A generalised method for measuring weak lensing magnification with weighted number counts

Published In:
Monthly Notices of the Royal Astronomical Society

OPEN ACCESS

A generalized method for measuring weak lensing magnification with weighted number counts

Bryan R. Gillis[★] and Andy N. Taylor

Institute for Astronomy, University of Edinburgh, Royal Observatory Edinburgh, Edinburgh EH9 3HJ, UK

ABSTRACT

We present a derivation of a generalized optimally weighted estimator for the weak lensing magnification signal, including a calculation of errors. With this estimator, we present a local method for optimally estimating the local effects of magnification from weak gravitational lensing, using a comparison of number counts in an arbitrary region of space to the expected unmagnified number counts. We show that when equivalent lens and source samples are used, this estimator is simply related to the optimally weighted correlation function estimator used in past work and vice-versa, but this method has the benefits that it can calculate errors with significantly less computational time, that it can handle overlapping lens and source samples, and that it can easily be extended to mass-mapping. We present a proof-of-principle test of this method on data from the Canada–France–Hawaii Telescope Lensing Survey, showing that its calculated magnification signals agree with predictions from model fits to shear data. Finally, we investigate how magnification data can be used to supplement shear data in determining the best-fitting model mass profiles for galaxy dark matter haloes. We find that at redshifts greater than $z \sim 0.6$, the inclusion of magnification can often significantly improve the constraints on the components of the mass profile which relate to galaxies' local environments relative to shear alone, and in high-redshift low- and medium-mass bins, it can have a higher signal-to-noise than the shear signal.

Key words: gravitational lensing: weak – methods: statistical.

1 INTRODUCTION

Weak gravitational lensing has proven to be a powerful probe of the distribution of dark matter in the Universe, allowing the dark mass within clusters to be mapped, providing useful measurements of the halo masses of galaxies at a variety of redshifts, and helping to constrain fundamental cosmological parameters, among other applications. However, these advances have relied only on shear, which is just one of the two effects weak gravitational lensing has on background galaxies; the other being magnification.

When a background object is affected weakly by lensing, its shape will be stretched tangentially to the direction of the lens, known as shear, and its size will change, known as magnification. For measuring lensing, shear presents two key advantages over magnification: it has an expectation value of zero in the absence of lensing, and it typically has a larger signal-to-noise ratio than magnification. The zero unlensed expectation value of shear has many benefits which make measuring it easier. It allows shear to be measured without calibration to determine the unlensed expectation value, and it also

indirectly leads to the result that most biases in measuring shear are multiplicative rather than additive. Any measurement of magnification, on the other hand, must deal with a comparison to a quantity that has a non-zero expectation value, which allows additive biases to play a role. As such, magnification requires much more care in analysis in order to determine an unbiased estimator.

A relatively simple calculation of signal-to-noise which assumes comparable shear and convergence (as done by e.g. Schneider, King & Erben 2000) can show that shear is generally expected to provide much stronger signal-to-noise than magnification at low redshift. However, the benefit from magnification is not entirely negligible, and no extra observing time is needed to gain the information required for a magnification measurement; any observation which can measure shear can also measure magnification. As such, if an unbiased measurement of it can be made, it will at least provide a slight improvement to weak lensing mass measurements. The benefits of magnification are more significant at higher redshifts, due to the fact that the number of galaxies sufficiently resolved such that their shapes can be measured drops off faster with redshift than the total number of observed galaxies. For lens galaxies at redshifts much greater than ~ 1, the shear signal is negligible, but the magnification signal is still sufficient to gain some information on the

masses of these galaxies (Van Waerbeke et al. 2010; Hildebrandt et al. 2013). Magnification can also provide additional information if the convergence and shear in a region of space are not of the same magnitude. For instance, as shear is only sensitive to changes in overdensity, a region of constant projected mass density could have positive convergence but no shear, allowing magnification to be used to break the 'mass-sheet degeneracy' that prevents shear alone from determining this overdensity.

It is thus well-motivated to further investigate the uses of weak lensing magnification in measuring the mass distributions of dark matter haloes. This has been shown to be possible in recent works (e.g. Scranton et al. 2005; Hildebrandt et al. 2009, 2013; Ménard et al. 2010; Ford et al. 2012, 2014), although the results so far are limited, particularly compared to the wealth of results which have been obtained from shear data. These recent measurements of magnification have frequently used an optimally weighted correlation function estimator (Ménard & Bartelmann 2002; Scranton et al. 2005), which compares the projected number density of source galaxies in the vicinity of lens galaxies to a random catalogue, properly weighted by a function of the luminosities of the source galaxies to optimize the magnification signal.[1] The measured correlation function is then compared to the correlation function predicted from models. While this methodology is sufficient for purposes such as galaxy–galaxy lensing, it cannot easily be applied to other lensing applications such as mass-mapping. It also has the drawbacks that it requires significant computational time to calculate errors through a jackknife or bootstrap approach, and it requires a clear separation between lens and source samples to avoid contamination from clustering.

In this paper, we start from the assumption that we desire an optimally weighted estimator for the magnification itself, as opposed to previous work which has used an optimized correlation function estimator, and we present a derivation of such an estimator. Using this estimator, we present a new method for measuring the magnification signal which avoids some of the drawbacks of the optimally weighted correlation function, and we show how their resulting estimators are related. We present a proof-of-principle of this new method, and we look at the results of this test to analyse the possible benefits magnification data may provide in supplementing shear data.

In Section 2 of this paper, we present a derivation of an estimator for magnification and a derivation of errors for it, and we discuss how it can be implemented in practice for galaxy–galaxy lensing. In Section 3, we discuss the data we use for testing our method, which we take from the Canada–France–Hawaii Telescope Lensing Survey. In Section 4, we demonstrate that our method can indeed detect magnification signals, and we look at the benefits of supplementing shear data with magnification data for galaxy–galaxy lensing. In Section 5, we discuss various aspects of the work in this paper, including a comparison of our method and the optimally weighted correlation function method, and we present a proposed implementation of our method for mass-mapping. In Section 6 we summarize and present our conclusions.

[1] There is also notable work using a conceptually different approach to magnification, relying on the observed properties of galaxies to predict their intrinsic luminosities (e.g. Huff & Graves 2014). It remains for future work to investigate if it is possible to combine the results of this approach with the type of approach presented in this paper.

In order to facilitate future work with magnification, we make the code used for this paper publicly available. Details of its access can be found in Appendix C.

For relevant calculations in this paper, we assume a cosmology with $H_0 = 70 \, \text{km s}^{-1} \, \text{Mpc}^{-1}$, $\Omega_m = 0.288$, $\Omega_r = 8.6 \times 10^{-5}$, and $\Omega_k = 0$. Unless stated otherwise, all mass values are in units of solar masses and all distance values are in units of kpc.

2 MAGNIFICATION MEASUREMENT METHODOLOGY

In this section we present a derivation of a generalized optimally weighted estimator for magnification and our proposed method for measuring it. In Section 2.1, we present the derivation of this estimator and show how its errors can be determined. In Section 2.2, we present our method for measuring magnification for the case of galaxy–galaxy lensing. Finally, in Section 2.3, we discuss the key step of determining the background, unmagnified source density.

2.1 Determining an optimal local estimator for magnification

Let us start with the assumption that we wish to measure the mean magnification μ in a patch of the sky where μ is expected to be approximately constant, such as in a circular annulus around a lens galaxy. We first bin the galaxies observed in this patch at redshifts greater than that of the lens by magnitude, using small bins of width dm, which we will later decrease to the limit of zero. Let

$$dn_i = n(m_i) \, dm \tag{1}$$

be the number of galaxies observed in the ith magnitude bin in this patch of sky, and let

$$dn_{0,i} = n_0(m_i) \, dm \tag{2}$$

be the expected number of galaxies in this magnitude bin, in a patch of sky of the same area, and in the same redshift range in the presence of no magnification. If we assume that the number count can be locally approximated by a power law with slope

$$\alpha_i = 2.5 \frac{d}{dm_i} \log n_0(m_i), \tag{3}$$

then we expect the observed number count of galaxies in this magnitude bin to be

$$\langle dn_i \rangle = dn_{0,i} \mu^{\alpha_i - 1} \tag{4}$$

(Narayan 1989; Broadhurst, Taylor & Peacock 1995).

In the weak lensing limit, we expect $(\mu - 1) \ll 1$, and so we can use the approximation

$$\langle dn_i \rangle = dn_{0,i} (1 + (\mu - 1))^{\alpha_i - 1}$$
$$\approx dn_{0,i}(1 + (\alpha_i - 1)(\mu - 1)). \tag{5}$$

Solving this equation for μ, we get

$$\mu = 1 + \frac{\langle dn_i \rangle - dn_{0,i}}{(\alpha_i - 1)dn_{0,i}}. \tag{6}$$

If we assume the statistical error on estimating $dn_{0,i}$ and α_i to be significantly less than the standard error of dn_i, we can then use

$$\mu_i = 1 + \frac{dn_i - dn_{0,i}}{(\alpha_i - 1)dn_{0,i}} \tag{7}$$

as an estimator for μ, with $\langle \mu_i \rangle = \mu$, and which is independent for each magnitude bin i.

In order to optimally combine the estimators μ_i for all magnitude bins, we apply inverse-variance weighting. We start with the standard deviation of each estimator, which, under the assumption that dn_i is Poisson, can be calculated to be

$$\sigma_{\mu_i} \approx \frac{\sqrt{1 + (\alpha_i - 1)(\mu - 1)}}{|\alpha_i - 1|\sqrt{dn_{0,i}}}. \tag{8}$$

The derivation of this equation is presented in equation (A1) in Appendix A. This then gives us the optimal inverse-variance weighting for bin i:

$$dw_{i,\mathrm{opt}} = \frac{1}{\sigma_{\mu_i}^2} = \frac{dn_{0,i}\,(\alpha_i - 1)^2}{1 + (\alpha_i - 1)(\mu - 1)}. \tag{9}$$

While the weighting scheme in equation (9) will provide the estimate of μ with the minimum standard error, it requires foreknowledge of the actual value of μ to determine the weights. For a likelihood-based approach, we can use the value of μ predicted by the model for the weight calculation, but this will require a separate summation over bins to calculate the likelihood for each model μ. Since we expect $(\mu - 1) \ll 1$, we instead use the weighting scheme:

$$dw_i = dn_{0,i}\,(\alpha_i - 1)^2, \tag{10}$$

which is model-independent. This will save significant computational time at the expense of a negligible decrease in the precision (but not accuracy) of our constraints.

We can then calculate our estimate of μ for this patch of sky through

$$\hat{\mu} = \frac{\sum_i dw_i \mu_i}{\sum_i dw_i}. \tag{11}$$

This will have standard error

$$\sigma_{\hat{\mu}} \approx \left(\sum_i dw_i\right)^{-1/2}, \tag{12}$$

the accuracy of which depends on the accuracy of the standard deviation we calculated for each μ_i in equation (8).

Let us now examine the case where the bin size dm approaches zero. We now express the parameters dn_i, $dn_{0,i}$, α_i, μ_i, and dw_i as $dn(m)$, $dn_0(m)$, $\alpha(m)$, $\mu(m)$, and $dw(m)$, respectively. Here, m is the magnitude of the differential bin, and we define $dn(m) = n(m)\,dm$, $dn_0(m) = n_0(m)\,dm$, and $dw(m) = w(m)dm$. $n_0(m)$ and $\alpha(m)$ have functional forms which can be constrained by observations, also giving us a functional form for $w(m)$. For a set of observed galaxies with magnitudes $m_1, m_2, \ldots, m_j$, we can express the differential number count $n(m)$ in terms of a sum of Dirac delta functions:

$$n(m) = \sum_j \delta(m - m_j). \tag{13}$$

This form gives the expected behaviour that integrating $n(m)$ from m to $m + dm$ results in the number of galaxies with magnitudes in this range.

Let a represent the minimum galaxy magnitude used in the survey, and let b represent the maximum observable magnitude for galaxies. We can then calculate our estimate of μ through

$$\hat{\mu} = \frac{1}{W}\int_a^b w(m)\mu(m)\,dm$$
$$= 1 + \frac{1}{W}\int_a^b [n(m) - n_0(m)]\,[\alpha(m) - 1]\,dm, \tag{14}$$

where

$$W = \int_a^b w(m)\,dm = \int_a^b n_0(m)\,(\alpha(m) - 1)^2\,dm. \tag{15}$$

Expanding $n(m)$ using equation (13), this gives us

$$\hat{\mu} = 1 + \frac{1}{W}\left(\sum_j \left[\alpha(m_j) - 1\right] - \int_a^b n_0(m)\,[\alpha(m) - 1]\,dm\right), \tag{16}$$

where m_j represent the magnitudes of observed galaxies in the range $a < m_j < b$.

As before, this estimator will have the standard error

$$\sigma_{\hat{\mu}} \approx W^{-1/2}, \tag{17}$$

and will be approximately Gaussian for large n. The accuracy of this depends on the validity of our assumptions on the distribution of number counts used in equation (8). In practice, an empirical error estimate (determined, for instance, from the standard deviation in the $\hat{\mu}$ for various similar regions) is likely to be more reliable.

From this, we can estimate the likelihood for a given value of μ to be

$$\mathcal{L}(\mu|\hat{\mu}) = \frac{1}{\sqrt{2\pi}\sigma_{\hat{\mu}}}\exp\left(-\frac{\mu - \hat{\mu}}{2\sigma_{\hat{\mu}}^2}\right)$$
$$\approx \sqrt{\frac{W}{2\pi}}\exp\left(-\frac{W(\mu - \hat{\mu})}{2}\right). \tag{18}$$

The weighting used for this estimator is in fact the same as that generally used for the optimally weighted correlation function, and was originally presented by Ménard & Bartelmann (2002) without derivation. We will later show in Section 5.1 that this estimator is simply related to that used for correlation-function-based methods in the scenario in which the lens and source samples are equivalent. In the following section, we will show how the estimator presented here can be used as a basis for an alternative method for calculating a magnification signal by directly comparing the observed source counts to the measured background density at a given redshift.

2.2 Implementation for galaxy–galaxy lensing

We start with the case of a galaxy–galaxy lensing analysis, which will allow us to readily test our implementation against the lensing signal measured with shear information. We bin our lens sample by both stellar mass and redshift so that the lenses in each bin will have approximately the same halo mass profiles.

For each lens, we bin the possible patch of sky in which sources could potentially appear based on projected distance at the redshift of the lens. This gives us a series of annuli for each lens, some of which may overlap masked regions of the image. We determine the unmasked fraction of each annulus through the ratio of unmasked pixel centres to all pixel centres inside the annulus. In practice, most masks are relatively coarse, and so this fraction can be reasonably estimated by regularly sampling pixels in order to conserve computational time.

The angular area of the kth annulus around the lth lens is then

$$\Omega_{kl} = D_l^{-2} f_{kl}\pi(R_{k+1}^2 - R_k^2), \tag{19}$$

where D_l is the angular diameter distance at the redshift of the lens, f_{kl} is the unmasked fraction of this annulus, and R_k and R_{k+1} are its inner and outer radii, respectively, in physical units.

To ensure minimal contamination of the source sample with galaxies which are actually in the lens plane, we require any sources used for our analyses to be separated from their associated lenses by a redshift buffer Δz, which we set to 0.2. As galaxies in the lens plane are clustered with each other, any which are misidentified as sources will be clustered with lenses as well, resulting in an artificial increase in the observed counts of sources around lenses, biasing the magnification signal. It is therefore necessary that this buffer be sufficiently large that this contamination has negligible effect on the measured magnification signal. See Section 3.1 for discussion of our choice of Δz.

Note that this separation between lenses and sources is applied to each lens individually. This allows for the possibility that the lens sample might overlap with the source sample. Previous methods for measuring magnification have relied on the two samples being well-separated. The loosening of this requirement for our method will allow for an increased signal-to-noise ratio in situations in which it is beneficial to use overlapping lens and source samples.

The expected unmagnified number count of sufficiently distant sources in a given magnitude bin projected in this annulus is

$$n_{0,kl}(m)\,\mathrm{d}m = \Omega_{kl}n(m, z_l + \Delta z)\,\mathrm{d}m, \tag{20}$$

where z_l is the redshift of the lens and $n_0(m, z_l + \Delta z)\mathrm{d}m$ is the number of sources per unit area with magnitude between m and $m + \mathrm{d}m$ at redshift $z_l + \Delta z$ or greater. This gives us the relevant $n_0(m)$ to use in equation (16), and $\alpha(m)$ can be determined through equation (3), allowing us to get the estimate $\hat{\mu}_{kl}$ for this annulus.

This process can be performed more efficiently through storing the integrals

$$\mathrm{d}(z) = \int_a^b n(m, z)\left[\alpha(m) - 1\right]\mathrm{d}m \tag{21}$$

and

$$w(z) = \int_a^b n(m, z)\left[\alpha(m) - 1\right]^2\mathrm{d}m \tag{22}$$

for all needed redshift values, and calculating $\hat{\mu}_{kl}$ through

$$\hat{\mu}_{kl} = 1 + \frac{1}{\Omega_{kl}w(z_l + \Delta z)}\left(\sum_j\left[\alpha(m_{jkl}) - 1\right] - \Omega_{kl}\mathrm{d}(z_l + \Delta z)\right), \tag{23}$$

where m_{jkl} represent the magnitudes of all sufficiently distant sources observed in this annulus, and a and b represent the lower and upper magnitude limits of the source sample, respectively.

By comparing equation (23) to equations (14) and (17), we can estimate the standard deviation of $\hat{\mu}_{kl}$ by

$$\sigma_{\hat{\mu}_{kl}} = \left[\Omega_{kl}w(z_l + \Delta z)\right]^{-1/2}. \tag{24}$$

When we stack lenses, we can use this estimate to apply inverse-variance weighting, giving each individual annulus a weight of

$$w_{kl} = \Omega_{kl}w(z_l + \Delta z). \tag{25}$$

Finally, we get the estimated magnification for the stacked lenses in this annulus from the weighted mean across all lenses:

$$\hat{\mu}_k = \frac{\sum_l w_{kl}\hat{\mu}_{kl}}{\sum_l w_{kl}}. \tag{26}$$

If our estimates of standard deviations in number counts were perfectly accurate, the standard error in this estimator would be

$$\sigma_{\hat{\mu}_k,\mathrm{ideal}} = \left(\sum_l w_{kl}\right)^{-1/2}, \tag{27}$$

but as this is unlikely to be the case (for instance, due to the clustering of sources making their distribution not perfectly Poisson), we instead determine it empirically. We start by empirically determining the weighted standard deviation in the $\hat{\mu}_{kl}$ estimates for the set of lenses,

$$\sigma_{\hat{\mu}_k} = C_k\sqrt{\frac{\sum_l w_{kl}\hat{\mu}_{kl}^2}{\sum_l w_{kl}} - \left(\frac{\sum_l w_{kl}\hat{\mu}_{kl}}{\sum_l w_{kl}}\right)^2}, \tag{28}$$

where

$$C_k = \frac{\sum_l w_{kl}}{\sqrt{\left(\sum_l w_{kl}\right)^2 - \sum_l w_{kl}^2}} \tag{29}$$

(ThePawn 2013).[2] We then estimate the standard error in the magnification estimate $\hat{\mu}_k$ for this set of annuli through

$$\sigma_{\hat{\mu},k} = \frac{\sigma_{\hat{\mu}_k}}{\sqrt{N_k - 1}}, \tag{30}$$

where N_k is the number of lenses with non-zero weight for this annulus.

This estimate of the magnification can then be converted into an estimate of the convergence:

$$\hat{\kappa}_k = 1 - \sqrt{\hat{\mu}_k^{-1} + \hat{\gamma}_{\mathrm{t},k}^2}, \tag{31}$$

where $\hat{\gamma}_{\mathrm{t},k}$ is the tangential-shear estimate for this annulus.[3] The standard error in this estimate is

$$\sigma_{\hat{\kappa},k} = \sqrt{\frac{\sigma_{\hat{\mu}}^2 + 2\hat{\gamma}_{\mathrm{t},k}\sigma_{\hat{\gamma}_{\mathrm{t},k}}^2}{4\left(\hat{\mu}_k^{-1} + \hat{\gamma}_{\mathrm{t},k}^2\right)}}. \tag{32}$$

The estimated projected overdensity in this annulus is then

$$\hat{\Sigma}_k = \Sigma_{\mathrm{crit},k}\hat{\kappa}_k, \tag{33}$$

where $\Sigma_{\mathrm{crit},k}$ is the critical density for this annulus:

$$\Sigma_{\mathrm{crit},k} = \frac{c^2}{4\pi G_{\mathrm{c}}}\frac{D_{\mathrm{s}}}{D_{\mathrm{l}}D_{\mathrm{ls}}}, \tag{34}$$

where D_{s} is the angular diameter distance from redshift zero to the mean source redshift, D_{l} from redshift zero to the mean lens redshift, and D_{ls} from the mean lens redshift to the mean source redshift. The standard error in the projected overdensity is then

$$\sigma_{\hat{\Sigma},k} = \Sigma_{\mathrm{crit},k}\sigma_{\hat{\kappa},k}. \tag{35}$$

---

[2] This formula for the standard error of a weighted mean was chosen based on a comparison of various estimators, as it was found to generally have the lowest amount of bias in various test cases, some of which involved biased weights being used.

[3] Equation (31) uses the assumption that the cross-shear $\gamma_{\mathrm{x},k}$ is expected to be zero, as is the case in galaxy–galaxy lensing. For applications where this is not the case, $\gamma_k = \sqrt{\gamma_{\mathrm{t},k}^2 + \gamma_{\mathrm{x},k}^2}$ should be used in place of $\gamma_{\mathrm{t},k}$ in this equation and equation (32).

2.3 Determining unmagnified source density

In measuring the magnification signal of a given region of space, we require knowledge of what the source density would be as a function of magnitude in the absence of any magnification. In equations (14) and (16), this represents $n_0(m)$, which is used to determine $\alpha(m)$ through equation (3) and W through equation (15). Given the large area covered by the Canada–France–Hawaii Telescope Lensing Survey, it is a relatively safe assumption that the survey region is on average unmagnified. There will be some overdense and underdense regions, but the region as a whole should have a projected density approximating that of the whole sky, and so we can thus use the measured source densities across the entire field to determine $n_0(m)$.

This approach has two key advantages over a theoretical model. First, it accounts for many observational effects which will identically affect the source samples we later use for measuring magnification, such as the difficulty in identifying faint galaxies. Secondly, as we are interested in all galaxies which appear behind a given annulus, we are in fact interested in the number density of all galaxies at a given redshift or greater, whereas models are typically fit only to galaxies within a single redshift slice. This resulted in us being unable to satisfactorily fit a model such as a Schechter function to the observed galaxy counts beyond a given redshift, even allowing for perturbations to the functional form.

This approach does result in noise being present in $n_0(m)$. As it is measured over a much greater area of the sky than even the stacked regions we wish to measure magnification for, however, this noise is likely to be sub-dominant to the Poisson noise in $n(m)$, the number count distribution of galaxies which actually appear in these regions. However, we must differentiate $n_0(m)$ to determine $\alpha(m)$, and differentiation amplifies this noise. In order to handle this, we first smooth the measured $n_0(m)$ using the Savitzky & Golay (1964) algorithm prior to differentiation.

As can be seen most easily in equation (14), the estimator for magnification is particularly sensitive to n_0. For instance, in the simplified scenario where $\alpha(m) = 2$, a bias of 0.1 per cent in n_0 leads directly to a bias of -0.001 in $\hat{\mu}$. As the typical magnification signal is of this order, we must take care to minimize bias in n_0 as much as possible.

If the positions of lenses and sources were completely independent, the approach we use would be expected to be unbiased when applied to galaxy–galaxy lensing. However, even aside from contamination of the source sample with galaxies in the lens plane misidentified as sources, some correlation or anticorrelation between the samples is expected. This can be caused by a number of position-dependent effects which can introduce biases into the redshift estimates for galaxies, including atmospheric seeing, depth variations due to the dithering strategy of the survey, and depth variations due to galactic dust.

The method used for estimating photometric redshifts in the Canada–France–Hawaii Telescope Lensing Survey by Hildebrandt et al. (2012) used a local model for atmospheric seeing to minimize the bias in galaxy colours due to it, so this is unlikely to be an issue. However, regions with poor seeing and regions with poor depth will have larger errors on the estimated colours of galaxies. If there is a gradient in the object density in colour space (e.g. more blue galaxies than red), then noise in colour determinations will result in asymmetric scatter along this gradient. Even if any particular galaxy is as likely to be scattered to a higher colour as a lower colour, the net result will be to smooth out the object density distribution in each axis of colour space. This will then result in the
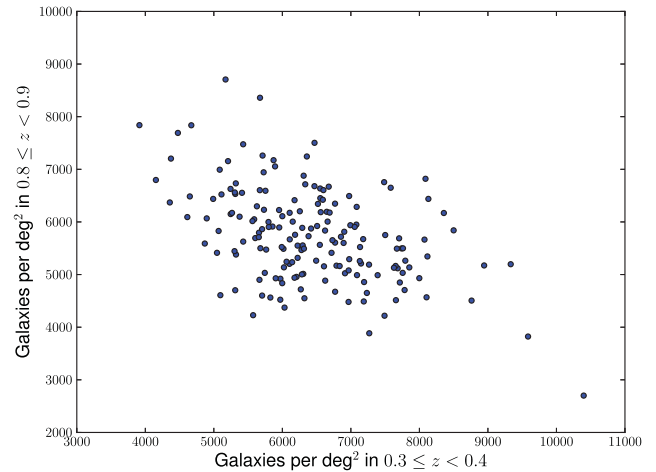

Figure 1. A comparison of the number densities of galaxies in two different photometric redshift slices in the Canada–France–Hawaii Telescope Lensing Survey. Each point represents an individual field. The anticorrelation between the two redshift slices implies the presence of a systematic error in redshift estimates which varies from field to field.

estimated redshifts of these objects being systematically biased. As the relation between redshift and colour is not linear, the direction and magnitude of this bias will vary with redshift. This means that some redshift bins may be spuriously underdense, while others are spuriously overdense.

Let us present an example of what might occur, and how this might lead to a systematic bias in the magnification signal. Consider an extreme example: a field is divided into two regions of equal size; one region has good seeing, and the other region has poor seeing. This poor seeing results in the measured luminosities of galaxies in this region having greater errors than the luminosities of galaxies in the other region. If we suppose that there are more blue galaxies than red and that, for the sake of this example, the distribution between the two extremes is smooth, then we would expect more galaxies to be scattered to redder colours than we would expect to see galaxies scattered to bluer colours. Thus, there would be a net bias in this region for all galaxies to appear redder. Let us assume a simple photometric redshift estimation which relies only on this colour. In this case, this would result in a systematic bias in this region for all galaxies to have higher redshifts. As there is more volume at higher redshifts, there are more galaxies there. Thus, a bias of all galaxies to higher redshifts would result in a lower density of galaxies in any given redshift bin.

We thus have one region with good seeing and approximately normal density of lenses and sources, and one with poor seeing and a lower density of both lenses and sources. In measuring n_0, we would weight both of these regions equally and calculate an expected source density which is the mean of the two. However, in measuring the magnification signal, we would be using the lower number of lenses in the region which also has a lower number of sources. Our measured n would thus be more akin to a weighted mean, which in this case would end up being weighted more by the region with a lower density of sources. Thus, even without any magnification, we have a scenario in which the measured count of galaxies differs from the expected count.

Evidence of this type of bias can be seen in Fig. 1, which shows the galaxy number densities of different fields from the Canada–France–Hawaii Telescope Lensing Survey in the redshift slices 0.3–0.4 and 0.8–0.9 plotted against each other. This plot shows
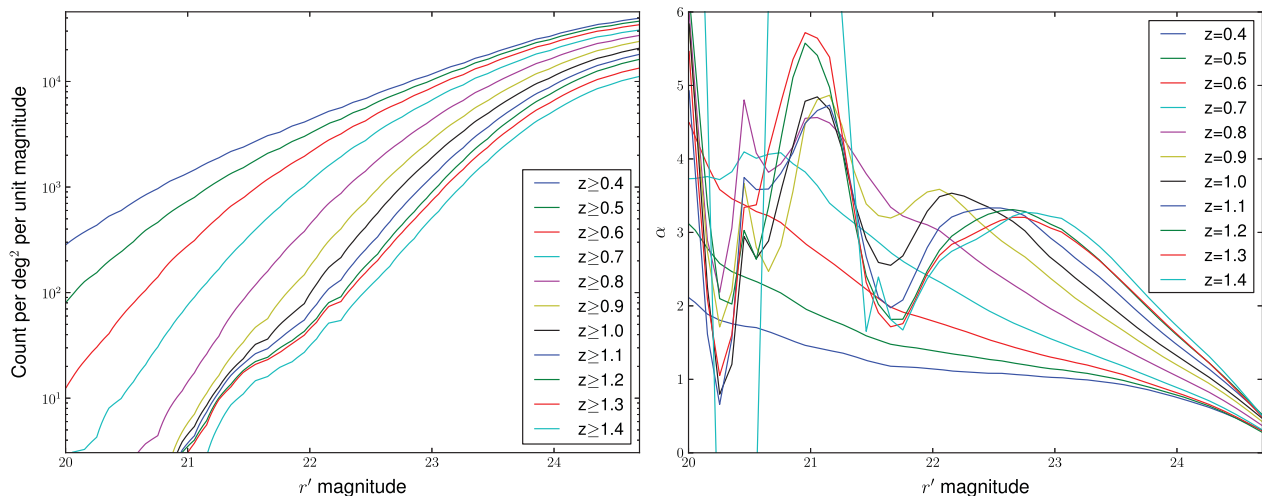
Figure 2. Number density of galaxies at a given redshift or greater as a function of r'-band magnitude (left) for the CFHTLenS, and the corresponding calculated values of α (right). The strong oscillations observed in α for high redshifts and low magnitudes are caused by Poisson noise, due to the low numbers of galaxies in these bins. This results in these bins being correspondingly downweighted, and so these oscillations are likely to have no appreciable effects on our results. In the left-hand panel, the lowest redshifts correspond to the highest lines, and in the right-hand panel, the lowest redshifts typically correspond to the lowest lines.

clearly that fields which have greater densities in the 0.3–0.4 slice have lower densities in the 0.8–0.9 slice, and vice-versa. There are also many cases of well-separated slices which have correlated number densities, but the anticorrelation here is most striking, as it is difficult to explain with any effect other than non-trivial redshift systematic errors such as we have discussed. More insight into this can be gained from the rightmost panel of fig. 4 of Hildebrandt et al. (2012), which shows a comparison of the photometric redshifts as used here against the spectroscopic redshifts of a matched sample of galaxies, and observing the complicated shape of the distribution. The field-by-field comparison we show here demonstrates that this distribution is not constant between fields.

A full discussion of this type of effect is beyond the scope of this paper. Interested readers are directed to Morrison & Hildebrandt (2015), which discusses this issue in detail.

In the example above of how this effect might lead to an error in the measured magnification signal, we pointed out that the result is similar to the difference between the unweighted and weighted means of source densities. We can thus correct for it by weighting our initial calculation of mean source density by the density of lenses in the same region of space. The more local this weight is, the noisier it will be, and the more it will become entangled with the actual magnification signal, but the better it will account for this effect.

For this paper, we have decided to weight on a field-by-field basis, as the discontinuity in both position and time of observation between fields is likely to result in much more variation in seeing and depth than the variation across the field of view within a single field. To do this, when we are counting all galaxies in an individual field at a redshift z_s or greater, we weight this by the number of lenses at redshift[4] $z_s - \Delta z$. See Section 5.4 for discussion of the impact of this weighting on our results.

3 DATA

In order to test our methodology, we use the publicly available galaxy catalogues from the Canada–France–Hawaii Telescope Lensing Survey (Heymans et al. 2012), hereafter referred to as 'CFHTLenS', which are available for download from http://www.cadc-ccda.hia-iha.nrc-cnrc.gc.ca/community/CFHTLens/query.html. The CFHTLenS is a 154 deg² survey (125 deg² after masking; Erben et al. 2013), based on the Wide component of the Canada–France–Hawaii Telescope Legacy Survey (CFHTLS), which was observed in the period from 2003 March 22 to 2008 November 1, using the MegaCam instrument (Boulade et al. 2003). It consists of deep, sub-arcsecond, optical data in the five optical Sloan Digital Sky Survey (SDSS)-like filters $u^*g'r'i'z'$. The data is divided into 171 individual patches of sky, which we refer to as 'fields' in this paper.

In order to avoid contamination from foreground stars and galaxies, cosmic rays, and other observational defects, masks were applied to the CFHTLenS fields using the automask tool (Dietrich et al. 2007; Erben et al. 2009). Shapes of background galaxies with $i' < 24.7$ in the unmasked regions were measured with the *lens*fit shape measurement algorithm (Miller et al. 2013), giving an effective density for sources with shapes of 11 arcmin⁻² in the redshift range $0.2 < z_{phot} < 1.3$ (Heymans et al. 2012). For the analysis in this paper, we use all fields in the survey, not simply those that passed the systematics tests for cosmic shear measurements (Heymans et al. 2012), as it has been demonstrated that fields with systematics that may affect cosmic shear have no effect on galaxy–galaxy lensing measurements (Velander et al. 2014).

Photometric redshifts for the survey were estimated with the BPZ code (Benítez 2000). The accuracy of photo-z measurements for the CFHTLS-Wide survey was improved by homogenizing the PSF through different bands in the CFHTLenS survey. Photo-zs were made available for the entire survey, with a typical redshift uncertainty of ~0.04(1 + z) in the redshift range $0.2 < z_{phot} < 1.3$ (Hildebrandt et al. 2012).

We use the stellar mass estimates described by Velander et al. (2014), obtained by fitting spectral-energy-distribution templates, following the method of Ilbert et al. (2010). These stellar masses

---

[4] The photometric redshifts we use are binned to discrete intervals of 0.01, hence the precise value being used here.

were found to be in general agreement with deeper data such as WIRDS, which includes NIR filters (Bielby et al. 2012), up to $z = 0.8$.

We use the magnitudes measured with the r' filter for all galaxies. We considered using the i'-band magnitudes, but this was found to be problematic, as the number density of galaxies as a function of i'-band magnitude in the CFHTLenS shows a prominent step at $m_{i'} = 23$, which would result in $\alpha(m)$, the logarithmic derivative of this function, being discontinuous at this value. This is possibly an indirect result of the fact that the original i' filter used for the CFHTLenS had to be replaced partway through observations. The r'-band magnitudes showed no such issue, and typically had the next best signal-to-noise after the i'-band magnitudes, and so we decided to use them. From here on, whenever we refer to magnitudes, we will thus be referring specifically to the r'-band magnitudes.

3.1 Sample selection

We select galaxies from the public catalogue using the following cuts on it:

   (i) Z_B ≤ 4
   (ii) MAG_R ≤ 24.7
   (iii) star_flag ≤ 0.01
   (iv) CHI_SQUARED_BPZ ≤ 2.

This provides us with a sample of galaxies with reasonable-quality redshifts. Other cuts were tested to improve redshift quality, such as imposing a cut on the 'ODDS' parameter, but they were found to cause too much bias in the resulting galaxy selection. We use Z_B as the photometric redshift estimate for each galaxy. The resulting number densities of galaxies at or beyond various redshift slices, as well as the calculated values of the logarithmic slope parameter α, can be seen in Fig. 2.

We do not impose a cut on the 'MASK' parameter. Rather, we download the mask files and apply them ourselves, removing any galaxy from the sample which resides in a region with MASK >1. We do this as we found that the catalogue's MASK values for a small number of galaxies did not agree with our determinations, possibly due to a different version of the mask files having been used to assign the values in the catalogue. Assigning the MASK values ourselves thus ensures consistency with the mask files we have available, which we also use for determining the unmasked fractions of annuli around lenses.

We use the entirety of these galaxies for our source sample, and all galaxies with Z_B ≤1.3 for our lens sample. In selecting lens–source pairs, we require that the source be in the background of the lens, separated by a minimum redshift of $\Delta z = 0.2$. This value is chosen to minimize the fraction of galaxies identified as appropriate sources which in fact reside in the same bin as the lenses. To determine this fraction, we matched the CFHTLenS galaxy catalogue to publicly available spectroscopic galaxy catalogues from the VIMOS VLT Deep Survey ('VVDS'; Le Fèvre et al. 2005; Garilli et al. 2008) and the DEEP2 galaxy survey (Davis et al. 2007; Newman et al. 2013), both of which overlap portions of the CFHTLenS. We then calculated the fractions of galaxies with sufficiently large photometric redshifts which have spectroscopic redshifts within a given bin. As shown in Fig. 3, a buffer of $\Delta z = 0.2$ ensures that fewer than 1 per cent of galaxies identified as appropriate sources in fact reside in the same bin as the lenses for lens redshift bins $z_{mid} \leq 0.65$, and the benefit of increasing this value further is indistinguishable from noise.
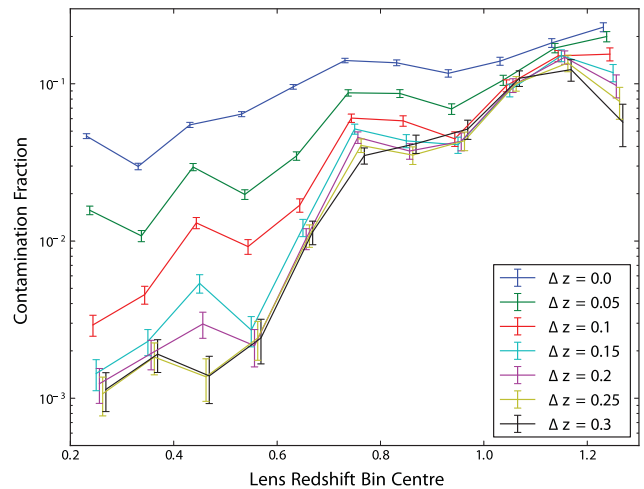
Figure 3. For a given lens redshift bin of width 0.1, the fraction of galaxies with photometric redshifts at least a given redshift Δz greater than the centre of the bin (corresponding to the source sample), which are likely to actually reside in the bin. We determine this fraction through matching galaxies in our sample with spectroscopic redshift catalogues from the VVDS and DEEP2, and determining the fraction of galaxies with sufficiently large photometric redshifts which have spectroscopic redshifts within a given bin. Larger Δz values consistently correspond to smaller contamination fractions, though this effect seems to converge to a maximum by $\Delta z \sim 0.2$. The contamination fraction values for each lens redshift bin are plotted near the centres of the bins, with a horizontal offset applied for visual clarity.

4 APPLICATION TO DATA

In order to provide a proof-of-principle that our proposed method can be successfully used to measure the magnification signal from weak gravitational lensing, we test it by measuring the galaxy–galaxy lensing signal in the CFHTLenS. We choose the galaxy–galaxy lensing application as it is computationally straightforward to measure the shear signal at the same time as the magnification signal. As weak lensing shear is a well-studied and robust method for determining halo mass profiles with signal-to-noise of the same order of magnitude as weak lensing magnification, it provides an ideal comparison.

One drawback in using shear to model halo mass profiles is that the shear signal is degenerate to the addition or subtraction of a sheet of constant projected mass density to the field of view. Magnification is not subject to this degeneracy, which leads to the result that a comparison with shear alone cannot be used to determine whether a flat offset in the magnification signal is due to a bias in the measurement or to the presence of an actual over- or underdensity in the region of interest. See the discussion on the potential for lens–source correlations or anticorrelations to cause such a bias in the magnification signal in Section 2.3.

In Section 4.1, we discuss the model we use for the expected form of the lensing signal, and in Section 4.2, we present the results of our tests.

4.1 Modelling lensing signal

To fit our measured lensing signals, we assume that the underlying mass distribution can be described by a profile which is on average radial, and which can be described with relatively few parameters. The most common model used for this purpose is known as the halo model (Seljak 2000; Mandelbaum et al. 2005), which has two

free parameters: the halo mass of the galaxies in the sample, and the satellite fraction, with other parameters directly fixed through a physical model. For this paper, however, we choose to use a different model, which has one additional parameter to characterize the surrounding environment of galaxies in the sample. This allows us to better investigate how magnification might aid in breaking degeneracies between separate factors which characterize the environment of galaxies.

Our model uses four free parameters: M_{1h}, the total one-halo mass of the lens; M_{gr}, the mean mass of the group or cluster a typical satellite lens resides in; f_{sat}, the fraction of lenses which are satellites of a group; and κ_{offset}, a flat offset to the measured κ due to observational biases. Other relevant parameters are fixed directly from observations or with fitting functions. We model the galaxies in each bin as having their mass dominated by a dark matter halo with a truncated Navarro, Frenk & White (1997, hereafter a 'tNFW' profile) profile, using the model of Baltz, Marshall & Oguri (2009). We leave the mass of this profile as a free parameter, determine the concentration through the fitting function of Neto et al. (2007), and set the truncation radius equal to double the virial radius.

Of these galaxies, we assume that a fraction f_{sat} are satellites of a group, where f_{sat} and M_{gr}, the total mass of the group, are free parameters. We assume the group's total mass can also be modelled with a tNFW profile, and we determine its concentration and truncation radius in the same manner as for the galaxy's halo. We assume that the distribution of satellites relative to group centre can be described by a scaled tNFW density profile with concentration $c = 2.5$, which is consistent with the results of McGee et al. (2009). The details of calculating the contribution of the group environment to the lensing signal can be found in Gillis et al. (2013), which uses the same model for shear except with a fixed f_{sat}. For computational simplicity, we do not include any contributions from large-scale structure to the lensing signal (the so-called two-halo term) at this stage of analysis, but we plan to do so in future work.

This model results in a mean mass profile for a given set of parameters M_{1h}, M_{gr}, f_{sat}. To determine the shear and magnification lensing signals, we project this mass profile along the line of sight to get the projected overdensity $\Sigma(\boldsymbol{R})$. The expected shear and convergence signals can then be obtained through

$$\langle \gamma_t \rangle = \frac{\overline{\Sigma(<R)} - \overline{\Sigma(R)}}{\Sigma_{crit}} \qquad (36)$$

and

$$\langle \kappa \rangle = \frac{\overline{\Sigma(R)}}{\Sigma_{crit}}, \qquad (37)$$

where $\overline{\Sigma(R)}$ is the mean projected density at a distance R from the galaxy, $\overline{\Sigma(<R)}$ is the mean projected density contained within a distance R of the galaxy, and Σ_{crit} is the critical projected density, using the same definition as equation (34).

Finally, we assume that the correlation and anticorrelation between the number counts of lenses and sources across the field of view due to observational effects will have the net effect of shifting the measured magnification signal by a constant value, which we leave here as a free parameter, κ_{offset}. While the other three free parameters affect both the measured shear and magnification signals, κ_{offset} only has an effect on the magnification signal.

We determine the best-fitting set of parameters and their associated errors for each bin of interest through a three-phase Markov chain Monte Carlo algorithm. In the first and second phases, we follow the Metropolis–Hastings algorithm (Metropolis et al. 1953), using a χ^2 comparison of the model lensing signal with the mea-

sured values to determine the acceptance ratio. The first phase is a 'burn-in' period, in which the algorithm is given time to reach the region with the lowest χ^2. In the second phase, we periodically record the points arrived at by the algorithm. The standard deviations in these values represent the 1σ errors in the best-fitting model parameters. In the final phase, we apply simulated annealing, successively decreasing the size of steps and the acceptance ratio, to arrive at the best-fitting set of parameters.

For each bin of interest, we fit the data in three manners: to magnification information only, to shear information only, and to both magnification and shear information.

4.2 Resulting measurements of magnification signal

In order to test that our model is performing as expected, we first perform a measurement of tangential shear in various lens stellar mass and redshift bins and test fitting our model to the resulting signal. The results of this test are shown in Fig. 4, where we present the shear signal in terms of the reduced tangential shear γ_t against angular separation for each lens bin as well as the predicted shear for the best-fitting model. In each bin, the model passes a χ^2 test in comparison with the data, as does the mean χ^2 of 17.6 for 15 degrees of freedom and 15 bins. Our model thus appears to provide a good fit to the shear data.

In Fig. 5, we show the measured magnification signal for the same lens stellar mass and redshift bins as in Fig. 4, presented in terms of the convergence κ, which is related to the measured magnification through equation (31). Also plotted on it are curves representing the predicted κ from the best-fitting model which uses only magnification data, the best-fitting model which uses only shear data (for which γ_t for each bin is plotted in Fig. 4), and the best-fitting model to the combination of the shear and magnification data sets. As the fit using shear data alone does not constrain κ_{offset}, the fitted value from using magnification alone is used for that model. The data generally shows the expected form and is well-fit by our model, though the errors are typically larger than for the shear measurements. Notice that in some bins the signal converges to a positive value at large angular separation between lenses and sources, and in others it converges to a negative value. This is most likely an effect of variable seeing and other position-dependent effects leading to correlation or anticorrelation between lenses and sources, as we discussed in Section 2.3, and which is the reasoning behind our use of the κ_{offset} model term.

We note that in many cases, the innermost radial bins show very high magnification signals compared to the models. We believe this to be most likely due to the obscuration of background galaxies by dust in the regions of lens galaxies (particularly those which reside within groups or clusters). Dust can most easily obscure the faintest magnitude galaxies, and it is at the faintest magnitudes that the slope of the luminosity function is shallowest (as can be seen in Fig. 2). At these magnitudes, $\alpha(m) - 1$ is negative, and so the obscuration of a galaxy would result in a spurious increase in the magnification signal. As we know that the innermost radial bins are likely to be problematic in this manner, we leave them out of our fitting procedure and calculation of χ^2 values. Methods have been proposed to handle dust extinction (e.g. Ménard et al. 2010, which finds that magnification is typically a stronger effect than dust extinction by a factor of ~ 3 in the V band, and Bauer et al. 2014), but as this paper is intended to serve as a first demonstration for the magnification methodology proposed here, we leave the full application of these for future study.
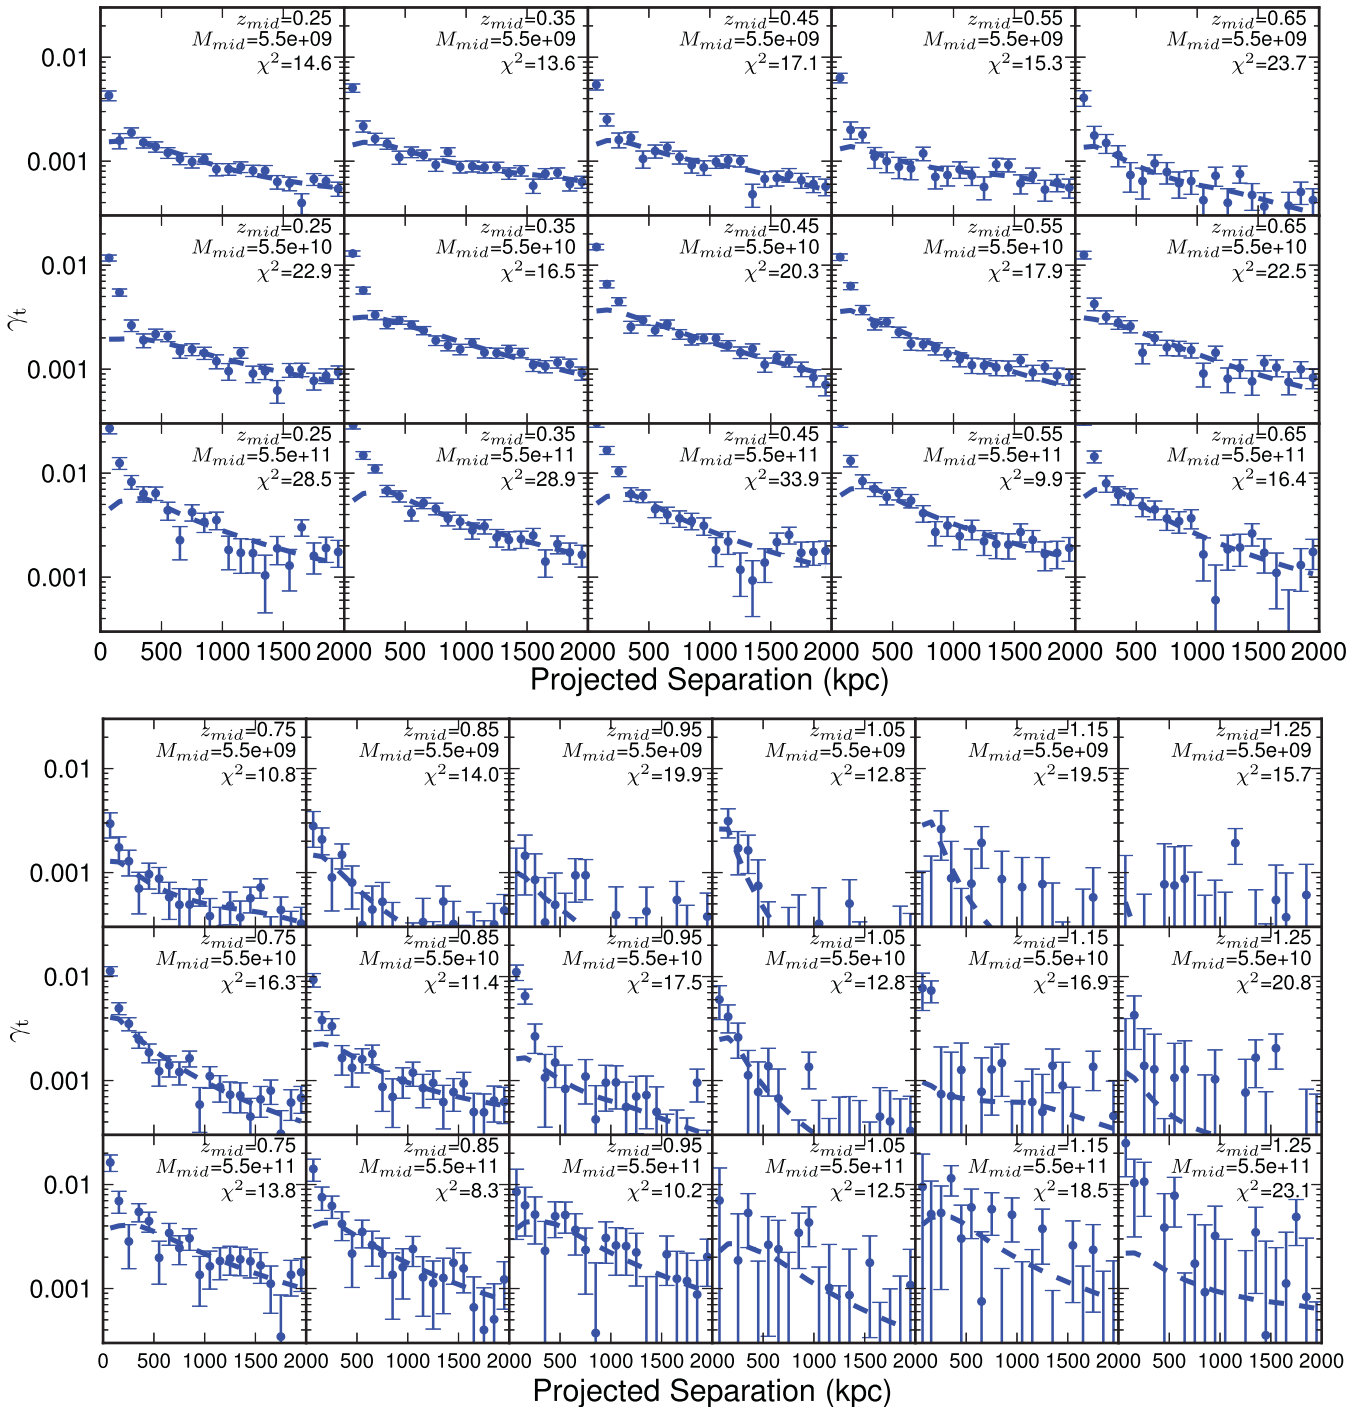
Figure 4. Measured shear signal γ (points and errorbars) for different lens redshift and mass bins. Also shown is the best-fitting model using shear data alone (dashed line). The top panel displays low-redshift bins, and the bottom panel displays high-redshift bins. The χ^2 values represent the χ^2 for a comparison of the shear data with the shear-only fit (15 degrees of freedom).

In the low-redshift bins, the model which is best fit to only shear data generally also does a good job at fitting the magnification data. It fails a χ^2 test at the 5 per cent threshold for only five of the 15 low-redshift bins, and the fit to the combined data set fails for only two of these bins: the $9 < \log M_{lens}/M_\odot < 10$ and $0.3 < z < 0.4$ bin, and the $11 < \log M_{lens}/M_\odot < 12$ and $0.6 < z < 0.7$ bin. This failure ratio of 13 per cent for the combined fit does imply that something beyond statistical error is to blame here, but this is not surprising given the simple model we use in this paper.

In the higher redshift bins, the best-fitting model to shear-only does a much poorer job at fitting the magnification signal. This is not surprising, as at these redshifts the signal-to-noise of the shear measurements is much lower, and so noise in the shear signal can propagate through into noise into the best-fitting measurements, resulting in a high χ^2 value when it is compared with the magnification data. The fit to the combined data set does a better job at fitting to both, though it still fails at the 5 per cent threshold for seven of the 18 high-redshift bins. It remains for future
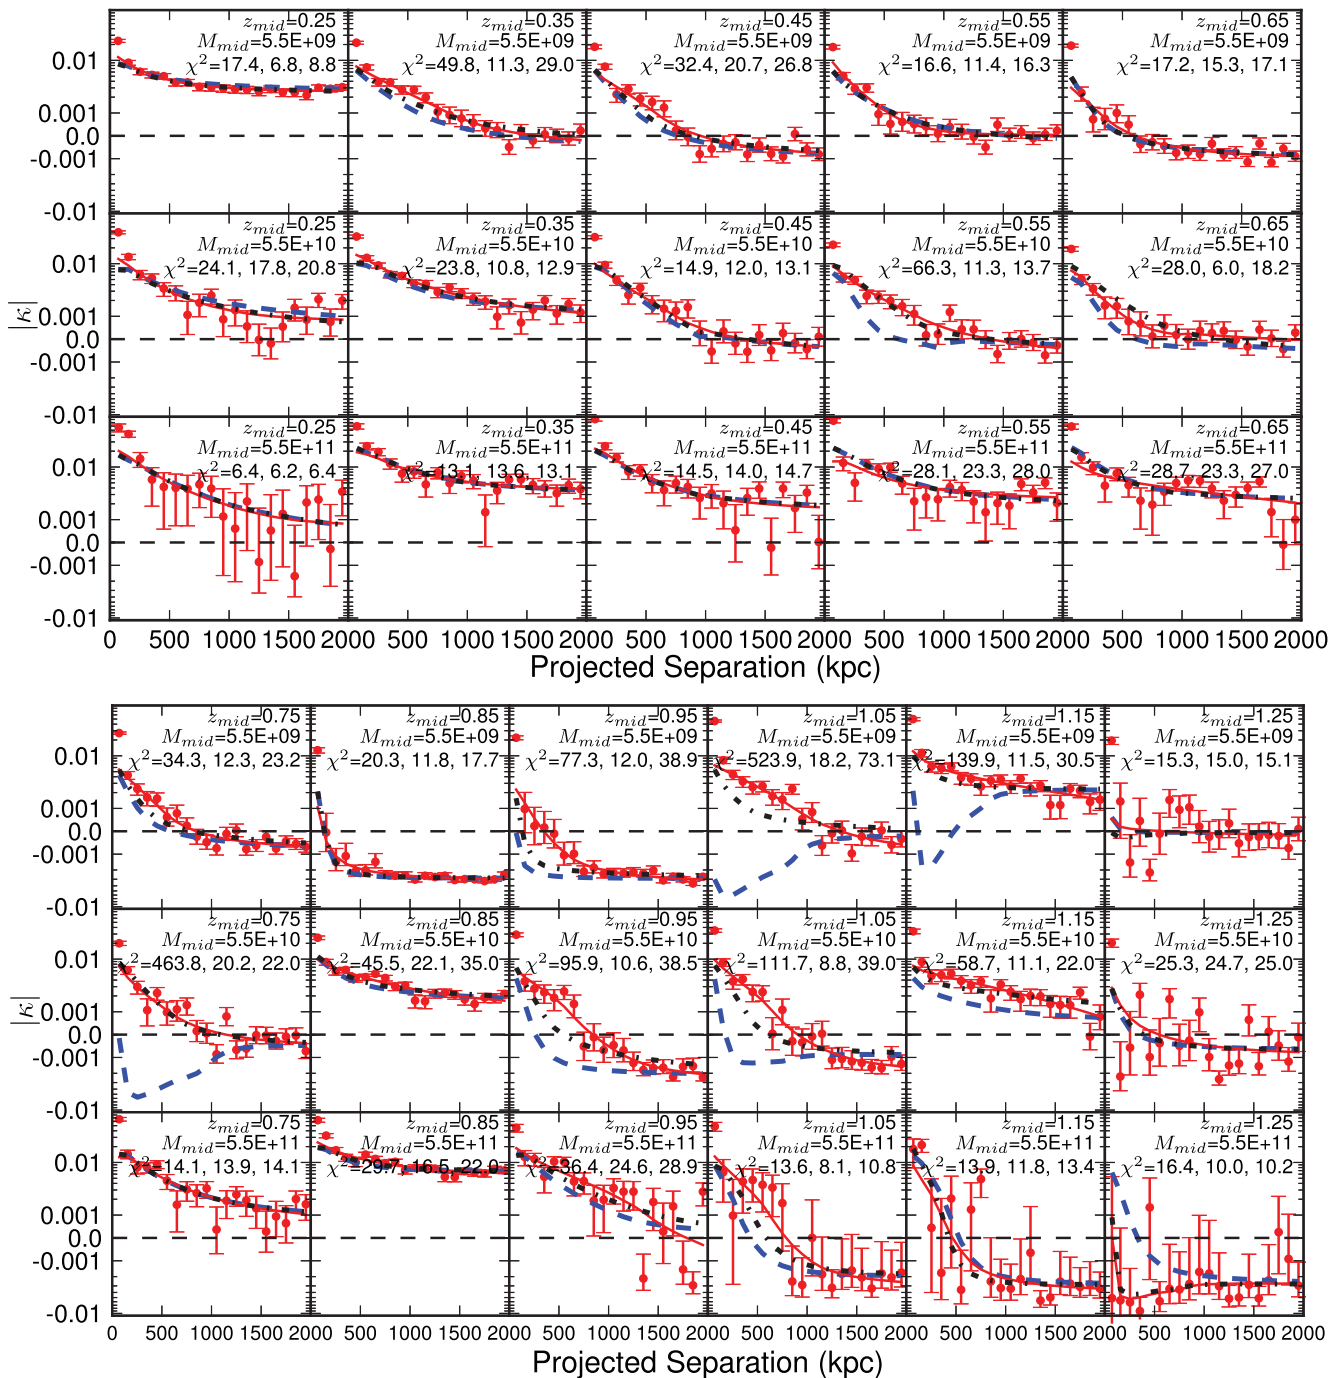
Figure 5. Measured magnification signal κ (points and errorbars) for different lens redshift and mass bins. Also shown are the best-fitting models using shear data alone (blue dashed line), using magnification data alone (red solid line), and for the combination of the shear and magnification data (black dot–dashed line). The top panel displays low-redshift bins, and the bottom panel displays high-redshift bins. As the fit using shear data alone does not constrain $\kappa_{\rm offset}$, the fitted value from using magnification alone is used for that model. The χ^2 values represent the χ^2 for a comparison of the magnification data with the shear-only fit (17 degrees of freedom), the magnification-only fit (14 degrees of freedom), and the combined fit (14 degrees of freedom), respectively.

investigation to determine whether this is due to issues with the shear signal, the magnification signal, the modelling process, or a combination of these factors.

Overall, the strong fit of our model implies that our method is successfully recovering a real magnification signal. This magnification signal may, however, still have biases in it due to effects we have not fully corrected for, such as dust obscuration or contamination

of the source sample. Thus at this stage, the specific values fit by our model may still be biased. Given the strong agreement with the shear signal, these biases are likely at a small level, however, and we can still gain insight into how magnification data can be used to improve the errors on fits to various parameters.

This can be seen in Fig. 6, where we show the best-fitting parameters for our model fits and their associated errors for all lens stellar
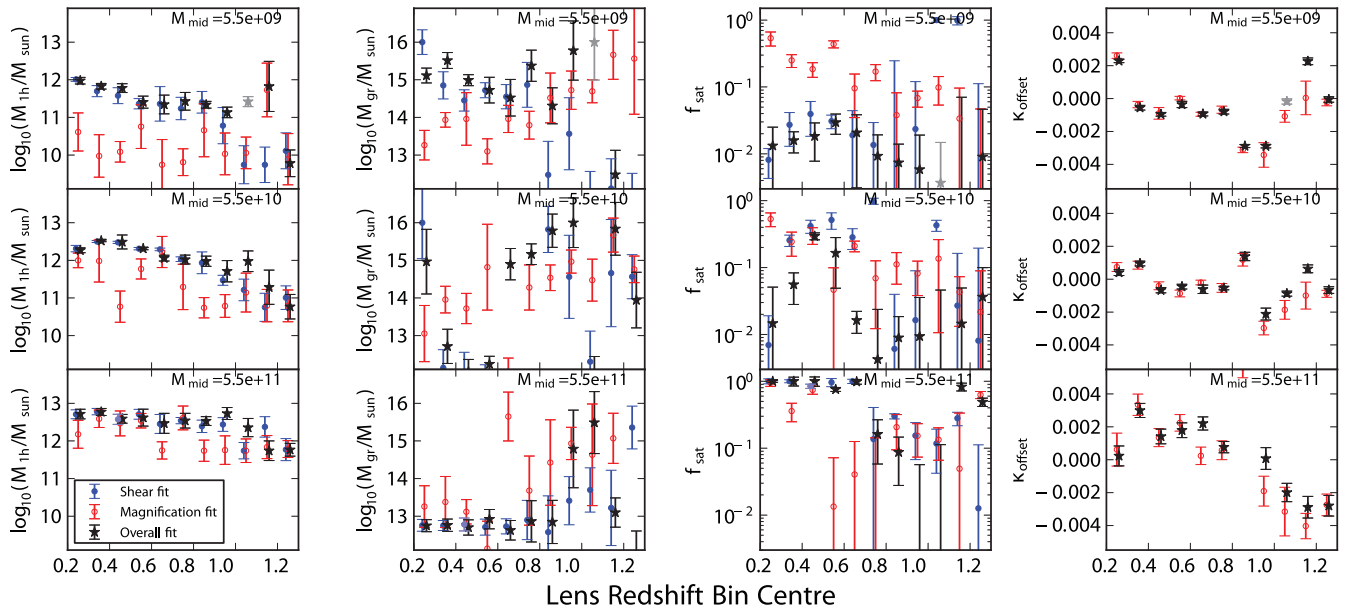
Figure 6. Best-fitting values and errors of our model parameters, for various lens redshift and mass bins. From left to right, the columns show the parameters: one-halo mass, group mass, satellite fraction, and κ_{offset}. The best-fitting values using shear data alone are shown in blue with filled circle markers, values from magnification data alone are shown in red with open circle markers, and values from the combined data are shown in black with star markers. As shear data alone cannot constrain κ_{offset}, the shear-only fit is not shown in that panel. Bins for which the χ^2 value indicates a poor fit (>30 for shear- and magnification-only fits, >50 for the fits to both) are indicated with a faded marker.

mass and redshift bins tested. This figure shows the fits to shear data alone, magnification data alone, and to the combined data set. As shown in Fig. 3, the contamination of the source sample with galaxies in the lens plane which were spuriously misidentified as sources becomes a significant issue at redshifts greater than $\sim$0.7, and so the fits including magnification data at these redshifts may be less reliable. The fitted M_{1h} and M_{gr} parameters from magnification alone seem to be in general agreement with those from shear alone. It is in the f_{sat} parameter that the two fits differ, with the magnification-only fit regularly favouring a larger f_{sat}.

Interestingly, some of the fitted κ_{offset} values show a significant change between the magnification-only fit and the fit to the combined data set. This seems to occur when the best-fitting group mass parameter M_{gr} for magnification alone was spuriously high but was corrected to a more-reasonable value when complemented with the shear data. There are also bins in which the shear-only fit appears to show a spuriously high or low value, but in which the fit to the combined data set is much more reasonable. This clearly shows the potential benefit of combining shear and magnification data.

The trend observed in the best-fitting values for the one-halo mass parameter M_{1h} at large redshift is potentially worrisome – the combined fit seems to consistently favour very low masses, often lower than either the shear or magnification data alone favours. This is possibly due to the fact that we are not fitting to the innermost radial bins, which means we are losing a significant amount of information relevant to the one-halo mass fit. This could also potentially be due to contamination of the source sample in high lens-redshift bins leading to spuriously low magnification signals. We discuss this possibility further in Section 5.3.

In Fig. 7, we show the ratios of the calculated errors on our fitted parameters to the errors obtained when shear information alone is used. Both the ratios of the magnification-only fit and the combined

data set fit are shown, as well as a calculation of the theoretical ratio, where the theoretical combined error is

$$\sigma_{th} = \left(\sigma_{shear}^{-2} + \sigma_{mag}^{-2} \right)^{-1/2}, \tag{38}$$

which is not necessarily valid for a fit of more than one parameter, depending on the relative degeneracies of the two fits.

This plot shows that when used alone, magnification typically does a poorer job at constraining the model parameters at low redshift. However, as expected, it is more beneficial at high redshift, particularly for the parameters relating to the host groups of the galaxies in the sample, M_{gr} and f_{sat}. Interestingly, the errors of the fit to the combined data set do not generally follow the theoretical calculations, in some cases being even smaller, and in some cases being even larger than in the shear-only case. This suggests that the situation is more complicated than the theoretical calculation assumes.

What is likely going on is a combination of two factors. The shear and magnification signals and the fits to them have different degeneracies, and the two fits disagree on the best-fitting set of parameters in some scenarios. The former factor can result in the fit to the combined data set having significantly smaller errors than the fit to either alone, as the combination of the two data sets breaks the degeneracies that are present in either alone, allowing much greater precision. The latter factor can result in the fit to the combined data set showing larger errors as the different regions preferred by each data set will spread out the most-reasonable region in the fit to the combined data set. The result is that when the regions preferred by the individual fits overlap, the fit to the combined data set is much smaller, while when they do not overlap, it can be larger.
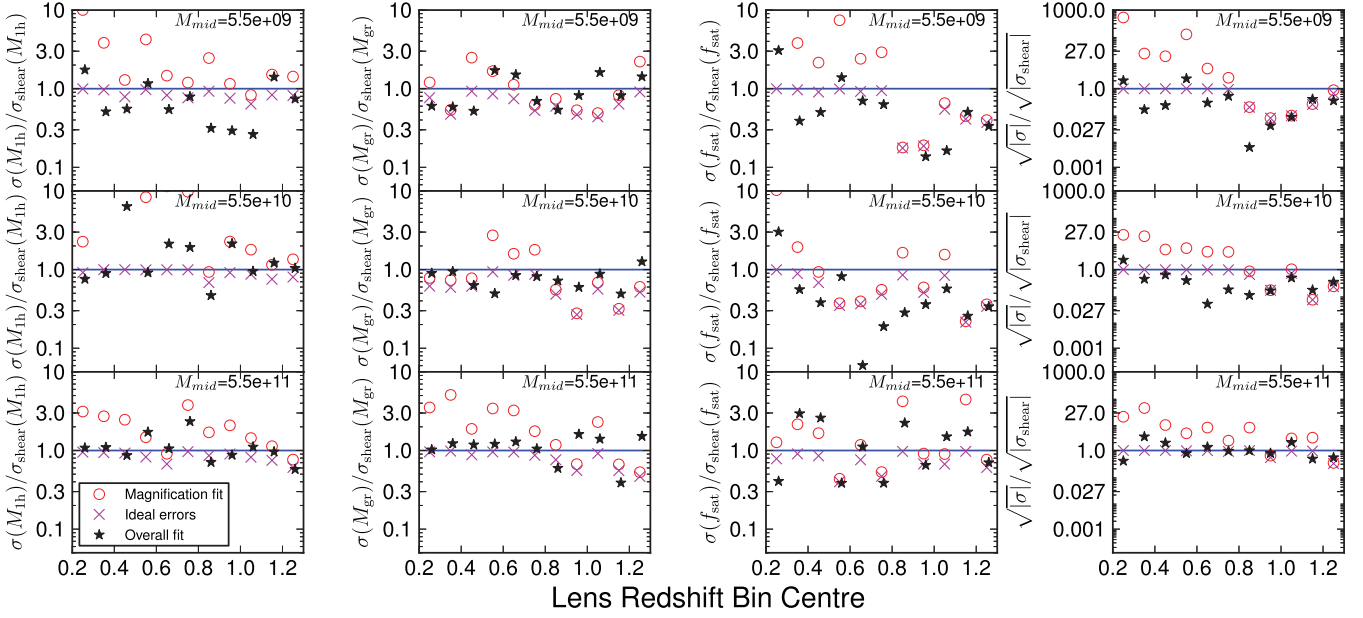
Figure 7. The ratios of errors in our model parameters to the errors obtained from using shear information alone, for various lens mass and redshift bins. From left to right, the columns show the parameters: one-halo mass, group mass, satellite fraction, and the square root of the determinant of the covariance matrix, which is representative of the total volume of the confidence regions. The red open circles show the ratios of errors from magnification alone to the errors from shear alone, the black stars show the ratios of errors from the combination of shear and magnification to shear alone, and the magenta crosses show the ratio of the theoretical combination of errors from magnification and shear to the errors from shear alone. Note that due to the multiparameter nature of the fit and the fact that the best-fitting values from shear and magnification alone sometimes disagree, the errors from the combined data set are occasionally greater than for shear alone.
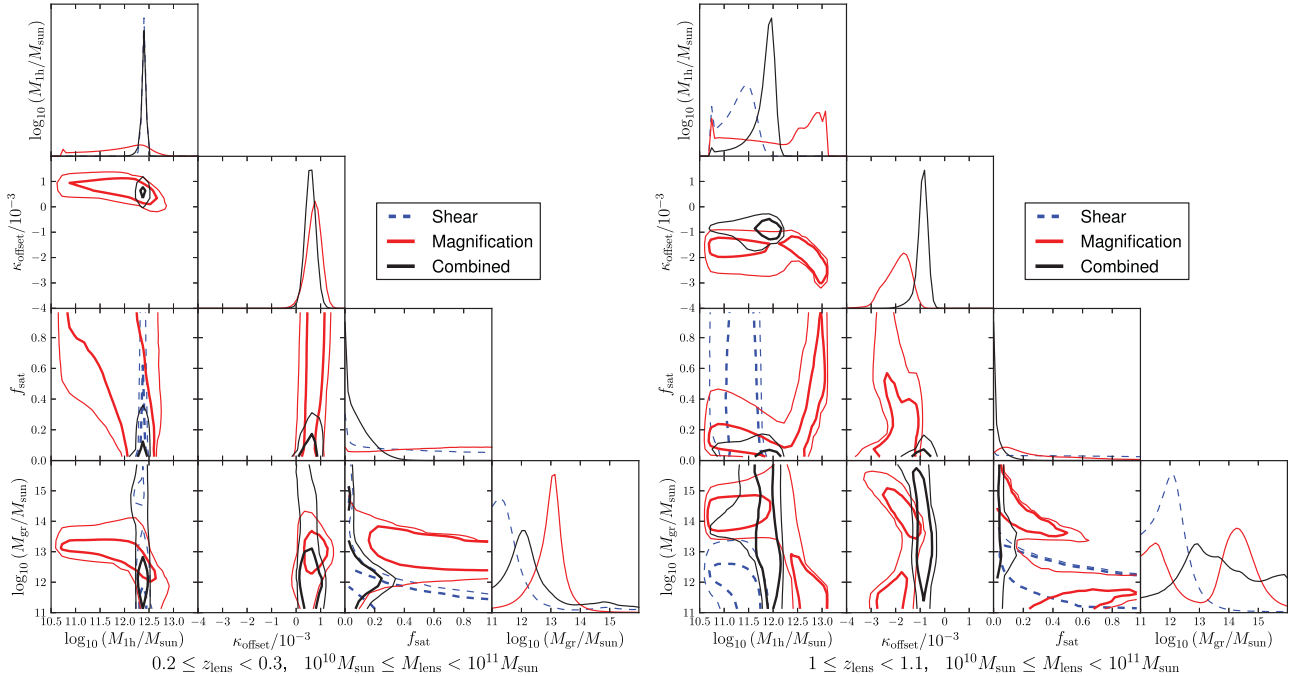
Figure 8. Confidence regions of fitted parameters for shear-only, magnification-only, and combined fits to the $0.2 \leq z < 0.3$, $10^{10}\,\mathrm{M}_\odot \leq M_\mathrm{lens} < 10^{11}\,\mathrm{M}_\odot$ bin (left) and the $1 \leq z_\mathrm{lens} < 1.1$, $10^{10}\,\mathrm{M}_\odot \leq M_\mathrm{lens} < 10^{11}\,\mathrm{M}_\odot$ bin (right). The thick and thin contours show 95 and 68 per cent confidence regions, respectively. Note that the shear-only fits are not shown for the panels which include the κ_offset parameter, as this parameter is unconstrained by the shear signal.

Notably, at high redshifts and low lens stellar masses, the shear signal approaches zero and is not reliably detectable, as can be seen in the bottom panel of Fig. 4. This results in the best-fitting model parameters being unreliable when shear alone is used. However, as can be seen in the bottom panel of Fig. 5, the magnification signal

can in fact still be measured for these bins, and so the best-fitting model parameters from using only magnification are likely to be more reliable.

In order to help illustrate the possible benefits of magnification data, Fig. 8 shows the confidence regions for the fitted parameters to

two lens bins, selected to illustrate cases in which magnification is particularly useful. These show that magnification provides the most benefit in the M_{gr}–f_{sat} parameter plane. Here, the shear signal is only able to provide an upper boundary to the combination of M_{gr} and f_{sat}. The magnification signal on its own is not much better, also covering larger regions of parameter space. However, the intersections of these confidence regions are relatively small, and so the fit to the combined data set has a much tighter confidence region.

It remains for future investigation to determine how much benefit magnification data will provide when all relevant biases are taken into account. The fact that even for this preliminary analysis, magnification occasionally provides large benefits in reducing the errors in fitted parameters is strongly encouraging and motivates further investigation to determine whether this will remain the case in a proper treatment.

5 DISCUSSION

In Section 5.1, we compare our method to the optimally weighted correlation function method, which has been used in past analyses. In Section 5.2, we present a proposed implementation of our methodology for mass-mapping. In Section 5.3, we present an estimate of the possible impact of source sample contamination on our results. In Section 5.4, we test and discuss the impact of our choice to weight our measurement of the unmagnified source density on a per-field basis. Finally, in Section 5.5, we discuss further possible work.

5.1 Comparison to other methods

When applied to galaxy–galaxy lensing, the results of this methodology are nearly equivalent to those obtained through an optimally weighted cross-correlation function estimator. This method was first proposed by Ménard & Bartelmann (2002), and in typical implementations (e.g. Scranton et al. 2005; Hildebrandt et al. 2009; Ford et al. 2012) uses a cross-correlation function in which all observed sources are weighted by $\alpha(m) - 1$. The measured correlation function can then be compared to the predicted correlation functions from mass models. To date, it has not been stated in the literature how this correlation function relates to the magnification, and so we investigate this here.

The angular cross-correlation represents the excess number density of galaxies in one sample in angular bins around galaxies in another sample. For scenarios in which only magnification is expected to influence the correlation between lens and source positions, which generally requires the two samples to be well separated in redshift, the expected cross-correlation between lenses and sources in a given angular and magnitude bin would be

$$\langle w(\theta, m) \rangle = \frac{\langle N(\theta, m) \rangle}{N_0(m)} - 1 = \mu(\theta)^{\alpha(m)-1} - 1, \quad (39)$$

where $N_0(m)$ is the unmagnified number count of source galaxies at a given magnitude and $\langle N(\theta, m) \rangle$ is the expected magnified number counts of galaxies at a given magnitude in a given angular bin relative to the lens sample, and using equation (4) to identify this ratio as corresponding to $\mu(\theta)^{\alpha(m)-1}$.

Approaches by e.g. Scranton et al. (2005), Hildebrandt et al. (2009), and Ford et al. (2012) calculate the following optimally weighted cross-correlation function, using a weighted version of the Landy & Szalay (1993) estimator,

$$w_{LS}^{opt} = \frac{LS^{\alpha-1} - RS^{\alpha-1} - \langle \alpha - 1 \rangle LR}{RR} + \langle \alpha - 1 \rangle, \quad (40)$$

where S stands for *sources*, L stands for *lenses*, and R stands for the *random* catalogue, which is at least 10 times as large as the lens and source catalogues. All sources are weighted by their $\alpha(m) - 1$ value, $\langle \alpha\text{-}1 \rangle$ is the mean value of $\alpha(m) - 1$ across the entire sample, and all pair counts are normalized by the product of the total numbers of galaxies in each catalogue.

This can be more easily compared to our method by instead presenting the cross-correlation function using the Davis & Peebles (1983) estimator, which converges to the same result as the Landy & Szalay (1993) estimator at high number counts. The optimally weighted cross-correlation function with this estimator is

$$w_{DP}^{opt} = \frac{LS^{\alpha-1}}{LR} - \langle \alpha - 1 \rangle, \quad (41)$$

using the same notation as in equation (40). The optimally weighted correlation can therefore be seen as a measurement of

$$w^{opt} = \frac{\sum_{j \in \text{annu}} \left[\alpha(m_j) - 1 \right]}{N_{0,\text{annu}}} - \frac{\sum_{j \in \text{samp}} \left[\alpha(m_j) - 1 \right]}{N_{0,\text{samp}}}, \quad (42)$$

where $j \in$ samp represents the indices of all source galaxies in the sample, $j \in$ annu represents only those source galaxies which lie in the appropriate annulus behind any lens in the sample, $N_{0,\text{samp}}$ is the total number of source galaxies in the sample, and $N_{0,\text{annu}}$ is the expected number of source galaxies lying in the annuli in the presence of no magnification. If we convert this sum to an integral in the same fashion as in equation (13), we get

$$w^{opt} = \frac{\int_a^b n(m) \left[\alpha(m) - 1 \right] dm - \int_a^b n_0(m) \left[\alpha(m) - 1 \right] dm}{\int_a^b n_0(m) \, dm}. \quad (43)$$

Compare this to the expansion of equation (14):

$$\hat{\mu} = 1 + \frac{\int_a^b n(m) \left[\alpha(m) - 1 \right] dm - \int_a^b n_0(m) \left[\alpha(m) - 1 \right] dm}{\int_a^b n_0(m) \left[\alpha(m) - 1 \right]^2 dm}. \quad (44)$$

In both equations, the number count of galaxies observed within an annulus of interest only appears within one term, $\int_a^b n(m) \left[\alpha(m_j) - 1 \right] dm$. The methods thus have equivalent statistical power, and it is possible to convert between their measurements through

$$\hat{\mu} = 1 + \frac{w^{opt}}{\langle [\alpha - 1]^2 \rangle}, \quad (45)$$

and

$$w^{opt} = \langle [\alpha - 1]^2 \rangle (\hat{\mu} - 1), \quad (46)$$

where

$$\langle [\alpha - 1]^2 \rangle = \frac{\int_a^b n_0(m) \left[\alpha(m) - 1 \right]^2 dm}{\int_a^b n_0(m) \, dm}. \quad (47)$$

This comparison allows the statistical error of w^{opt} to be easily estimated without expending computational power on a bootstrap or jackknife approach. It can instead be estimated as

$$\sigma_{w^{opt}} = \sqrt{\frac{\langle [\alpha - 1]^2 \rangle}{N_0 A_{\text{samp}}}}. \quad (48)$$

The derivation of this equation is presented in equation (A2). This estimate relies on the assumption that background galaxies follow an ideal Poisson distribution. Since clustering is expected among these galaxies, their actual distribution will have a larger standard

deviation than that of a Poisson distribution, and so this estimate will be biased low.

Computationally, both our method and the correlation function method scale as at best $O(N_l \log N_s)$ with the number of lens galaxies N_l and number of source galaxies N_s. This is due to the fact that they both include a step in which the source galaxies residing in a given annulus around lens galaxies must be determined, a process which can have this scaling at best. The difference comes in the coefficient of this term, which is dominated by the calculation of errors. Our method can calculate errors with only a few extra calculations per lens, as detailed in Section 2.2. However, for the correlation function method, a bootstrap or jackknife method is typically used to determine the errors, which requires rerunning the analysis on subsets of the data multiple times (e.g. 50 times in the case of Ford et al. 2012).

The comparison we show here allows calculations of the magnification signal through a correlation function estimator to use equation (48) as an estimate of the error. However, this estimate will be biased low due to the clustering of source galaxies. The algorithm we use to calculate a more accurate error estimate for our method (see equations 28 through 30) cannot be readily translated to a correlation function estimator, as it relies on the unmasked fractions of the annuli surrounding each lens in the sample, which are not calculated when a correlation-function-based approach is used. This results in our method being significantly faster than a correlation-function-based approach when accurate errors on the magnification estimate must be calculated, thanks to the fact that it does not have to resort to a computationally expensive bootstrap or jackknife approach.

It is worth noting that many correlation function estimators, when measuring the function for large separations, bin objects into cells and determine the unmasked fraction of each cell (e.g. Scranton et al. 2005). The correlation function between these cells is then calculated through a summation over pairs of cells. In this case, the error can be determined empirically through an analogous calculation to the one used for our method. However, such binning is only practical at large separations. Modifications can be made to the small-scale estimation to allow for error calculations, but this will require the calculation of the unmasked fraction of each annulus. If this is done, it is then no longer necessary to use a random catalogue to determine the overdensity within each annulus; since its area is known, its overdensity can be more efficiently calculated by simply comparing its density to the density of the full observation. At this point, the method is in fact more similar to our proposed method than to the traditional method of estimating a correlation function, lacking only the generalization done in our method to allow different redshifts for lenses. It is thus apparent that our proposed method is the most practical means to accurately determine the error in a measurement of a magnification signal.

Finally, our method also provides the advantage that it can handle overlapping lens and source samples. Correlation-function-based approaches require the lens and source samples to be well separated in redshift space so that there is no expected clustering between the two samples. Our method has the benefit that each lens is treated individually, and so it is only necessary that each lens be well separated from the subset of the source galaxy sample which we use to measure magnification. This results in the ability for lenses at the low-redshift end of the lens sample to use a greater amount of redshift space to measure the magnification signal, regardless of the fact that higher redshift lenses in the sample cannot use this space. This will result in an increased signal-to-noise of the magnification measurement.

5.2 Proposed implementation for mass-mapping

As we showed in Section 5.1, our method has equivalent statistical power to cross-correlation function methods already in use when applied to galaxy–galaxy lensing. One important advantage this method has is that it is capable of providing an estimate of the mean magnification of an arbitrary region of space, not limited to a circular annulus, or even to a region of constant weight.

As the statistical power of magnification varies with the square root of the area of the region measured, mass-mapping with magnification alone must necessarily involve a trade-off between resolution and accuracy. The simplest approach to mass-mapping would involve measuring the magnification, and thus projected overdensity, in a circular aperture around any point in the region of interest. While this method can be easily applied with the methodology outlined in the previous sections, it suffers from the drawback that the resulting mass-map will show discontinuous behaviour whenever a source galaxy enters or exits the aperture.

This discontinuous behaviour can be avoided through instead applying a filter function which approaches zero at large separation, such as a two-dimensional Gaussian weight function. This also provides the benefit that source galaxies closer to the point of interest will have more influence on its measured signal than distant galaxies. Let us call this function $F(\theta)$.

The magnification at a given angular position $\boldsymbol{\theta}$ and redshift z can then be estimated as

$$\hat{\mu}_{kl} = 1 + \frac{1}{A_F w(z + \Delta z)}$$
$$\times \left(\sum_j F\left(\left|\boldsymbol{\theta}'_j - \boldsymbol{\theta}\right|\right) \left[\alpha(m_j) - 1\right] - A_F d(z + \Delta z) \right), \quad (49)$$

where $d(z)$ and $w(z)$ are defined as in Section 2.2, j represents a summation over all source galaxies at redshift $z + \Delta z$ or greater, with $\boldsymbol{\theta}'_j$ representing the angular coordinates of these galaxies, and

$$A_F = \iint F(\theta) \, d\Omega. \quad (50)$$

In practice, source galaxies sufficiently distant from the position of interest such that $F\left(\left|\boldsymbol{\theta}'_j - \boldsymbol{\theta}\right|\right) \to 0$ can be ignored in this calculation.

One notable issue with mass-mapping with magnification is the fact that position-dependent biases in the redshift determinations of source galaxies can result in a spurious magnification signal, as discussed in Section 2.3. In the scenario of galaxy–galaxy lensing, this effect is of large enough scale that it can be treated as a flat offset and included as a fitted parameter in the model, but this is not possible for mass-mapping. The effect is somewhat reduced in mass-mapping, as the positions of lenses are not used, and so this effect should only appear in tomographic analyses in which a redshift cut-off is imposed on sources. This cut-off would allow for position-dependent effects, as the number of sources beyond the cut-off will depend on the errors in their redshift determinations, which is itself position-dependent. Care will thus have to be taken in such an analysis to account for this effect.

5.3 Estimating source sample contamination

As shown in Fig. 3, the fraction of galaxies whose photometric redshifts make them appear to be sources but which in fact lie within the lens plane is of order 0.1 per cent at low redshifts, and can be
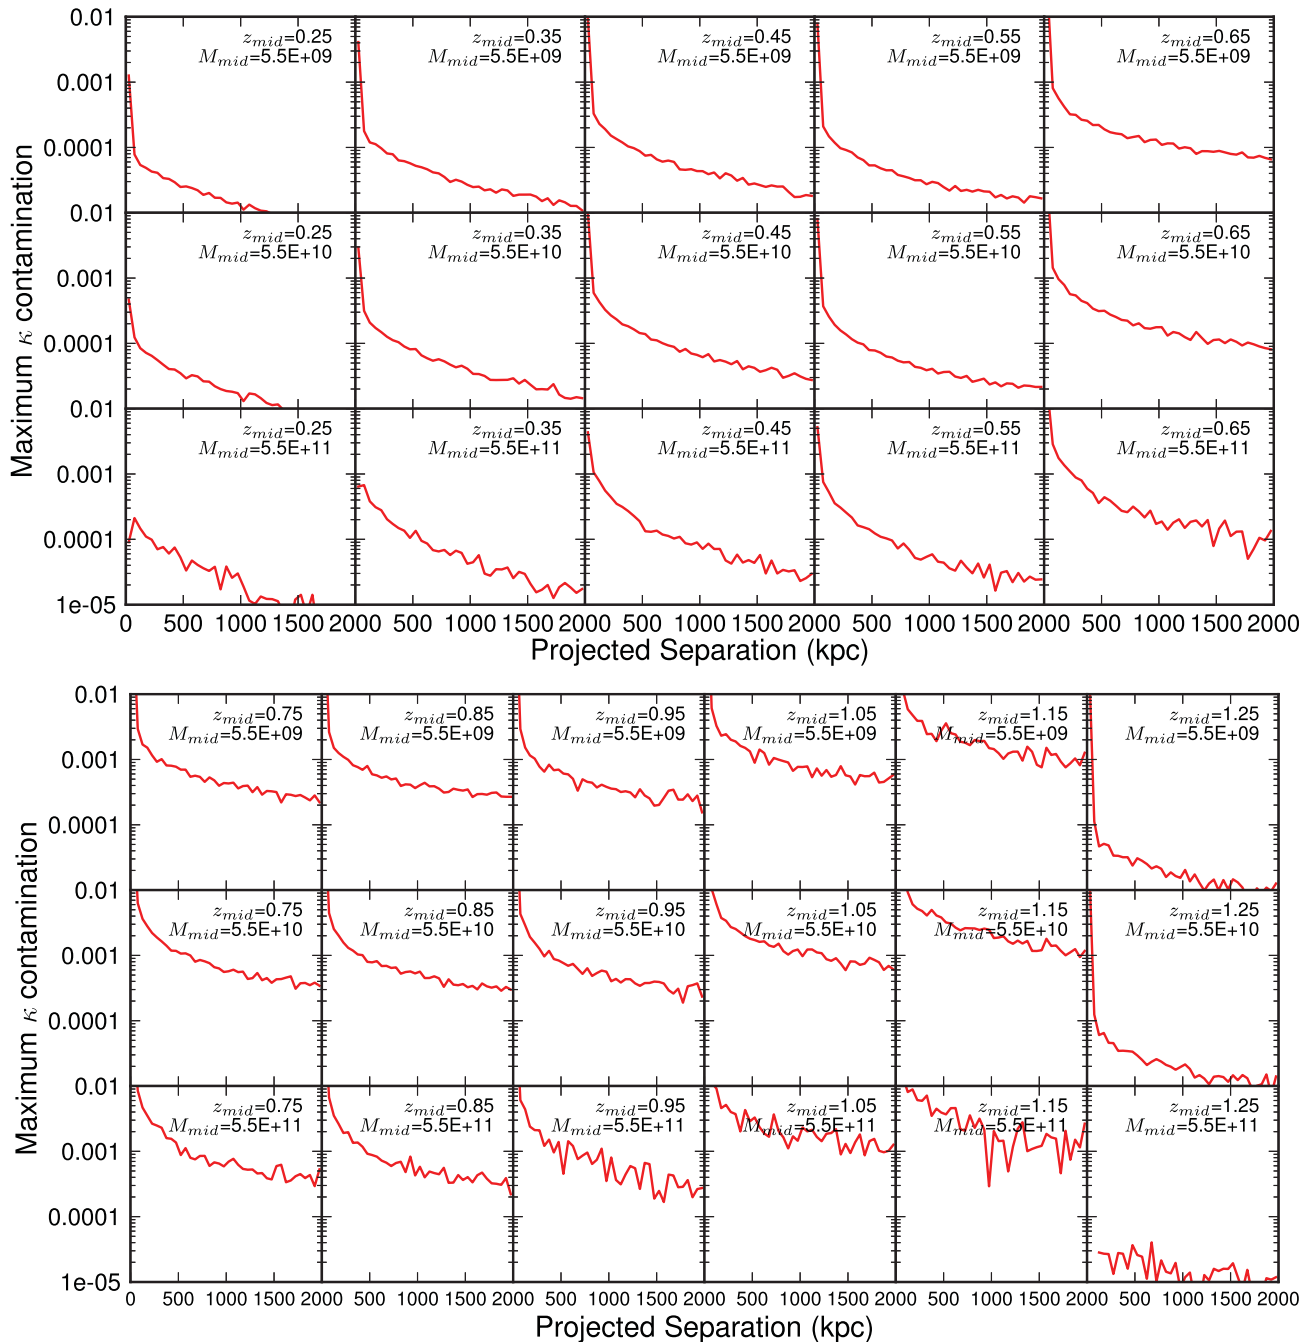
Figure 9. The maximum change in the measured κ due to contamination of the source sample with lenses. Low-redshift bins are shown in the top panel, and high-redshift bins are shown in the bottom panel. This maximum change is calculated from the product of the correlation function shown in Fig. B1, the contamination at the lens redshift, which is shown in Fig. 3, and the factor 0.25, which represents a conservative upper bound on the slope of κ/w.

as high as ~10 per cent at high lens redshifts, even with a large buffer redshift between lenses and sources imposed. Galaxies in the lens plane will typically be correlated in position with each other due to their mutual gravitational attraction. Therefore, if galaxies in this plane are misidentified as sources, we would expect them to generally appear closer to lens galaxies than they would through random chance alone.

We can make an order-of-magnitude estimate of the impact this will have on our measured magnification signal through multiply-

ing the correlation function of galaxies in the lens plane with the contamination and a factor representing the conversion from the correlation function to κ. For this conversion factor, we choose 0.25, which would correspond to a mean $|\alpha(m) - 1|$ of ~0.5. Typically, the mean is less than this, so this results in an upper bound on the contamination of κ. We plot the resulting predicted contamination for selected lens redshift and stellar mass bins in Fig. 9. As can be seen through comparison with Fig. 5, this value is typically negligible for $z < 0.6$, becoming comparable in magnitude to κ in
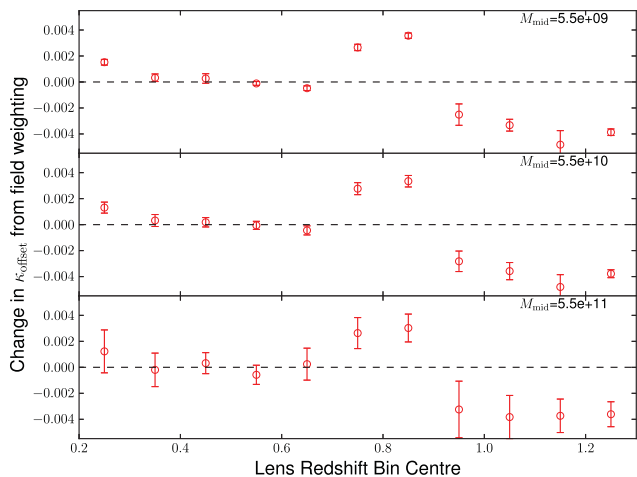
Figure 10. The change in the fitted values of κ_{offset} when weighting is applied on a per-field basis in the source-counting procedure. This weighting accounts in part for observational effects such as variable seeing which indirectly result in the number densities of lenses and sources being correlated or anticorrelated. The magnitude of the change shown here provides an estimate of the uncertainty in the value of κ_{offset} due to smaller-scale variations.

the $0.6 < z < 0.7$ bin. In higher redshift bins, the effect is even worse.

Further investigation of correcting for this effect is beyond the scope of this paper, but it is unlikely to be present significant difficulties. By virtue of the fact that the galaxies in the lens plane which are assigned photometric redshifts appropriate for sources, likely have a similar magnitude distribution to legitimate source galaxies, correction for their differing magnitude distribution may not be necessary. As such, a simple calculation based on the correlation function of galaxies in the lens plane and the mean value of $\alpha - 1$ for sources may be all that is needed.

5.4 Effects of field weighting

In Section 2.3, we discussed how correlations and anticorrelations of lenses and sources in regions of the sky can result in an effect which mimics a magnification signal. We chose to partially correct for this effect by applying weights on a field-by-field basis in determining the unmagnified source density. We can demonstrate the impact of this weighting through a comparison of the best-fitting κ_{offset} parameters in the scenarios in which we do and do not apply field-by-field weighting, as shown in Fig. 10.

As expected, the application of weighting has the same effect on all lens stellar mass bins, but the effect varies with redshift. This is consistent with our hypothesis that variable seeing could result in varying amounts of scatter between redshift bins from field to field, as this effect would not be expected to be constant across all redshifts due to the complicated nature of determining photometric redshifts.

In general, the magnitude of change in κ_{offset} due to field-by-field weighting is $\lesssim 0.004$. Any remaining effect due to variable seeing on smaller scales is likely to result in a smaller change, and so we would expect the proper fitted values for κ_{offset} for all bins to be of smaller magnitude. From the fitted values shown in Fig. 6, it appears that the scatter in fitted values of κ_{offset} is comparable to this magnitude. Additionally, many bins of the same redshift but differ-

ent stellar masses show significantly different fitted values of κ_{offset}. This suggests that the fitted values for κ_{offset} are not sensitive to this effect alone, but are also affected by other aspects of the measured lensing signals. This is possibly due to the fact that the model we use does not include contributions for larger-scale structure, which is dominant at the same large radial separations that κ_{offset} is.

Alternatively, let us consider that this may be due not to an observational effect, but due to actual over- or underdensities in the regions of these lenses. The convergence maps generated by Van Waerbeke et al. (2013) from shear data alone show that, when smoothed on scales of 2 arcmin and correcting for the contribution of noise, the convergence has typical standard deviation of ~ 0.006.[5] The typical area covered by all lenses within one of our stellar mass and redshift bins is of order $\sim 10^5$ arcmin^2. As this is roughly four orders of magnitude greater than the smoothing scale applied by Van Waerbeke et al. (2013), we would expect the standard deviation of κ due to large-scale structure within this region to drop by two orders of magnitude, to $\sim 0.000\,06$. Clearly, this is too small to account for the scatter we observe in the fitted values of κ_{offset}, and so they are more likely due to observational effects or fitting errors than cosmic variance.

5.5 Further research

In our attempts to fit a mass profile to shear and magnification data simultaneously, we found several indications that our fit was not adequate, including poor χ^2 comparisons, extreme and unlikely values fit for certain bins, and inconsistent trends over changing lens redshift and stellar mass. It is possible that these issues might be resolved through the use of a different model mass profile, or through the inclusion of the effects of large-scale structure on the lensing signal. This will allow us to better determine the potential benefits of magnification in breaking degeneracies between parameters characterizing the environment surrounding galaxies in our samples.

As mentioned in Section 4.2, dust obscuration can result in the measured magnification signal to be spuriously high if it is not corrected for. In this paper, we handled this effect through discarding the two innermost radial bins when we fit our mass profiles to the magnification data, as these are the bins which are likely to show the most significant contribution from dust. A proper analysis would require modelling the obscuration from dust, as done by e.g. Ménard et al. (2010). This will be necessary before firm conclusions can be drawn from our model fits.

6 SUMMARY AND CONCLUSIONS

In this paper, we have presented a new method to estimate weak lensing magnification which can be applied to an arbitrary region of space, as expressed in equation (16) and in a more efficient form for computations in equation (23). We have shown that the estimator resulting from this method has equivalent statistical power to the optimally weighted correlation function estimator which has

---

[5] We determined this value through mass-map data provided by the authors of Van Waerbeke et al. (2013). Under the assumption that the variance in the predicted convergence from their mass-maps is the sum of the variances due to noise and from the actual variance in κ, we took the difference between the variances in κ for their data maps and noise maps to estimate the variance, and thus standard deviation, in κ due to the actual distribution of matter.

been used for past analyses, and we provide conversions between them in equations (45) and (46). This conversion also allows a simple way to estimate the errors on measurements made with the correlation function estimator without the use of bootstrap or jackknife methods, which we present in equation (48). However, this error estimate is likely to be biased low due to the clustering of source galaxies. Our method can calculate a more accurate estimate of error with only a few extra calculations, which presents a large benefit over the requirement of correlation-function-based methods to use bootstrap or jackknife approaches to accurately calculate error. Our method is also capable of handling overlapping lens and source samples, which can increase the signal-to-noise of the magnification signal when this is the case.

The estimator we have presented here provides a suitable mathematical basis to use magnification for mass-mapping, for instance through the form in equation (49). However, applying this will be difficult in practice, as the contributions to the magnification signal from actual fluctuations in the projected mass density field are degenerate with localized observational biases, such as due to variable seeing across the field of view. It remains for future research to determine whether this issue can be resolved, through advances in photometric redshift estimation to correct for seeing-induced biases or other techniques.

We have shown a proof-of-principle of our method applied to galaxy–galaxy lensing in the CFHTLenS fields, and shown that the convergence profiles determined with it are consistent with the convergent predictions from mass profiles which are fit to shear data in the fields, as can be seen in Fig. 5. As expected, in most cases magnification data provides poorer signal-to-noise than does shear, but its contribution is not always negligible. In fact, in certain cases its combination with shear seems to provide significantly smaller errors on model parameters than either alone, likely due to the magnification data helping break degeneracies between multiple parameters. This finding is tentative, however, as our present algorithm was only designed as a proof-of-principle and does not yet include provisions for effects such as dust obscuration, which can result in biases in the magnification signal.

REFERENCES

Baltz E. A., Marshall P., Oguri M., 2009, J. Comsol. Astropart. Phys., 1, 15
Bauer A. H., Gaztañaga E., Martí P., Miquel R., 2014, MNRAS, 440, 3701
Benítez N., 2000, ApJ, 536, 571
Bielby R. et al., 2012, A&A, 545, A23
Boulade O. et al., 2003, Proc. SPIE, 4841, 72
Broadhurst T. J., Taylor A. N., Peacock J. A., 1995, ApJ, 438, 49
Davis M., Peebles P. J. E., 1983, ApJ, 267, 465
Davis M. et al., 2007, ApJ, 660, L1
Dietrich J. P., Erben T., Lamer G., Schneider P., Schwope A., Hartlap J., Maturi M., 2007, A&A, 470, 821
Erben T. et al., 2009, A&A, 493, 1197
Erben T. et al., 2013, MNRAS, 433, 2545
Ford J. et al., 2012, ApJ, 754, 143
Ford J., Hildebrandt H., Van Waerbeke L., Erben T., Laigle C., Milkeraitis M., Morrison C. B., 2014, MNRAS, 439, 3755
Garilli B. et al., 2008, A&A, 486, 683
Gillis B. R. et al., 2013, MNRAS, 431, 1439
Heymans C. et al., 2012, MNRAS, 427, 146
Hildebrandt H., Pielorz J., Erben T., van Waerbeke L., Simon P., Capak P., 2009, A&A, 498, 725
Hildebrandt H. et al., 2012, MNRAS, 421, 2355
Hildebrandt H. et al., 2013, MNRAS, 429, 3230
Huff E. M., Graves G. J., 2014, ApJ, 780, L16
Ilbert O. et al., 2010, ApJ, 709, 644
Landy S. D., Szalay A. S., 1993, ApJ, 412, 64
Le Fèvre O. et al., 2005, A&A, 439, 845
McGee S. L., Balogh M. L., Bower R. G., Font A. S., McCarthy I. G., 2009, MNRAS, 400, 937
Mandelbaum R., Tasitsiomi A., Seljak U., Kravtsov A. V., Wechsler R. H., 2005, MNRAS, 362, 1451
Ménard B., Bartelmann M., 2002, A&A, 386, 784
Ménard B., Scranton R., Fukugita M., Richards G., 2010, MNRAS, 405, 1025
Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, J. Chem. Phys., 21, 1087
Miller L. et al., 2013, MNRAS, 429, 2858
Morrison C. B., Hildebrandt H., 2015, MNRAS, 454, 3121
Narayan R., 1989, ApJ, 339, L53
Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493 (tNFW)
Neto A. F. et al., 2007, MNRAS, 381, 1450
Newman J. A. et al., 2013, ApJS, 208, 5
Savitzky A., Golay M. J. E., 1964, Anal. Chem., 36, 1627
Schneider P., King L., Erben T., 2000, A&A, 353, 41
Scranton R. et al., 2005, ApJ, 633, 589
Seljak U., 2000, MNRAS, 318, 203
ThePawn 2013, Bias correction in weighted variance. Available at: http://stats.stackexchange.com/a/47379
Van Waerbeke L., Hildebrandt H., Ford J., Milkeraitis M., 2010, ApJ, 723, L13
Van Waerbeke L. et al., 2013, MNRAS, 433, 3373
Velander M. et al., 2014, MNRAS, 437, 2111

APPENDIX A: DERIVATIONS OF SELECTED EQUATIONS

In equation (8), we presented the expected standard deviation of the estimator μ_i for the magnification of an individual magnitude bin. This equation can be derived as follows, starting from equation (7):

$$
\begin{aligned}
\sigma_{\mu_i} &= \frac{\sigma_{\mathrm{d}n_i}}{|\alpha_i - 1| \, \mathrm{d}n_{0,i}} \\
&= \frac{\sqrt{\mathrm{d}n_i}}{|\alpha_i - 1| \, \mathrm{d}n_{0,i}} \\
&= \frac{\sqrt{\mathrm{d}n_{0,i}\,\mu^{\alpha-1}}}{|\alpha_i - 1| \, \mathrm{d}n_{0,i}} \\
&\approx \frac{\sqrt{1 + (\alpha_i - 1)(\mu - 1)}}{|\alpha_i - 1| \, \sqrt{\mathrm{d}n_{0,i}}}.
\end{aligned}
\tag{A1}
$$

In equation (48), we presented an equation for estimating the standard error of the optimally weighted correlation function estimator for magnification, without relying on a bootstrap or jackknife approach. This equation can be derived as follows, starting from equation (46):

$$
\begin{aligned}
\sigma_{w^{\mathrm{opt}}} &= \left\langle [\alpha - 1]^2 \right\rangle \sigma_{\hat{\mu}} \\
&= \frac{\int_a^b n_0(m)\,[\alpha(m) - 1]^2 \, \mathrm{d}m}{\int_a^b n_0(m)\,\mathrm{d}m} \\
&\quad \times \left(\int_a^b n_0(m)\,[\alpha(m) - 1]^2 \, \mathrm{d}m \right)^{-1/2} \\
&= \frac{\left(\int_a^b n_0(m)\,[\alpha(m) - 1]^2 \, \mathrm{d}m \right)^{1/2}}{\int_a^b n_0(m)\,\mathrm{d}m} \\
&= \sqrt{\frac{\left\langle [\alpha - 1]^2 \right\rangle}{N_0 A_{\mathrm{samp}}}},
\end{aligned}
\tag{A2}
$$

where N_0 is defined as in Section 5.1.

APPENDIX B: SUPPLEMENTAL FIGURES

Fig. B1 shows the correlation function between galaxies in the lens and source samples within the same redshift slices for low-redshift bins. This can be used to estimate the contamination fraction of lenses in the source sample, which we showed in Fig. 3.
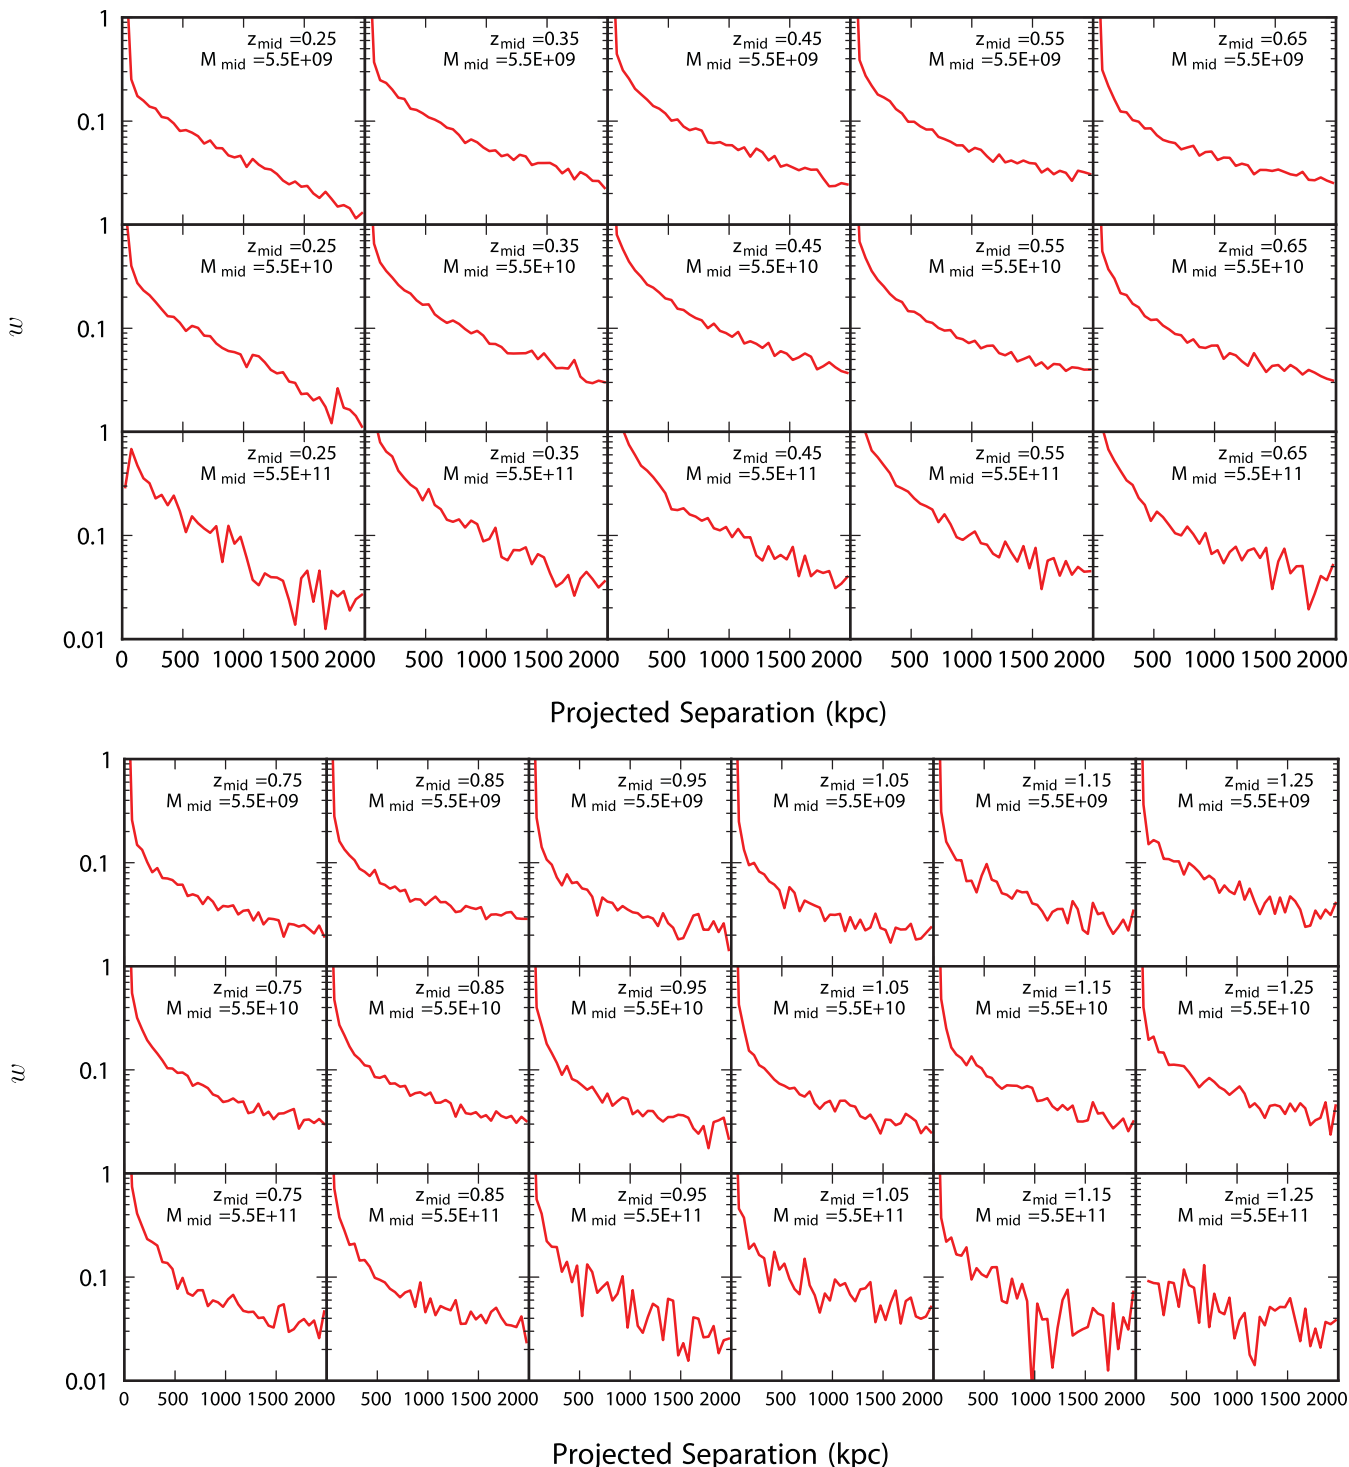
Figure B1. Correlation function between galaxies from the lens and source samples in various redshift and stellar mass bins. Low-redshift bins are shown in the top panel, and high-redshift bins in the bottom panel.

APPENDIX C: CODE ACCESS

In order to facilitate further work with magnification, we make the magnification and shear measurement code used for this paper publicly available. The code may be accessed online through a repository hosted at https://bitbucket.org/brgillis/magnification_public. If this repository is no longer available and a web search fails to locate a new host, interested parties are invited to contact the lead author, Bryan Gillis, at brg@roe.ac.uk to request the code.

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