

GAMA+KiDS: Empirical correlations between halo mass and other galaxy properties near the knee of the stellar-to-halo mass relation

Edward N. Taylor,^{*}¹ Michelle E. Cluver,^{1,2} Alan Duffy,¹ Pol Gurri,¹ Henk Hoekstra,³ Alessandro Sonnenfeld,³ Malcolm N. Bremer,⁴ Margot M. Brouwer,⁵ Nora Elisa Chisari,⁶ Andrej Dvornik,⁷ Thomas Erben,⁸ Hendrik Hildebrandt,⁷ Andrew M. Hopkins,⁹ Lee S. Kelvin,¹⁰ Steven Phillipps,⁴ Aaron S. G. Robotham,¹¹ Cristobál Sifón,¹² Mohammadjavad Vakili,³ and Angus H. Wright⁷

¹Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn 3122, Australia

²Department of Physics and Astronomy, University of the Western Cape, Robert Sobukwe Road, Bellville, South Africa

³Leiden Observatory, Leiden University, PO Box 9513, Leiden, 2300 RA, The Netherlands

⁴HH Wills Physics Laboratory, University of Bristol, Tyndall Avenue, Bristol, BS8 1TL, UK

⁵Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, the Netherlands

⁶Institute for Theoretical Physics, Utrecht University, Princetonplein 5, 3584 CC Utrecht, The Netherlands

⁷Ruhr-University Bochum, Astronomical Institute, German Centre for Cosmological Lensing, Universitätsstr. 150, 44801 Bochum, Germany

⁸Argelander-Institut für Astronomie, Auf dem Hügel 71, 53121 Bonn, Germany

⁹Australian Astronomical Optics, Macquarie University, 105 Delhi Rd, North Ryde, NSW 2113, Australia

¹⁰Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA

¹¹ICRAR, M468, University of Western Australia, Crawley, WA 6009, Australia

¹²Instituto de Física, Pontificia Universidad Católica de Valparaíso, Casilla 4059, Valparaíso, Chile

18 June 2020

ABSTRACT

We use KiDS weak lensing data to measure variations in mean halo mass as a function of several key galaxy properties (namely: stellar colour, specific star formation rate, Sérsic index, and effective radius) for a volume-limited sample of GAMA galaxies in a narrow stellar mass range ($M_* \sim 2\text{--}5 \times 10^{10} M_\odot$). This mass range is particularly interesting, inasmuch as it is where bimodalities in galaxy properties are most pronounced, and near to the break in both the galaxy stellar mass function and the stellar-to-halo mass relation (SHMR). In this narrow mass range, we find that both size and Sérsic index are better predictors of halo mass than colour or SSFR, with the data showing a slight preference for Sérsic index. In other words, we find that mean halo mass is more strongly correlated with galaxy structure than either stellar populations or star formation rate. These results lead to an approximate lower bound on the dispersion in halo masses among $\log M_* \approx 10.5$ galaxies: we find that the dispersion is $\gtrsim 0.3$ dex. This would imply either that offsets from the mean SHMR are very tightly coupled to size/structure, or that the dispersion in the SHMR is larger than past results have suggested. Our results thus provide new empirical constraints on the relationship between stellar and halo mass assembly at this particularly interesting mass range.

Key words: galaxies: formation and evolution – galaxies: mass functions – galaxies: statistics – galaxies: stellar content – galaxies: fundamental parameters

1 INTRODUCTION

As the quantitative link between the observed galaxy population and the cosmological population of dark matter halos, the stellar-to-halo mass relation (SHMR) represents a crucial interface between observation and theory (see Wechsler & Tinker 2018,

for a recent review). The messy baryonic processes of galaxy formation and evolution are understood to be seeded by the dissipationless collapse of their larger dark matter halos. The ongoing accretion onto and dynamics within galaxies are thus driven by the gravitational potential well at the centre of the halo, and regulated by shocks, outflows, and other astrophysical processes of feedback within and at the outskirts of the halo. Secular evolutionary processes like gas accretion, dynamical insta-

* entaylor@swin.edu.au

bility, star formation, and feedback would lead to the expectation of self-similar evolution of galaxies of a given mass, with the possibility of second order effects tied to formation time, and so to large-scale environment. In addition, stochastic, perturbative effects like interactions and/or mergers between galaxies produce significant differences in the evolutionary trajectories of individual galaxies. The outstanding challenge of galaxy formation and evolution is to identify and articulate the relative importance of these many different processes and mechanisms by connecting the cosmological population of dark matter halos to the correlated distributions of galaxy parameters as observed in the real Universe.

Techniques like abundance matching (*e.g.* Conroy, Wechsler & Kravtsov 2006; Guo et al. 2010; Moster et al. 2010) have been used to derive the average SHMR by forcing consistency between a halo mass function from theory (*e.g.* Tinker et al. 2010) and the observed galaxy stellar mass function (*e.g.* Bell et al. 2003; Marchesini et al. 2009; Baldry et al. 2012; Driver et al. 2018). Extensions or refinements like halo occupation modelling (*e.g.* Berlind & Weinberg 2002; Yang, Mo & van den Bosch 2003) also require consistency with clustering statistics and other observational constraints (but see Moster et al. 2010, who argue that clustering statistics do not have a significant influence on the inferred SHMR). A number of studies have now extended this formalism to infer the SHMR based on the combination of weak lensing measurements with number counts and/or spatial clustering (*eg* Leauthaud et al. 2012; Velander et al. 2014; Coupon et al. 2015; van Uitert et al. 2015). While lensing provides an avenue for direct measurement of the SHMR for $\log M_* \lesssim 10.5$, van Uitert et al. (2015) found that current weak lensing data are not sufficient on their own to constrain the SHMR for $\log M_{\text{halo}} \gtrsim 12$: at these high masses, it is the stellar mass function that provides the tighter constraint on the SHMR.

There is a qualitative and quantitative consensus on the form of the SHMR, at least in terms of a population average, that emerges from these analyses. The generic result is that the knee in the galaxy stellar mass function at $\log M_* \sim 10.5$ is tied to a break in the SHMR, with an associated peak in the stellar-to-halo mass ratio of $\sim 2\text{--}3\%$, at a halo mass $\log M_{\text{halo}} \sim 12$ (*e.g.* Moster et al. 2010; Behroozi, Conroy & Wechsler 2010; van Uitert et al. 2015). On either side of this peak, the SHMR is reasonably described as a power law, with the low- and high-mass slopes usually taken as reflecting the suppression of star formation by supernova (*e.g.* Larson 1974; McKee & Ostriker 1977; Joung & Mac Low 2006) and by AGN feedback (*e.g.* Croton et al. 2006; Bower et al. 2006), respectively, in the low- and high-mass regimes (see also, *e.g.* Mitchell 2016). But it is worth emphasising that this generic result is a virtually inescapable consequence of trying to reconcile the observed Schechter-like galaxy stellar mass function with a close-to-power-law halo mass function (different versions of this argument can be found in, *e.g.*, Marinoni & Hudson 2002; Moster et al. 2010; Behroozi, Conroy & Wechsler 2010; van Uitert et al. 2015). Any model that gets both the stellar and halo mass functions right will necessarily give a similar form for the SHMR.

Weak gravitational lensing (see reviews by Bartelmann & Schneider 2001; Hoekstra & Jain 2008), and more specifically galaxy–galaxy weak lensing, is one of the most successful observational avenues to obtaining direct halo mass measurements for large and representative galaxy samples (*e.g.* Brainerd, Blandford & Smail 1996; Hudson et al. 1998; Hoekstra, Yee & Gladders 2004; Mandelbaum et al. 2006). In order to more directly challenge models of galaxy formation and evolution in a cosmological context, our goal in this paper is to take a more

empirical approach to exploring the role of halo mass, as measured by galaxy–galaxy weak lensing, in influencing or determining the observable properties of galaxies.

Our specific interest is to probe correlations between halo mass and galaxy properties *at fixed stellar mass* — or, in other words, to identify which galaxy property or properties are most directly correlated with the dispersion around the average SHMR.¹ By attempting a systematic (if not exhaustive) exploration of second-order correlations around the SHMR, our goals are similar to, but distinct from, studies by Mandelbaum et al. (2006), Hudson et al. (2015), Charlton et al. (2017), and others who have identified differences in the SHMRs inferred for galaxy subsamples separated by colour and/or size. Our goals are similarly complementary to, *e.g.*, van Uitert et al. (2015), who have considered whether halo mass correlates more strongly with stellar mass or velocity dispersion (see also, *e.g.* Li et al. 2013).

One novel aspect of this paper is that, as a means to control for the mass dependence of the SHMR, we focus on a narrow mass range in stellar mass: $10.3 < \log M_* < 10.7$, or $M_* \approx 2\text{--}5 \times 10^{10} M_\odot$. This mass range is particularly interesting for several reasons. First, it is close to the knee of both the galaxy stellar mass function and the SHMR. It is thus where the stellar-to-halo mass ratio peaks, and so (in the canonical view) the point of transition where stellar feedback gives way to AGN feedback as the dominant regulator of star formation. Robotham et al. (2014) have also shown, based on galaxy pair counts and star formation rates, that this mass range is where galaxy mergers take over from star formation as the dominant channel for galaxy stellar mass growth.

Perhaps most significantly for this work, this mass range is also where the bimodality (or bimodalities) in galaxy properties is most pronounced, and where there are approximately equal numbers of canonically ‘early’ and ‘late’ type galaxies — whether that distinction be made on the basis of broadband colour (see, *e.g.* Baldry et al. 2006; Peng et al. 2012; Taylor et al. 2015), specific star formation rate (Renzini & Peng 2015), morphological classification (Bamford et al. 2009; Kelvin et al. 2015; Moffett et al. 2016), structure (van der Wel 2008), size (Shen et al. 2003; Lange et al. 2015), etc. (see also, *e.g.*, Robotham et al. 2013). By having a relatively large spread in galaxy properties across our sample, we obtain the best lever arm on any dependence on halo mass with these properties. Further, because this is the mass regime where there is the greatest diversity in galaxy properties, this is also where the influence of halo mass is potentially the most interesting: can one or more manifestations of galaxy bimodality be linked to differences in halo mass?

The structure of our discussion is as follows. We lay out our experimental design in §2, including our sample selection (§2.1) and subdivision (§2.2), weak lensing measurements (§2.3), and halo mass modelling (§2.4). We present proof of concept for our novel approach in §3, including demonstrated consistency with existing results (§3.1) and a variety of null results (§3.2), before presenting our main results in §4. In §5.1, we discuss our results and their implication for the role of halo mass in galaxy formation and

¹ In more theory-minded approaches like abundance matching and halo occupation modelling, the dispersion in the SHMR is usually framed in terms of the diversity of galaxy stellar mass values at fixed halo mass; *i.e.*, $p(M_*|M_{\text{halo}})$. For this paper, we consider instead the complementary quantity: the distribution of halo mass at fixed stellar mass, $p(M_{\text{halo}}|M_*)$. The two quantities are related, but distinct, with the latter being a more natural observable.

the dispersion in the SHMR. We also consider potential confounding effects and biases in §5.2 and §5.3. A full summary of our quantitative results is given in Table 1, and we summarise our main results and conclusions in §6. For our stellar mass estimates, we assume a Chabrier (2003) stellar initial mass function (IMF), and we have adopted a concordance cosmology $(\Omega_m, \Omega_\Lambda, h) = (0.3, 0.7, 0.7)$ throughout.

2 EXPERIMENTAL DESIGN

2.1 Lens galaxy sample selection

Our lens sample is selected from the Galaxy And Mass Assembly (GAMA) survey (Driver et al. 2011; Liske et al. 2015; Baldry et al. 2018), which has obtained near-total ($\gtrsim 99.5\%$) spectroscopic completeness for $r < 19.8$ galaxies over three equatorial fields totalling 180 sq. degrees (plus two Southern fields we do not consider here). We make use of a number of data products that have been described elsewhere, and have been made public with GAMA Data Releases 2 and 3 (Liske et al. 2015; Baldry et al. 2018), including stellar mass estimates and stellar population parameters (Taylor et al. 2011), group identifications (Robotham et al. 2011), and Sérsic profile fits (Kelvin et al. 2012). We also make use of ultraviolet-plus-total infrared star formation rates (SFRs) described in Davies (2016).

Our primary sample selection is in terms of stellar mass: namely, $10.3 < \log M_* < 10.7$. The stellar mass estimates are based on stellar population synthesis modelling of optical–near infrared spectral energy distributions (SEDs), using the Bruzual & Charlot (2003) simple stellar population models with the Chabrier (2003) prescription for the stellar initial mass function (IMF), and single-screen dust following Calzetti et al. (2000). The GAMA SEDs are all measured in (large) matched apertures on seeing-matched imaging, in order to obtain the best characterisation of SED shape (Hill et al. 2011; Wright et al. 2016). Because the apertures are finite, we re-normalise the SEDs to match the Sérsic model photometry in the r -band, as a measure of total flux.

We limit our sample to the redshift range $0.10 < z < 0.18$. The upper redshift limit, which is to ensure that we have a properly volume-limited sample, has been determined following the arguments given in Taylor et al. (2015); see also Fig. 14 of Baldry et al. (2018). The lower redshift limit is to minimise the impact of photometric redshift errors for nearby sources in the overall lensing measurements², which are not fully formally propagated. To identify and exclude catastrophic errors in the photometry (which sometimes happen when a very bright star disrupts the segmentation and aperture definition) we throw away 185 cases where the Sérsic and aperture photometry do not agree within 0.3 mag. Together with basic quality control cuts for redshift reliability, etc., this gives a parent sample of 11392 galaxies.

In order to ensure that we are truly looking at the primary halos of the particular galaxies in our sample, we have also done our best to select only central galaxies based on the GAMA group catalogues (G³C; Robotham et al. 2011). First, if a galaxy is ungrouped in the G³C then it is taken to be a central by default. Then, the G³C gives two quantities that can be used to identify the central galaxy within a group: RankBCG, which ranks galaxies within each group according to brightness, and RankIterCen, which ranks galaxies

according to their distance from an iteratively re-calculated centre of light. For our analysis, we require either RankBCG = 1 and RankIterCen < 4 or vice versa; that is, we require approximate consistency between the two measures. This gives us our main sample of 7593 central galaxies. Note that both values are precisely 1 for $\approx 93\%$ of our main sample, and that our final sample is $\approx 3\%$ larger than it would be if we defined our sample based on just one of the two measures. In any case, we have repeated our analysis using different central galaxy selections and verified that none of our main results or conclusions change.

2.2 Sample subdivision

Fig. 1 shows the range of properties spanned by the galaxies in our sample — namely: stellar mass, $\log M_*$; intrinsic (*i.e.*, dust corrected) stellar colour, $(g - i)_*$; specific star formation rate (SSFR); Sérsic index, n ; and Sérsic effective radius, R_e . In the diagonal panels of this Figure, the lighter gray histogram shows the distributions for the parent sample, and the black histograms refer to our main sample of central galaxies only. We note that while the effect of excluding satellite galaxies in this mass range is to slightly reduce the relative numbers of generically red, passive, and de Vaucouleurs-like ($n \gtrsim 2$) galaxies, it also leads to a more nearly flat distribution of $\log M_*$ for our main sample.

One of the ways that we will explore halo mass variations will be to split our main sample into subsamples, according to a particular property. The dashed lines in each panel of Fig. 1 show the quartiles for each property within our main sample; *i.e.*, these lines show how to split our main sample into four equally-sized subsamples on the basis of any one property. One nice aspect of selecting this particular mass range is that dividing this sample into quartiles closely aligns with the peaks and saddle of the bimodalities in colour, SSFR, shape, size, etc.

The principal difficulty that we will grapple with in this paper is that, even at fixed mass, many of these properties are closely correlated. There are good astrophysical reasons why a sample of galaxies with blue colours is also likely, in general, to have higher star formation rates, and to have a diskier morphologies (and hence lower values for the Sérsic index, n). This can be seen from the bivariate distributions for the main parameters of interest within our main sample, which are shown in the off-diagonal panels of Fig. 1. As a quantitative description of the overlap between subsamples divided in different ways, the upper-right panels of Fig. 1 remap each property to a dimensionless rank or percentile. Each cell in these panels thus shows how the quartile subdivision in one property projects onto a similar subdivision in the other property; the numbers refer to how many galaxies are found in each cell. For example, it can be seen that while the bluest quarter of our main sample (the lefthand column in the upper $(g - i)_* - \log M_*$ panel) do span the full mass range that we consider, it is nevertheless biased slightly to lower stellar masses (more galaxies in the lower cells than in the upper cells); conversely, the reddest quarter of the sample (righthand column) has slightly higher stellar masses (more galaxies in higher cells). Similarly, the lowest quarter of our sample in stellar mass has on average bluer stellar colours, and the highest quarter in stellar mass has on average slightly redder stellar colours.

By virtue of our decision to focus on only a narrow stellar mass range, the interdependence is stronger between parameters other than mass. This shows how we can effectively control for the mass dependence of the SHMR to isolate second order correlations between halo mass and the dispersion in the SHMR. For the other properties we consider, the interdependencies mean that

² For nearby lenses, the relevant error in the lensing measurement scales with the square of the source redshift; see Eq. (1) and Eq. (2).

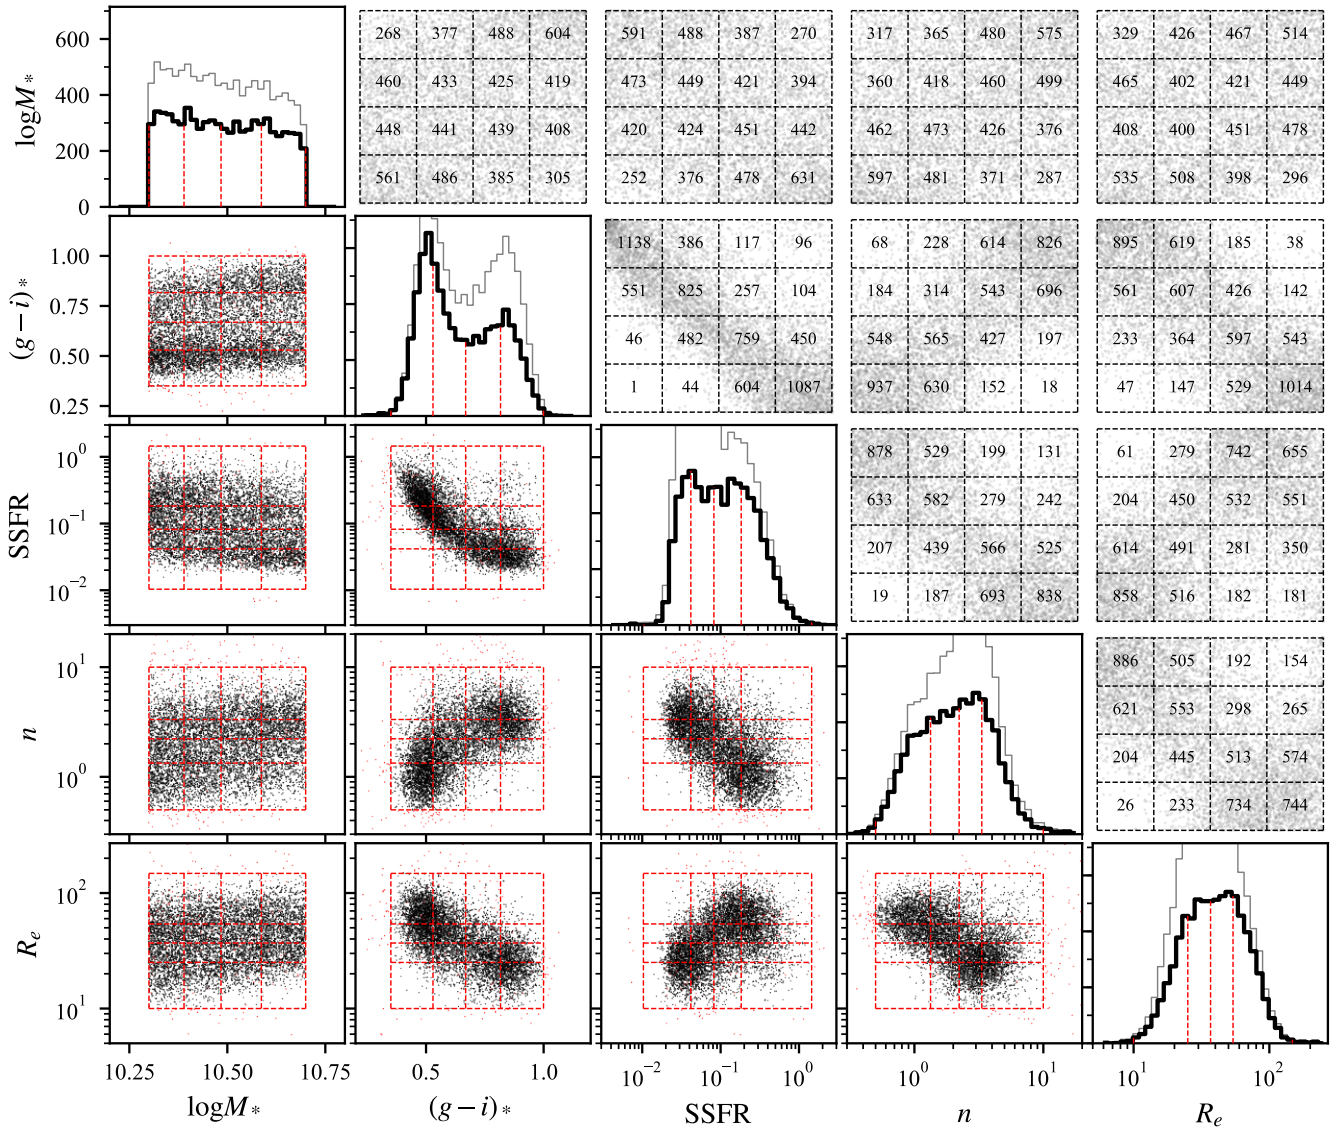


Figure 1. Illustrating the distributions of and correlations between galaxy parameters within our sample of $\log M_* \sim 10.5$ galaxies.— Along the diagonal, panels show histograms of the distributions of: stellar mass, $\log M_*$; intrinsic (dust-corrected) stellar colour, $(g-i)_*$; specific star formation rate, SSFR; Sérsic index, n ; and half-light radius, R_e . The black histograms show our sample of central galaxies only; the lighter gray histogram shows the parent sample, including satellites. In the lower left panels, the black points show bivariate distributions for these parameters. The red lines show how the sample is divided into four equally-sized samples according to each property, after excluding a few outlying values (red points). The upper right panels explicitly show the overlap between these different sample divisions. As discussed in §2.2, the main points to take from this figure are 1.) across this limited mass range, the distribution of other galaxy properties as a function of mass is roughly constant (each cell in the top row is roughly uniformly populated); and 2.) although there is strong covariance between other galaxy properties (there is clear structure in the gray points), there are considerable differences when dividing the sample according to different properties (many galaxies are found in off-diagonal cells).

any ‘true’, causal relation (or relations) between halo mass and one (or more) of these properties will induce ‘spurious’, coincident correlations with other properties. On the other hand: while these interdependencies are significant, they are not total – there are real differences between subdividing this sample according to different properties. This gives us cause to hope that we may be able to distinguish between a genuine correlation between lensing signal and some galaxy property and any ‘spurious’ or tertiary correlations. We return to this point in §5.1.

2.3 Weak lensing measurements from KiDS

The distorting effect of weak gravitational lensing, called shear, is to slightly change the observed ellipticity and position angle of a background source seen in close projection to a nearer lens. The aggregate effect for an ensemble of many background sources is that their sizes are seen to be, on average, very slightly compressed radially and stretched tangentially around the lens.

In the thin lens approximation, and assuming spherical symmetry for the lensing mass, the degree of shear is related to the geometry of the observer–lens–source configuration, and to the mass

distribution of the lens via:

$$\gamma(R) = \Delta\Sigma(R) / \Sigma_{\text{crit}}. \quad (1)$$

The critical surface density, Σ_{crit} , describes the characteristic scale for lensing, and is given by:

$$\Sigma_{\text{crit}} = \frac{c^2}{4\pi G} \frac{D_l}{D_s D_{ls}}, \quad (2)$$

where D_l , D_s , and D_{ls} are the angular diameter distances between the observer and the lens, between the observer and the source, and between the lens and the source, respectively. Σ_{crit} is thus determined entirely by physics and geometry, rather than any intrinsic properties of the lens. The action of the lensing mass distribution is determined by the *excess surface density* (ESD), which can be expressed as:

$$\Delta\Sigma(R) = \widehat{\Sigma}(R) - \Sigma(R). \quad (3)$$

In words, this is the difference between the mean projected surface density within the radius R :

$$\widehat{\Sigma}(R) = \frac{1}{\pi R^2} \int_0^R dR' \Sigma(R') \quad (4)$$

and the projected surface density at that radius, $\Sigma(R)$. All else being equal, the degree of shear thus depends on the *density contrast*, and not the density (or mass) *per se*.

The method for deriving weak lensing measurements from KiDS imaging is described in detail in, e.g., Viola et al. (2015), Hildebrandt et al. (2017), and Dvornik et al. (2018). In brief: the lensing signal is measured as a weighted average of the tangential projections of the background source ellipticities, measured in annuli, to build up a lensing profile for each lens; *viz.*:

$$\Delta\Sigma_{R,i} = \left(\frac{\sum_j \tilde{w}_{ij} \varepsilon_j \tilde{\Sigma}_{\text{crit},ij}}{\sum_j \tilde{w}_{ij}} \right) \frac{1}{1 + K(R)}. \quad (5)$$

Here, $\Delta\Sigma_{R,i}$ is the ESD profile for the i th lens measured in an annulus with radius, R ; ε_j is the tangential projection of the shape tensor for the j th source; \tilde{w}_{ij} is the weight given to each source according to its ellipticity and the lens–source geometry; $\tilde{\Sigma}_{\text{crit},ij}$ is the effective critical surface density for the lens/source pair ij ; and K is a small ($\lesssim 10\%$) scalar correction to account for the multiplicative ‘noise bias’ and ‘weight bias’ in the overall shear inferred from the optimally weighted shapes of small and/or low signal-to-noise galaxies. The Fenech Conti et al. (2017) calibration of the K s used here has recently been updated by Kannawadi et al. (2019), but the differences are negligible for our purposes.

We use the KiDS galaxy–galaxy weak lensing pipeline (see, e.g. Dvornik et al. 2018) to obtain ESD profiles for each of the galaxies in our lens sample, based on the KiDS-450 catalogues (Hildebrandt et al. 2017). In constructing ESD profiles for the galaxies in our lens sample, we measure the ESD in 20 concentric annuli around each lens, with a logarithmic spacing between annulus edges in the range 12–2000 kpc, and only considering background sources with $z_s > 0.2$ and $(z_s - z_l) > 0.1$. These ESD profiles — one for each of the lens galaxies we consider — represent the weak lensing dataset that we analyse in this paper. While the signal-to-noise ratio in any one ESD profile is very low, by considering them aggregate, we can hope to derive information about the global lensing properties of the ensemble.

In principle, because the same background source galaxy can contribute to the ESD profile for multiple lenses, the errors in the

ESD profiles of different lenses can be significantly correlated, particularly where there is a high density of lenses on the sky. In practice, for our specific and relatively small sample of $\log M_* \sim 10.5$ galaxies, this covariance is negligible: the Pearson correlation coefficients for ESD profile measurements at different radii for different lenses are $\lesssim 0.001$. This should not be surprising given the relatively low sky densities of our lens sample: $\sim 40 / \text{deg}^2$, over three independent fields. We therefore ignore the covariances between the ESD measurements for the lenses in our sample, which makes the computation described in the next section tractable.

2.4 Lens modelling and mass estimation

Eq. (1) shows how, knowing both the geometry of the lens–source system and the projected mass distribution of the lens, one can predict the shear. Conversely, given observational constraints on both the geometry and the shear, one can infer the lens mass distribution.

We assume that the halos can be described using the NFW mass distribution (Navarro, Frenk & White 1997), which is fully described by its mass³, M_{halo} , and a shape parameter, c , which is usually referred to as the concentration. Analytic expressions for the values of $\widehat{\Sigma}$, Σ , and/or $\Delta\Sigma$ for the NFW profile are given in, e.g., Wright & Brainerd (2000) or Coe (2010). We find that, in general, we cannot place strong constraints on the values of c using our data (see §5.3 where we explore this issue in greater detail). For this reason, and following van Uitert et al. (2015), we adopt a fiducial mass-concentration relation based on Duffy et al. (2008):

$$c = f_{\text{conc}} \cdot 10.14 \left(\frac{M_{\text{halo}}}{2 \cdot 10^{12} h^{-1} M_{\odot}} \right)^{-0.089} (1+z)^{-1.01}, \quad (6)$$

where the scaling factor $f_{\text{conc}} = 0.70$ comes from the lensing-plus-stellar mass function modelling of van Uitert et al. (2015).

For small values of R , the stars within the galaxy itself can make a non-negligible contribution to the observed shear. In the thin lens approximation, the shear from the stars is simply added linearly to that from the halo. Knowing each galaxy’s total stellar mass, circularised effective radius, and Sérsic shape parameter, we can approximately account for this using a circularised Sérsic profile to describe the stellar mass distribution of the galaxy. The relevant assumptions are 1.) that the Sérsic parameters derived from fits to the r -band images can be used to describe the stellar mass distribution, and 2.) that a circularised model is sufficient for the purposes of computing the lensing shear. Analytic expressions for the values of $\widehat{\Sigma}$ and Σ for the circularly-symmetric Sérsic profile can be found in, e.g., Graham & Driver (2005).

Putting these two pieces together — *i.e.*, knowing the Sérsic parameters for the galaxy, and given a trial value for the halo mass — we can then generate a model ESD profile to compare to the data. The goodness of fit for the ensemble is given by the summa-

³ Following the convention for N -body dark matter simulations, M_{200} is defined as the mass enclosed within a radius R_{200} from the halo centre, such that the mean mass density within R_{200} is 200 times the cosmological mean matter density at that redshift, *i.e.*, $\bar{\rho} = 200 \Omega_m(z) \rho_{\text{crit}}$. With this definition, the shape parameter c can be related to a scale radius R_s via $c = R_{200}/R_s$. Note that for clarity, elsewhere in the paper we will use the symbol M_{halo} and just ‘the halo mass’; this should properly be understood as referring to this proxy measurement of the total or virialised halo mass.

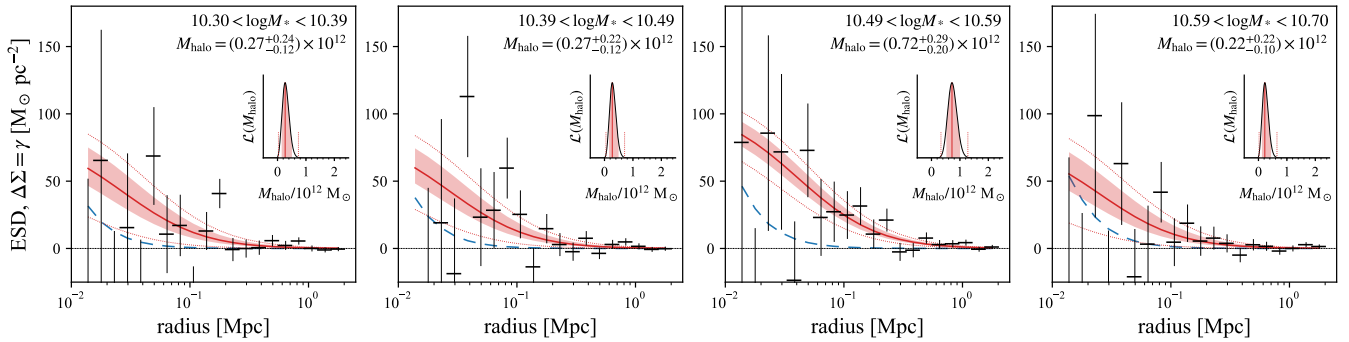


Figure 2. Stacked lensing profiles and halo mass modelling, having subdivided our sample according to stellar mass, $\log M_*$. — In each panel, the mean lensing signal from the halo is shown in black; in each case, these are based on the KiDS lensing measurements for approximately 1600 lenses. The blue dashed line shows the lensing contribution from the stars, based on the GAMA Sérsic fit parameters for the lens galaxies. These values have been subtracted from the observed ESD profiles to isolate the effect of the halo. The red shows the inferred NFW halo model: the heavy solid lines show the maximum likelihood fits; the shaded regions show the equivalent of the $\pm 1\sigma$ uncertainties; the thin dotted lines bound the 95 % confidence region. Inset within each panel is the posterior PDF for the mean halo mass for each bin. Note that for each subsample, the uncertainty on the halo mass estimate is $\approx 0.2 \times 10^{12} M_\odot$.

tion over all radii and for all lenses:

$$\chi^2(M_{\text{halo}}) = \sum_{R,i} \left[\Delta\Sigma_{R,i} - \Delta\Sigma_{\text{Sersic}}(R | M_{*,i}, R_{e,i}, n_i) - \Delta\Sigma_{\text{NFW}}(R | M_{\text{halo}}) \right]^2 / \sigma_{R,i}^2, \quad (7)$$

where $\sigma_{R,i}$ is the formal uncertainty associated with the measured $\Delta\Sigma_{R,i}$. The (log) likelihood function for the ensemble is then just:

$$\ln \mathcal{L}(M_{\text{halo}}) = -\frac{1}{2} \chi^2(M_{\text{halo}}). \quad (8)$$

This can be evaluated for the sample as a whole, or for any specific subset of the sample.

Within the framework of Bayesian statistics, the quantity of interest is not the likelihood *per se*, but instead the posterior probability distribution function (PDF) for the value of M_{halo} , which is given by the product of the likelihood and an assumed prior on M_{halo} . Here, we adopt a prior that is flat in M_{halo} (and not, say, flat in $\log M_{\text{halo}}$) which means that the PDF is directly proportional to \mathcal{L} as defined above. The effect of choosing a prior that is flat in $\log M_{\text{halo}}$ would be to make the PDF proportional to $\mathcal{L}/M_{\text{halo}}$. With the data used here, this $1/M_{\text{halo}}$ up-weighting blows up more rapidly than the value of \mathcal{L} drops for values of $M_{\text{halo}} \ll 10^{11} M_\odot$. The result is a pathological effect where the PDF diverges for small values of M_{halo} . Quite aside from this point, this choice of prior is also natural since, all else being equal, our observables (*i.e.*, the ESD and the shear) scale linearly with mass, and our fits are framed in terms of linear ESD. In other words, our decision to use a flat prior in M_{halo} is both important and reasonable.

Looking at Eq. (7) and Eq. (8), it can be seen that if all i galaxies are assumed to have the same halo mass, then the summations over i need only be done once. If this summation is done in advance then the computation of the likelihood function only involves the single comparison of one co-added ESD profile, representing the ensemble in aggregate, and one for the model. This is the standard approach of ‘stacking’, which has the well known limitation that it is necessarily limited to considering the mean properties of the larger ensemble (and always with the implicit assumption of Gaussian statistics).

For this reason, we will also pursue a different approach: we

will fit each lens in the full ensemble simultaneously, while allowing each lens to have its own particular mass. We do this by adopting a simple (linear) parametric prescription to predict halo mass from some other property (*e.g.*, colour, size, etc.), here denoted x :

$$M_{\text{halo},i}(x_i | A, b) = A(x_i - x_0) + b, \quad (9)$$

Being free to choose any value of the arbitrary reference value, x_0 , we choose x_0 to be the mean value of x for the ensemble, so as to minimise the covariance between the parameters A and b . We neglect observational errors on the x quantity, essentially because properly accounting for these errors is computationally expensive. The effect of this decision will be to (weakly) systematically bias our results towards apparently weaker correlations; *i.e.*, lower values of A .

The only other additional complication is whether and how to accommodate negative values for $M_{\text{halo},i}$. Even though, physically, mass is a strictly positive quantity, from an experimental standpoint it is perfectly reasonable to obtain a negative measurement where the errors are comparable to the actual value. In an attempt to mitigate biases in our results, we allow the parameter M_{halo} to take negative values. Our scheme for treating negative values of M_{halo} is very simple: we use the absolute value of M_{halo} to determine the shape of the ESD profile, and then reverse the sign of the shear in the case of a negative mass, so that $\Delta\Sigma(-M_{\text{halo}}) \equiv -\Delta\Sigma(M_{\text{halo}})$. The motivation and rationale for this decision are discussed further in §3.2.

Rather than fitting for the value of a single parameter, M_{halo} , for the ensemble, we now allow each lens to have its own unique halo mass, $M_{\text{halo},i}$, which is derived from the value of x for that lens, x_i , and the free parameters A and b . We adopt the standard ‘non-informative’ or reference prior, which is flat in both A and b . Our likelihood function then becomes:

$$\ln \mathcal{L}(A, b) = -\frac{1}{2} \chi^2(A, b), \quad (10)$$

with the only other formal modification being to replace the M_{halo} that appears in Eq. (7) with $M_{\text{halo},i}(x_i | A, b)$ as defined in Eq. (9). Formally, the approach outlined above is equivalent to fitting for the ensemble of values $M_{\text{halo},i}$ all lenses in the sample (which are in general poorly constrained on their own), and then fitting a linear relation to those results. Seen through this lens, our approach is not

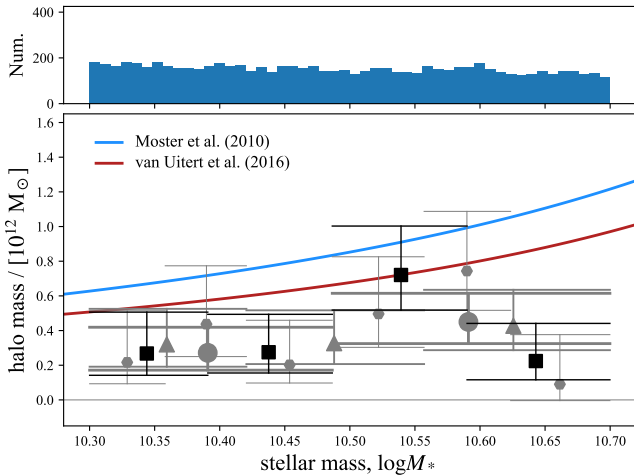


Figure 3. Inferred halo masses from stacked ESD profile fitting for our sample sub-divided by stellar mass.— The black points show the inferred halo masses for the four equally populated bins shown in Fig. 2; the gray points show what we get when splitting our sample into two (circles), three (triangles), or six (hexagons) equally sized bins. The horizontal length of the errorbar ‘caps’ show the extent of each bin, and the size of each point reflects the number of lenses that have gone into each stack. The points are the maximum likelihood values, and the (asymmetric) error bars reflect the 68% confidence interval in the PDF for M_{halo} for each stack.

to track the constraints on the individual $M_{\text{halo},i}$ values, and instead marginalise over these unknowns as nuisance parameters.

The practical consequence of this change is that we cannot pre-compute the summation over i in the definition of the χ^2 : it is necessary to compute a separate model–observed ESD comparison for each individual lens, and for every set of (A, b) trial values. But the benefit justifies the cost: by allowing each lens to have its own halo mass, it becomes possible to make inferences about the distribution of halo masses across the ensemble — which, after all, is our primary goal in this paper.

3 PROOF OF CONCEPT AND SOME SIMPLE SANITY CHECKS

3.1 The stellar-to-halo mass relation across our sample

As a demonstration of our ability to identify and measure variations in halo mass across our sample, we first consider the correlation between stellar mass and halo mass. A secondary motivation here is to check that the results we obtain are broadly consistent with existing results. From the outset, however, we make the disclaimer that past determinations have been derived following very different methods, and our ability to constrain the relation between stellar and halo mass with our sample is limited by our deliberate decision to focus on a narrow mass range.

3.1.1 Stacked lensing profiles

Fig. 2 shows the mean halo lensing profiles for our lens sample, subdivided by stellar mass into four equally populated bins. In each panel of Fig. 2, the blue dashed line shows the mean lensing profile for the stellar component, which is derived from the combination of the SED-fit mass-to-light ratios and Sérsic fit parameters from

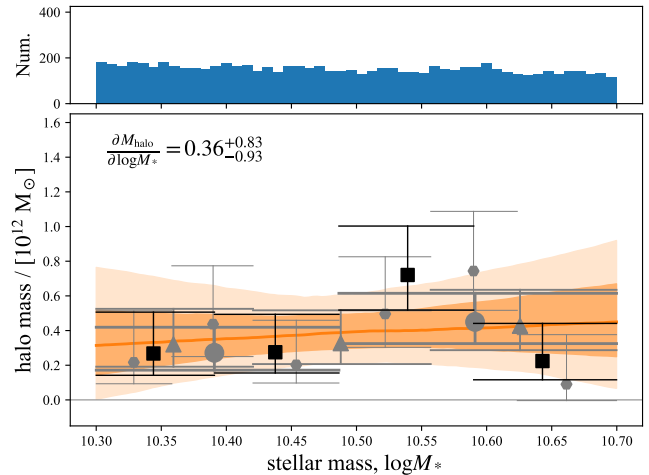


Figure 4. Inferred linear correlation between halo mass and stellar mass across our sample.— The solid yellow line shows the best fit linear relation between $\log M_*$ and M_{halo} , based on simultaneous fits to the KiDS ESD profiles for our sample of lenses. The darker and lighter shaded regions show the percentile equivalents of the ± 1 and $\pm 2\sigma$ regions for the fit. The points and error bars are simply repeated from Fig. 3; the line shown is *not* simply a fit to these points. Instead, the line and the various points are both different visualisations of the same underlying lensing behaviour of the galaxies in our sample.

GAMA. These values have been subtracted from the KiDS lensing measurements, shown in black, to isolate the lensing effect of the halo. We note that accounting for lensing by the stars reduces the inferred lensing mass by ~ 10 – 15% ; *i.e.*, by ~ 3 – 5 times the actual stellar mass.⁴ Based on this, we estimate that residual systematic errors in the values of the halo mass due to the influence of the stars within the lens galaxies’ is limited to the few percent level: even if the errors in our accounting for lensing by the stars are as bad as $\sim 20\%$, the impact on the inferred halo mass would be just 2–3%; *i.e.*, negligible in comparison to the statistical uncertainties.

The maximum likelihood estimates for each halo ESD profile are shown in Fig. 2 as the red lines, with the red shaded regions showing the equivalent of the $\pm 1\sigma$ uncertainties (*i.e.*, the 68% confidence interval), and the red dotted lines bounding the 95% confidence region.

The inset panels show the posterior PDF for the mean halo mass for the sample, with the same confidence limits marked in a similar fashion. Note that simple Gaussian statistics are not always a good way to represent our results, particularly where there is small but non-zero likelihood that $M_{\text{halo}} \leq 0$. (Recall that we do allow the parameter M_{halo} to be negative; see §2.4 and §3.2, below.) For example, in the first panel, the PDF for the lensing halo mass can be seen to be slightly skewed towards higher masses. For this reason, we will always show the (generally asymmetric) percentile equivalents of the $\pm 1\sigma$ errors.

⁴ The reason for this discrepancy comes down to the different ESD profile shapes for the stars and the halo, and the different lensing-signal-to-mass ratios, $\Delta\Sigma(R)/M$, of the two components at small radii. In other words, for the same mass, the lensing signal at small radii from the stars is a factor of a few stronger than from the halo.

3.1.2 Measuring mean halo mass in bins of fixed stellar mass

Fig. 3 shows the halo masses we derive from our stacked ESD profile fits for our sample sub-divided according to stellar mass. The black points highlight the four equally-sized mass bins shown in Fig. 2. But we could just as well have split our sample into fewer or more bins: the gray points show what we get splitting our sample into two, three, or six equally populated bins. (For reference, the histogram in the upper panel shows the observed $\log M_*$ distribution for our sample.) Naturally, the uncertainties are larger for the smaller and less populated bins. Since each set of points is depicting the same data, it is not surprising that they all show a more or less consistent picture. That is, these binned results are a useful means for visualising the average relation between halo mass and stellar mass across our sample.

For comparison, we show the SHMR determination from van Uitert et al. (2015), which is derived from halo modelling of the a joint GAMA+KiDS dataset over 100 square degrees, and also the one by Moster et al. (2010), which is derived from halo occupation modelling. Our directly inferred halo mass measurements are slightly but systematically low compared to the two SHMRs shown here: where we might have expected a mean halo mass of $\sim (1 \pm 0.1) \times 10^{12} M_\odot$, the mean measured value for our sample is $(0.4 \pm 0.1) \times 10^{12} M_\odot$. It is beyond the scope of this paper for us to resolve this apparent disagreement, but we do note that both of the SHMR determinations shown are based primarily on the stellar mass function constraints, rather than direct observational constraints on galaxies' halo masses in this mass range. While the uncertainties are significant, and especially given the very different approaches to measuring these quantities, it is encouraging that our results are at least broadly consistent, to within a factor of 2.5.

3.1.3 Quantifying the relation between stellar mass and halo mass by fitting to the ensemble

If we wanted to quantify the relation between stellar and halo mass using our data, we could imagine fitting a linear relation to any one of the different sets of binned and stacked halo mass measurements shown in Fig. 3. In practice, however, the answer we would get would depend somewhat on what binning we chose to adopt. This problem is compounded by the fact that the uncertainties on each point is manifestly non-Gaussian, so the error propagation would be non-trivial. It is for these reasons that we do not base our analysis on the binned and stacked ESD profiles, but instead perform simultaneous fits to the many independent ESD profiles for the lenses in our sample, as described in §2.4.

The results of this fit are shown in Fig. 4. We find a slightly shallower slope to the SHMR across the narrow range of our sample than either Moster et al. (2010) or van Uitert et al. (2015): we find $\partial M_{\text{halo}}/\partial \log M_* = 0.36^{+0.83}_{-0.93}$, as opposed to a value nearer 3. This tension can be largely alleviated by allowing for the factor 0.4 difference in mean halo mass, which would bring our two results into agreement at the 1σ level. Again, we take it as encouraging that we find consistency to within a factor of a few between these very different approaches.

3.2 Some simple sanity checks, and the potential for bias

Before we move on to searching for galaxy parameters that show statistical correlations with halo mass, it is worthwhile to demonstrate a null result. In Fig. 5, we show the inferred dependence of

halo mass on several parameters that we would expect to be completely unrelated to halo mass: namely, declination, and a random value. We also show the inferred relation between halo mass and both redshift and apparent magnitude. If there is no significant evolution across our $0.10 < z < 0.18$ redshift window, and if our sample is properly volume limited (*i.e.*, that the stellar populations across our sample do not vary with redshift), then we would expect to find $\partial M_{\text{halo}}/\partial z = 0$, and $\partial M_{\text{halo}}/\partial m_r = 0$. In each of the four cases shown, while the uncertainties are large, the measured values do conform to these simple expectations.

Fig. 5, and particularly the panel showing the inferred correlation between M_{halo} and apparent magnitude, is also useful to illustrate how our scheme for allowing the value of M_{halo} to be negative mitigates potential biases in our results. As shown in this Figure, the data are sometimes consistent with linear relations that would imply negative halo masses. Taken at face value, negative masses would seem to be unphysical, but this is totally consistent with a low signal-to-noise measurement of a strictly positive quantity.

We have considered two alternatives to our preferred approach for accommodating negative values for M_{halo} (which is simply to define $\Delta\Sigma(-M_{\text{halo}}) \equiv -\Delta\Sigma(M_{\text{halo}})$; see §2.4). The first would be to restrict the allowed region of (A, b) parameter space to forbid any non-positive values of $M_{\text{halo},i}$. Looking at Fig. 5, the problem with this approach becomes clear: the allowed range of fit parameters would become very sensitive to the furthest outlying point and/or the precise limits over which the fitting is done. Looking at the trend in halo mass as a function of apparent magnitude, for example, we would get a very different answer if we required M_{halo} to be positive for $r < 17$, or 16, or 12. The net result would be a potentially strong bias against large values of A and/or small values of b .

At the other extreme, we could simply place a floor of zero on the halo mass values, or, equivalently, say that $\Delta\Sigma(M_{\text{halo}}) = 0$ for $M_{\text{halo}} \leq 0$. With this decision, any observed ESD profile would be equally well consistent with $M_{\text{halo}} = -10^{14} M_\odot$ as 0, and only data in the $M_{\text{halo}} > 0$ regime would contribute to the constraints on the values of A and b . To the extent that approach would allow us to consider steeper gradients, it would mean that we would be more likely to overestimate any correlations than to underestimate them: that is, the net result would be a potential bias towards larger values of A and/or small values of b .

With these simple arguments, we motivate our specific approach to accommodating negative values of M_{halo} as being intermediate between these two extremes, and with the hope that any bias in our results that arise as a consequence of this decision are small.

4 RESULTS: EXPLORING CORRELATIONS BETWEEN HALO MASS AND GLOBAL GALAXY PARAMETERS

Our basic results are shown in Fig. 6. These panels show the inferred variation in halo mass as a function of several key galaxy properties: intrinsic stellar colour, $(g - i)_*$, which is a proxy for light-weighted mean stellar age; specific star formation rate, SSFR; Sérsic shape parameter, n ; and half-light radius, R_e . We note that we have also considered a number of other galaxy properties, including effective colour (*i.e.*, without correcting for internal extinction), ellipticity, star formation rate, $H\alpha$ equivalent width, mass-to-light ratio, and others, which are not shown. We have chosen to focus our attention on the particular properties shown here either because we consider these properties to be more ‘fundamental’ than

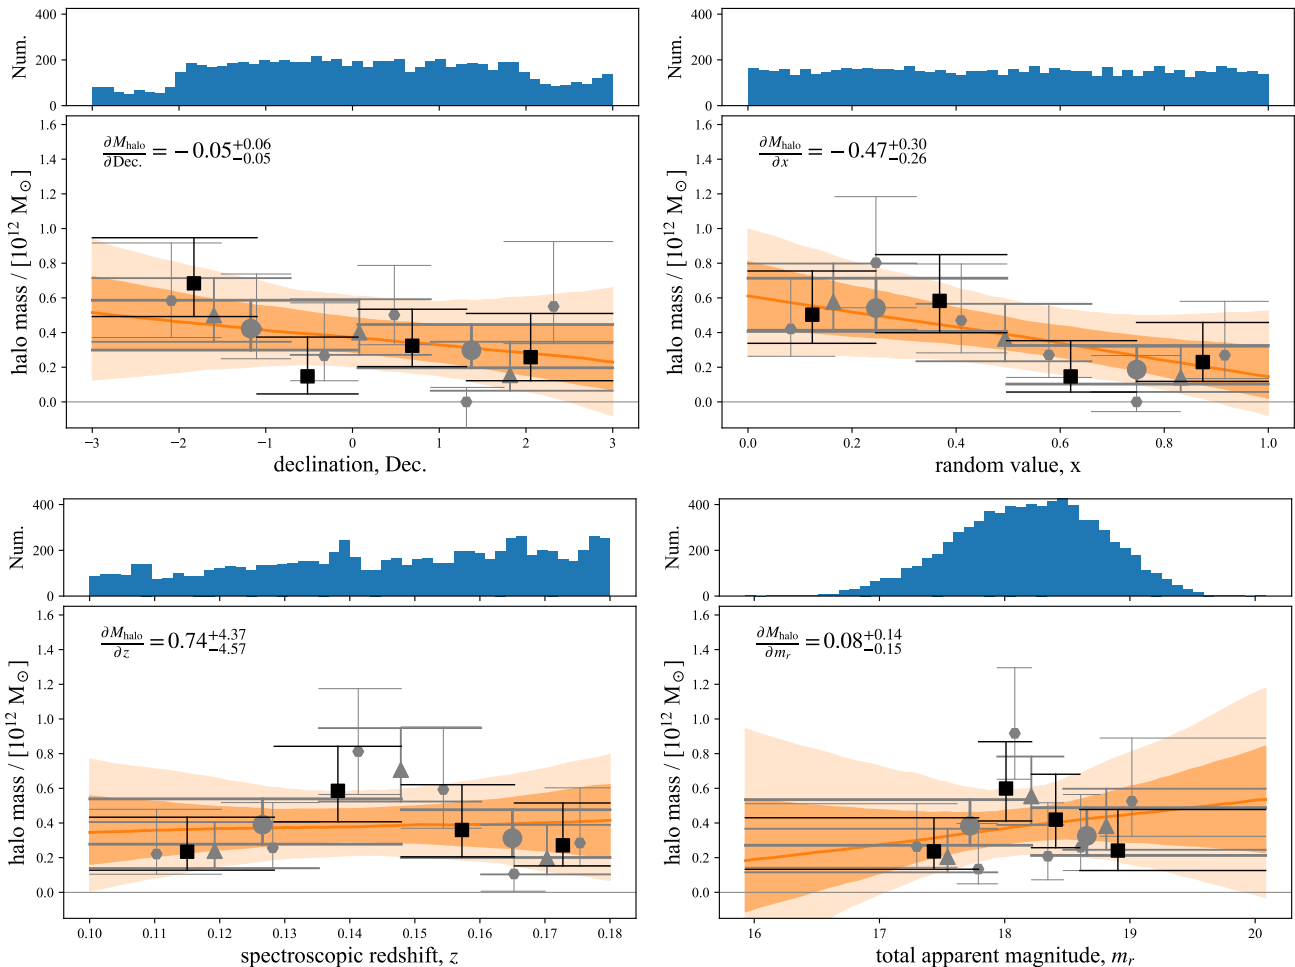


Figure 5. Some simple sanity checks for our analysis.— Each plot shows the apparent variation in halo mass as a function of some property that is expected to be unrelated to halo mass. As in Fig. 3 and Fig. 4, the histograms in the upper panels show the distributions of each property across our sample. The fact that the observed dependence of halo mass on each of these quantities is consistent with zero is reassuring.

other related properties (*e.g.*, preferring intrinsic stellar colour to effective colour or Dn_{4000} , or SSFR to SFR or $H\alpha$ equivalent width), or because we see no evidence for a correlation with halo mass.

The panels in Fig. 6 are analogous and directly comparable to Fig. 4 and Fig. 5. The upper panels show the distribution for each property across the full sample. The points show the maximum likelihood values for the mean halo mass, in bins of the quantity in question; the horizontal error bars show the width of each bin; the vertical error bars show the (asymmetric) $\pm 1\sigma$ confidence intervals for the inferred halo mass. The lines show the inferred correlation between halo mass and the quantity in question; the heavier and lighter shaded regions show the equivalent of the 1 and 2 σ range for these fits, as a function of the property in question; the inferred values for the slope of the linear fit are also given with 1 σ uncertainties in each panel.

Our first and most basic observation is that *even over the narrow range of stellar masses covered by our sample, there is significant variation in galaxies’ halo masses*. Taking the stacked-by-quartile results in Fig. 6 at face value, the observed minimum-to-maximum variation in stellar-to-halo mass fractions across the sample is at least 1 dex ($0.1\text{--}1.2 \times 10^{12} M_{\odot}$).

Further to this, *at (approximately) fixed stellar mass, there are clear correlations between global galaxy properties and halo mass*.

More specifically, we see that canonically ‘early type’ galaxies (*i.e.*, red, quiescent, or elliptical galaxies) have larger halo masses than ‘late types’ (*i.e.*, blue, star forming, or disk galaxies). In general terms, this result is consistent with, *e.g.*, Hoekstra et al. (2005), Mandelbaum et al. (2006), and others, who find offset SHMRs for generically red/early versus blue/late samples. Our results offer additional detail and insight, by beginning to map the relations between these properties and halo mass at fixed stellar mass.

The obvious next question is: which property or properties are most closely correlated with halo mass? This is the question that occupies the rest of this paper.

Considering first the empirical correlation between halo mass and intrinsic stellar colour, $(g - i)_*$, the binned results appear broadly consistent with all blue ($(g - i)_* \lesssim 0.75$) galaxies in the sample having $M_{\text{halo}} \approx 2\text{--}4 \times 10^{11} M_{\odot}$, and red ($(g - i)_* \gtrsim 0.075$) galaxies having $M_{\text{halo}} \approx 8 \times 10^{11} M_{\odot}$.

Let us entertain this scenario for a moment to tease out an important aspect of our approach. While there is nothing preventing us from fitting a linear relation between any two properties (*e.g.*, to quantify the strength of the correlation), it should also be recognised that there is no guarantee that the result of such a fit will provide a good description of reality. In order to have confidence that the linear fit provides a faithful description of the underlying

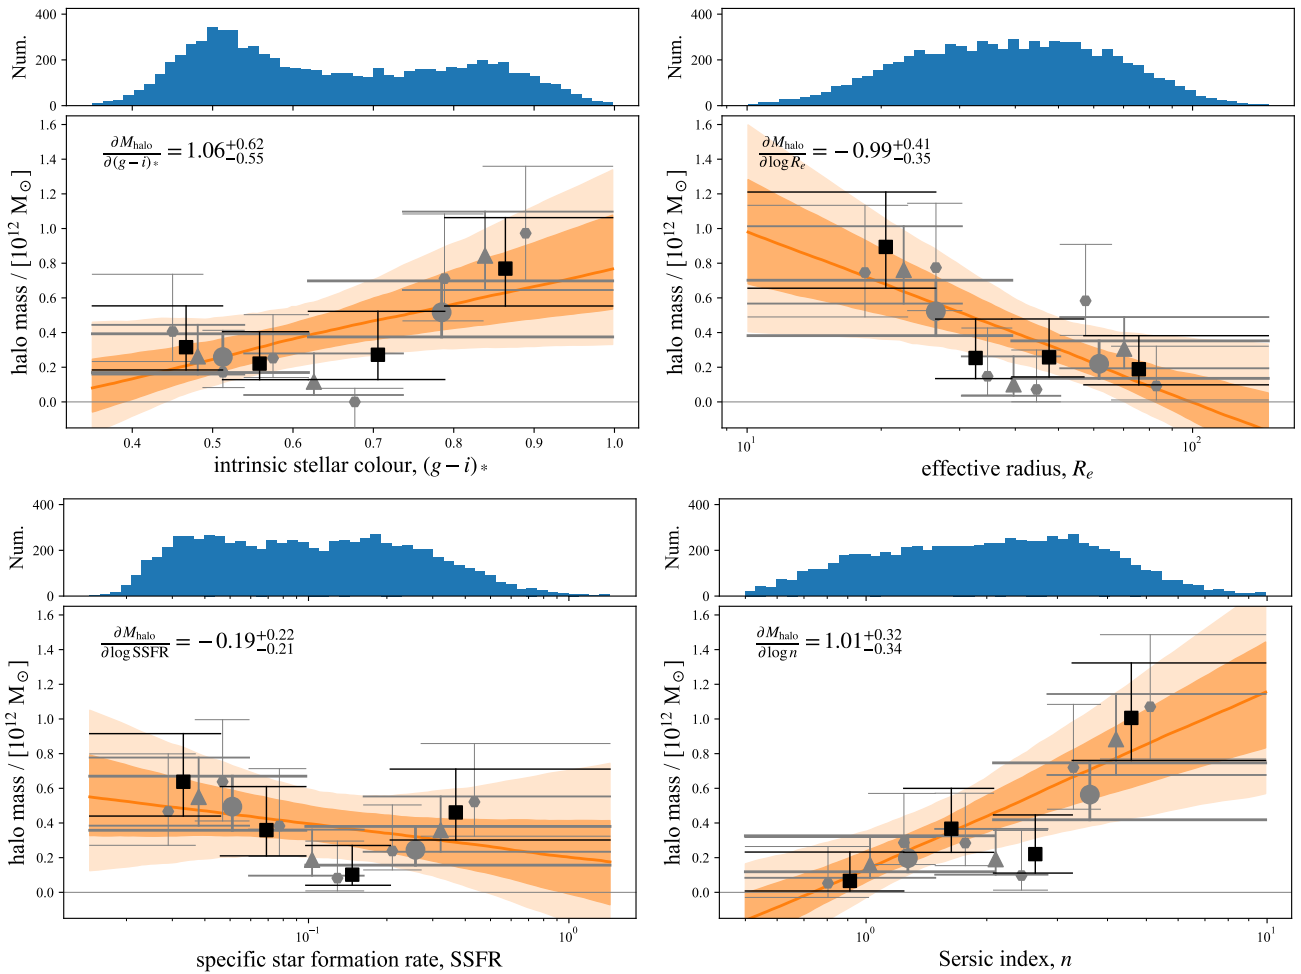


Figure 6. Exploring variations in halo mass as a function of other key galaxy observables.— The different panels show the apparent trend in the mean halo mass as a function of stellar colour, specific star formation rate, Sérsic index, and effective radius. Focussing on the points, which show the mean inferred halo mass in bins, there is a clear and apparently linear correlation between halo mass and Sérsic index, that is in close accord with the linear fit to the sample as a whole. For the other properties shown here, we do see significant correlations with halo mass across the sample, but the binned results are less consistent with a simple linear dependence.

data, therefore, what we are looking for is consistency between the the fits and the binned results. Conversely, where the binned results are not consistent with the linear fits, this is a sign that the relationship between these two properties is more complex than a simple linear correlation. That is why we show both the binned values as well as the linear fits: as complementary representations of the same underlying data. At the same time, recognising the size of the formal uncertainties in these panels, we acknowledge the real risk of over-interpreting the data in this regard.

Looking next at the empirical correlation between halo mass and effective radius, R_e , the indication is that galaxies with $R_e \gtrsim 4$ kpc (*i.e.*, the canonical ‘late types’, which have larger sizes, bluer colours, and lower Sérsic indices) have approximately constant halo masses $\approx 3\text{--}5 \times 10^{11} M_\odot$. There is possibly the suggestion from the binned results that there is a strong correlation between effective radius and halo mass for canonical ‘early types’, such that more compact galaxies have higher halo masses, but this cannot be taken as anything more than suggestive.

Parenthetically, we note that the inverse correlation between M_{halo} and R_e seen in Fig. 6 would seem, on its face, to be counter to the results in Charlton et al. (2017), who find a positive correla-

tion between the offsets from the SHMR and size–mass relations: averaging over $9 \lesssim \log M_* \lesssim 11.5$ and $0.2 < z_{\text{phot}} < 0.8$, their result is $\Delta \log M_{\text{halo}} = (0.42 \pm 0.16) \Delta \log R_e$. Noting the major differences in how our results are derived (volume limited vs. magnitude selected lens samples; spectroscopic vs. photometric redshifts used for lens selection and characterisation; narrow vs. broad redshift windows; exclusion vs. inclusion of satellites), we cannot hope to uniquely identify the cause for this apparent tension. That said, we do highlight that the Charlton et al. (2017) results are derived in bins of absolute magnitude *and* effective colour, with both offsets determined separately for the red and for blue galaxies subsamples. In this sense, the Charlton et al. (2017) results are perhaps better viewed as probing *third-order* correlations between M_{halo} and R_e at fixed mass and colour — and assuming that colour is the correct choice for the second-order term.

Returning to the main discussion, the empirical correlation between halo mass and specific star formation rate is less impressive than the other parameters we consider. There is maybe a hint that there is a correlation between halo mass and SSFR for actively star forming galaxies, and maybe also a slight inverse correlation for more quiescent galaxies. The fact that the variation in halo mass

seen in this panel is less than in some others imply that SSFR does not play a primary role in predicting halo mass (and vice versa). Instead, we judge it more likely that the empirical correlation seen is a spurious correlation induced by correlation between both SSFR and halo mass and some other parameter(s).⁵

Turning finally to the empirical correlation between Sérsic index and halo mass, here the binned data are most consistent with a simple linear correlation between $\log n$ and halo mass. This is also the most significant (in a statistical sense) correlation that we see in our dataset.

In summary, then, we have three possible ways of understanding our data:

- (i) a direct correlation between halo mass and galaxy concentration, as quantified by Sérsic index, n .
- (ii) a correlation between halo mass and size for early type galaxies, with late types having approximately constant halo masses; or
- (iii) a bimodal distribution of halo masses, which is tied to the bimodality in intrinsic stellar colour, $(g - i)_*$.

Or, of course, any combination of these three. While we cannot unambiguously discriminate between these scenarios with the present data, we discuss potential interpretations of these results further in the next section.

5 DISCUSSION: PROBING THE APPARENT CONNECTION BETWEEN GALAXY PROPERTIES AND HALO MASS

5.1 Quantifying the connection between halo mass and galaxy properties

The essential idea that we explore in this section is the extent to which, at fixed mass, the scatter around the mean/median stellar-to-halo mass relation (SHMR) is directly coupled to one (and only one) particular galaxy property. Our motivation here is that by finding the galaxy property that is most directly associated with halo mass, we then have circumstantial evidence for the mechanism(s) by which halo mass influences galaxy evolution. Such a finding would also have implications for, *e.g.*, cosmological studies where assembly bias is relevant.

Given the tight correlations between many different galaxy observables, how could we hope to identify such a property? The simplest approach—and the only one that we might hope our data to support—is to find the property that implies the largest spread of halo masses across the sample. To see this, imagine that there is a tight correlation between halo mass and some galaxy property, x , and that this property x is itself correlated with some other property, y , but with some scatter. To the extent that binning by y results in mixing galaxies with different values for x —and by extension, M_{halo} —in each bin, the effect will be to drive the mean value of M_{halo} in each bin towards the average value for the sample as a

whole. (Think of the central limit theorem.) The result is, therefore, that the observed (mean) trend in halo mass as a function of y will be less than that as a function of x .⁶

Another way of framing this question is: what galaxy property is the best predictor of halo mass across our sample of $\log M_* \approx 10.5$ galaxies? By writing $M_{\text{halo}} = A(x - x_0) + b + \epsilon$ where ϵ represents the offset from the mean relation between M_{halo} and x , we can see the net dispersion in the SHMR as the sum (in quadrature) of two terms. The first term is tied directly to the gross trend between M_{halo} and x : *viz.*, $A\sigma_x$, where σ_x represents the distribution of x across the sample via the RMS value. The second term represents some ‘random’ or unknown dispersion around that mean relation via the RMS value of the (unknown) ϵ s. Since the net dispersion is a finite quantity determined by astrophysics, it can be considered fixed. The best predictor of M_{halo} is the quantity for which the ϵ s are minimised, or, conversely, where the value of $A\sigma_x$ is maximised. The quantity $A\sigma_x$ can also be seen as providing an approximate lower limit on the true dispersion in the SHMR, inasmuch as any intrinsic scatter around the relation between M_{halo} and x will propagate through to a greater net dispersion.

With this in mind, we use the implied dispersion in the SHMR as a quantitative basis for comparing the strength or significance of the correlation between halo mass and other observables. In our motivating remarks above, which assume Gaussian statistics, we have used M_{halo} . Since the SHMR dispersion is expected to be log-normal, it makes more astrophysical sense to consider the Gaussian dispersion in $\log M_{\text{halo}}$; *i.e.*:

$$\sigma_{\log M_{\text{halo}}} \gtrsim \frac{\partial \log M_{\text{halo}}}{\partial x} \sigma_x = \ln 10 \frac{\partial M_{\text{halo}}}{\partial x} \frac{\sigma_x}{\langle M_{\text{halo}} \rangle}, \quad (11)$$

where values for $\partial M_{\text{halo}}/\partial x = A$ and $\langle M_{\text{halo}} \rangle = b$ come directly from the MCMC chains for our modelling. We note that since this quantity is pure scalar (*i.e.*, is dimensionless), it is robust to any gross systematic errors in our halo mass measurements that might arise due to choice of cosmology, halo modelling, etc. For this reason, we strongly recommend the use of this parameter for comparisons between our observational results and the results of other studies or models.

The inferred values for the SHMR dispersion, derived in this way, are given in Table 1 for all of the properties we have considered. Taking these estimated dispersions at face value, we would conclude that the property most directly related to halo mass is Sérsic index, with an implied SHMR dispersion $\gtrsim 0.28$ dex. The strength of the correlations between halo mass and effective radius and intrinsic stellar colour are only slightly weaker, with an implied SHMR dispersion $\gtrsim 0.24$ dex and $\gtrsim 0.20$ dex, respectively.

5.2 Is it really not just stellar mass?

Our analysis specifically focuses on a narrow range of stellar mass, in an attempt to identify correlations with halo mass at fixed stellar mass. The formal random errors in the stellar mass estimates that our selection is based on are typically 0.12 dex, which is not negligible in comparison to our 0.4 dex selection window. The particular concern here would come from Eddington-like biases:

⁵ This view is informed by a series of numerical experiments where we explore the potential for ‘spurious’ correlations induced by more ‘fundamental’ correlation between one particular parameter and halo mass. More specifically, we find that the empirical correlation between halo mass and SSFR can be largely or fully explained as a natural consequence of the empirical correlation between halo mass and any of Sérsic index, effective radius, or intrinsic stellar colour. The converse—that a correlation between halo mass and SSFR can explain the observed correlations with other properties—is untrue.

⁶ By the same token, we cannot exclude the possibility that any (or even all) of the four apparent correlations shown in Fig. 6 is ‘spurious’, in the sense that they are simply a consequence of a properly ‘fundamental’ astrophysical correlation between halo mass and some property, unidentified, which is itself correlated with each the observables we consider—but this is inescapable.

quantity, x	mean value, x_0	$b = \langle M_{\text{halo}} \rangle$	$A = \partial M_{\text{halo}} / \partial x$	RMS value, σ_x	implied SHMR dispersion
<i>Principal quantities of interest:</i>					
intrinsic stellar colour, $(g - i)_*$	0.6192	$0.384^{+0.095}_{-0.093}$	$1.065^{+0.623}_{-0.549}$	0.158	$0.20^{+0.10}_{-0.10}$
effective radius, $\log R_e$ / [kpc]	1.5942	$0.406^{+0.101}_{-0.091}$	$-0.989^{+0.411}_{-0.352}$	0.226	$0.24^{+0.06}_{-0.09}$
specific star formation rate, $\log \text{SSFR}$ / [Gyr^{-1}]	-1.0103	$0.398^{+0.106}_{-0.094}$	$-0.188^{+0.223}_{-0.214}$	0.400	$0.09^{+0.08}_{-0.06}$
Sérsic index, $\log n$	0.3179	$0.457^{+0.122}_{-0.095}$	$1.007^{+0.320}_{-0.335}$	0.287	$0.28^{+0.06}_{-0.07}$
<i>Other astrophysical quantities:</i>					
stellar mass, $\log M_*$ / [M_\odot]	10.4865	$0.388^{+0.105}_{-0.094}$	$0.357^{+0.834}_{-0.928}$	0.115	$0.09^{+0.09}_{-0.06}$
star formation rate, $\log \text{SFR}$ [$M_\odot \text{ yr}^{-1}$]	0.4616	$0.407^{+0.100}_{-0.105}$	$-0.153^{+0.244}_{-0.230}$	0.392	$0.08^{+0.08}_{-0.06}$
effective stellar colour, $(g - i)$	1.0449	$0.401^{+0.104}_{-0.098}$	$0.388^{+0.517}_{-0.564}$	0.145	$0.08^{+0.06}_{-0.05}$
<i>Null results and controls:</i>					
Sérsic magnitude (apparent), m_r	18.2147	$0.389^{+0.096}_{-0.092}$	$0.084^{+0.136}_{-0.154}$	0.462	$0.06^{+0.06}_{-0.04}$
redshift, z	0.1479	$0.387^{+0.107}_{-0.101}$	$0.744^{+4.367}_{-4.575}$	0.022	$0.08^{+0.07}_{-0.06}$
Declination, Dec. / [deg.]	0.0631	$0.370^{+0.113}_{-0.086}$	$-0.048^{+0.061}_{-0.050}$	1.502	$0.10^{+0.08}_{-0.07}$
random value, x	0.4973	$0.388^{+0.101}_{-0.086}$	$-0.467^{+0.300}_{-0.265}$	0.289	$0.16^{+0.08}_{-0.09}$
axis ratio, q	0.4105	$0.409^{+0.095}_{-0.100}$	$-0.466^{+0.457}_{-0.424}$	0.224	$0.12^{+0.10}_{-0.08}$
position angle, θ / [deg]	0.0813	$0.406^{+0.101}_{-0.104}$	$-0.003^{+0.002}_{-0.002}$	51.485	$0.19^{+0.10}_{-0.10}$

Table 1. Summarising the results of our linear fits to the correlation between halo mass and various quantities. — For each quantity, x , the first three columns give the values for the linear relation $M_{\text{halo}} = A(x - x_0) + b$, as defined in Eq. (9). As described in §5.1, we use $\ln 10 \sigma_x A/b$ as a metric to compare the relative strength of the correlation between each property and halo mass. These values, given in the final column, can be interpreted as the amount of dispersion in the SHMR that is directly coupled to the property in question, and as such, they provide an approximate lower bound on the dispersion around the SHMR.

while individual galaxies are as likely to scatter to higher or lower masses, the fact that lower mass galaxies are more common than higher mass galaxies means that more low mass galaxies scatter up into the sample than high mass galaxies scatter down. The consequence of this would be that we might be overestimating stellar masses nearer to our lower mass limit, and underestimating those nearer to our upper mass limit. Simple numerical experiments suggest that despite the fact that $\sim 20\%$ more galaxies scatter across the lower selection boundary than the upper one, the scale of this kind of bias on the mean mass is < 0.006 dex at the low mass end, and even smaller at the high mass end. In short, the systematic impact that random errors in the stellar mass estimates might have on our results really is negligible.

The bigger concern might be differential systematic errors in the stellar mass estimates, whereby there is some bias in galaxies’ stellar mass estimates that correlates directly with one or more of the observable properties we have looked at. While we cannot unambiguously rule out this possibility, we can ask how big such a bias would have to be in order to fully explain our results, following a simple argument based on propagation of errors. In order to explain the apparent correlation between halo mass and Sérsic index, the size of the differential systematic bias would have to be such that we are underestimating the stellar masses of high n galaxies by ~ 0.4 dex (*i.e.*, a factor of ≈ 2.5) relative to low n galaxies; this would be equivalent to missing 1 magnitude of flux. For size, the bias would have to be such that we are overestimating the masses of large galaxies by ≈ 0.25 dex, or a factor of ≈ 1.8 , relative to small galaxies. For colour, the masses of red galaxies would have to be underestimated by ≈ 0.35 dex, or a factor of ~ 2.25 , relative to blue galaxies. While biases of this size are not inconceivable, they would certainly be extreme and undermine the vast majority of work on galaxy formation and evolution over the past several decades (see also Taylor et al. 2010, for constraints on differential systematic errors on stellar mass estimates).

5.3 Is it really halo mass?

At the most basic level, what we have shown is that there are differences in the observed ESD profiles for lens samples split by different lens galaxy properties, which we have interpreted these differences as being due to variations in mean halo mass across the sample. Another possibility is that the observed differences in the ESDs profile are in fact the result of variations of some halo property other than mass; *e.g.*, halo concentration.

Recall from §2.4 that we are forced to assume a strict prior on halo concentration, c , based on the Duffy et al. (2008) prescription for c as a function of halo mass. This is because we cannot properly constrain the values of either c or M_{halo} for our lens sample without such a prior. In this section, we do away with this prior, and look at the joint or bivariate halo mass–concentration constraints that we can derive using our data.

The change is just to allow c to be a free parameter in the NFW description of the halo ESD profile, $\Sigma_{\text{NFW}}(R|M_{\text{halo}}, c)$ so that Eq. (7) and Eq. (8) become bivariate expressions of both M_{halo} and c , rather than M_{halo} alone. What we have done is to map the bivariate likelihood function, $\mathcal{L}(M_{\text{halo}}, c)$, based on fits to the stacked ESD profiles for each of our quartile subsamples.

The results of this exercise are shown in Fig. 7. While the uncertainties are large, there is no clear evidence for significant variation in halo concentration across the sample: the c values for each subsample are generally consistent with one another. This is not true for halo mass, where the values of M_{halo} do clearly vary across the sample. These observations allow us to rule out the possibility that the effects that we are seeing are being driven primarily by variations in halo concentration, rather than mass.

Further to this point, we also note that, in general, the inferred variation as a function of mass is *greater* when we relax our strong prior on c . The implication, therefore, is that a less restrictive prior on c would only *increase* the apparent correlation between halo mass and galaxy properties. This is another reason why our inferred values for the dispersion in the SHMR should be viewed as lower limits.

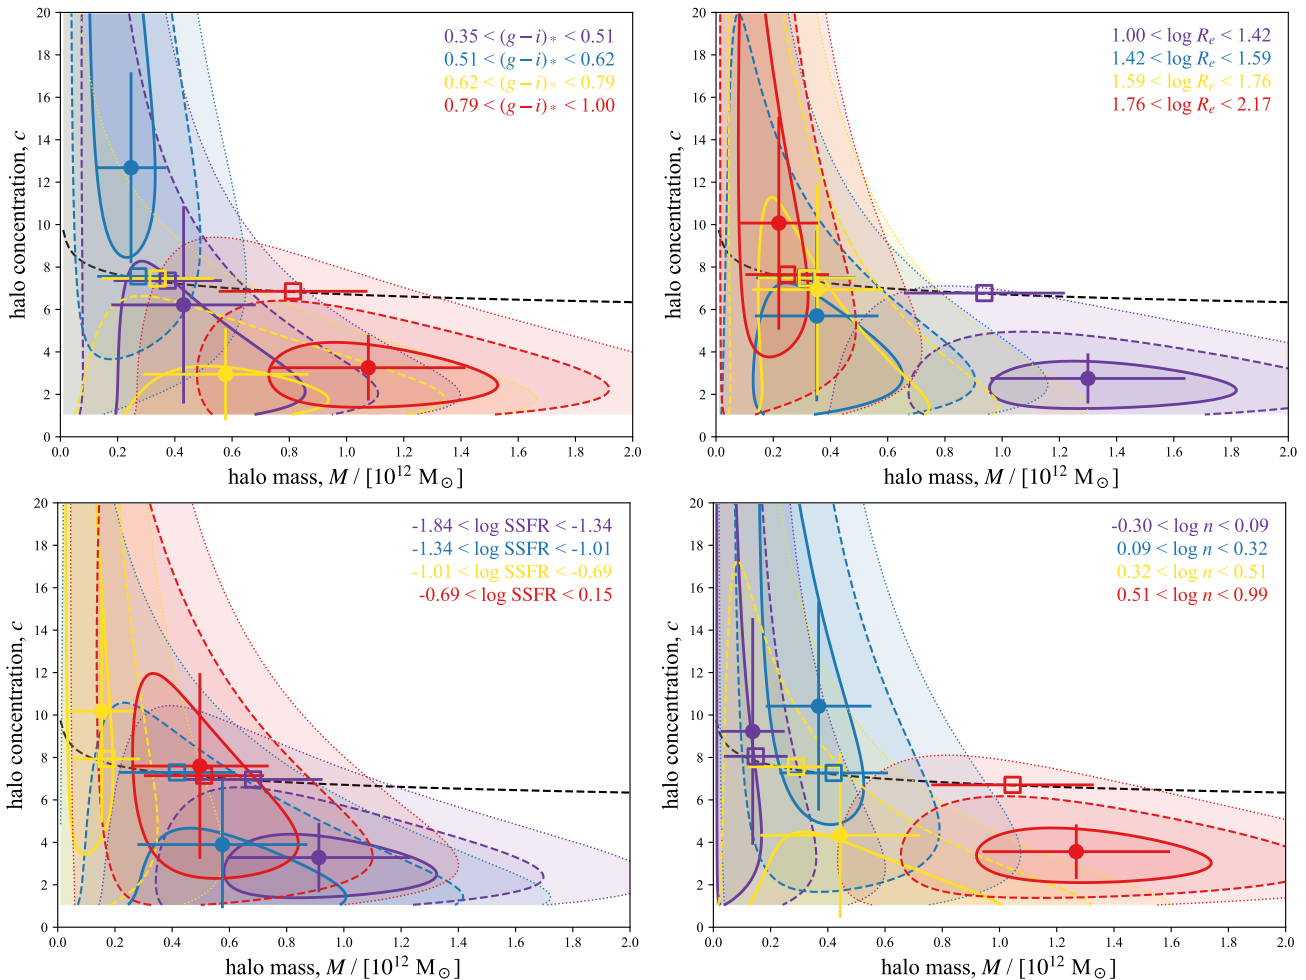


Figure 7. The joint constraints on halo mass and concentration, when binning by stellar colour, SSFR, effective radius, and Sérsic index. — These figures provide a means to test the proposition that it is really halo mass that is varying across the sample, rather than concentration. The solid, dashed, and dotted contours in each panel show the region bounding the 50%, 90%, and 99% confidence regions in the halo mass-concentration plane when fitting to stacked ESD profiles for the sample split according to different properties. The filled points with error bars show the usual least-squares estimator and (symmetric) standard error estimates, which are what one would obtain by marginalising over either M_{halo} or c . The black dashed line shows the Duffy et al. (2008) $M_{\text{halo}}-c$ relation that we adopt as our prior elsewhere in the text; the open squares show the values of M_{halo} we obtain with this prior. In each case, the results for each subsample are consistent with no systematic variation in concentration across the sample as a whole, where there is clear evidence for significant variation in halo mass. This shows that we really are seeing variations in halo mass, rather than halo shape, across our sample.

6 SUMMARY AND CONCLUSIONS

Our broad purpose with this paper has been to explore correlations between halo mass, as measured by galaxy-galaxy weak lensing, and other observable galaxy properties. Our work is based on a volume-limited sample of central galaxies (see Fig. 1) spanning a narrow range in stellar mass ($10.3 < \log M_* < 10.7$) and redshift ($0.10 < z < 0.18$). This particular mass range is interesting because 1.) it is near the knee in the galaxy stellar mass function and the stellar-to-halo mass relation (SHMR), and so represents the transition between the low- and high-mass regimes of galaxy evolution; and 2.) it is where there is the greatest diversity in the general galaxy population — in particular, it is where the various bimodalities in galaxy properties are most pronounced.

From a technical standpoint, the first novel aspect of this work is that as well as deriving halo mass measurements from stacked ESD profiles, we have also explicitly modelled the full set of unstacked lensing profiles for all of the lenses in our sample. Specifically, as described in §2.4, we have made linear fits to the relations

between halo mass and galaxy observables, where we have considered all of the individual lenses in our sample simultaneously. We have validated this approach in §3 through comparison between our measurement of the SHMR over this narrow mass range and previous results, and by demonstrating that we see no correlation between halo mass and a number of unrelated variables; *viz.* declination, redshift, or a random variable. We discuss the potential for systematic biases in our results in §3.2.

When considering the results shown in Figs 4, 5, and 6, it is important to recognise that the lines should not be understood as fits to the points. Where the points show the mean halo mass inferred from stacked ESD profiles, after binning by a particular galaxy property, the lines represent the inferred mean relation between halo mass and the property in question, as inferred from our modelling of the full ensemble of unstacked ESD profiles. The points and the lines are therefore best understood as complementary representations of the general trend across the sample.

Fig. 6 shows our essential results, which are the empirical cor-

relations between halo mass, and several key galaxy properties; *viz.* intrinsic stellar colour (as an tracer of stellar populations and particularly light-weighted mean stellar age), specific star formation rate, effective radius (as a proxy for size and/or density), and Sérsic index (as an indicator of concentration, which can also be taken as a proxy for bulge-to-disk ratio). We see evidence for variation in halo mass as a function of each of the galaxy properties that we have considered. In general terms, our main observational result is thus that, for the same stellar mass, canonically ‘early type’ galaxies have larger halo masses than canonical ‘late types’.

Our results are qualitatively and quantitatively consistent with the stellar-to-halo mass measurements (in the relevant mass range) of red versus blue galaxies by, *e.g.*, Hoekstra et al. (2005), Mandelbaum et al. (2006), and Hudson et al. (2015). As discussed in §4, there is at least superficial tension between our qualitative interpretation of the relation between halo mass and galaxy size and that of Charlton et al. (2017). Our results should be contrasted to complimentary studies by, *e.g.*, Sonnenfeld, Wang & Bahcall (2019) and Huang et al. (2020), who have investigated stellar-to-halo mass ratios as a function of size/structure for very- to super-massive galaxies ($\log M_* > 11$ and > 11.7 , respectively).

At a basic level, what we have shown is that there are differences in the lensing signal from halos (more specifically, the ESD profiles; see Fig. 2), which we are then interpreting as being due to variations in the mean halo masses, as a function of different observables. It is conceivable what we are seeing are really variations in some other halo property, which we are mistakenly attributing to variations in mass. An important factor in our process for inferring halo masses is a strong prior constraint on halo concentration as a function of mass, since we cannot place strong constraints on halo mass without such a prior. In order to address the possibility that what we may be seeing is variation in halo concentration, rather than mass, as a function of galaxy properties, we have looked at the joint halo mass–concentration likelihood surface for various subsamples of our data. As shown in Fig. 7, while the data are consistent with all subsamples having approximately the same concentration, there is clear evidence for variations in halo mass across the sample. It remains possible that there is also some variation in halo concentration across the sample, but it is clear that there is certainly significant variation in halo mass. It also seems likely that relaxing the prior on concentration as a function of halo mass would only strengthen the observed correlations, rather than reduce them.

As a way to derive new insights into the influence that halo mass has on the formation and evolution of individual galaxies, our particular goal is to identify galaxy properties that are most directly related to halo mass. Saying the same thing in another way, we are looking for what galaxy properties are most closely correlated with offsets from or dispersion around the main SHMR. Compared to past studies, the novel aspects of this work are 1.) use of a narrow stellar mass range, to control for stellar mass dependence as best we can; and 2.) a systematic consideration of multiple galaxy observables.

The observed variation in mean halo mass as a function of galaxy properties within our sample demonstrates that $\log M_* \sim 10.5$ galaxies span a range of halo masses; that is, at fixed mass, there is significant dispersion in the SHMR. Moreover, the fact that the observed variations in halo mass across our sample as a function of colour, SSFR, size, and shape are larger than as a function of stellar mass clearly demonstrates that we are directly probing the dispersion in the SHMR. We can thus use our sample to get an approximate lower bound on the SHMR dispersion, under the assumption that the offset from the mean SHMR is fully and directly

tied to one given observable (see §5.1). In this way, we can limit the dispersion in the SHMR at $\log M_* \sim 10.5$ to be $\gtrsim 0.3$ dex.

While the expectation from simulations is that the dispersion should be large ($\gtrsim 0.4$ dex Mitchell 2016), there are not yet many strong, direct observational constraints. Abundance matching and halo occupation modelling approaches typically find an inferred dispersion *in stellar mass at fixed halo mass* of order 0.2 ± 0.02 dex (*e.g.* Moster et al. 2010; van Uitert et al. 2015; Tinker et al. 2017). This propagates through to an expected dispersion *in halo mass at fixed stellar mass* of order 0.24 dex for a $10.3 < \log M_* < 10.7$ sample like ours. If there is any additional variation in stellar-to-halo mass ratios beyond what is directly correlated with galaxy structure, then the dispersion in the SHMR must be larger than previously suggested. Alternatively, if past results are correct, then our results would suggest that the dispersion in the SHMR is essentially perfectly coupled to structure; in other words, that galaxies follow a sort of ‘fundamental plane’ as a function of stellar mass, halo mass, and structure, with essentially no scatter.

It is not possible with the present dataset to definitively address the question of which parameter is (or parameters are) most directly and fundamentally tied to halo mass. But in general terms, we find that the structural properties of concentration and size are better predictors of halo mass than stellar population properties like intrinsic colour or star formation rate. The suggestion from the data is that Sérsic index is slightly preferred over effective radius as the ‘more fundamental’ parameter: using the metric of the inferred dispersion in the SHMR to compare the relative significance of the trends with different parameters, the values are $0.28^{+0.06}_{-0.07}$ for Sérsic index and $0.24^{+0.06}_{-0.09}$ for effective radius, compared to $0.20^{+0.10}_{-0.10}$ for intrinsic stellar colour. The value for SSFR is just $0.09^{+0.08}_{-0.06}$, but with the caveats that 1.) the binned-and-stacked results indicate that a simple linear fit does not provide a faithful description of the mean relation between M_{halo} and SSFR, and 2.) the observed trend with SSFR can be explained as a spurious or secondary correlation.

A naive interpretation of our results would be that, *at fixed stellar mass*: halo mass determines structure (but with some scatter); then, structure determines stellar populations (but with some scatter). One implication of this would be that galaxy structure (as traced by Sérsic index) responds more quickly to changes in halo mass than does stellar colour, which would imply that the structural transition from disk-dominated to bulge-dominated precedes the colour transition from blue to red. By contrast, it would seem that halo mass does not play a primary role in determining the instantaneous star formation rate — or at least, SSFR is not a good predictor of M_{halo} — in this mass range.

By adopting stellar mass as the independent variable or regressor, the implicit assumption in the above is that stellar mass can be taken as a proxy for something like ‘degree of evolution’. A more theory-minded view would instead have halo mass play this prime role. In this framing, our stellar mass selection might be viewed as mixing halo populations where a higher halo mass is offset by a lower stellar-to-halo mass ratio, or vice versa. Where we observe a higher halo mass, this would therefore imply that the processes of star formation and/or stellar assembly have gone slower (through differences in environment, merger history, stellar feedback, or internal dynamics) and/or ended sooner (through differences in environment, merger history, AGN feedback, or internal dynamics), with the result being a relatively smaller stellar mass. Taking this view, the interpretation would be that, *at fixed halo mass*: a more concentrated stellar structure in the present day is associated with a faster and/or shorter star formation/stellar assembly history; conversely, a more extended and/or diskier stellar distribution is asso-

ciated with a longer and/or slower star formation/stellar assembly history.

The crucial question remains: what is the nature of the astrophysical causal connection(s) underpinning the observed statistical correlations? From the observers' side, one avenue for further study is the degree to which the dispersion around the SHMR is 'random'—in the sense that it is the product of stochastic processes that are not closely correlated with other halo particulars like formation time or large scale environment—or if it instead reflects some form of assembly bias (see, *e.g.*, Wang, De Lucia & Weinmann 2013). We have no direct means of probing this question with the present data/analysis—instead, what is needed are models that reproduce our results, which can then be interrogated to see how this behaviour comes about in the models.

In cosmological models of galaxy formation, the mechanism for long-term quenching of star formation in massive galaxies is by heating the gas in the outer halo, and thereby preventing or disrupting further gas accretion onto the galaxies. The subgrid prescription for this 'maintenance mode' feedback is usually, but not uniquely, associated with kinetic and/or energetic feedback from the central black hole (see, *e.g.*, Davé, Thompson & Hopkins 2016). To the extent that AGN feedback scales with black hole mass, and black hole mass scales with bulge mass, and bulge-to-total ratio scales with Sérsic index, our results might be used as an indirect test of AGN feedback prescriptions. For a different perspective, our results would seem at least superficially consistent with, *e.g.*, Snyder et al. (2015), who show that within $10^{12} M_{\odot}$ halos from the IllustrisTNG simulations, diskier galaxies tend to have higher stellar masses, and also Tachella et al. (2019), who argue that present day structure is set during the star-forming phase, where stellar feedback is more important. Following this line of argument, our results might be a more sensitive test of stellar feedback prescriptions. In this spirit, we present our observations as targets for modellers to aim to reproduce.

With this first exploratory study, we have demonstrated the feasibility and utility of unstacked lensing profiles to probe variations in halo mass across an ensemble. Looking ahead, the obvious next question is whether similar trends exist for lower and higher masses. This will need more work, and possibly also larger samples to obtain sufficient signal in the lensing measurements. Taking a broader perspective, this study also shows the value of having galaxy demographic survey data as a foreground screen for wide area lensing surveys. In this, we particularly highlight the opportunities that will be afforded by the combination between KiDS and WAVES-Wide (Driver et al. 2019) in the next few years.

ACKNOWLEDGEMENTS

This research is partially funded by the Australian Government through an Australian Research Council Future Fellowship (FT150100269) awarded to ENT. H. Hoekstra acknowledges support from Vici grant 639.043.512, financed by the Netherlands Organisation for Scientific Research (NWO). AS acknowledges funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 792916. This work is part of the Delta ITP consortium, a program of the Netherlands Organisation for Scientific Research (NWO) that is funded by the Dutch Ministry of Education, Culture and Science (OCW). H. Hildebrandt is supported by a Heisenberg grant of the Deutsche Forschungsgemeinschaft (Hi 1495/5-1) as well as an ERC Consolidator Grant (No. 770935). CS acknowledges support from the

Agencia Nacional de Investigación y Desarrollo (ANID) through FONDECYT Iniciación grant no. 11191125.

REFERENCES

- Bell E F, McIntosh D H, Katz N, Weinberg M D, 2003, *ApJ* 585, L117
 Bell E F et al., 2004, *ApJ* 608, 752
 Baldry I K, Balogh M L, Bower R G, Glazebrook K, Nichol R C, Bamford S P, Budavari T, 2006, *MNRAS* 373, 469
 Baldry I K et al., 2012, *MNRAS* 421, 634
 Baldry I K, 2018, *MNRAS* 474, 3875
 Bamford S et al., 2009, *MNRAS* 393, 1324
 Bartelmann, M & Schneider P, 2001, *Phys Rep*, 340, 291
 Behroozi P S, Conroy C, Wechsler R H, 2010, *ApJ* 717, 379
 Berlind A A, Weinberg D H, 2002, *ApJ* 575, 587
 Bower R G et al., 2006, *MNRAS* 370, 645
 Brainerd T G, Blandford R D, Smail I, 1996, *ApJ* 466, 623
 Bruzual G & Charlot S, 2003, *MNRAS* 344, 1000
 Calzetti D, Armus L, Bohlin R C, Kinney A L, Koorneef J, Storchi-Bergmann T, 2000, *ApJ* 533, 682
 Chabrier G, 2003, *ApJ* 586, L133
 Charlton P J L, Hudson M J, Balough M L, Khatri S, 2017, *MNRAS* 472, 2367
 Croton D J et al., 2006, *MNRAS* 365, 11
 Coe D, 2010, arXiv:1005.0411
 Conroy C, Wechsler R H, Kravtsov A V, 2006, *ApJ* 647, 201
 Coupon J et al., 2015, *MNRAS* 449, 1352
 Davé R, Thompson R, Hopkins P F, 2016, *MNRAS* 462, 3265
 Davies L J M et al. 2016, *MNRAS* 461, 458
 Driver S P et al., 2011, *MNRAS* 413, 971
 Driver S P et al., 2018, *MNRAS* 475, 2891
 Driver S P et al., 2019, *Messenger* 175, 46
 Duffy A R, Schaye J, Kay S T, Dalla Vecchia C, 2008, *MNRAS* 390, L64
 Dvornik A et al., 2018, *MNRAS* 479, 1240
 Fenech Conti I, Herbonnet R, Hoekstra H, Merten J, Miller J, Viola M, 2017, *MNRAS* 467, 1627
 Graham A W & Driver S P, 2005, *PASA* 22, 118
 Guo Q, White S D M, Li C, Bolyan-Kolchin M, 2010, *MNRAS* 404, 1111
 Hildebrandt H et al., 2017, *MNRAS* 465, 1454
 Hill D T et al., 2011, *MNRAS* 412, 765
 Hoekstra H & Jain B, 2008, *ARNPS* 58, 99
 Hoekstra H, Yee H C K, Gladders M D, 2004, *ApJ* 606, 67
 Hoekstra H, Hsieh B C, Yee H C K, Lin H, Gladders M D, 2005, *ApJ* 635, 73
 Huang S et al. *MNRAS* 492, 3685
 Hudson M J, Gwyn S D J, Dahle H, Kaiser N, 1998, *ApJ* 503, 531
 Hudson M J et al., 2015, *MNRAS* 447, 298
 Joung M K R, Mac Low M-M, 2006, *ApJ* 653, 1266
 Kannawadi A et al., 2019, *A&A* 624, A92
 Kelvin L S et al., 2012, *MNRAS* 421, 1007
 Kelvin L S et al., 2014, *MNRAS* 444, 1647
 Larson R B, 1974, *MNRAS* 169, 229
 Lange et al., 2015, *MNRAS* 447, 2603
 Leauthaud A et al., 2012, *ApJ* 744, 159
 Li C, Wang L, Jing Y P, 2013, *ApJL* 726, L7
 Liske J et al., 2015, *MNRAS* 452, 2087
 Madelbaum R, Seljak U, Kauffman G, Hirata C, Brinkmann J, 2006, *MNRAS* 368, 715
 Marchesini D, van Dokkum P G, Forster Schreiber N, Franx M, Labb'e I, Wuyts S, 2009, *ApJ* 701, 1765
 Marinoni C & Hudson M J, 2002, *ApJ* 569, 1
 McKee C F & Ostriker J P, 1977, *ApJ* 218, 148
 Mechior P & Viola M, 2012, *MNRAS* 424, 2757
 Miller L et al., 2013, *MNRAS* 429, 2858
 Mitchell P D, Lacey C G, Carlton C M, Cole S, 2016, *MNRAS* 462, 878
 Moffett A et al., 2016, *MNRAS* 457, 1308
 Moster B et al., 2010, *ApJ* 710, 903

- Navarro J F, Frenk C S & White S D M, 1997, ApJ 490, 493
Peng Y, Lilly S, Renzini A, Carollo M, 2012, ApJ 757, 23
Renzini A & Peng Y, 2015, ApJL 801, L29
Robotham A S G et al., 2011, MNRAS 416, 2640
Robotham A S G et al., 2013, MNRAS 431, 167
Robotham A S G et al., 2014, MNRAS 444, 3986
Shen S et al., 2003, MNRAS 343, 978
Sonnenfeld A, Wang W, Bahcall N, 2019, A&A 622, A30
Snyder G F et al., 2015, MNRAS 454, 1886
Tachella S et al., 2019, MNRAS 487, 5416
Taylor E N, Franx M, Brinchmann J, van der Wel A, van Dokkum P G,
2010, ApJ 722, 1
Taylor E N et al., 2011, MNRAS 418, 1587
Taylor E N et al., 2015, MNRAS 446 2144
Tinker J L, Robertson B E, Kravtsov A V, Klypin A, Warren M S, Yepes G,
Gottlöber S, 2010 ApJ 724, 878
Tinker J L et al., 2017, ApJ 839, 121
van Uitert E, Hoekstra H, Franx M, Gilbank D G, Gladders M D, Yee H K
C, 2013, A&A 539, A7
van Uitert E et al., 2016, MNRAS 459, 3251
Velander M et al., 2014, MNRAS 437, 2111
Viola M et al., 2015, MNRAS 452, 3529
Wang L, De Lucia G, Weinmann S M, 2013, MNRAS 431, 600
Wechsler R H & Tinker J L, 2018, ARAA 56, 435
van der Wel A, 2008, ApJL 675, L13
Wright A H et al., 2016, MNRAS 460, 765
Wright C O & Brainerd T G, 2000, MNRAS 534, 34
Yang X, Mo H J & van den Bosch F C, 2003, MNRAS 339, 1057