M_{200}^2$. This fall off in stellar mass at low-halo masses means that stars quickly become unimportant for lensing when $M_{200} < M_1$. These low-mass haloes then only have a DM component, which should be reasonably well fit by a Navarro, Frenk, and White (NFW) density profile (Navarro, Frenk & White 1997). The strong lensing cross-section of an NFW halo decreases exponentially with decreasing halo mass below $10^{13} M_{\odot}$ (Hilbert et al. 2007), so that these low-mass haloes have very low lensing cross-sections, which – even combined with their large abundance – leads to them making a negligible contribution to the total optical depth.

While the median $\sigma_{\text{lens}}^{\text{S}}$ is a monotonically increasing function of M_{200} , inspection of Fig. 3 reveals a feature around $10^{12} M_{\odot}$, where lensing is most efficient compared with the SIS prediction. This is the same mass scale at which M^*/M_{200} peaks (Moster et al. 2013), and this feature in $\sigma_{\text{lens}}^{\text{S}}(M_{200})$ can be ascribed to the turnover of M^*/M_{200} .

4.3 Comparison with Hilbert et al. (2008)

Compared with the H08 $\partial\tau/\partial\log M$, hydrodynamical simulations predict slightly more strong lensing in total, with the increase primarily due to galaxies with halo masses of $10^{12}-10^{13} M_{\odot}$. The cosmology and mass definition⁸ are different between H08 and the hydrodynamical simulations. To investigate whether this explains the different lensing predictions, we calculated two different $\partial\tau/\partial\log M$ from our SIS model – one using the cosmology and mass definition from EAGLE and BAHAMAS, the other those adopted by H08. These two $\partial\tau/\partial\log M$ are plotted as dashed lines in the right-hand panel of Fig. 4, and the differences between them are smaller than the differences between the lines from the different simulations.

For haloes with $M_{200} > 10^{14} M_{\odot}$, i.e. those from BAHAMAS, $\partial\tau/\partial\log M$ is similar between the hydrodynamical simulations and H08. Robertson et al. (2019) looked at the density profiles of BAHAMAS clusters, both DM-only and including baryons, and found that the total density at the centre of clusters increases when simulations include baryons, not just because of the contribution from stars (which should be captured by H08 who included analytical potentials associated with a stellar disc and bulge), but also because the DM profile itself becomes more centrally concentrated due to adiabatic contraction (Gnedin et al. 2004). The fact that this does not push the BAHAMAS curve above that from H08 probably reflects the fact that in H08 the analytical baryonic potential is added to the DM distribution from a DM-only simulation. Such a simulation does not ignore the baryonic material in the universe, rather it simulates it as if it were also made of DM. As such, there is more DM in a DM-only simulation than in a hydrodynamical simulation, such that when H08 add in a stellar component there is now more total mass in each halo. This appears to mimic the effects of adiabatic contraction on

⁸H08 use $M_{200, \text{mean}}$ whereas we use $M_{200, \text{crit}}$.

strong lensing cross-sections such that the hydrodynamical and DM-only plus analytical galaxies predictions are similar in the galaxy cluster regime.

The largest discrepancy between our results and those from H08 is in the 10^{12} – $10^{13} M_{\odot}$ mass range. At these halo masses the DM-only prediction is for negligible amounts of strong lensing (Hilbert et al. 2007), so the lensing signal for these systems is dominated by the stellar component. In EAGLE, the lensing cross-sections of these galaxies are well resolved (Fig. 3), and in H08 they use analytic expressions to calculate the ray distortions induced by stellar mass distributions, so they are not limited by any sort of resolution effects. The difference must therefore come down to differences between the stellar distributions found in EAGLE and those used by H08. H08 modelled the stellar component as the sum of an exponential disc

$$\Sigma_{\text{disc}} = \frac{M_{\text{disc}}}{2\pi r_{\text{disc}}^2} \exp\left(-\frac{r}{r_{\text{disc}}}\right) \quad (17)$$

and a bulge

$$\Sigma_{\text{bulge}} = \frac{94.5 M_{\text{bulge}}}{r_{\text{bulge}}^2} \exp\left[-7.67 \left(\frac{r}{r_{\text{bulge}}}\right)^{1/4}\right]. \quad (18)$$

The disc mass, bulge mass, and disc radius (M_{disc} , M_{bulge} , and r_{disc} , respectively) were taken from the De Lucia & Blaizot (2007) semi-analytic model that had been run on merger trees generated from the same Millennium simulation that H08 used for the DM component. They used an observationally derived $M_{\text{bulge}}-r_{\text{bulge}}$ relation to get the bulge radius.

To verify that these parametric mass distributions, with values taken from De Lucia & Blaizot (2007), produce less strong lensing than EAGLE galaxies, we use the Millennium data base⁹ to find the De Lucia & Blaizot (2007) galaxy parameters for haloes with $10^{12} < M_{200}/M_{\odot} < 10^{13}$ in the $z = 0.362$ snapshot. For each halo, we calculate $\sigma_{\text{lens}}^{\text{S}}$ for $|\mu| > 10$, which we can then compare with the EAGLE haloes in Fig. 3. We find that at fixed halo mass, the median $\sigma_{\text{lens}}^{\text{S}}$ from this procedure is one order of magnitude lower than from EAGLE, which is true for both 10^{12} and $10^{13} M_{\odot}$ haloes. This comparison is slightly unfair, because we have not included the DM component in the Millennium lensing. At $10^{13} M_{\odot}$ (where EAGLE and H08 differ by a factor of two to three in $\partial\tau/\partial\log M$), the DM fraction within the Einstein radius is typically significant, whereas at $10^{12} M_{\odot}$ the lensing is dominated by the stars. As such, the increasing discrepancy in $\partial\tau/\partial\log M$ going from $10^{13} M_{\odot}$ to $10^{12} M_{\odot}$ is explained by the increasing importance of the stars for the lensing, and the stellar components used by H08 being less efficient lenses than those found in EAGLE.

Given that EAGLE has a stellar mass – stellar size relation that is a good match to observations (Schaye et al. 2015), and a similar comparison with the stellar mass–halo mass relation suggests that EAGLE has too few stars in haloes around $10^{12} M_{\odot}$, it seems unlikely that EAGLE is significantly overestimating the lensing contribution from $10^{12} M_{\odot}$ haloes. A thorough analysis of the De Lucia & Blaizot (2007) semi-analytic galaxies and their lensing signal as implemented by H08 is beyond the scope of this work, but here we mention possible reasons for the lensing signals being lower than in EAGLE. The simplest possibility is that there is simply not enough mass in stars, or that the galaxies are too large (and therefore

more diffuse, and so less efficient strong lenses). Another possibility is that departures from circular symmetry may be important. In particular, H08 assume the stellar disc is always seen face on. Most of their galaxies in the 10^{12} – $10^{13} M_{\odot}$ halo mass range are disc-dominated, and viewing these discs face-on their surface densities rarely exceed Σ_{crit} . An edge-on disc reaches much higher surface densities and is therefore a more powerful lens (Bartelmann & Loeb 1998; Blain, Moller & Maller 1999). The choice to place all discs face-on may therefore cause H08 to underestimate the contribution to the strong lensing optical depth from lower mass haloes.

4.4 Comparison with lensed submillimeter galaxies

As mentioned in the introduction, one of the arguments for strong lensing being dominated by galaxy lenses, is that these are the primary lensing population when galaxies detected as being the brightest at submillimeter wavelengths are followed up to allow identification of a potential lens (e.g. Wardlow et al. 2013). Here, we sketch out a qualitative argument for why this is expected, and does not contradict our finding that lenses with $M_{200} > 10^{13} M_{\odot}$ make a dominant contribution to the strong lensing optical depth for high magnifications of point sources.

The very brightest objects observed in the submillimeter are almost entirely gravitationally lensed, which can be understood from the steepness of the bright-end of the intrinsic luminosity function of dust-obscured star-forming galaxies (Perrotta et al. 2002; Lima et al. 2010). This means that a small fraction of less-bright objects (of which there are many) being highly magnified, can dominate over the intrinsically bright objects with the same observed flux (Negrello et al. 2010).¹⁰

The flux limits employed in submillimeter surveys to find likely lensed objects, primarily select for objects with only modest magnifications. For example, Wardlow et al. (2013) expect magnifications of around 9 given their sample cuts. Given the universal form of $P(> \mu) \propto \mu^{-2}$ for large μ , if a class of objects dominates lensing for large magnifications (greater than 100 say), then it should also dominate lensing for more modest magnification (such as those relevant for submillimeter galaxies). However, this relation is true only for point sources, with extended sources having more complicated magnification distributions.

It is perhaps intuitive that with increasing source size, the maximum magnification achievable decreases, as less of the source can lie close to a caustic. A less obvious fact is that this decrease in maximum magnification is accompanied by an increase in the probability of being moderately magnified. This can be understood from noting that for an extended source, the magnification of the source as a whole is a weighted mean of the source plane magnification for point sources over the surface brightness profile of the extended source. As such, the mean magnification of sources randomly distributed on the source plane must be independent of their size. A large source cannot all be close to a caustic, but there is an increased chance that at least some of it will be. This effect was calculated explicitly by de Freitas, Makler & Dúmet-Montoya (2018), who showed that for a circular source with constant surface brightness lensed by an SIS, the cross-section for $|\mu| > 10$ is

⁹<http://gavo.mpa-garching.mpg.de/Millennium/>, described in Lemson & Virgo Consortium (2006). Strictly speaking we use the milli-Millennium data base, similar to the Millennium data base but openly accessible, and for a simulation with a volume 1/512 of the full Millennium simulation.

¹⁰This is the same argument being made by Broadhurst et al. (2018), who, by assuming that the mass function of black holes exponentially decreases at masses $> 10 M_{\odot}$, find that objects detected as having masses $\sim 30 M_{\odot}$ would in fact be dominated by lensed objects that are intrinsically less massive.

maximized when the angular radius of the source is 30 per cent of the Einstein radius. As galaxies have smaller Einstein radii than clusters, the relative size of a submillimeter source is larger when lensed by a galaxy than by a cluster. This decreases the importance of galaxy-lenses for very high magnifications of extended sources, but increases their prevalence as lenses with moderate magnifications, which is what dominates observed submillimeter samples.

4.5 Implications for searches for lensed gravitational waves

Aside from its importance for trying to estimate the probability of strong lensing from observed galaxies or galaxy clusters, knowing which halo masses are responsible for strong lensing is important for strategies to find optical counterparts to gravitationally lensed gravitational waves. At present, gravitational wave detections have large positional uncertainties, with 90 per cent confidence regions typically covering a few hundred square degrees (Abbott et al. 2018). Surveying this whole area with optical telescopes requires a large number of exposures to *tile* the sky localization region.

Reducing the telescope time required to find optical counterparts to gravitational waves requires novel observing strategies. For example, the optical counterpart to the binary neutron star merger detected by LIGO and Virgo (GW170817, Abbott et al. 2017b) was first discovered (Coulter et al. 2017) using an observing strategy that targeted known galaxies in the three-dimensional LIGO-Virgo localization (Gehrels et al. 2016).

For high-redshift gravitational wave sources, narrowing down the search using plausible host galaxies will be difficult, first because a larger fraction of possible hosts will be undetected and secondly because of the large number of galaxies per unit solid angle at high redshift compared with low redshift. However, if there is evidence that a particular source may have been lensed, then a search strategy that targets lenses rather than source-hosts can be used. Such a strategy has been employed by Smith et al. (2019b), who observed two known strong lensing clusters within the sky localization of a binary black hole seen by LIGO-Virgo (GW170814, Abbott et al. 2017a).

Assuming that an observed gravitational wave has been strongly lensed, and that there is an electromagnetic counterpart to detect, for the Smith et al. (2019b) strategy to have a high chance of success requires that massive clusters dominate the strong lensing optical depth. From $\partial\tau/\partial\log M$ in Fig. 4, we can see that the most massive clusters contribute only a small fraction of the total optical depth. In fact, using the BAHAMAS + EAGLE prediction, if we take a $z = 2$ point source that we know to have been highly magnified, the probability that it was lensed by a halo with $M_{200} > 10^{15} M_{\odot}$ is only around 2 per cent, rising to 25 per cent for $M_{200} > 10^{14} M_{\odot}$ and 50 per cent for $M_{200} > 10^{13} M_{\odot}$. However, a single powerful strong-lensing cluster can have a source-plane cross-section for $|\mu| > 10$ of ~ 500 arcsec².¹¹ This would account for around 5 per cent of the highly magnified source plane solid angle within the sky localizations of the best localized GW detections (90 per cent localization uncertainty of ~ 50 degree²), if indeed such a cluster lens is located within the sky uncertainties. In general, our results indicate that finding optical counterparts to high-redshift strongly lensed GWs will most likely require very wide-field optical follow-up as discussed by Smith et al. (2019a). Full exploration of such wide-field follow-up will

¹¹As a concrete example, MACS 0717 has $\sigma_{\text{lens}}^{S}(|\mu| > 30) = 65$ arcsec² (Vega-Ferrero, Diego & Bernstein 2019). Converting this to $|\mu| > 10$ using $P(|\mu| > \mu_0) \propto 1/\mu_0^2$ gives a cross-section of 580 arcsec².




Figure 5. The source-plane optical depth for magnifications greater than 10, as predicted by different models. The solid blue line shows our prediction from hydrodynamical simulations, while the blue dashed line shows the prediction from our SIS model. The red line shows results from H08, who also used cosmological simulations, but using a semi-analytic model of galaxy formation rather than hydrodynamical simulations. The yellow line shows the results from the SIS model used by Hannuksela et al. (2019), in which the mass function of lenses does not change with redshift. The different predictions agree to within a factor of two over most of the redshift range, except for the Hannuksela et al. (2019) line which is a factor of four to six larger than the others.

benefit from highly complete lists of strong-lensing systems down to halo masses of $10^{13} M_{\odot}$ (Rydzanowski et al. 2020).

4.6 Strong lensing probabilities at different source redshifts

The mass and redshift distribution of gravitational lenses is important for designing strategies to find the likely lenses of a lensed gravitational wave, as we have discussed. However, for calculating the rate of detectable lensed events, or for finding the probability that a given event has been strongly lensed (assuming one does not have an accurate sky localization), the important quantity is simply the probability of strong lensing as a function of source redshift. Ng et al. (2018) have made predictions for the rate at which lensed gravitational waves should be detected, while Hannuksela et al. (2019) recently found no evidence that any observed gravitational wave signals have been strongly lensed. Both groups used a strong lensing optical depth based on the assumption that lensing was done by a population of SIS mass profiles, whose mass function does not evolve, and has a normalization determined from galaxy surveys (Fukugita & Turner 1991).

In Fig. 5, we plot the evolution of the source-plane optical depth for high-magnification lensing as predicted from our simulations. This was done by calculating $\partial^2\tau/\partial\log M \partial z$, as shown for $z^S = 2$ in Fig. 4, but at many different source redshifts. At each z^S , we integrate over lens mass and lens redshift to obtain $\tau(z^S)$. As we only generated mass maps for $z^L \leq 2$, we can only calculate the optical depth out to $z^S = 2$. For comparison we show this same quantity from using our SIS model and from H08. We find that the two simulation predictions agree well with one another, while the SIS model predicts slightly more strong lensing, particularly at low source redshifts. We also used the SIS model to understand the impact of a change in cosmology from that adopted in our simulations (Planck Collaboration II 2014) to the latest *Planck* cosmology (Planck Collaboration VI 2018). This latest cosmology has a slightly increased Ω_m and a decreased σ_8 compared with the

earlier *Planck* cosmology. These changes approximately cancel one another to produce a $\tau(z^S)$ that differs from the Planck Collaboration II (2014) SIS result in Fig. 5 by less than the line width (the Planck Collaboration VI 2018 line is not plotted in fig. 5).

To compare with work that was used to calculate the probabilities that observed gravitational waves have been strongly lensed, we also plot the optical depth used by Hannuksela et al. (2019). They used a model where the population of lenses was comprised of singular isothermal spheres, with a mass function that does not evolve with time. This leads to an optical depth as a function of source redshift

$$\tau(z^S, |\mu| \geq \mu_0) = F \left(\frac{d_C(z^S)}{cH_0^{-1}} \right)^3 \left(\frac{2}{\mu_0} \right)^2, \quad (19)$$

where $d_C(z)$ is the comoving distance to redshift z , and F is a dimensionless constant for which they use 0.0017. This is plotted in Fig. 5 for $\mu_0 = 10$, along with the same quantity predicted from our simulations, our SIS model and from H08.

In Fig. 5, the different models agree reasonably well on $\tau(z^S)$, except for that used by Hannuksela et al. (2019) which predicts more strong lensing than the others by a factor of approximately five. This means their value for F is probably optimistic in terms of the amount of strong lensing, but a lower value for F would only strengthen their conclusion that it is unlikely that any of the already detected gravitational wave signals have been highly magnified by gravitational lensing.

5 CONCLUSIONS

In this paper, we have calculated the optical depth for high-magnification gravitational lensing of high-redshift point sources, and its contribution from different lens masses and redshifts, as predicted by two recent hydrodynamical simulations, EAGLE and BAHAMAS. We combine these two simulations because EAGLE has sufficient resolution to resolve the lowest mass haloes important for strong lensing, while the large volume of BAHAMAS allows for an adequate sampling of the high-mass end of the halo mass function. The predicted contribution from group mass (10^{13} – $10^{14} M_\odot$) haloes is rather uncertain, because the strong lensing regions of these haloes are not well resolved in BAHAMAS, and the EAGLE simulation has insufficient volume for a robust prediction. Future simulations, with sufficiently large volume and resolution to accurately estimate the lensing contribution from galaxy groups, will improve our predictions further.

The relative contribution of different lens masses to the total optical depth for large magnifications has been a topic of recent debate, particularly in the community studying the effects of gravitational lensing on observations of gravitational waves. The importance of galaxies versus galaxy clusters is a ‘tug-of-war’ between the high number density of galaxies, and the large lensing cross-sections of galaxy clusters. We first studied a simple model, in which haloes are modelled as singular isothermal spheres. In this model, the lensing cross-section of haloes grows as $M^{4/3}$, while the number density of haloes per decade in halo mass is approximately proportional to M^{-1} . As such, the high lensing cross-sections of more massive haloes wins over the increasing number of less massive haloes, and it is the more massive haloes that contribute more to the strong lensing cross-section. This argument holds up to the mass scale at which the mass function is no longer a power law, but is exponentially suppressed, which reduces the importance of the most massive haloes for strong lensing.

The hydrodynamical simulations confirm this picture, while altering slightly the importance of different halo masses. The

primary difference between the simulation predictions and those from our SIS model is that the simulations pick out a particular mass scale ($M_{200} \sim 10^{12} M_\odot$) as being more efficient at lensing than predicted by the SIS model. This scale corresponds to the scale at which galaxy formation is most efficient, in that the stellar mass to halo mass ratio is the highest there (Moster et al. 2013). A second result of the hydrodynamical simulations is that they predict a minimum halo mass below which lensing becomes inefficient. Again this is related to the stellar-to-halo mass relationship, as in low-mass haloes it is the stars rather than dark matter that dominate the strong lensing region. The stellar mass falls off quickly with decreasing halo mass below a halo mass of $10^{12} M_\odot$, and so lower mass haloes quickly become unimportant for lensing. Overall we find that around half of all high-magnification lensing is done by haloes with $M_{200} > 10^{13} M_\odot$ (i.e. galaxy groups and clusters), with the other half coming from less massive systems (galaxies). This result differs somewhat from a previous simulation-based result (H08), which used a dark matter only simulation combined with analytical gravitational potentials for the stellar components of galaxies, in that we find an enhanced contribution from galaxies living in 10^{12} – $10^{13} M_\odot$ haloes.

We also discussed the implications of this work for strategies to hunt for optical counterparts to gravitationally lensed gravitational waves. If the bulk of high magnification lines of sight resulted from the most massive haloes, then a credible strategy for finding the optical counterpart to a gravitational wave that was known to be strongly lensed, would be to look in the strong lensing region of the most massive haloes within the gravitational wave sky localization. However, massive clusters (with $M_{200} > 10^{15} M_\odot$) contribute only 2 per cent of the total optical depth for large magnifications, with the bulk of the signal coming from massive galaxies or low-mass clusters. The number densities of these objects is much higher, such that there will be many within the sky localization of a detected gravitational wave. This means that finding the optical counterparts to high-redshift gravitationally lensed gravitational waves will most likely require tiling the credible area on the sky as determined from the gravitational waves. That said, if a particularly powerful strong lensing cluster was in the sky localization of a well-constrained gravitational wave ($\sim 50 \text{ deg}^2$) then it could account for around 5 per cent of the strong lensing within the sky localization.

Finally, we presented our prediction for the optical depth for high magnification as a function of source redshift, a key ingredient in calculating the expected rate of gravitationally lensed gravitational waves, as well as the abundance of other gravitationally lensed point sources. We found that our result was in reasonable agreement both with H08’s previous simulation-based result, as well as simple models that treat lenses as SISs. At low redshift, high magnifications are extremely unlikely. Each $z^S = 0.5$ source produces, on average, 5×10^{-7} images with magnification greater than 10; this increases to about 2×10^{-5} images for $z^S = 2$. These lensing probabilities are lower than assumed by Hannuksela et al. (2019) in recent work on the probability that observed gravitational waves have been strongly lensed. They found that strong lensing was unlikely to have affected the current sample of observed gravitational waves, and our lower intrinsic lensing probabilities strengthen this result.

ACKNOWLEDGEMENTS

We give special thanks to Stefan Hilbert for producing new lensing results from the Millennium simulation, and to Ian McCarthy for granting us access to the particle data from the BAHAMAS simulations. Thanks also to Christopher Berry for helping with

the early stages of this work and for providing comments on a draft version of this paper, and to James Nightingale, David Lagattuta, Joop Schaye, Otto Hannuksela and Thomas Callingham for helpful discussions.

AR is supported by the European Research Council (ERC-StG-716532-PUNCA) and the Science and Technology Facilities Council (STFC) (ST/N001494/1). GPS and MB acknowledge support from STFC grant number ST/S000305/1. RM acknowledges the support of a Royal Society University Research Fellowship. VRE acknowledges support from STFC grant ST/P000541/1. MJ is supported by the United Kingdom Research and Innovation (UKRI) Future Leaders Fellowship ‘Using Cosmic Beasts to uncover the Nature of Dark Matter’ (MR/S017216/1).

This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

REFERENCES

- Abbott B. P. et al., 2016, *Phys. Rev. Lett.*, 116, 061102
 Abbott B. P. et al., 2017a, *Phys. Rev. Lett.*, 119, 141101
 Abbott B. P. et al., 2017b, *Phys. Rev. Lett.*, 119, 161101
 Abbott B. P. et al., 2017c, *Ann. Phys.*, 529, 1600209
 Abbott B. P. et al., 2018, *Living Rev. Relat.*, 21, 3
 Bahé Y. M. et al., 2017, *MNRAS*, 470, 4186
 Barnes D. J., Kay S. T., Henson M. A., McCarthy I. G., Schaye J., Jenkins A., 2017a, *MNRAS*, 465, 213
 Barnes D. J. et al., 2017b, *MNRAS*, 471, 1088
 Bartelmann M., Loeb A., 1998, *ApJ*, 503, 48
 Behroozi P. S., Wechsler R. H., Conroy C., 2013, *ApJ*, 762, L31
 Blain A. W., Moller O., Maller A. H., 1999, *MNRAS*, 303, 423
 Bolton A. S., Burles S., Koopmans L. V. E., Treu T., Gavazzi R., Moustakas L. A., Wayth R., Schlegel D. J., 2008, *ApJ*, 682, 964
 Broadhurst T., Diego J. M., Smoot George I., 2018, preprint ([arXiv:1802.05273](https://arxiv.org/abs/1802.05273))
 Broadhurst T., Diego J. M., Smoot George F. I., 2019, preprint ([arXiv:1901.03190](https://arxiv.org/abs/1901.03190))
 Bussmann R. S. et al., 2013, *ApJ*, 779, 25
 Chirivì G., Suyu S. H., Grillo C., Halkola A., Balestra I., Caminha G. B., Mercurio A., Rosati P., 2018, *A&A*, 614, A8
 Choi Y.-Y., Park C., Vogeley M. S., 2007, *ApJ*, 658, 884
 Coulter D. A. et al., 2017, *Science*, 358, 1556
 Crain R. A. et al., 2015, *MNRAS*, 450, 1937
 Cui W. et al., 2018, *MNRAS*, 480, 2898
 Davé R., Thompson R., Hopkins P. F., 2016, *MNRAS*, 462, 3265
 Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
 de Freitas V. P., Makler M., Dúmet-Montoya H. S., 2018, *MNRAS*, 481, 2189
 De Lucia G., Blaizot J., 2007, *MNRAS*, 375, 2
 Fukugita M., Turner E. L., 1991, *MNRAS*, 253, 99
 Gaudi B. S., Petters A. O., 2002, *ApJ*, 574, 970
 Gavazzi R., Treu T., Rhodes J. D., Koopmans L. V. E., Bolton A. S., Burles S., Massey R. J., Moustakas L. A., 2007, *ApJ*, 667, 176
 Gehrels N., Cannizzo J. K., Kanner J., Kasliwal M. M., Nissanke S., Singer L. P., 2016, *ApJ*, 820, 136
 Gnedin O. Y., Kravtsov A. V., Klypin A. A., Nagai D., 2004, *ApJ*, 616, 16
 Hannuksela O. A., Haris K., Ng K. K. Y., Kumar S., Mehta A. K., Keitel D., Li T. G. F., Ajith P., 2019, *ApJ*, 874, L2
 Haris K., Mehta A. K., Kumar S., Venumadhav T., Ajith P., 2018, preprint ([arXiv:1807.07062](https://arxiv.org/abs/1807.07062))
 Hilbert S., White S. D. M., Hartlap J., Schneider P., 2007, *MNRAS*, 382, 121v
 Hilbert S., White S. D. M., Hartlap J., Schneider P., 2008, *MNRAS*, 386, 1845 (H08)
 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
 Kaviraj S. et al., 2017, *MNRAS*, 467, 4739
 Kelly P. L. et al., 2015, *Science*, 347, 1123
 Kelly P. L. et al., 2018, *Nat. Astron.*, 2, 334
 Le Brun A. M. C., McCarthy I. G., Schaye J., Ponman T. J., 2014, *MNRAS*, 441, 1270
 Lemson G., Virgo Consortium T., 2006, preprint ([arXiv:astro-ph/0608019](https://arxiv.org/abs/astro-ph/0608019))
 Li S.-S., Mao S., Zhao Y., Lu Y., 2018, *MNRAS*, 476, 2220
 Lima M., Jain B., Devlin M., Aguirre J., 2010, *ApJ*, 717, L31
 McAlpine S. et al., 2016, *Astron. Comput.*, 15, 72
 McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, *MNRAS*, 465, 2936
 McCarthy I. G., Bird S., Schaye J., Harnois-Deraps J., Font A. S., van Waerbeke L., 2018, *MNRAS*, 476, 2999
 Miller M. C., Yunes N., 2019, *Nature*, 568, 469
 Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
 Murray S. G., Power C., Robotham A. S. G., 2013, *Astron. Comput.*, 3, 23
 Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
 Negrello M. et al., 2010, *Science*, 330, 800
 Ng K. K. Y., Wong K. W. K., Broadhurst T., Li T. G. F., 2018, *Phys. Rev. D*, 97, 023012
 Oguri M., Taruya A., Suto Y., Turner E. L., 2002, *ApJ*, 568, 488
 Oguri M. et al., 2013, *MNRAS*, 429, 482
 Perrotta F., Baccigalupi C., Bartelmann M., De Zotti G., Granato G. L., 2002, *MNRAS*, 329, 445
 Pillepich A. et al., 2018a, *MNRAS*, 473, 4077
 Pillepich A. et al., 2018b, *MNRAS*, 475, 648
 Planck Collaboration II, 2014, *A&A*, 571, A1
 Planck Collaboration VI, 2018, preprint ([arXiv:1807.06209](https://arxiv.org/abs/1807.06209))
 Press W. H., Schechter P., 1974, *ApJ*, 187, 425
 Richard J. et al., 2010, *MNRAS*, 404, 325
 Robertson A., Harvey D., Massey R., Eke V., McCarthy I. G., Jauzac M., Li B., Schaye J., 2019, *MNRAS*, 488, 3646
 Ryczanowski D., Smith G. P., Bianconi M., Massey R., Robertson A., Jauzac M., 2020, *MNRAS*, 495, 1666
 Schaye J. et al., 2015, *MNRAS*, 446, 521
 Schneider P., Ehlers J., Falco E. E., 1992, *Gravitational Lenses*. Springer-Verlag, New York
 Sembolini F., Yepes G., De Petris M., Gottlöber S., Lamagna L., Comis B., 2013, *MNRAS*, 429, 323
 Sharon K. et al., 2005, *ApJ*, 629, L73
 Sharon K. et al., 2017, *ApJ*, 835, 5
 Sheth R. K., Mo H. J., Tormen G., 2001, *MNRAS*, 323, 1
 Singer L. P., Goldstein D. A., Bloom J. S., 2019, preprint ([arXiv:1910.03601](https://arxiv.org/abs/1910.03601))
 Smith G. P., Jauzac M., Veitch J., Farr W. M., Massey R., Richard J., 2018, *MNRAS*, 475, 3823
 Smith G. P., Robertson A., Bianconi M., Jauzac M., 2019a, preprint ([arXiv:1902.05140](https://arxiv.org/abs/1902.05140))
 Smith G. P. et al., 2019b, *MNRAS*, 485, 5180
 Springel V. et al., 2005, *Nature*, 435, 629
 The EAGLE team, 2017, preprint ([arXiv:1706.09899](https://arxiv.org/abs/1706.09899))
 Tinker J., Kravtsov A. V., Klypin A. et al., 2008, *ApJ*, 688, 709
 Vega-Ferrero J., Diego J. M., Bernstein G. M., 2019, *MNRAS*, 486, 5414
 Vogelsberger M. et al., 2014, *Nature*, 509, 177
 Walsh D., Carswell R. F., Weymann R. J., 1979, *Nature*, 279, 381
 Wardlow J. L. et al., 2013, *ApJ*, 762, 59

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