

A hierarchical clustering method for quantifying satellite abundance

Chengyu Xi^{1,2★} and James E. Taylor^{1,2★}¹*Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, ON N2L 3G1, Canada*²*Waterloo Centre for Astrophysics, University of Waterloo, 200 University Avenue West, Waterloo, ON N2L 3G1, Canada*

Accepted 2021 February 19. Received 2021 January 19; in original form 2020 August 7

ABSTRACT

We present a new method for quantifying the abundance of satellites around field galaxies and in groups. The method is designed to work with samples such as photometric redshift catalogues, that do not have full spectroscopic coverage, but for which some redshift or distance information is available. It consists of identifying the galaxies most likely to be centrals, and using the clustering signal around them as a template to iteratively decompose the full population into satellite and central populations. Thus, it is similar to performing crowded-field photometry by using isolated stars to determine the point spread function of the image. The method does not identify individual satellites or centrals conclusively, but assigns a probability to each galaxy of being one or the other. Averaged over a large sample, it provides a statistical estimate of satellite abundance, even in crowded fields with large redshift uncertainties. We test the method using data from the COSMOS field, which includes a large set of local objects with accurate photometric redshifts. We measure satellite abundance as a function of central stellar or halo mass, as well as the satellite luminosity function, and find results consistent with previous studies, but extending over a broader range of central masses. We also consider a number of possible systematic uncertainties in the method, and show that they are generally smaller than our random errors. Having presented the method in this paper, we will use it to study the properties of the satellite populations in a forthcoming one.

Key words: galaxies: dwarf – galaxies: formation – galaxies: groups: general – Local Group – galaxies: luminosity function, mass function – dark matter.

1 INTRODUCTION

In the current picture of hierarchical structure formation, cold dark matter (CDM) haloes merge together to form progressively larger systems as the Universe evolves. Smaller dark matter haloes often survive accretion on to larger systems, leaving a population of distinct ‘subhaloes’ within CDM haloes. While the gas that cools and settles into the centre of the main halo will contribute to the growth of a central galaxy (CG), smaller galaxies that formed within subhaloes before they merged can survive as distinct satellite galaxies, subject to a broad range of environmental effects that may transform or destroy them, including tidal heating or stripping, encounters, and internal or external feedback. The observed abundance of satellites in the local Universe provides a detailed test of this complex picture and gives important insights into the overall effect of environment on galaxy formation.

The dominant galaxies of the Local Group (LG), the Milky Way (MW) and M31, have the best studied satellite populations in the Universe. Recent surveys have discovered many new, faint members of the LG (e.g. Bechtol et al. 2015; Drlica-Wagner et al. 2015; Kopolov et al. 2015), such that the total abundance of LG satellites can be estimated with increasing confidence (Newton et al. 2018). Over the past two decades, however, several points of tension have arisen between the observed population of LG satellites and that

expected from theory. The most famous is the ‘missing satellite problem’, which contrasts the small number of observed satellites with the large number of dark structures predicted by theory (Klypin et al. 1999; Moore et al. 1999). A second, ‘too-big-to-fail’ problem contrasts the low central densities estimated in the massive satellites of the MW with the much higher densities expected from theory (Boylan-Kolchin, Bullock & Kaplinghat 2011). There may be other tensions as well, in the radial clustering (e.g. Kravtsov, Gnedin & Klypin 2004; Taylor, Silk & Babul 2004) or 3D spatial distribution (e.g. Pawlowski et al. 2015). We refer the reader to Bullock & Boylan-Kolchin (2017) for a detailed review of these challenges.

Many solutions have been proposed to resolve the tensions between theory and observations of the LG satellites, including internal feedback due to star formation (e.g. Dekel & Silk 1986; Mashchenko, Wadsley & Couchman 2008; Governato et al. 2010; Wetzel et al. 2016), the effects of global (e.g. Bullock, Kravtsov & Weinberg 2000; Gnedin & Kravtsov 2006) and/or inhomogeneous (e.g. Lunnan et al. 2012) reionization, tidal or other environmental effects (e.g. Taylor & Babul 2001; Mayer et al. 2006; Łokas, Kazantzidis & Mayer 2012), or modifications to the underlying dark matter model such as warm dark matter (e.g. Macciò & Fontanot 2010; Anderhalden et al. 2013; Kennedy et al. 2014; Lovell et al. 2014), self-interacting dark matter (e.g. Spergel & Steinhardt 2000; Elbert et al. 2015; Fry et al. 2015), or fuzzy dark matter (e.g. Nadler et al. 2019).

There remains, however, the important question of whether the MW and/or LG satellites are representative of all satellite populations. Many observational studies have shown that the MW is not

* E-mail: cxi@uwaterloo.ca (CX); taylor@uwaterloo.ca (JET)

typical in having two Large/Small Magellanic Cloud (LMC/SMC)-like satellites (Guo et al. 2011; Liu et al. 2011; Robotham et al. 2012; Strigari & Wechsler 2012; Speller & Taylor 2014), and similar conclusions have been suggested by numerical simulations (Boylan-Kolchin et al. 2010; Busha et al. 2011; Kang, Wang & Luo 2016; Zhang, Luo & Kang 2019). The ongoing Satellites Around Galactic Analogs survey (SAGA – Geha et al. 2017) has also shown that there is a large variation in satellite populations from system to system. Theoretical models predict that the abundance of halo substructure should vary more than expected from Poisson statistics alone, and should be correlated with the formation redshift of the system (e.g. Chua et al. 2017; Jiang & van den Bosch 2017). These complications caution us from relying too heavily on the properties of a single system to constrain models of galaxy formation. To determine whether the LG is representative, and to understand satellite properties across a broad range of environments, we should seek out satellites around as large a sample as possible of CGs.

Identifying satellites and distinguishing them from foreground or background systems require some form of distance information. The main approaches in the literature include the following: (1) the use of existing complete spectroscopy to identify satellites around the nearest and brightest systems (e.g. Yang et al. 2007); (2) dedicated spectroscopic campaigns to obtain spectroscopy for fainter targets around a smaller number of selected systems (Geha et al. 2017); (3) the use of photometric distance estimates from techniques such as the tip of the red giant branch (e.g. Carlin et al. 2016; Danieli et al. 2017; Cohen et al. 2018; Danieli et al. 2019), or surface brightness fluctuations (e.g. van Dokkum et al. 2018; Carlsten et al. 2019); (4) statistical abundance measurements based on clustering (e.g. Guo et al. 2011, 2012a; Liu et al. 2011; Strigari & Wechsler 2012; Wang & White 2012; Sales et al. 2013; Speller & Taylor 2014; Wang et al. 2014; Xi et al. 2018). The four approaches have different strengths and weaknesses. Method (1) requires only existing data, but is restricted to the brightest satellites in the nearest systems, and may also suffer from incompleteness due to fibre positioning limitations in dense fields (e.g. Guo, Zehavi & Zheng 2012b; Smith et al. 2019). Method (2) is extremely expensive in terms of observing time, and thus limited to small numbers of systems. Method (3) is restricted to very nearby systems (<20 Mpc), whose virial radii subtend large angles on the sky, making complete coverage difficult. Method (4) cannot confirm individual galaxies as satellites or centrals; it has been very successful, however, in making measurements of the average satellite abundance, and is the least resource-intensive method of the four a priori.

Clustering-based methods have generally been applied to samples at redshifts of ~ 0.05 – 0.2 , selected from the Sloan Digital Sky Survey (SDSS – York et al. 2000; e.g. Guo et al. 2011, 2012a; Liu et al. 2011; Strigari & Wechsler 2012; Wang & White 2012; Sales et al. 2013; Wang et al. 2014). A different strategy was adopted by Speller & Taylor (2014), who focused on very nearby systems (<42 Mpc). This allowed them to estimate the abundance of intrinsically faint satellites, at the expense of significant background contamination. They used a selection technique based on galaxy structural properties (mainly apparent size) to reduce the background contamination and boost the signal-to-noise ratio (SNR) of the clustering measurement. The technique was further developed and tested in Xi et al. (2018), using a broader range of morphological cuts. An optimized version was shown to be effective up to $z \sim 0.15$, far beyond the range considered in Speller & Taylor (2014).

These previous clustering-based studies have generally considered samples of primaries that are clearly isolated, in the sense that they have no brighter companion within fixed projected and line-

of-sight separations. This approach works well for bright, massive primaries, but becomes inefficient for less luminous ones. By dropping isolation cuts, Xi et al. (2018) were able to detect a clear clustering signal and constrain satellite abundance using only observations from the fairly small Cosmic Evolution Survey (COSMOS) field, but this resulted in a broad selection of primaries, including many systems with overlapping virial regions. As a result, the interpretation of their results remains slightly unclear, relative to previous studies, as not all of their primaries are true CGs.

In this work, we introduce a new method to deal with the complications of overlapping systems and crowded fields. We start by identifying the subset of galaxies in a sample most likely to be true CGs, using a hierarchical search in which galaxies are checked for isolation in order of decreasing stellar mass, with isolation criteria that scale with the estimated virial radius of the system. The cross-correlation function of the sample with respect to this set of most likely primaries provides an initial template for the clustering signal. This template is used to estimate the probability that *any* member of the sample is a primary or a secondary. Finally, we can iterate through the last two steps, recalculating a probability-weighted cross-correlation function and the modified primary/secondary probabilities until convergence. The final primary/secondary probabilities for the whole sample then allow us to estimate satellite abundance, luminosity functions, and other distributions of secondary properties. We note that we have developed and optimized our method for low redshift samples; some of our assumptions may need modification, in order to apply the method at higher redshifts.

In this paper, we present the method and give some simple estimates of satellite abundance; in a forthcoming paper, we will study the properties of the detected satellite populations in more detail. The paper is structured as follows. In Section 2, we describe our data selection, including the basic cuts that define our initial sample. In Section 3, we measure the clustering signal and use it to define a ‘region of interest’ (ROI) around each primary likely to contain most genuine satellites. In Section 4, we describe our iterative method for estimating primary and secondary probabilities for each galaxy. In Section 5, we present our main results on satellite abundance. In Section 6, we test the method for possible systematic uncertainties. Finally, in Section 7 we summarize our results and discuss future prospects for this new method.

2 DATA – COSMOS

The satellite galaxies we can hope to detect around a low-redshift primary (at most a few tens per system, based on abundances in the LG) will be seen in projection with a much larger number of foreground and background galaxies (of the order of thousands) that are not physically associated with the primary. Precise distance information is essential for separating true satellites from this foreground/background population. Spectroscopic redshifts are ideal for this purpose, but impractical for large samples. For instance, if we want to search for satellites brighter than -18 in absolute magnitude out to a redshift of 0.2 , this requires distance information for galaxies down to an apparent magnitude of roughly 22 . However, with a few exceptions (e.g. Geha et al. 2017), wide-field spectroscopic catalogues are usually only complete down to an apparent magnitude of 17 to 18 , far from the depth required. Thus, using photometric redshifts (‘photo- z s’) is the only realistic solution. The COSMOS field features high-quality photo- z s generated from 30+ deep bands (Scoville et al. 2007; Ilbert et al. 2013; Laigle et al. 2016), making it an ideal place to test our method.

2.1 The COSMOS photometric redshift catalogue

COSMOS is a deep ($AB \sim 25\text{--}26$), multiwavelength ($0.25\text{--}24\ \mu\text{m}$) survey covering a $2\ \text{deg}^2$ equatorial field (Scoville et al. 2007). The multiwavelength imaging includes *Hubble Space Telescope* imaging with the Advanced Camera for Survey and follow-up observations from many other facilities across a wide range of wavelengths – X-ray, UV, optical/IR, FIR/submillimetre, and radio (Scoville et al. 2007). In this paper, we will use a recently updated photometric redshift catalogue (Laigle et al. 2016, ‘COSMOS 2015’ hereafter) for our analysis. The main improvement of this catalogue compared to the previous releases is the addition of new, deeper NIR and IR data from the second data release (DR2) of the UltraVISTA and SPLASH (*Spitzer* Large Area Survey with Hyper-Suprime-Cam; Miyazaki et al. 2012) projects. Compared to the first data release (DR1) of UltraVISTA, the exposure time of DR2 was significantly longer (McCracken et al. 2012), providing the deeper IR and NIR data as well as better SNRs (Laigle et al. 2016). On the other hand, the DR2 data only covers a part (namely the ‘ultra-deep stripes’, roughly $0.6\ \text{deg}^2$) of the COSMOS field. This causes a slight inconsistency in depth and SNR across the field, which we will address below by applying a magnitude cut.

The COSMOS photo- z s were derived using χ^2 template fitting, as described in Mobasher et al. (2007) and Ilbert et al. (2009). The spectral energy distribution (SED) templates used in the COSMOS 2015 catalogue include a set of 31 spiral and elliptical galaxies from Polletta et al. (2007) and a set of templates for young blue star-forming galaxies generated using Bruzual & Charlot (2003) models. Given the updated NIR and IR data and two additional star-forming galaxy templates, Laigle et al. (2016) further improved on photo- z quality relative to previous COSMOS catalogues (Capak et al. 2007; Ilbert et al. 2009, 2013). The accuracy of the photo- z s has been verified by comparing them to a large number of highly reliable (97 per cent confidence) spectroscopic redshifts (Lilly et al. 2007) that are available in the COSMOS field. For the objects of magnitude $i_{AB}^+ < 22.5$ and redshift (z) range of 0–1.2, the photo- z s have an r.m.s scatter of $\sigma = 0.7$ per cent with respect to the spectroscopic redshifts, while the occurrence of ‘catastrophic failures’ with relative errors $|z_p - z_s|/(1 + z_s) > 0.15$ is only 0.51 per cent. For this work, we choose the median of photo- z likelihood distribution from the template fitting (‘ZPDF’ in the catalogue) as the base redshift. From this redshift, we calculate angular-diameter and luminosity distances, and corresponding luminosities and projected separations, assuming all galaxies follow the Hubble flow. In the process of template fitting and photo- z estimation, Laigle et al. (2016) also calculated stellar masses and star formation rates for the galaxy samples, which will be used in our analysis below. Specifically, we use ‘MASS_MED’ and ‘SFR_MED’, the medians of the stellar mass and star formation rate probability distribution functions (PDFs).

2.2 Additional spectroscopic redshifts

We can further improve on our distance estimates by supplementing the COSMOS photo- z s with spectroscopic redshifts, where these are available. While there is no single public spectroscopic redshift catalogue for the whole COSMOS field, most of the measured redshifts in the region are now accessible through the NASA Extragalactic Database.¹ In addition to these redshifts, we also obtained a few other unpublished redshifts from the COSMOS collaboration (M.

¹<https://ned.ipac.caltech.edu>

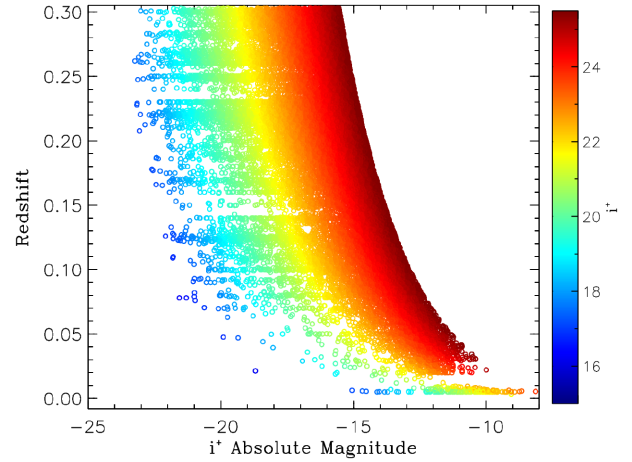


Figure 1. The redshift versus i^+ -band absolute magnitude distribution of our base catalogue, coloured by the i^+ -band apparent magnitude, as indicated in the right-hand colour scale.

Salvato, private communication). The redshifts used in this work will be mainly photo- z s from the COSMOS 2015 catalogue, but replaced with spectroscopic redshifts where possible. Given the numerous literature sources and slightly different qualities of the spectroscopic redshifts, a universal redshift uncertainty of 0.0001 is assigned to each galaxy whose photometric redshift is replaced with a spectroscopic value. Absolute magnitudes and stellar masses for those objects are also corrected, based on the resulting change in the distance modulus.

2.3 The base sample

As mentioned above, the depth of the COSMOS 2015 catalogue varies across the field, depending on whether the new ‘ultra-deep’ (UltraVISTA DR2) imaging is available or not. In general, the catalogue appears to be relatively complete down to a magnitude of $i^+ < 25.5$ (MAG_AUTO), but becomes incomplete beyond this. [The 3σ depths in the i^+ band are 26.2 and 26.9 for 3 and 2 arcsec apertures, respectively (Laigle et al. 2016).] To ensure reasonable completeness over the redshift range of interest, we apply the following initial cuts on the catalogue, which are the same cuts used in Xi et al. (2018):

- (i) $i^+ < 25.5$
- (ii) $0 < z_{\text{pdf}} < 6.9$
- (iii) $z - 2\sigma_z < 0.3$
- (iv) $\sigma_z < 0.5$

where z_{pdf} refers to the median of photo- z likelihood distribution measured using galaxy template fitting, and σ_z refers to the photo- z error, estimated by using the 68 per cent confidence level upper and lower limits of the photo- z likelihood distribution provided in the catalogue [i.e. $\sigma_z = (z_{\text{pdf}}^{U68} - z_{\text{pdf}}^{L68})/2$]. Note that we include the broad redshift cut $0 < z_{\text{pdf}} < 6.9$ to exclude stars and X-ray sources in the catalogue, as well as objects in the masked regions, as these objects do not have robust z_{pdf} estimates; we include the upper limit cut of redshift $z - 2\sigma_z < 0.3$ to focus on the local volume while keeping a reasonable completeness over a target redshift range of 0–0.25; finally, we use a redshift error cut $\sigma_z < 0.5$ to exclude those galaxies with poor-quality redshifts (mainly faint galaxies) from the further analysis. These cuts produce a base catalogue of 41 559 galaxies (37 578 after excluding galaxies with large redshift errors). Fig. 1 shows the redshift versus i^+ absolute magnitude distribution

for our base catalogue after applying the cuts above. Given our cut in apparent magnitude, the sample galaxies generally have absolute magnitudes between -24 and -10 for the redshift range ($z = 0-0.25$) we will consider below.

2.4 Stellar mass completeness

Laigle et al. (2016) estimated the stellar mass completeness of their catalogue; for redshift range of $0 < z < 0.35$, they suggested a 90 per cent completeness limit of $M_* = 10^{8.6} M_\odot$. We are considering systems at a slightly lower redshift range ($z < 0.25$). Examining the stellar mass function and stellar mass errors for this redshift range, we conclude that we are complete down to at least $M_* < 10^{8.2} M_\odot$, where the differential mass function peaks, and cut the main sample at this value. For very low redshifts ($z < 0.07$), we appear to be complete down as low as $M_* < 10^{7.2-10^{7.5}} M_\odot$; we will discuss local satellite abundance at these lower stellar masses below and in subsequent work.

3 DETERMINING THE CLUSTERING SCALE

We assume a model in which the galaxy with the largest stellar mass is the dominant galaxy within each halo, and resides at or close to its geometric and dynamical centre. To tell whether a given galaxy is the dominant CG (or ‘primary’ hereafter) within its own halo, or a potential satellite of another primary, we need to search its surroundings to see if there is another more massive galaxy close by. If there are more massive neighbours nearby, this suggests that the galaxy may be satellite, whereas if all nearby galaxies are less massive, it suggests that the galaxy is a primary. To quantify the characteristic scale on which satellites are associated with their primaries, we measured the clustering pattern of all the galaxy pairs in our base sample. Based on these clustering results, we will quantify ‘spatially nearby’ and specify the ROI for the primary–secondary classification.

3.1 Halo mass assignment

We expect the characteristic extent of the satellite distribution in a halo to scale with its virial radius. We calculate a fiducial halo mass and virial radius for each galaxy, assuming that it is the CG of its host halo. These masses and virial radii will be used to characterize the clustering throughout this work.

To estimate halo masses, we could in principle use abundance matching, assuming a monotonic relation between stellar mass and halo mass that we derived empirically by comparing the observed stellar mass function and the predicted halo mass function within the observation volume. However, the effective area of the COSMOS field (after correcting for masking as discussed below) is only 1.46 deg^2 , giving an effective comoving volume of $1.67 \times 10^5 \text{ Mpc}^3$ up to $z = 0.25$. The cosmic variance in the mean density for a volume this size is a factor of approximately 0.4 (Somerville et al. 2004); if we consider 10 independent, equal-volume redshift slices, the relative cosmic variance of each slice increases to 0.7. Thus, there is a large systematic uncertainty in the normalization of the halo mass function within this volume. Instead, we use the stellar-to-halo mass relation (SHMR) derived by Behroozi, Wechsler & Conroy (2013) (B13 hereafter). They provide a formula (equation A1) for the inverse halo-to-stellar mass ratio (HSMR), with parameters as listed in Appendix A.

We note, however, that a combination of observational errors in the stellar mass estimates, intrinsic scatter in the SHMR, and the non-linearity of the halo mass function will bias halo masses estimated directly from the B13 formula. As there are many more

low-mass haloes than high-mass ones, errors in stellar mass will more frequently scatter objects into a given mass range from below than from above. Thus, directly applying the HSMR to estimate the halo mass corresponding to an observed stellar mass will lead to systematic overestimates, especially at the high-mass end. To quantify and correct this bias, we have performed a Monte Carlo simulation of the effect of errors to obtain a bias-corrected HSMR based on B13. We explain this correction in more detail in Appendix A.

Given an estimate of the halo mass for each galaxy, we also assign a corresponding virial radius and virial velocity. The virial radius is taken to be $\left(\frac{3M_h}{4\pi\rho_c\Delta_c}\right)^{1/3}$, where ρ_c is the critical density of the universe and Δ_c is the mean overdensity of the halo within the virial radius, with respect to the critical density. We use the fitting formula from Bryan & Norman (1998) for Δ_c . Once the virial radius has been calculated, the virial velocity is given by $\sqrt{GM_h/R_{\text{vir}}}$.

The stellar mass completeness limit is estimated empirically, following the method described in Laigle et al. (2016). They calculated the K_s -band magnitude limits for the COSMOS catalogue to be 24.0 and 24.7 for the deep and ultra-deep fields, respectively. In this work, we choose a limit of 24.0 in order to have uniform depth across the whole field. Given this magnitude limit, the limiting mass a galaxy would need to have to be observed at a given redshift is calculated as

$$\log M_{\text{lim}} = \log M - 0.4(K_{s\text{lim}} - K_s). \quad (1)$$

Next, a stellar mass limit is estimated for each redshift bin, within which 90 per cent of the galaxies lie, given the stellar mass errors. We also calculate a corresponding halo mass limit in each redshift bin using the bias-corrected HSMR (although in this case we ignore the effect of scatter on the completeness threshold).

3.2 Defining an ROI

Around each CG, we define an ROI in which we will search for potential satellites. The size of ROI is determined by two considerations: First, the clustering signal should be consistent across systems with different CG masses, and secondly, the ROI should include most of the ‘one-halo’ clustering signal associated with the main halo around the CG, while excluding the regions that are dominated by the ‘two-halo’ or background terms. We have explored different possible choices of ROI boundaries by measuring the two-dimensional (line-of-sight and projected) clustering signal of all pairs in our base catalogue.

For each pair in the catalogue, the galaxy with the larger stellar mass is assumed to be the primary. We then count pairs as a function of separation in projected distance and line-of-sight velocity offset, scaled by the halo virial radius and the velocity error, respectively. These scaling choices are found to give us relatively consistent clustering signals for primaries of different masses (as shown in Fig. 2 below). Note that some fraction of an annulus around a given primary may be missing from the catalogue, as it overlaps with field boundaries and masked regions. We carefully measure the shape of the survey boundaries and masked regions to produce a single template for the whole field. Monte Carlo sampling of this template is then used to determine the area completeness η around each primary as a function of radius, as explained in Appendix B, and the counts are corrected by this factor.

Fig. 2 shows the density of pair counts in 2D phase space for all galaxies (top panel), and binned by primary halo mass (bottom panels), where the primary is defined to be the member of the pair with the larger stellar mass. A clear overdensity of pairs can be

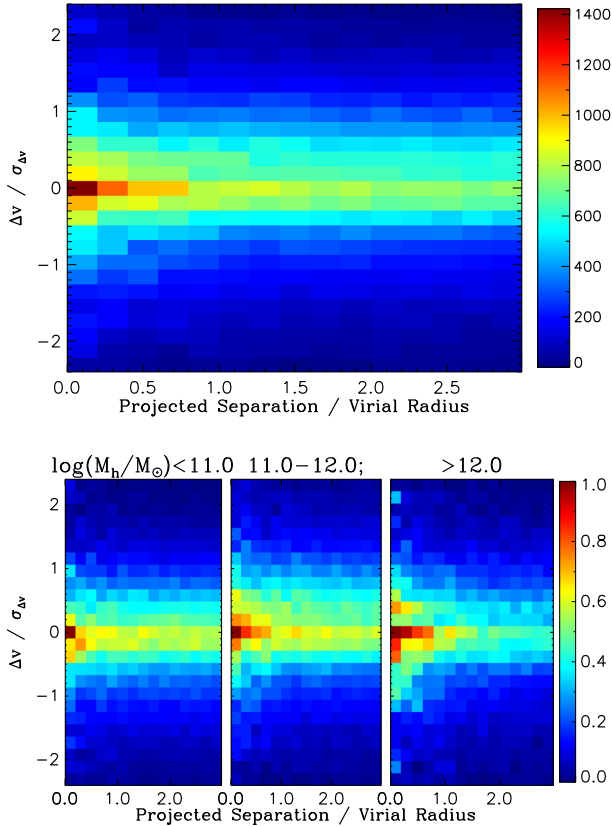


Figure 2. Primary–secondary clustering signal in the base catalogue. The colour scale shows the density of galaxy pairs $\Sigma = N/\Delta R\Delta v$ as a function of projected separation and velocity offset, where these have been scaled by the estimated virial radius and velocity error, respectively. Top panel: all primaries; bottom panels: same quantity normalized and binned by primary mass.

seen both in the projected separation and in the velocity separation directions. Overall, primaries with larger masses show a stronger clustering signal. At the same time, the clustering patterns in the different mass ranges have a similar dependence on separation scale: They all have the strongest clustering within $0.5R_{\text{vir}}$. The signals all extend to fairly large radii, but start to drop significantly after $1.5\text{--}2R_{\text{vir}}$. Along the velocity axis, which is scaled by redshift error, the signals in all three mass bins drop off at a similar rate, reaching the background level at $\Delta v \sim 1\text{--}1.5\sigma_{\Delta v}$. It is worth noting the slight asymmetry of the clustering signal along the velocity separation axis, with slightly more negative velocity separations than positive ones. This is due to the incompleteness at the faint, high-redshift end of the survey volume (as shown in Fig. 1). We calculate the velocity offset with respect to the more massive (and thus more luminous) member of the pair, which as a result of Eddington bias due to incompleteness tends to be further away on average. Thus, it produces a negative velocity offset more often than a positive one.

Based on these clustering patterns, around each primary galaxy we will define an ROI for potential satellites using the following cuts:

- (i) A cut in projected separation (assuming both galaxies are at the line-of-sight distance of the primary), scaled by the virial radius of the primary: $R_p/R_{\text{vir}} < A$.
- (ii) A cut in velocity difference, scaled by the circular velocity of the primary: $\Delta v/v_{\text{vir}} < B$.

- (iii) A cut in velocity difference relative to the uncertainty in velocity difference between the primary and secondary: $\Delta v/\sigma_{\Delta v} < C$.

Note the final cut depends on the secondary properties as well as those of the primary, so this is determined for each galaxy pair individually. Any secondary galaxy in a pair that meets conditions (i) and (ii or iii) is considered a potential satellite of the more massive member. We will choose the values $(A, B, C) = (3.0, 2.0, 1.5)$ as our default, but test the effect of changing these definitions of the ROI in Section 6 below.

4 THE ITERATIVE CLUSTERING METHOD

Our goal is to quantify satellite abundance using the strength of the clustering signal. An obvious complication is that any given galaxy may be a central or a satellite; without further selection, the ‘raw’ clustering signal measured in the previous section consists of a complicated sum of central-satellite and off-set-satellite terms (e.g. Leauthaud et al. 2012). Depending on the stellar mass and range of separations considered, the latter can significantly bias estimates of the true satellite abundance.

To avoid this complication, our strategy will be to identify the galaxies most likely to be centrals, and use the clustering signal around these objects as a template to separate out central and satellite contributions. In this section, we will first describe the initial, ‘first-run’ primary selection; then, we will model the clustering of secondaries around these first-run primaries to determine primary and secondary probabilities for all galaxies in the sample, and finally we will test the results of iterating over this process, by adding new potential primaries to our initial sample, weighted by their primary probability, and remeasuring a weighted clustering signal around the enlarged primary sample.

4.1 Initial primary selection

We want to select primaries that dominate an ROI that scales with their halo mass, as described in the previous section. Since the ROI is larger for the more massive systems, a smaller system that has no more massive companions within its own ROI can still lie within the ROI of a larger system. This asymmetry naturally leads us to a hierarchical search algorithm, where we start searching around the most massive systems first. The detailed steps are as follows:

- (i) All galaxies in the catalogue are assigned a halo mass and virial velocity dispersion based on their stellar mass, as if they were the CG in their own host halo.
- (ii) We then go through the catalogue in ranked order of stellar mass, selecting the most massive galaxy in the catalogue as the first primary.
- (iii) All galaxies that lie in the ROI around the first primary (as defined in the previous section) are classified as its secondaries and removed from the list of potential primaries.
- (iv) The next most massive unclassified galaxy is then selected as the next primary candidate.
- (v) We check the stellar masses of all galaxies within the ROI around this next candidate. If the candidate is the most massive galaxy in its ROI, then it is classified as a primary and the other galaxies in the ROI are classified as its secondaries.
- (vi) We iterate over the last two steps until all the galaxies in the catalogue are classified as primaries or secondaries.

This produces our first-run primary sample. To guarantee reasonable completeness, we make two additional cuts on this initial sample as follows:

- (i) Primaries with ROIs that are heavily affected by survey boundaries or masking are removed from the initial primary list.
- (ii) Primaries with redshifts higher than 0.25 are removed from the initial primary list.

These cuts are necessary to remove primaries close to the sample boundaries, either on the sky or in redshift. These galaxies may have more massive companions that lie just outside the field or beyond our redshift cut. Thus, there is a higher probability that they are not actual CGs, but are in fact satellites of another, more massive galaxy.

Overall, this process is very conservative in selecting primaries, producing a sample of 1490 galaxies that is incomplete (in the sense of missing many genuine CGs), but relatively uncontaminated by satellites. Cutting out systems with redshifts exceeding 0.25 reduces the number of first-run primaries to 873, while excluding those with area completeness less than 0.65 (i.e. those with ROIs that are masked or cut off by field boundaries by more than 35 per cent) reduces the number to 815.

4.2 Clustering of the first-run primary and secondary samples

To study satellite abundance and its dependence on primary properties, we first need to model and separate the contributions to the clustering signal from the satellite population and the background galaxy population. Having classified potential primaries and secondaries, we measure the surface number density of secondaries within the ROI and in an extended region around it (with the same velocity offset limits, but extended out to $3.2R_{\text{vir}}$ in order to have a better estimate of the background surface density).

Around each primary, we count secondaries in annuli spaced evenly in $\log(R_p/R_{\text{vir}})$. The annuli range from 0.1 to $3.2R_{\text{vir}}$, in steps of 0.25 dex. We exclude secondaries with projected separations of less than $0.1R_{\text{vir}}$ (~ 25 kpc, for the MW) to avoid outlying H II regions or other components of the CG. We calculate the surface density of secondary galaxies around each primary in physical units (Mpc^{-2}), assuming all secondaries lie at the same distance as the primary. Given counts $N^{i,j}$ in radial bin i of projected area A_i around primary j , the surface density is

$$\Sigma_{i,j} = \frac{N^{i,j}}{\eta_i A_i} = \frac{N^{i,j}}{\eta_i 2\pi R^i \Delta R}, \quad (2)$$

where η_i is the mean area completeness in radial bin i (estimated as described in Appendix B), R^i is the mean radius of the bin, and ΔR is the width of the bin. As we want to scale the surface density of secondary galaxies in units of R_{vir} , it is useful to define a separation variable $X_i \equiv R_p^{i,j}/R_{\text{vir}}^j$. Thus, equation (2) can be written as

$$\Sigma_{i,j} = \frac{N^{i,j}}{\eta_i 2\pi (X_i \cdot R_{\text{vir}}^j) \Delta X \cdot R_{\text{vir}}^j} = \frac{N^{i,j}}{\eta_i 2\pi X_i \Delta X (R_{\text{vir}}^j)^2}. \quad (3)$$

In what follows, we will fit the surface density in bins of primary redshift and mass. Where the primary sample contains more than five objects, we use the bootstrap method to estimate the uncertainties in the surface densities, by re-sampling the primary sample 120 times. As the bootstrap method does not work well when the sample size is smaller than five, we have also calculated Poisson uncertainties for each bin. The final uncertainties for the bin are taken to be the larger of the two values.




Figure 3. The redshift–halo mass distribution of the first-run primary sample. The coloured boxes show the boundaries for the 20 subsets used to fit the mass and redshift dependence of the clustering signal.

The secondary surface density consists of two main parts: the contribution from clustered satellites and the contribution from background or foreground galaxies.

$$\Sigma(X|M_{\text{halo}}, z) = \Sigma_{\text{sat}}(X|M_{\text{halo}}) + \Sigma_{\text{bg}} \quad (4)$$

The first component Σ_{sat} should correlate with the halo mass of the primary, but should be roughly independent of redshift over the narrow redshift range considered here, while the second component should be roughly independent of halo mass, but should depend on redshift. There should also be a more extended clustered component due to large-scale structure (the ‘two-halo’ term; see e.g. Cooray & Sheth 2002, for a detailed review), but the characteristic scale of this component (cf. 4–8 Mpc) is much larger than the scales considered here. Thus, we treat it as a constant with respect to radius, and include it in the foreground/background term. To fit the two terms, we split the primaries into 5 fixed redshift slices and 2–5 halo mass bins per slice, with adaptive boundaries as shown in Fig. 3.

We assume that the satellite distribution roughly matches the subhalo distribution, which in turn approximately follows a Navarro–Frenk–White (NFW; Navarro, Frenk & White 1996) density profile. Over the range of radii we are most sensitive to, a projected NFW profile scales as r^{-2} in the outer parts of the halo, and is somewhat shallower in the inner parts. We could fit the density profile of the satellite component with the exact form of a projected NFW profile, assuming some mean concentration–mass relation. Given the low SNR of the satellite component, however, we choose to fit it with the simplified form

$$\Sigma_{\text{sat}}(X) = S_{\text{halo}} F(X) = \frac{S_{\text{halo}}}{X^2 + \alpha X + \beta}, \quad (5)$$

where S_{halo} is an overall normalization (in units of Mpc^{-2}) that is independent of X but depends on the primary halo mass, while α and β are parameters describing the radial dependence. We have compared this simplified form to a projected NFW profile over the range $X = 0.05$ – 3 , for concentration parameters $c = 5$ – 15 , and found that using fixed values $\alpha = 0.2$ and $\beta = 0$ gives a good fit in all cases. For this choice of parameters, $\int_0^1 2\pi F(X) X dX = \log[1 + \alpha] - \log[\alpha] = 11.18$, so the total number of satellites within the virial radius is $N_{\text{sat}} = 11.18 S_{\text{halo}} R_{\text{vir}}^2$.

The resulting fits for the secondary surface density (fitting S_{halo} and Σ_{bg} jointly for each individual primary bin, with α and β set to fixed values 0.2 and 0, respectively) are shown in Fig. 4. Summing

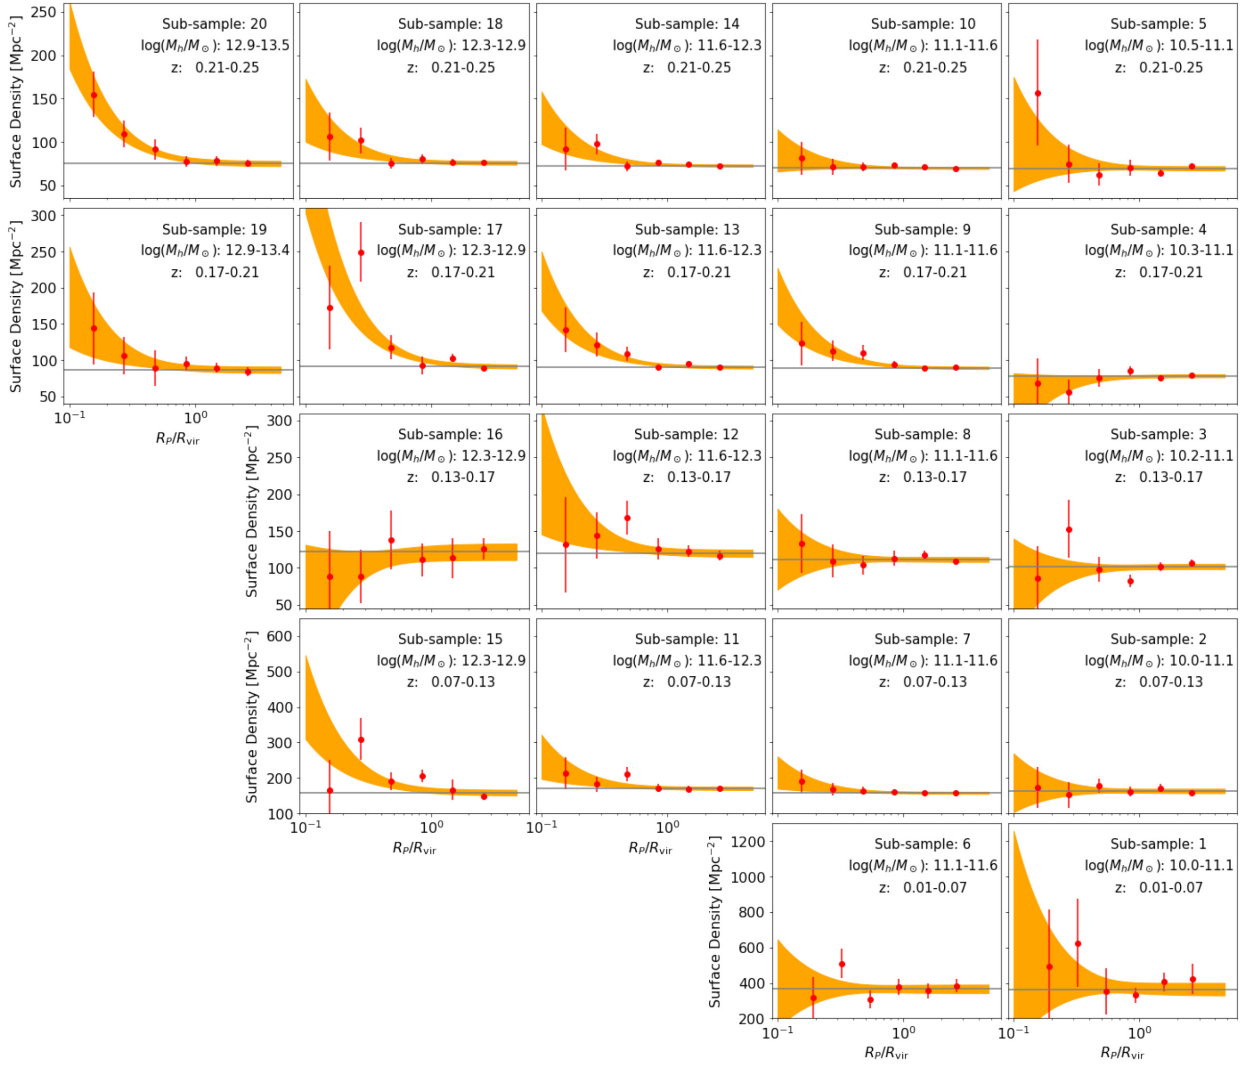


Figure 4. The surface density of secondaries around each of the primary subsamples shown in Fig. 3 (separate panels, with redshift increasing from bottom to top and halo mass increasing from right to left). Red points with error bars show the measured surface number density in radial bins, while the orange shading shows the 1σ region around the best fit from equation (5). The horizontal grey lines indicate the background level in the best-fitting model.

over all primary bins, a clustered excess in Σ is detected at an SNR of approximately 11. We can see that the overall background surface density (as measured in physical units) decreases with primary redshift, while the clustered satellite component increases with primary mass. In the next section, we will explore these correlations in more detail.

4.2.1 Satellite/halo component

We fit the halo mass dependence of S_{halo} over the 20 primary bins using a linear relation in log–log space:

$$\log_{10} \left(\frac{S_{\text{halo}}}{1 \text{ Mpc}^{-2}} \right) = a \times \log_{10} \left(\frac{M_{\text{halo}}}{10^{12} M_{\odot}} \right) + b, \quad (6)$$

where a and b are free parameters. We choose $10^{12} M_{\odot}$ as the pivot mass in our fit, as this is roughly the median halo mass of our 20 primary bins. The projected area of a halo will scale as the virial radius squared, that is as $M^{2/3}$. If we assume systems have a fixed number of satellites per unit halo mass (as expected

from subhalo abundance; e.g. Gao et al. 2004), then the projected surface density should go as $M^{1/3}$, so we expect $a \sim 0.33$. The value of b (the normalization at $M_{\text{halo}} = 10^{12} M_{\odot}$) will depend on the depth of the catalogue, as discussed below. From our fits, we find $a = 0.30^{+0.11}_{-0.10}$ and $b = 0.26^{+0.08}_{-0.11}$, so the scaling with halo mass seems fairly consistent with the expected value.

Fig. 5 shows our fit for S_{halo} as a function of mean halo mass, over all 20 bins in primary mass and redshift. Given its sensitivity to smaller radial bins with larger errors, the fitted value of the parameter has an SNR ~ 8 , significantly lower than the SNR for the whole clustering signal. To illustrate the dependence on halo mass more clearly, we also show in Fig. 6 a version combining bins with similar mean halo masses. We note that at the low-mass end [$\log(M_h/M_{\odot}) < 11$], the fitted value is actually negative (although consistent with zero, given the uncertainties). This may be partly due to completeness problems at low mass, which affect the high redshift bins in particular. Repeating the fitting process for low-redshift primaries ($z < 0.15$) only, we obtain a less negative value, which is once again consistent with zero (as shown in the subpanel of Fig. 6).




Figure 5. Average S_{halo} versus halo mass $[\log_{10}(M_h/M_\odot)]$, with colour indicating the mean redshift of the primary subsample. The blue curve and grey shaded area show the best-fitting model of the mass dependence (equation 6), together with the 1σ uncertainty range.

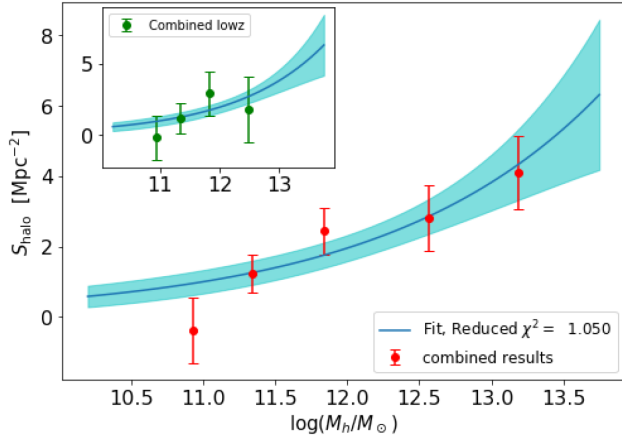


Figure 6. As Fig. 5, but with subsamples of similar mean halo mass combined into single bins for clarity. The top-left subplot shows the results for primaries with redshift $z < 0.15$.

4.2.2 Background component

Our surface densities are calculated in physical units (Mpc^{-2}) at the distance of the primary. Thus, if the foreground/background component consisted of a fixed field population with a broad redshift distribution and thus a fixed number per square degree, we would expect its inferred physical surface density to scale as d_A^{-2} , where d_A is the angular diameter distance of the primary. (It is worth noting that this assumption will not necessarily hold if extending the method to higher redshift.) Fig. 7 shows the fitted value of the background surface density Σ_{bg} in each bin versus the average value of d_A^{-2} for that bin. We fit the trend with a simple linear model:

$$\Sigma_{\text{bg}} = c + d \times \left[\left(\frac{d_A}{10^3 \text{ Mpc}} \right)^{-2} - 2.0 \right], \quad (7)$$

where c and d are free parameters.

Fitting the 20 bins gives tight constraints on the parameters: $c = 79.5^{+0.6}_{-0.6}$ and $d = 21.6^{+0.5}_{-0.5}$. Note that if the background scaled exactly as the inverse of the angular diameter distance, we would expect the constant term $c - 2d \sim 0$ to be small; in practice, various minor effects, notably the varying width of the redshift range Δz over which we measure the background, will cause the background density to




Figure 7. The background surface density Σ_{bg} versus d_A^{-2} . The grey area shows the 1σ uncertainty range around the best-fitting model (equation 7).

deviate from the simple scaling. As it is, for our fitted values of the parameters we find $c - 2d = 36.3$, which is small relative to the typical values of Σ_{bg} .

4.2.3 Single-step fit

While the two-step fitting procedure outlined above is useful to illustrate the features the model, it is more robust to fit the entire four-parameter model for both terms in the surface density in a single step, given the potential covariance between the model parameters. We use the function `Minimizer.emcee2` from the PYTHON module LMFIT (least-squares minimization fitting) to do MCMC sampling of the likelihood (Foreman-Mackey et al. 2013). The marginalized results of this fit are shown in Fig. 8 and are also included in Table 1. We note that the parameters a and b are strongly (anti)correlated; indeed, with higher SNR data we could imagine a more detailed, halo occupation distribution (HOD)-based fit to the halo-mass dependence of the satellite abundance. There is also some correlation between parameters c and b (or to a lesser degree c and a), indicating that satellite abundance estimates do require careful accounting for the background term.

4.2.4 Assigning probability

Given our model fit to the clustering measurements, we can estimate the amplitude and radial distribution of the satellite component and background components around each primary. We define the probability of a secondary galaxy in the ROI being an actual satellite as

$$P_{i,j}^{\text{sat}}(X, M_h, z^p) = \frac{\Sigma_{\text{sat}}}{\Sigma_{\text{tot}}} = \frac{\Sigma_{\text{sat}}(X, M_h)}{\Sigma_{\text{sat}}(X, M_h) + \Sigma_{\text{bg}}(z^p)}, \quad (8)$$

where M_h and z^p are the halo mass and redshift of the primary, respectively.

This equation can result in very small probabilities at large radii. Since real satellites (objects that have crossed the virial radius at least once) almost all lie within $3R_{\text{vir}}$ (e.g. Wetzel et al. 2014), we truncate the probability around $X = 3$ as follows:

$$P'(X) = P(X) \cdot \frac{1}{1 + 1000^{X-3}}. \quad (9)$$

²See <https://lmfit.github.io/lmfit-py/fitting.html> for a detailed description of LMFIT and `Minimizer.emcee`.

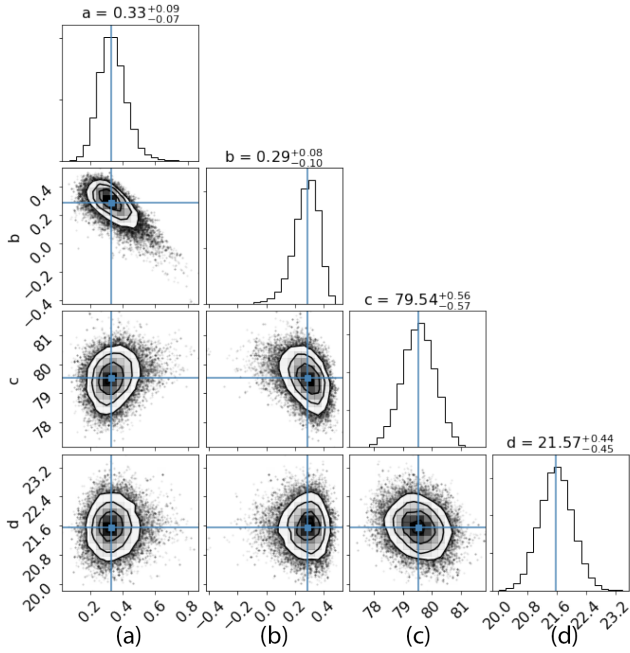


Figure 8. Likelihood distributions for the clustering model parameter values, derived by fitting the full model to all 20 primary subsamples simultaneously. Panels show the distribution marginalized over the two (or, on the diagonal, all three) other model parameters.

Table 1. The model fitting results from the first run and first iteration.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
F0	$0.33^{+0.09}_{-0.07}$	$0.29^{+0.08}_{-0.10}$	$79.5^{+0.6}_{-0.6}$	$21.6^{+0.4}_{-0.5}$
F1	$0.26^{+0.07}_{-0.07}$	$0.38^{+0.05}_{-0.06}$	$79.3^{+0.5}_{-0.5}$	$20.8^{+0.4}_{-0.4}$

4.3 Iterating over the fit

Our initial primary selection ignores many galaxies that may well be primaries, but appear close to more massive systems when seen in projection. To attempt to correct for this, we iterate over our clustering measurement, including a weighted contribution from all galaxies, proportional to their probability of being a primary.

For each iteration, we run the top-down selection again. During the new selection process, every galaxy starts with a 100 per cent probability of being a primary. Starting with the most massive galaxy as the first primary, we assign probabilities of nearby galaxies being its satellites, using the method described in the last section. The probability of each of these galaxies being an independent primary is reduced accordingly. We then proceed through the catalogue in order of decreasing stellar mass. If a galaxy has a probability of being a primary between zero and 1, we estimate that nearby galaxies have a probability of being its satellites that is the *product* of its probability of being a primary and the satellite probability given in Section 4.2.4. Proceeding through the entire catalogue in decreasing order of stellar mass ensures that the calculation is well defined and that every galaxy is assigned a final probability of being a primary, equal to 100 per cent minus the sum of all probabilities that it is a satellite of nearby systems.

Running through this process once, we find 8920 primaries with probabilities greater than 0.99, of which 3567 lie at redshifts of 0.25 or less. If we further remove systems with areal completeness of 0.65 or less, the number of high-confidence primaries is 3246, versus

815 in our original sample; repeating this exercise for systems with primary probabilities of 0.999 or more reduces the number to 1478.

If we include primaries with probabilities greater than 0.99 (and weight all satellites by their CG primary probability), the effect of iteration on the fit to the clustering signal is shown in Table 1. Overall, the parameter values after iteration show good consistency with our initial estimates, shifting by less than 2σ in all cases. The uncertainties in the fitted parameter values drop, but only slightly. This suggests that in a data set like the COSMOS catalogue that has extremely precise photometric redshifts, even the first-run sample of isolated primaries can provide a good estimate of satellite abundance. Since the use of lower probability primaries may dilute the clustering signal and introduce some bias, in what follows we will use our initial, first-run estimates of the fitted parameters to derive satellite abundance. We anticipate that in data sets with less accurate redshifts, iteration will become more important in deriving accurate estimates of the clustering signal.

5 ESTIMATING SATELLITE ABUNDANCE

In this section, we will make some basic estimates of the overall abundance of satellites, as well as their abundance as a function of properties such as stellar mass or luminosity. In each case, our estimate is based on the clustering signal, which is typically small compared to the background. The simplest way to estimate satellite abundance is to count every galaxy in the ROI, weighted by the satellite probability calculated in Section 4.2.4, so we will use this approach first in Section 5.1, referring to it as ‘method A’.

As explained below, method A assumes that the clustering amplitude is uncorrelated with stellar mass, luminosity, or any of the other secondary properties considered. More generally, we expect the fraction of galaxies in the ROI that are true satellites to depend on these other properties. In Section 5.2, we develop a more sophisticated approach, ‘method B’, similar to the one introduced in Speller & Taylor (2014), that attempts to correct for possible correlations in the limit of a weak clustering signal. Future surveys with stronger detections of clustering should be able to bypass these complications by dividing the galaxy sample directly into bins of secondary property value before they measure the clustering amplitude, simplifying the analysis considerably; we call this ‘method C’.

5.1 Abundance estimates using method A

In method A, to estimate satellite abundance N_{sat}^i around primary galaxy i we simply add up the probabilities P_j^i of each galaxy j in its ROI being a true satellite:

$$N_{\text{sat}}^i = \sum_j P_j^i. \quad (10)$$

Note that although P_j^i can remain non-zero at large distances from the primary, to compare to previous results from literature we set a radial limit of $1.5R_{\text{vir}}$ by default, and only count towards the total satellite abundance secondaries that lie within this projected separation of the primary.

We can study the dependence of satellite abundance on primary mass by stacking systems with similar stellar or halo masses, as shown in Fig. 9 (black points with error bars). As expected, there is a strong trend in satellite abundance with halo mass. As a consistency check, we also split our primary sample in two by redshift, and calculate satellite abundance separately for each of the two subsamples (red and blue points). The results for both subsamples show good consistency with those for the whole sample.




Figure 9. Satellite abundance as a function of primary stellar mass, estimated using method A. Red and blue points show the results for independent primary samples at higher redshifts ($z = 0.17\text{--}0.25$) and lower redshifts ($z = 0.07\text{--}0.17$), respectively.

We note that N_{sat} increases faster with primary stellar mass at the high-mass end of the range ($M_h > 10^{10.5} M_\odot$). This is consistent with the pattern seen in HOD modelling (e.g. Peacock & Smith 2000; Seljak 2000; Berlind & Weinberg 2002) and is a result of the changing slope of the SHMR; halo mass increases faster with stellar mass at large stellar masses, resulting in a faster increase in satellite numbers.

5.2 Abundance estimates using method B

Method A provides the simplest estimate of satellite abundance, and is guaranteed to be correct when averaging over all galaxies in the secondary sample used to calculate the satellite probability (equations 8 and 9). It is not necessarily correct, however, for subsamples of secondaries selected by luminosity, stellar mass, colour, or other properties, if these properties are correlated with the clustering amplitude. As a simple example, one can imagine a galaxy population that consisted of two distinct types, labelled ‘red’ and ‘blue’. If the red galaxies were completely clustered, but the blue galaxies were completely unclustered, then we would measure some intermediate average clustering strength for the combined population, and give every galaxy a satellite probability based on this average value. If we weighted all galaxies by this average satellite probability, but then split them back into subsamples by colour, we would conclude that the satellite and background populations both had the same net colour distribution. In effect, the true colour distribution of the satellite population (100 per cent red, in this example) would be contaminated by the colour distribution of the field population. More generally, whenever the background population is significantly different from the satellite population, the satellite properties inferred using method A will be biased towards those of the background population.

With sufficient SNR in the clustering signal, we could avoid this problem by selecting subsamples with a limited range of the desired satellite property (stellar mass, luminosity, colour, etc.) *before* measuring the clustering amplitude and calculating the satellite probability (we refer to this as ‘method C’). In the limit of low SNR, however, splitting the galaxies into narrow bins in a given property will increase the shot noise in the background estimate until it is unacceptably large. Instead, we have developed an intermediate solution, ‘method B’, based on the approach in Speller & Taylor (2014).

In method B, first the number of the background galaxies within a given radius is estimated for each primary, by summing up the non-satellite probabilities of each pair:

$$N_{\text{bg}}^i = \sum_j (1 - P_j^i). \quad (11)$$

We then measure the fraction of *all* background galaxies (i.e. summing over the ROIs of all primaries or over the whole field) with a property of interest (e.g. stellar mass, luminosity, and colour) in a given range, and scale the total number of background galaxies in the ROI by this fraction. This gives the expected background contribution to a particular subsample, which we then subtract to calculate satellite abundance. For instance, if we want to measure the luminosity function of satellites, Φ_{sat} , we first need to measure the total luminosity function within the ROI, Φ_{TTL} , and the total luminosity function for all background galaxies, Φ_{bg} . Then, we remove the background contribution from the total abundance in each bin k in luminosity, such that

$$\Phi_{\text{sat},k} = \Phi_{\text{TTL},k} - \frac{N_{\text{bg}}}{N_{\text{TTL}}} \Phi_{\text{bg},k}. \quad (12)$$

Here, N_{bg} and N_{TTL} refer to the number of background galaxies and the total number of galaxies within the radial cut around each primary, respectively (before any cut in luminosity). Note that this approach can be used for individual primaries, except for the lowest mass systems, where the galaxy counts are so small that Poisson fluctuations dominate. To correct for this, if the number of objects within the ROI is less than three times the number of luminosity bins, then we stack results for multiple primaries at similar redshifts, and use the average signal.

We used this method to calculate satellite abundance for different luminosity ranges (below), as well as full satellite luminosity functions (see Section 5.4). Fig. 10 shows abundance for various cuts in M_{i+} . Overall, the dependence of abundance on primary stellar (top panel) or halo (bottom panel) mass has a similar form for different luminosity cuts, although there may be a truncation at lower stellar masses that depends on the luminosity cut. Here too, this pattern is very similar to those seen for brighter galaxies in HOD modelling. Plotted as a function of stellar mass, satellite abundance shows a change in slope around $10.5 M_\odot$ at all luminosities. Plotting as a function of halo mass, this feature disappears, confirming that it is a result of the non-linear SHMR. We also indicate on the plot the abundance of MW satellites brighter than $V = -14.5$ with a galactocentric distance greater than 20 kpc (brown diamond – three satellites meet these criteria), assuming an MW halo mass of $12.1 M_\odot$, as discussed below.

5.3 Comparison to previous work

For massive galaxies, a number of other estimates of satellite abundance exist in the literature. In Fig. 11, we compare our abundance estimates to the results of Conroy, Wechsler & Kravtsov (2006, C06). These are based on HOD modelling (e.g. Peacock & Smith 2000; Seljak 2000; Zehavi et al. 2002) of the luminosity functions and correlation functions of samples from the SDSS (York et al. 2000) and DEEP2 (Newman et al. 2013) surveys. HOD modelling provides an estimate of the average number $\langle N_{\text{gal}} \rangle$ of galaxies within a halo of a given mass, so in the limit where $\langle N_{\text{gal}} \rangle \gg 1$, $\langle N_{\text{gal}} \rangle - 1$ should match our measured satellite abundance. As $\langle N_{\text{gal}} \rangle$ decreases, some haloes will contain no galaxies over a given magnitude limit, so we can only compare results in the large $\langle N_{\text{gal}} \rangle$ regime, i.e. for large halo masses. In addition, the results in C06 are given in bins

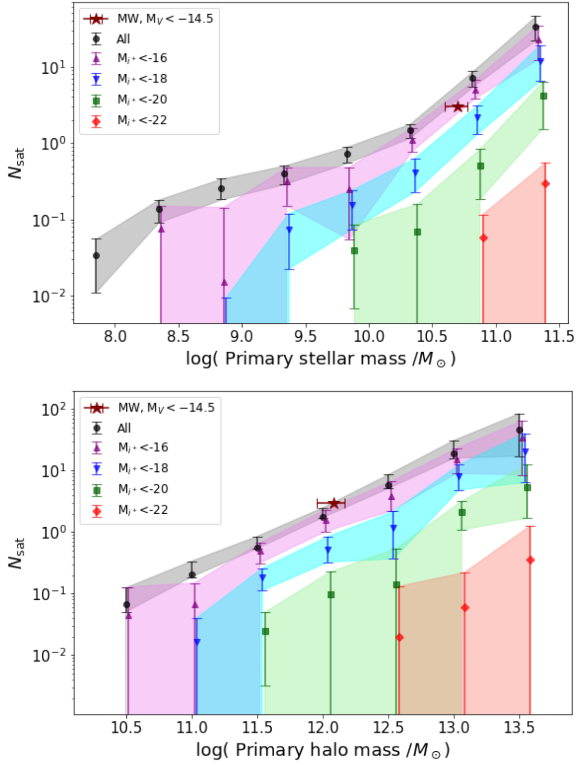


Figure 10. Satellite abundance as a function of primary stellar/halo mass (top/bottom panels, respectively), estimated using method B. The individual shaded regions show results for different cuts in satellite luminosity, as indicated in the legend. The brown diamond shows the number of MW satellites brighter than $V = -14.5$ with a galactocentric distance greater than 20 kpc (three satellites meet these criteria).

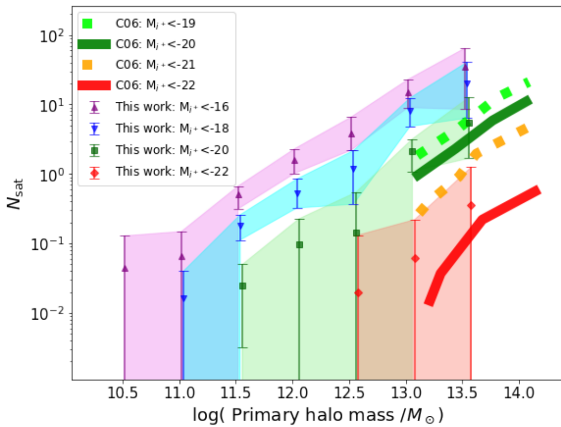


Figure 11. Comparison to the HOD-based results of Conroy et al. (2006, fig. 5). Their luminosity bins have been converted to i^+ -band magnitudes, assuming a mean colour of $(M_r - M_{i^+}) = 0.25$, and have had 1 subtracted from them to account for the CG.

of $M_r - 5 \log(h)$; we convert to our i^+ band assuming a fixed mean colour $(M_r - M_{i^+}) \sim 0.25$ (roughly the value measured for our sample in the COSMOS catalogue), which produces a shift of $(M_r - M_{i^+}) - 5 \log(h) \sim 1$ mag exactly in the C06 bin boundaries.

Given these conversions, examining Fig. 11, we see that there is excellent agreement between our results and those of C06. Our estimated satellite abundance matches that measured by C06 to within half a standard deviation, for all four of the magnitude cuts

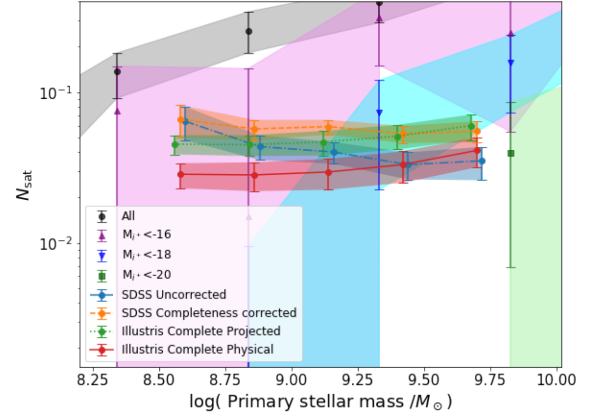


Figure 12. Comparison to the observed and mock results of Besla et al. (2018; see their figs 11 and 12). (Note that the points are shifted to left or right by up to 0.02 dex for clarity.)

in that study (two lie between our luminosity cuts, but are clearly consistent with our results). The slope of the $N_{\text{sat}}-M_h$ relation is harder to judge given the limited baseline in C06, but generally it appears to be consistent with our inferred slope at halo masses of $\log(M/M_\odot) > 13$ or more. The agreement between these two sets of results is particularly striking, given that they employ completely different samples and methods, and that there is no parametric freedom in adjusting our results. Overall, SDSS provides a more robust estimate of satellite abundance for massive haloes, but as a deeper survey with more accurate redshift estimates, COSMOS is better able to probe the low-halo-mass regime.

We have also compared our results to more recent work by Besla et al. (2018, B18), which is one of the few studies to estimate satellite abundance at lower halo mass. Fig. 12 compares our results to theirs, for the primary stellar mass range $10^8-10^9 M_\odot$. The B18 results are based on an SDSS spectroscopic sample at a redshift of 0.013–0.0252, with r -band magnitudes between 14 and 17.77, and primary stellar masses in the range of $(0.2-5) \times 10^9 M_\odot$. (The apparent magnitude limit of 17.77 corresponds to absolute magnitude limits of -16.03 or -17.5 at $z = 0.013$ or 0.0252 respectively, equivalent to -16.3 and -17.8 in the i^+ band.) They also compare these to a mock catalogue generated using the Illustris hydrodynamical simulations (Vogelsberger et al. 2014; Nelson et al. 2015). Four sets of results are shown. The ‘uncorrected SDSS’ results are from raw counts of nearby companions; the ‘completeness corrected’ version is after correcting for observational selection effects, using the mock catalogues. The ‘physical’ simulation results show the abundance of real satellites, while the ‘projected’ counts show the result including a background contribution introduced by projection effects. Completeness and projection corrections move the measured abundance up and down, respectively; overall, the best estimate of satellite abundance from B18 is the completeness corrected (orange) curve, decreased by a factor of ~ 50 per cent to account for projection effects.

Over the mass range of $(0.2-5) \times 10^9 M_\odot$ covered by B18, our estimates for magnitude cuts $M_{i^+} < -16$ or $M_{i^+} < -18$ are consistent with their measured values. Given their effective magnitude limits range from -16.3 and -17.8 in the i^+ band, there once again seems to be excellent consistency in the overall abundance estimated by the two methods. In contrast, B18 measure almost no trend in satellite abundance with primary stellar mass. It is worth pointing out, however, that the mean redshift of the sample in B18 is lower at smaller primary masses (see their fig. 3). Thus, the satellite luminosity function may be measured to greater depth for these

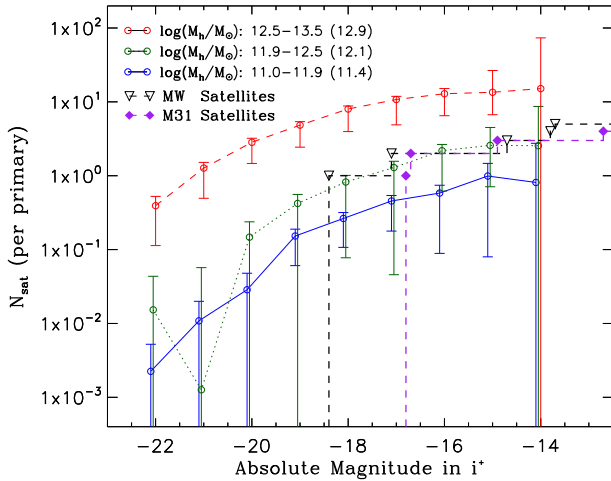


Figure 13. Satellite luminosity functions, for three ranges of primary halo mass, estimated using method B (see Appendix C for the same results derived using method A). The mean halo mass for each range is given in parentheses. The observed luminosity functions of the MW and M31 are shown for comparison (black and purple points and lines). These magnitudes have been converted from the V band, assuming a fixed average colour of $V - i^+ = 0.3$.

systems, explaining the flatter slope. B18 also point out that they may be biased towards preferentially identifying multiple systems at low stellar mass, due to their bluer colours. We conclude that our results are consistent with B18, once again despite very different methods and samples.

5.4 The satellite luminosity function

Finally, we can use method B to estimate satellite luminosity functions directly (for comparison, results using method A are shown in Appendix C). In Fig. 13, we present the luminosity function of satellites for three subsets of primaries. The subsets were chosen such that the mean halo mass of the middle bin, $\langle M_h \rangle = 12.1 M_\odot$, is close to the estimated mass of the MW or M31 (Bland-Hawthorn & Gerhard 2016; Posti & Helmi 2019), such that we can compare to the observed luminosity functions for these systems. For method B, we are able to measure the satellite luminosity functions reliably down to absolute magnitudes of -14 . Within this magnitude range, the observed abundance of satellites around the MW and M31 is close to the average value. One exception is at the bright end of the MW satellite luminosity function, where the presence of the LMC and SMC represents a slight (1σ – 2σ) departure from the average. This unusual feature of the MW’s satellite population has been noted and discussed extensively elsewhere (e.g. Robotham et al. 2012 – see Speller & Taylor 2014, for further references). We will consider satellite luminosity and stellar mass functions in more detail in a subsequent paper.

6 TESTING FOR SYSTEMATIC UNCERTAINTIES IN THE METHOD

While we have shown that our method produces estimates of satellite abundance consistent with previous studies using larger samples, there remain a number of choices, assumptions, or free parameters in the method that could take on different values. In this section, we will perform a set of tests to understand the effects of the various assumptions and free parameters in the clustering method.




Figure 14. Results of a null test where the halo component of the clustering signal is measured using galaxies outside the velocity limits of the ROI. A black dashed line corresponding to no clustering is included for comparison.

6.1 Null test

First, as a null test, we calculated the clustering signals around our initial sample of primaries, using only secondaries that lie within the radial cuts, but outside the redshift cuts we defined in Section 3.2. Following the procedure in Section 4, we measured the surface number density of the secondaries for the 20 bins in primary mass and redshift, and refit our surface density model. The satellite component from the fit is shown in Fig. 14.

Most of secondaries in this test should not be physically associated with the primaries, except for a small number of real satellites that are scattered into the foreground or background by the redshift errors. Thus, we expect the clustering signal to be close to zero, and to show little dependence on the primary halo mass. This is confirmed in Fig. 14.

6.2 Parameter and systematic tests

We have tested for systematic effects and uncertainties in our method by varying the parameters that define the initial primary selection and the calculation of the satellite probability. The tests include

- (i) Using radial cuts of $2R_{\text{vir}}$ or $3.5R_{\text{vir}}$ to define the radial extent of the ROI.
- (ii) Using velocity cuts of $1\sigma_{\Delta v}$ or $3\sigma_{\Delta v}$ to define the velocity extent of the ROI.
- (iii) Increasing/decreasing the stellar mass by 0.16 dex, which is comparable to or larger than the typical stellar mass uncertainties in our data.
- (iv) Increasing/decreasing the halo mass derived from our fiducial SHMR by 50 per cent.
- (v) Varying the slope of bias-corrected SHMR (see Appendix A) at the high-mass end to 1.5 and to 2.5, with respect to the original value of ~ 2.1 from B13.
- (vi) Varying the definition of the virial radius, increasing or decreasing it by 20 per cent.
- (vii) Adding 1σ scatter to the initial stellar masses before ranking them. (We perform this test three times to check the consistency of the potential effects.)
- (viii) Keeping the primary–secondary selection fixed, but adding $0.2R_{\text{vir}}$ scatter to the coordinates of the primaries, to test the effects on potential miscentring. (We repeat this test three times to check for the consistency of the effects.)

In each case, all other parameters and steps in the method are kept fixed. The results of these tests are summarized in Table 2.

Table 2. The fitting results of all tests.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
F0	0.33 ^{+0.09} _{-0.08}	0.29 ^{+0.08} _{-0.10}	79.5 ^{+0.6} _{-0.6}	21.6 ^{+0.5} _{-0.5}
LMh	0.29 ^{+0.08} _{-0.07}	0.24 ^{+0.09} _{-0.10}	79.0 ^{+0.5} _{-0.5}	21.7 ^{+0.5} _{-0.5}
SMh	0.28 ^{+0.08} _{-0.07}	0.37 ^{+0.06} _{-0.08}	79.0 ^{+0.5} _{-0.6}	21.6 ^{+0.5} _{-0.5}
LMs	0.30 ^{+0.09} _{-0.08}	0.12 ^{+0.09} _{-0.12}	78.8 ^{+0.5} _{-0.5}	21.7 ^{+0.5} _{-0.5}
SMS	0.45 ^{+0.12} _{-0.10}	0.33 ^{+0.08} _{-0.09}	80.2 ^{+0.6} _{-0.6}	22.0 ^{+0.5} _{-0.5}
$\sigma_{\Delta v}$ 1	0.32 ^{+0.08} _{-0.07}	0.22 ^{+0.07} _{-0.09}	78.2 ^{+0.5} _{-0.5}	21.2 ^{+0.4} _{-0.4}
$\sigma_{\Delta v}$ 3	0.21 ^{+0.10} _{-0.09}	0.29 ^{+0.08} _{-0.10}	79.9 ^{+0.7} _{-0.7}	21.1 ^{+0.6} _{-0.5}
R2	0.30 ^{+0.06} _{-0.06}	0.32 ^{+0.06} _{-0.06}	78.7 ^{+0.5} _{-0.5}	21.4 ^{+0.4} _{-0.4}
R3.5	0.23 ^{+0.09} _{-0.08}	0.28 ^{+0.06} _{-0.08}	82.2 ^{+0.5} _{-0.5}	21.0 ^{+0.4} _{-0.4}
LR	0.28 ^{+0.09} _{-0.08}	0.14 ^{+0.08} _{-0.10}	78.1 ^{+0.5} _{-0.5}	21.2 ^{+0.4} _{-0.4}
SR	0.26 ^{+0.10} _{-0.09}	0.28 ^{+0.09} _{-0.12}	80.7 ^{+0.7} _{-0.7}	22.1 ^{+0.5} _{-0.5}
H SHMR	0.20 ^{+0.06} _{-0.06}	0.31 ^{+0.05} _{-0.07}	80.2 ^{+0.5} _{-0.5}	21.3 ^{+0.3} _{-0.4}
L SHMR	0.34 ^{+0.10} _{-0.09}	0.28 ^{+0.07} _{-0.09}	79.4 ^{+0.6} _{-0.6}	21.6 ^{+0.5} _{-0.5}
RS 1	0.26 ^{+0.08} _{-0.07}	0.28 ^{+0.07} _{-0.09}	78.6 ^{+0.6} _{-0.5}	21.6 ^{+0.5} _{-0.5}
RS 2	0.25 ^{+0.08} _{-0.07}	0.29 ^{+0.07} _{-0.08}	78.7 ^{+0.6} _{-0.5}	21.6 ^{+0.5} _{-0.5}
RS 3	0.25 ^{+0.08} _{-0.07}	0.28 ^{+0.07} _{-0.09}	78.6 ^{+0.5} _{-0.5}	21.6 ^{+0.5} _{-0.6}
CS 1	0.34 ^{+0.12} _{-0.10}	0.05 ^{+0.11} _{-0.15}	79.4 ^{+0.5} _{-0.5}	21.8 ^{+0.5} _{-0.5}
CS 2	0.25 ^{+0.10} _{-0.08}	0.11 ^{+0.9} _{-0.12}	79.3 ^{+0.5} _{-0.5}	21.8 ^{+0.5} _{-0.5}
CS 3	0.26 ^{+0.08} _{-0.08}	0.21 ^{+0.08} _{-0.10}	79.6 ^{+0.6} _{-0.6}	21.6 ^{+0.5} _{-0.5}

Notes. [Key to the tests:

F0 = original fit result;

LMh = larger halo mass;

SMh = smaller halo mass;

LMs = larger stellar mass;

SMS = smaller stellar mass;

$\sigma_{\Delta v}$ 1 = using $1\sigma_{\Delta v}$ velocity cut to define ROI;

$\sigma_{\Delta v}$ 2 = using $3\sigma_{\Delta v}$ velocity cut for ROI;

R2 = using $2R_{\text{vir}}$ projected separation cut for ROI;

R3.5 = using $3.5R_{\text{vir}}$ projected separation cut for ROI;

LR = larger R_{vir} (increasing the primary virial radius by 20 per cent);

SR = smaller R_{vir} (decreasing the primary virial radius by 20 per cent);

H SHMR = using a higher slope of HSMR at high-mass end;

L SHMR = using a lower slope of HSMR at high-mass end;

RS = ranking shuffle (adding 1σ scatter to the primary stellar masses before ranking them);

CS = centring shift (adding $0.2R_{\text{vir}}$ scatter to the coordinates of the primaries before measuring clustering).]

6.3 Discussion

Considering the results of the individual tests in detail, changing the halo mass mainly just shifts the points horizontally on the $S_{\text{halo}}-M_h$ plot, so the fitted parameters a and b remain relatively constant. Changing the stellar mass has quite a different effect, however. As the HSMR is shallow at the low-mass end but steep at the high-mass end, increasing or decreasing the stellar mass does not change the inferred halo mass much at the low-mass end, but can produce significant change at the high-mass end. As a result, the slope (a) of the $S_{\text{halo}}-M_h$ relation is more strongly affected.

Similarly, changing the slope of the HSMR at the high-mass end will mainly affect the halo mass estimates in this range. As a result, the slope a is shifted systematically to higher or lower values. As for the mass ranking test, in addition to the random scatter in individual mass estimates and resulting variations in detailed primary selection, there is a net change in the mass function. Since there are more low-mass galaxies than high-mass ones, adding the random scatter tends to increase the number of massive galaxies relative to the fiducial

catalogue. This leads to slight shifts in the fitted satellite abundance, although they are less important than in the case of the mass tests.

The parameter values obtained in each test are given in Table 2. Tests that produce a variation of more than 2σ in the fitted parameters are highlighted in bold. We note that only one test (our first re-centring test) produces a significant change in the parameters of the satellite component. Three of the 38 tests produce significant deviations in the background fit, but this is only slightly higher than the expected rate of 2σ deviations given the random errors (8 per cent versus 5 per cent). Thus overall, the systematic uncertainties associated with our tests do not appear to significantly increase the random errors quoted in the fiducial model fit.

7 CONCLUSIONS

In this work, we have developed and tested a method for quantifying satellite abundance, using galaxy clustering. The method establishes a basic template for the radial dependence and amplitude of the satellite component of the clustering signal by using a subsample of isolated (or at least locally dominant) systems, but then applies this template iteratively to estimate the probability that any given galaxy in the field is a satellite of a nearby system. (Note that the form of the template assumes that the surface number density of background galaxies is inversely proportional to the square of angular-diameter distance; this assumption works well at low redshift, but may need modification if applying the method at higher redshift.) In that sense it is similar to crowded-field photometry, where an initial sample of isolated stars is used to determine the point spread function of the image, and that point spread function is then applied iteratively to the entire field. Using our method, we have estimated satellite abundance as a function of primary stellar and halo mass, and also measured the satellite luminosity function, over a very broad range of primary halo mass ($10^{10}-10^{13.5} M_{\odot}$). We have also tested the method for systematic uncertainties by varying the model parameters, and found variations in the final results that are generally smaller than our random error estimates.

We have compared the results of this new technique to several previous estimates of satellite abundance from the literature, which were derived using larger catalogues. Our results are fully consistent with those of Conroy et al. (2006) at the high-mass end, and of Besla et al. (2018) at the low-mass end, while covering a much larger range in primary mass overall. We have also compared our measured luminosity functions to those of the dominant LG galaxies, assuming an average halo mass of $12.1 M_{\odot}$ for these systems. The LG satellite populations seem fairly typical, with the exception of the bright satellites of the MW (the LMC and the SMC), as noted previously in the literature. The main purpose of this paper was to describe and validate our method; in subsequent work, we will consider in more detail the properties of the detected satellite populations, including their spatial distribution, colours, star formation rates, and dependence on primary properties.

The COSMOS catalogue was chosen for this work for its deep photometry and extremely accurate photo- z s. Other deep surveys with accurate distance information will also be good candidates to apply our method. One potentially important survey is planned with SPHEREx,³ an all-sky survey satellite with a wide-field spectral imager. SPHEREx is currently scheduled to launch in 2024, and will produce, during its 2-yr mission, four all-sky maps, with hundreds of millions of near-infrared stellar and galactic spectra

³See this <https://spherex.caltech.edu/> for more details.

(0.75–5.0 micron) (Spangelo et al. 2015; Bock & SPHEREx Science Team 2018). The redshifts in the SPHEREx surveys are estimated by fitting template SEDs to observations, similarly to COSMOS. While COSMOS used photometry in 30 bands to derive its phot-zs, SPHEREx will produce low-resolution ($R \sim 20\text{--}100$) spectra, with a similar final redshift accuracy, as discussed in Stickley et al. (2016). While the main survey will be shallower than COSMOS, two regions at the polar caps will be visited multiple times, providing ~ 100 square deg of coverage to a depth similar to COSMOS. Thus, SPHEREx should provide a redshift catalogue of similar accuracy to the COSMOS catalogue used here, but covering an area 50 times larger. The resulting increase in the SNR of the clustering signal would allow much finer binning in primary or secondary properties, giving a much more detailed view of the relationship between satellites and their CGs.

ACKNOWLEDGEMENTS

We thank N. Afshordi, A. Broderick, M. Balogh, M. Hudson, and our friends and collaborators from the COSMOS survey for their comments and advice. JET acknowledges support from the Natural Sciences and Engineering Research Council of Canada, through a Discovery Grant. The COSMOS 2015 catalogue is based on data products from observations made with European Southern Observatory (ESO) Telescopes at the La Silla Paranal Observatory under ESO programme ID 179. A-2005 and on data products produced by TERAPIX and the Cambridge Astronomy Survey Unit on behalf of the UltraVISTA consortium.

DATA AVAILABILITY

Most of the data underlying this article are publicly available. The COSMOS 2015 catalogue (Laigle et al. 2016) can be accessed from the COSMOS website, at <http://cosmos.astro.caltech.edu/page/photom>. A few spectroscopic redshifts from the COSMOS collaboration are not yet published (M. Salvato, private communication), but will be shared on reasonable request to the corresponding author with permission of the COSMOS collaboration. The derived data generated in this research will also be shared on reasonable request to the corresponding author.

REFERENCES

Anderhalden D., Schneider A., Macciò A. V., Diemand J., Bertone G., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 014
 Bechtol K. et al., 2015, *ApJ*, 807, 50
 Behroozi P. S., Wechsler R. H., Conroy C., 2013, *ApJ*, 770, 57
 Berlind A. A., Weinberg D. H., 2002, *ApJ*, 575, 587
 Besla G. et al., 2018, *MNRAS*, 480, 3376
 Bland-Hawthorn J., Gerhard O., 2016, *ARA&A*, 54, 529
 Bock J., SPHEREx Science Team, 2018, American Astronomical Society Meeting Abstracts #231, #354.21
 Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., 2010, *MNRAS*, 406, 896
 Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2011, *MNRAS*, 415, L40
 Bruzual G., Charlot S., 2003, *MNRAS*, 344, 1000
 Bryan G. L., Norman M. L., 1998, *ApJ*, 495, 80
 Bullock J. S., Boylan-Kolchin M., 2017, *ARA&A*, 55, 343
 Bullock J. S., Kravtsov A. V., Weinberg D. H., 2000, *ApJ*, 539, 517
 Busha M. T., Wechsler R. H., Behroozi P. S., Gerke B. F., Klypin A. A., Primack J. R., 2011, *ApJ*, 743, 117
 Capak P. et al., 2007, *ApJS*, 172, 99
 Carlin J. L. et al., 2016, *ApJ*, 828, L5

Carlsten S. G., Beaton R. L., Greco J. P., Greene J. E., 2019, *ApJ*, 879, 13
 Chua K. T. E., Pillepich A., Rodriguez-Gomez V., Vogelsberger M., Bird S., Hernquist L., 2017, *MNRAS*, 472, 4343
 Cohen Y. et al., 2018, *ApJ*, 868, 96
 Conroy C., Wechsler R. H., Kravtsov A. V., 2006, *ApJ*, 647, 201
 Cooray A., Sheth R., 2002, *Phys. Rep.*, 372, 1
 Danieli S., van Dokkum P., Merritt A., Abraham R., Zhang J., Karachentsev I. D., Makarova L. N., 2017, *ApJ*, 837, 136
 Danieli S., van Dokkum P., Abraham R., Conroy C., Dolphin A. E., Romanowsky A. J., 2020, *ApJ*, 895, L4
 Dekel A., Silk J., 1986, *ApJ*, 303, 39
 Drlica-Wagner A. et al., 2015, *ApJ*, 813, 109
 Elbert O. D., Bullock J. S., Garrison-Kimmel S., Rocha M., Oñorbe J., Peter A. H. G., 2015, *MNRAS*, 453, 29
 Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
 Fry A. B. et al., 2015, *MNRAS*, 452, 1468
 Gao L., White S. D. M., Jenkins A., Stoehr F., Springel V., 2004, *MNRAS*, 355, 819
 Geha M. et al., 2017, *ApJ*, 847, 4
 Gnedin N. Y., Kravtsov A. V., 2006, *ApJ*, 645, 1054
 Governato F. et al., 2010, *Nature*, 463, 203
 Guo Q., Cole S., Eke V., Frenk C., 2011, *MNRAS*, 417, 370
 Guo Q., Cole S., Eke V., Frenk C., 2012a, *MNRAS*, 427, 428
 Guo H., Zehavi I., Zheng Z., 2012b, *ApJ*, 756, 127
 Ilbert O. et al., 2009, *ApJ*, 690, 1236
 Ilbert O. et al., 2013, *A&A*, 556, A55
 Jiang F., van den Bosch F. C., 2017, *MNRAS*, 472, 657
 Kang X., Wang L., Luo Y., 2016, *MNRAS*, 460, 2152
 Kennedy R., Frenk C., Cole S., Benson A., 2014, *MNRAS*, 442, 2487
 Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, *ApJ*, 522, 82
 Koposov S. E., Belokurov V., Torrealba G., Evans N. W., 2015, *ApJ*, 805, 130
 Kravtsov A. V., Gnedin O. Y., Klypin A. A., 2004, *ApJ*, 609, 482
 Laigle C. et al., 2016, *ApJS*, 224, 24
 Leauthaud A. et al., 2012, *ApJ*, 744, 159
 Lilly S. J. et al., 2007, *ApJS*, 172, 70
 Liu L., Gerke B. F., Wechsler R. H., Behroozi P. S., Busha M. T., 2011, *ApJ*, 733, 62
 Lokas E. L., Kazantzidis S., Mayer L., 2012, *ApJ*, 751, L15
 Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, *MNRAS*, 439, 300
 Lunnan R., Vogelsberger M., Frebel A., Hernquist L., Lidz A., Boylan-Kolchin M., 2012, *ApJ*, 746, 109
 Macciò A. V., Fontanot F., 2010, *MNRAS*, 404, L16
 McCracken H. J. et al., 2012, *A&A*, 544, A156
 Mashchenko S., Wadsley J., Couchman H. M. P., 2008, *Science*, 319, 174
 Mayer L., Mastropietro C., Wadsley J., Stadel J., Moore B., 2006, *MNRAS*, 369, 1021
 Miyazaki S. et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 8446, Ground-Based and Airborne Instrumentation for Astronomy IV. SPIE, Bellingham, p. 84460Z
 Mobasher B. et al., 2007, *ApJS*, 172, 117
 Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, *ApJ*, 524, L19
 More S., van den Bosch F. C., Cacciato M., Mo H. J., Yang X., Li R., 2009, *MNRAS*, 392, 801
 Nadler E. O., Gluscevic V., Boddy K. K., Wechsler R. H., 2019, *ApJ*, 878, L32
 Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
 Nelson D. et al., 2015, *Astron. Comput.*, 13, 12
 Newman J. A. et al., 2013, *ApJS*, 208, 5
 Newton O., Cautun M., Jenkins A., Frenk C. S., Helly J. C., 2018, *MNRAS*, 479, 2853
 Pawlowski M. S., Famaey B., Merritt D., Kroupa P., 2015, *ApJ*, 815, 19
 Peacock J. A., Smith R. E., 2000, *MNRAS*, 318, 1144
 Polletta M. et al., 2007, *ApJ*, 663, 81
 Posti L., Helmi A., 2019, *A&A*, 621, A56

- Reddick R. M., Wechsler R. H., Tinker J. L., Behroozi P. S., 2013, *ApJ*, 771, 30
- Robotham A. S. G. et al., 2012, *MNRAS*, 424, 1448
- Sales L. V., Wang W., White S. D. M., Navarro J. F., 2013, *MNRAS*, 428, 573
- Scoville N. et al., 2007, *ApJS*, 172, 1
- Seljak U., 2000, *MNRAS*, 318, 203
- Smith A. et al., 2019, *MNRAS*, 484, 1285
- Somerville R. S., Lee K., Ferguson H. C., Gardner J. P., Moustakas L. A., Giavalisco M., 2004, *ApJ*, 600, L171
- Spangolo S. C., Katti R. M., Unwin S. C., Bock J. J., 2015, *J. Astron. Telesc. Instrum. Syst.*, 1, 037001
- Speller R., Taylor J. E., 2014, *ApJ*, 788, 188
- Spergel D. N., Steinhardt P. J., 2000, *Phys. Rev. Lett.*, 84, 3760
- Stickley N. R., Capak P., Masters D., de Putter R., Doré O., Bock J., 2016, preprint (arXiv:1606.06374)
- Strigari L. E., Wechsler R. H., 2012, *ApJ*, 749, 75
- Taylor J. E., Babul A., 2001, *ApJ*, 559, 716
- Taylor J. E., Silk J., Babul A., 2004, in Ryder S., Pisano D., Walker M., Freeman K., eds, Proc. IAU Symp. 220, Dark Matter in Galaxies. Kluwer, Dordrecht, p. 91
- Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, *ApJ*, 688, 709
- van Dokkum P., Danieli S., Cohen Y., Romanowsky A. J., Conroy C., 2018, *ApJ*, 864, L18
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
- Wang W., White S. D. M., 2012, *MNRAS*, 424, 2574
- Wang W., Sales L. V., Henriques B. M. B., White S. D. M., 2014, *MNRAS*, 442, 1363
- Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2014, *MNRAS*, 439, 2687
- Wetzel A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, *ApJ*, 827, L23
- Xi C., Taylor J. E., Massey R. J., Rhodes J., Koekemoer A., Salvato M., 2018, *MNRAS*, 478, 5336
- Yang X., Mo H. J., van den Bosch F. C., Pasquali A., Li C., Barden M., 2007, *ApJ*, 671, 153
- Yang X., Mo H. J., van den Bosch F. C., 2009, *ApJ*, 695, 900
- York D. G. et al., 2000, *AJ*, 120, 1579
- Zehavi I. et al., 2002, *ApJ*, 571, 172
- Zhang D., Luo Y., Kang X., 2019, *MNRAS*, 486, 2440

APPENDIX A: BIAS IN HALO MASSES DERIVED FROM THE SHMR

Throughout this work, we assume the SHMR proposed by Behroozi et al. (2013):

$$\log_{10}(M_*(M_h)) = \log_{10}(\epsilon M_1) + f\left(\log_{10}\left(\frac{M_h}{M_1}\right)\right) - f(0), \quad (\text{A1})$$

where the function $f(x)$ is defined as

$$f(x) = -\log_{10}(10^{\alpha x} + 1) + \delta \frac{(\log_{10}(1 + \exp(x)))^\gamma}{1 + \exp(10^{-x})}. \quad (\text{A2})$$

The free parameters vary with redshift as follows:

$$\begin{aligned} \log_{10}(M_1) &= M_{1,0} + (M_{1,a}(a-1) + M_{1,z}z) \exp(-4a^2) \\ \log_{10}(\epsilon) &= \epsilon_0 + (\epsilon_a(a-1) + \epsilon_z z) \exp(-4a^2) + \epsilon_{a,z}(a-1) \\ \alpha &= \alpha_0 + (\alpha_a(a-1)) \exp(-4a^2) \\ \delta &= \delta_0 + (\delta_a(a-1) + \delta_z z) \exp(-4a^2) \\ \gamma &= \gamma_0 + (\gamma_a(a-1) + \gamma_z z) \exp(-4a^2), \end{aligned} \quad (\text{A3})$$

where $a = 1/(1+z)$ is the scale factor. [The 1σ uncertainty range in these parameter values is listed on p. 9 of Behroozi et al. (2013).] In our case, since our primary sample covers the fairly narrow redshift (z) range of 0–0.25, little variation is predicted in the SHMR. Thus,




Figure A1. Stellar mass errors versus stellar mass in the mock sample (for clarity, only a third of the data points are shown), together with the mean relation in bins of stellar mass (red points and error bars).

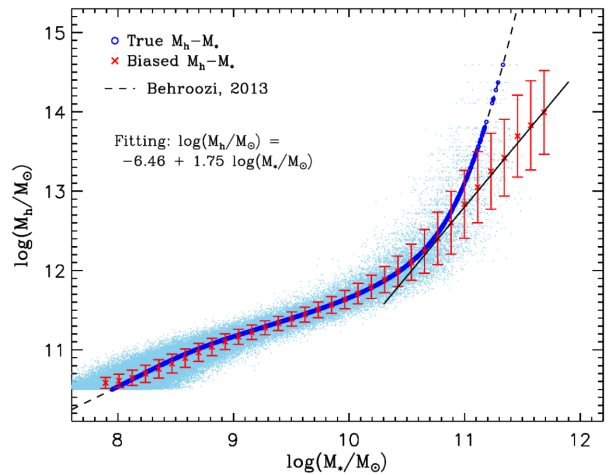


Figure A2. Halo-to-stellar mass relation of the Monte Carlo samples. The dark blue dots are our Monte Carlo halo mass sample with stellar masses assigned using B13. The light blue points show the ‘observed’ stellar masses of the sample after adding intrinsic scatter and observational errors. (Note that only 15 per cent of the mock sample is shown for clarity.) The red dots with error bars are the average halo mass in each bin of ‘observed’ stellar mass. The solid black line is a linear fit to the red dots [in $\log(M_*/M_\odot)$], over the range of 10.7–11.8.

we simply use an intermediate redshift of $z = 0.15$ for the analysis below.

The Behroozi et al. (2013) relations are theoretical, unbiased mean values of M_* , given a specific halo mass M_h . In any real survey, this relationship will be biased by intrinsic scatter and observational errors (B13; Leauthaud et al. 2012). To quantify this bias for the COSMOS catalogue, we generated a Monte Carlo sample of 12 000 haloes selected at random from the halo mass function given in Tinker et al. (2008). ‘True’ stellar masses were calculated for these objects using the SHMR given above. We then added intrinsic scatter and random errors to each stellar mass 50 times independently, to simulate an ‘observed’ stellar mass sample. The intrinsic scatter in the SHMR is about 0.14–0.2 dex at a redshift of $z = 0$ (More et al. 2009; Yang, Mo & van den Bosch 2009; Reddick et al. 2013), and there is no evidence of any trend with mass, at least down to halo

masses of $10^{12} M_{\odot}$ (Behroozi et al. 2013; Reddick et al. 2013). Thus, we added an intrinsic scatter of 0.15 to all stellar masses derived for our ‘observed’ sample. Average observational errors, as a function of stellar mass, were estimated directly from the COSMOS 2015 catalogue, as shown in Fig. A1. We added these in quadrature to determine the final stellar masses of the mock sample.

The light blue points in Fig. A2 show the ‘observed’ stellar masses of the mock sample, after adding the intrinsic scatter and observational errors. The dark blue points show the underlying ‘true’ stellar masses, while the black dashed line shows the theoretical SHMR from B13. The red points show the mean halo mass in each ‘observed’ stellar mass bin. We can see that the ‘observed’ SHMR follows the theoretical SHMR reasonably well at low masses, but departs from it at the high masses. Given this pattern, we will use a two-part SHMR to assign halo masses in our work. For stellar masses of $10^{10.7} M_{\odot}$ or less, we use an unmodified B13 SHMR, while for masses greater than $10^{10.7}$, we use the linear fit to the average ‘observed’ values listed and shown as a black solid line on the plot.

APPENDIX B: MEASURING THE MASKING AND FIELD BOUNDARIES

Regions of the COSMOS field have poor photometry in one or more bands, due to contamination from bright stars, internal reflections, or other artefacts. Data from these regions are tagged with a ‘masking’ flag (‘FLAG_PETER’ in the COSMOS 2015 catalogue), which can be used to exclude those data from further analysis. The shape of these masked regions, together with the field boundaries, needs to be measured to determine the area completeness η around any given primary. Although detailed mask files are available for the COSMOS field, we found it less computationally demanding for our work to use a single, approximate mask image with coarser spatial sampling. We use the method described in Xi et al. (2018) to generate this global mask. First, a coarse map consisting of 390×390 cells is constructed, covering the whole COSMOS field. We search for objects in each cell of this map, to determine whether it should be included or masked out. In the first round, any cell with one or fewer objects counts is selected as a potential masked region. In the second round, these candidates are confirmed as masked if they have one or more neighbouring cells with no counts. This two-step selection process reduces to 0.0026 per cent the probability of artificially eliminating cells due to Poisson fluctuations in their object counts. The map resolution and count threshold were determined empirically after testing various resolutions from 200×200 to 600×600 , with different thresholds in each case. We found that the effect on the clustering signal of variations in the masking parameters is small, producing variations in S_{halo} of roughly 5 per cent or less. The final resolution was selected to provide the most accurate overall mask, relative to the full images.

Given a single global mask for the COSMOS field, we then generated a large, random sample of points, and used the distribution of the points around each primary to estimate its area completeness as a function of projected separation. For each galaxy, we counted the number of random samples in projected radial bins with and without applying the masks and boundaries. Each bin had a size of $0.2R_{\text{vir}}$ of the galaxy, up to $3.6R_{\text{vir}}$. The area completeness is then

$$\eta(R_{\text{P}}) \equiv \frac{A_{\text{M}}(R_{\text{P}})}{A_{\text{T}}(R_{\text{P}})} = \frac{N_{\text{M}}(R_{\text{P}})}{N_{\text{T}}(R_{\text{P}})}, \quad (\text{B1})$$

where A_{M} and A_{T} are the masked and total areas, and N_{M} and N_{T} are the random counts with and without masking, respectively.

Besides the area completeness in individual radial bins, we also measured the total area completeness of each primary within $3.0R_{\text{vir}}$. Galaxies with poor completeness were excluded from the primary sample, as described in the main text.

APPENDIX C: COMPARING BACKGROUND ESTIMATION METHODS A AND B

While method A is simple to implement, it may introduce systematic biases in the inferred satellite properties, as described in Section 5.1. Method B removes the contribution from the background statistically, and should produce less biased, albeit noisier, results.

Fig. C1 shows the (cumulative) satellite luminosity function for three sets of primaries with different halo mass ranges, using methods A (dashed lines) and B (solid lines). Overall, method A produces a luminosity function with a steeper slope, which continues to rise at faint magnitudes, whereas for method B, the cumulative luminosity function flattens. As shown in Fig. 13, the method B results are in better agreement with LG data.




Figure C1. Cumulative satellite luminosity functions estimated using methods A (dashed) and B (solid), for primaries in three halo mass bins (as in Fig. 13).

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.