

Satellite mass functions and the faint end of the galaxy mass–halo mass relation in LCDM

Isabel M. E. Santos-Santos ^{1,2}★, Laura V. Sales ³, Azadeh Fattahi ² and Julio F. Navarro¹

¹*Department of Physics and Astronomy, University of Victoria, Victoria, BC V8P 5C2, Canada*

²*Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*

³*Department of Physics and Astronomy, University of California, Riverside, 900 University Avenue, Riverside, CA 92507, USA*

Accepted 2022 July 18. Received 2022 June 13; in original form 2021 October 29

ABSTRACT

The abundance of the faintest galaxies provides insight into the nature of dark matter and the process of dwarf galaxy formation. In the LCDM scenario, low-mass haloes are so numerous that the efficiency of dwarf formation must decline sharply with decreasing halo mass in order to accommodate the relative scarcity of observed dwarfs and satellites in the Local Group. The nature of this decline contains important clues to the mechanisms regulating the onset of galaxy formation in the faintest systems. We explore here two possible models for the stellar mass (M_*)–halo mass (M_{200}) relation at the faint end, motivated by some of the latest LCDM cosmological hydrodynamical simulations. One model includes a sharp mass threshold below which no luminous galaxies form, as expected if galaxy formation proceeds only in systems above the hydrogen-cooling limit. In the second model, M_* scales as a steep power law of M_{200} with no explicit cut-off, as suggested by recent semi-analytical work. Although both models predict satellite numbers around Milky Way-like galaxies consistent with current observations, they predict vastly different numbers of ultrafaint dwarfs and of satellites around isolated dwarf galaxies. Our results illustrate how the satellite mass function around dwarfs may be used to probe the M_* – M_{200} relation at the faint end and to elucidate the mechanisms that determine which low-mass haloes ‘light up’ or remain dark in the LCDM scenario.

Key words: galaxies: dwarf – galaxies: haloes – galaxies: luminosity function, mass function.

1 INTRODUCTION

Ultrafaint dwarfs, defined here as dwarf galaxies with stellar masses $M_* < 10^5 M_\odot$ (Bullock & Boylan-Kolchin 2017), are typically systems whose extremely low surface brightness ($\mu_v \geq 27$ mag arcsec^{−2}) hinders their discovery and makes follow-up studies extremely difficult. Indeed, although recent efforts have led to the discovery of dozens of ultrafaints in the Milky Way (MW) halo (see Simon 2019, and references therein), it remains unclear how many more of them may still lurk undetected in the vicinity of our Galaxy.

The ultrafaint population also remains largely unexplored in external galaxies, with only loose constraints available on the massive end of the ultrafaint regime in M31 (see the dwarf galaxy catalogue compiled and maintained by McConnachie 2012). Identifying isolated ultrafaints in the field is even more difficult, with few, if any, reported so far outside the Local Group.

Because of their extreme intrinsic faintness, few bright stars are available for spectroscopic study in ultrafaints, even when using some of the largest ground-based telescopes. This implies that the characterization of some of their basic properties, such as their metallicity distribution, elemental abundances, or velocity dispersion, is subject to large uncertainty. Poorly determined velocity dispersions, in particular, affect our ability to estimate halo masses

and to constrain the relation between stellar mass (M_*) and halo virial¹ mass (M_{200}) at the very faint end of the galaxy luminosity function.

Indeed, our best constraints on the M_* – M_{200} relation at the faint end arguably come from abundance-matching techniques (Conroy, Wechsler & Kravtsov 2006; Guo et al. 2010; Behroozi, Wechsler & Conroy 2013; Moster, Naab & White 2013). Because the galaxy stellar mass function around $M_* \sim 10^8 M_\odot$ (the faintest luminosities for which it is well constrained; see e.g. Baldry et al. 2012) is substantially shallower than the LCDM halo mass function in that regime (Springel et al. 2008; Boylan-Kolchin et al. 2009), it is clear that galaxy formation must become increasingly inefficient towards decreasing halo masses. Characterizing this decline in galaxy formation efficiency at the low-mass end is difficult, and there is so far no consensus on how steep the decline is, on what the scatter in M_* at fixed M_{200} might be, and on whether there is a characteristic ‘threshold’ halo mass below which no luminous galaxy forms in LCDM.

The lack of consensus concerns not only abundance-matching studies, but also direct cosmological simulations of the formation of the faintest galaxies. For example, Local Group simulations

¹We shall use halo ‘virial’ properties defined within a radius, r_{200} , enclosing a mean density 200 times the critical density for closure. A subscript ‘200’ identifies quantities defined within or at that radius.

* E-mail: isabel.santos@durham.ac.uk

from the APOSTLE project (Fattahi et al. 2016; Sawala et al. 2016) suggest a relation with a fairly sharp cut-off at low halo masses, where few, if any, isolated haloes with V_{\max} below $\sim 15 \text{ km s}^{-1}$ host a galaxy (Fattahi et al. 2018). At least qualitatively, this is the behaviour expected in scenarios where luminous galaxy formation only proceeds in haloes with masses exceeding the ‘hydrogen-cooling limit’ (HCL) set by the primordial abundance cooling function after accounting for the presence of an evolving, ionizing ultraviolet (UV) background (see e.g. Gnedin 2000; Okamoto, Gao & Theuns 2008; Benitez-Llambay & Frenk 2020).

On the other hand, some cosmological simulations suggest that even haloes below the HCL may be able to form stars, so that no clear minimum ‘threshold mass’ for galaxy formation exists. For example, FIRE-2 simulations (Wetzel et al. 2016; Hopkins et al. 2018; Wheeler et al. 2019) seem better described by a power-law M_*-M_{200} relation similar to that reported by Brook et al. (2014) and which extends well below the HCL mass.

This argument has been strengthened by semi-analytical models that attempt to reproduce simultaneously the MW satellite mass function and its radial distribution. Because tides may, in principle, disrupt subhaloes near the MW disc, accounting for the large number of ultrafaint satellites discovered in the inner ~ 40 kpc of the MW halo has led to the suggestion that populating subhaloes well below the HCL with luminous galaxies may be needed (Graus et al. 2019; Kelley et al. 2019).

However, there is still substantial uncertainty about whether Galactic tides are actually able to fully disrupt cuspy LCDM subhaloes (van den Bosch et al. 2018; Errani & Navarro 2021) and no cosmological simulation has actually reached the ultrafaint regime probed by observations. Despite these uncertainties, it is clear that simulation predictions for the faintest dwarfs appear to differ, depending on the resolution and subgrid physics adopted in the simulations (see e.g. Munshi et al. 2019). This is problematic, as the steep halo mass function in LCDM implies that even small differences in the stellar mass–halo mass relation should result in large differences in the expected number of faint galaxies.

We explore here how the abundance of ultrafaint satellites may be used to place constraints on the behaviour of the M_*-M_{200} relation at the faint end. Their abundance around isolated dwarf primaries is particularly constraining. This is because the subhalo mass function is well approximated by a power law (Springel et al. 2008) and therefore a power-law stellar mass–halo mass relation would result in ‘self-similar’ satellite mass functions independent of primary mass (Sales et al. 2013). This is a clear prediction that can be used to gain insight into the shape of the stellar mass–halo mass relation for primaries at the faint end.

We explore these issues here, and argue that the ultrafaint satellites of isolated dwarf galaxies are a promising way to elucidate how the faintest galaxies form and populate dark haloes at the low-mass end. This paper is organized as follows. We begin by motivating in Section 3 two particular analytical forms of the faint-end M_*-M_{200} relation (a power law and one with an explicit low-mass cut-off) based on results from recent cosmological hydrodynamical simulations (Section 2). We validate the ‘cut-off’ model in Section 4.1 by reproducing results from the APOSTLE runs. We then compare the results from both our models for the ultrafaint satellite population of dwarf primaries spanning a wide range of stellar mass (Section 4.2), and then contrast these results with available data for the Local Group in Section 5. We conclude with predictions for future satellite surveys of ultrafaint dwarfs around primaries such as the Large Magellanic

Cloud (LMC) in Section 5.3 and summarize our main results in Section 6.

2 NUMERICAL METHODS

We shall use results from a number of recent cosmological hydrodynamical simulations of dwarf galaxy formation in LCDM. These include simulations of individual galaxies from the NIHAO project (Wang et al. 2015; Buck et al. 2019), an ensemble of simulations using the FIRE (Hopkins et al. 2018; Garrison-Kimmel et al. 2019; Wheeler et al. 2019) and CHANGA (Munshi et al. 2021) codes, as well as simulations of constrained Local Group environments from the APOSTLE project (Fattahi et al. 2016, 2018; Sawala et al. 2016). Since we shall use the latter to calibrate our modelling procedure, we describe the APOSTLE simulations in some detail below. Results from the other runs are taken directly as reported in those publications, to which we refer the interested reader for details.

2.1 The APOSTLE simulations

The APOSTLE project is a set of 12 ‘zoom-in’ cosmological volumes tailored to reproduce the main properties of the Local Group. Each volume is selected from a large cosmological box to contain a pair of haloes with masses, relative radial and tangential velocities, and surrounding Hubble flow, consistent with the corresponding values observed for the MW–Andromeda pair (Fattahi et al. 2016).

The APOSTLE runs used the EAGLE galaxy formation code (Crain et al. 2015; Schaye et al. 2015), using the so-called reference parameters. This code includes subgrid physics recipes for radiative cooling, star formation in gas exceeding a metallicity-dependent density threshold, stellar feedback from stellar winds, radiation pressure and supernovae explosions, homogeneous X-ray/UV background radiation, supermassive black hole growth, and active galactic nucleus feedback (note that the latter has negligible effects on dwarf galaxies and is therefore unimportant in APOSTLE).

The EAGLE model was calibrated to approximate the observed $z = 0.1$ galaxy stellar mass function in the $M_* = 10^8-10^{12} M_{\odot}$ range. Simulated galaxies thus roughly follow the abundance-matching M_*-M_{200} relation of Behroozi et al. (2013) and Moster et al. (2013). No extra calibration is made in APOSTLE, and therefore the stellar–halo mass relation that results for fainter galaxies may be regarded as the extrapolation of the same subgrid physics to lower mass haloes.

The APOSTLE volumes have been run at three different levels of resolution. In this paper, we use the five highest resolution volumes (labelled ‘AP-L1’ in Fattahi et al. 2016). These runs have initial dark matter and gas particle masses of $m_{\text{DM}} \sim 5 \times 10^4 M_{\odot}$ and $m_{\text{gas}} \sim 1 \times 10^4 M_{\odot}$, respectively, and a gravitational softening length of 134 pc at $z = 0$. The APOSTLE volume simulated at highest resolution fully contains a sphere of radius $r \sim 3.5$ Mpc from the midpoint of the MW and M31 analogue haloes.

The friends-of-friends (FoF) groupfinding algorithm (Davis et al. 1985) (with linking length equal to 0.2 times the mean interparticle separation) and the SUBFIND halo finder (Springel et al. 2001; Dolag et al. 2009) were used to identify haloes and subhaloes. We shall refer to the galaxies formed in the most massive subhaloes of each FoF group as ‘centrals’, and to the rest of galaxies within the virial radius of each FoF central as ‘satellites’. Throughout the paper, we shall use the term ‘primary’ to refer to a central galaxy that may have satellites.

APOSTLE assumes a flat Λ CDM cosmological model following WMAP-7 parameters (Komatsu et al. 2011): $\Omega_m = 0.272$; $\Omega_\Lambda = 0.728$; $\Omega_{\text{bar}} = 0.0455$; $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$; $\sigma = 0.81$; $h = 0.704$.

3 MODELLING THE SATELLITE STELLAR MASS FUNCTION

The satellite mass function of a primary of given stellar mass, M_*^{pri} , depends mainly on (i) the mass function of subhaloes present in the halo of that system, on (ii) the relation between stellar mass and subhalo mass, and on (iii) the possible reduction of stellar mass due to tidal stripping after infall. The first item depends mainly on the primary halo virial mass, or, equivalently, on V_{200}^{pri} , and has been extensively studied through cosmological N -body simulations.

For the second item, which, in the case of satellites, applies before first infall into the primary halo, it is customary to express the stellar mass not as a function of (sub)halo mass, but rather in terms of its maximum circular velocity, V_{max} , a quantity more resilient to tidal effects than virial mass.

The third item is the most difficult to treat analytically, since it depends strongly on the pericentric distance of the orbit, the number of orbits completed, and the radial segregation of stars within each subhalo. Fortunately, as we shall see below, the fraction of stellar mass lost to tides is, on average, small, and we shall neglect it in our modelling in the interest of keeping the model as simple as possible.

We describe below the parametrizations we adopt to build an analytical model for the satellite mass function of a primary of mass M_*^{pri} . These parametrizations are motivated by the results of cosmological N -body and hydrodynamical simulations, as discussed in detail in the remainder of this section. We note that the two satellite mass function models explored here differ only in the assumptions made for the M_* – V_{max} relation.

3.1 Subhalo mass function

The substructure mass function of Λ CDM haloes scales in direct proportion to the virial mass of the primary halo and has been shown to be fairly well approximated by a power law. Following Wang et al. (2012), the average number of subhaloes within the virial radius of an isolated (central) Λ CDM halo may be expressed as

$$\langle N_{\text{sub}} \rangle (> \nu) = 10.2 (\nu/0.15)^{-3.11}, \quad (1)$$

where $\nu = V_{\text{max}}/V_{200}^{\text{pri}}$. This function applies to all Λ CDM haloes regardless of mass, and has been tested well over the $0.1 < \nu < 0.5$ range. The scatter around the average number at given ν is well approximated by Poisson statistics.

We compare in Fig. 1 the results of three sets of cosmological simulations with the predictions from equation (1) (thick black line). The simulations include the average of all MW-sized haloes of the Aquarius project (dotted black line; Springel et al. 2008), that of the cluster-sized haloes of the Phoenix project (dot–dashed black line; Gao et al. 2012), as well as that of all haloes with $V_{200} > 40 \text{ km s}^{-1}$ in the APOSTLE project (solid grey line). As is clear from this figure, equation (1) reproduces quite well the subhalo mass function of haloes spanning a wide range of virial mass.

However, this function is expressed in terms of the present-day subhalo maximum circular velocity, V_{max} , which may have been affected by tidal stripping after infall. Since the stellar content of a subhalo is more closely tied to V_{peak} , the maximum circular velocity prior to infall, equation (1) must therefore be corrected to yield the distribution of V_{peak} values needed in the modelling.

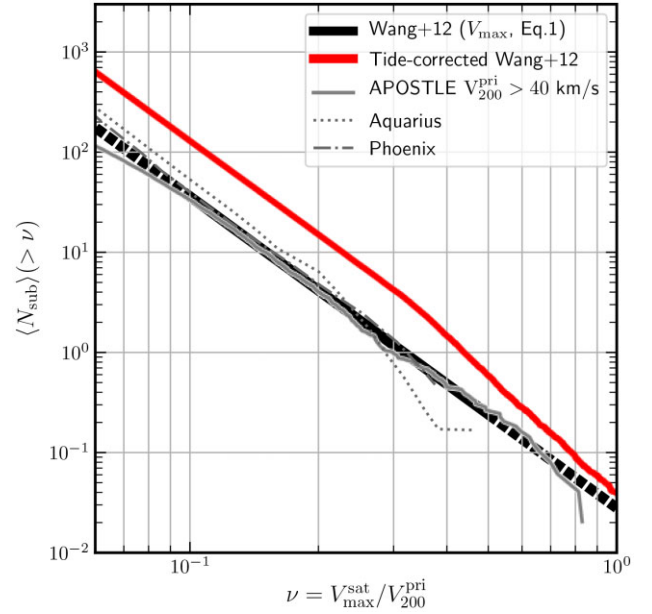


Figure 1. The subhalo velocity function (i.e. average number of subhaloes with $\nu = V_{\text{max}}/V_{200}$ above a certain value). The solid black line shows the function proposed by Wang et al. (2012, equation 1), which describes well the substructure mass function of Λ CDM haloes of all masses in the Millenium DM-only cosmological simulation, particularly in the $0.1 < \nu < 0.5$ range. Extrapolations of this line beyond such range are shown in dotted line style. The average subhalo velocity function for APOSTLE haloes with $V_{200}^{\text{pri}} > 40 \text{ km s}^{-1}$ is shown with a solid grey line. For comparison, thin grey dotted and dash–dotted lines show the average subhalo functions found for haloes in simulations of the Aquarius and Phoenix projects, respectively. A thick red line shows the ‘tide-corrected’ version of equation (1), resulting from statistically converting V_{max} values to V_{peak} as explained in Section 3.

To this end, we explore the relation between V_{max} and V_{peak} in APOSTLE haloes. This is shown in the left-hand panel of Fig. 2, scaled by the virial velocity of the primary, V_{200} , at $z = 0$. As expected, APOSTLE subhaloes had V_{peak} values systematically larger than V_{max} . The distribution of the ratio $V_{\text{peak}}/V_{\text{max}}$ is shown (in cumulative form) in the right-hand panel of Fig. 2, for various bins in V_{max}/V_{200} . This panel shows that, on average, the reduction in subhalo maximum circular velocity that results from tides is fairly modest and largely independent of subhalo mass.

Only the most massive subhaloes (i.e. $V_{\text{max}}/V_{200} > 0.3$) deviate from this trend, and appear substantially less affected by tides than less massive subhaloes. The median $V_{\text{max}}/V_{\text{peak}}$ is ~ 0.82 for low-mass subhaloes, but climbs to ~ 0.93 at the massive end. Why do tides seem to affect more massive haloes less? This is most likely a result of the rapid dynamical friction-driven evolution of massive haloes, which tend to merge with the primary halo quickly after accretion. In other words, the (few) very massive haloes present at any given time result from recent accretion events where tides have not had any substantive effect yet.

The results shown in Fig. 2 can be used to statistically correct the distribution of V_{max}/V_{200} measured in cosmological simulations and to estimate the V_{peak} subhalo distribution of a given halo. The result is illustrated by the thick red line in Fig. 1, which shows the tide-corrected form of equation (1). We shall hereafter adopt the tide-corrected version of equation (1) (with Poisson scatter) to model the V_{peak} distribution of a halo of given V_{200} .




Figure 2. V_{\max}/V_{200} versus V_{peak}/V_{200} (left) and normalized cumulative distribution of V_{\max}/V_{peak} (right) for APOSTLE subhaloes. Lines of different colours show results for subhaloes in different V_{\max}/V_{200} bins, as indicated in the legend.

3.2 The stellar mass–halo mass relation

What is the stellar mass expected for a subhalo of given V_{peak} ? Fig. 3 motivates our choice of models for the stellar mass–halo mass relation. This figure shows the M_* – V_{\max} relation reported for central galaxies at $z = 0$ selected from recent cosmological hydrodynamical simulations, as indicated in the legend. For central galaxies V_{peak} is in general achieved at $z = 0$ (except for those centrals that have tidally interacted in the past, which have been removed from our sample), and therefore the peak maximum circular velocity coincides with the maximum circular velocity at present-day, V_{\max} . When necessary, we have transformed quoted halo masses into V_{\max} assuming they follow a Navarro–Frenk–White density profile (hereafter NFW; Navarro, Frenk & White 1996, 1997) with a mass–concentration relation as given by Ludlow et al. (2016).

These simulations suggest two different behaviours for the M_* – V_{\max} relation. On the left-hand panel of Fig. 3, we have grouped simulations where M_* and V_{\max} seem better described by a simple power law that extends down from $V_{\max} \sim 200 \text{ km s}^{-1}$ to less than $\sim 10 \text{ km s}^{-1}$, deep into the ultrafaint regime ($M_* \sim 10\text{--}10^2 M_\odot$).

Interestingly, the power law follows closely the extrapolated M_* – V_{\max} relation from Moster et al. (2013) (dashed red line),

$$M_* = \frac{0.0702}{[(m_1)^{-\beta} + (m_1)^\gamma]} M_{200}, \quad (2)$$

where $m_1 = M_{200}/M_1$, $M_1 = 10^{11.59} M_\odot$, $\beta = 1.376$, and $\gamma = 0.608$. As above, M_{200} in this relation can be easily transformed into V_{\max} assuming an NFW density profile and a mass–concentration relation.

On the other hand, the panel on the right in Fig. 3 groups simulations whose results seem better described by a rapidly steepening relation between M_* and V_{\max} towards decreasing V_{\max} , suggesting the presence of a cut-off in the relation. Following Fattahi et al. (2018), this ‘cut-off’ relation may be parametrized as:

$$M_* = \eta^\alpha \exp(-\eta^\mu) M_0, \quad (3)$$

with $\eta = V_{\max}/50 \text{ km s}^{-1}$, and $(M_0, \alpha, \mu) = (3 \times 10^8 M_\odot, 3.36, -2.4)$, shown by the solid black line in Fig. 3.

Although those authors fitted only results for central galaxies, an indistinguishable fit is obtained when adding the M_* – V_{peak} data for APOSTLE satellites, which justifies the use of equation (4) to model the stellar content of a satellite of given V_{peak} . (For centrals $V_{\text{peak}} \approx V_{\max}$ at $z = 0$, by construction.)

The ‘power-law’ and ‘cut-off’ relations between stellar mass and peak velocity are the sole difference between the two models we explore in this paper. We emphasize that it is not our intention to categorize the different simulations into one or the other behaviour, but instead to motivate these two different analytical models that seem to describe well the current predictions from several simulations.

Moreover, both models explored here are meant to be purely empirical, without being strongly linked to particular choices of subgrid physics or numerical resolution. The fact that most of current predictions from state-of-the-art numerical simulations align well with one of the two models, independent of their assumed galaxy formation physics and resolution, is reassuring and provides support to the approach presented in this work.

These two M_* – V_{peak} relations are plotted in both panels of Fig. 3 for ease of comparison (red dashed curve for ‘power law’ and solid black for ‘cut-off’). The main differences between them are the behaviour at low V_{peak} and the slope of the relation at intermediate V_{peak} , between ~ 30 and 80 km s^{-1} . Because their predictions differ for systems like the MW, which we shall use to calibrate our models, we adopt a single M_* – V_{\max} relation for $V_{\max} > 84 \text{ km s}^{-1}$ (where the Fattahi et al. 2018 and the Moster et al. 2013 lines cross each other). This is shown by the power-law solid grey line depicted in Fig. 3 and may be expressed as

$$M_*/M_\odot = 3.29 \times 10^9 (V_{\max}/84 \text{ km s}^{-1})^{4.52}, \quad (4)$$

applicable only for $V_{\max} > 84 \text{ km s}^{-1}$.

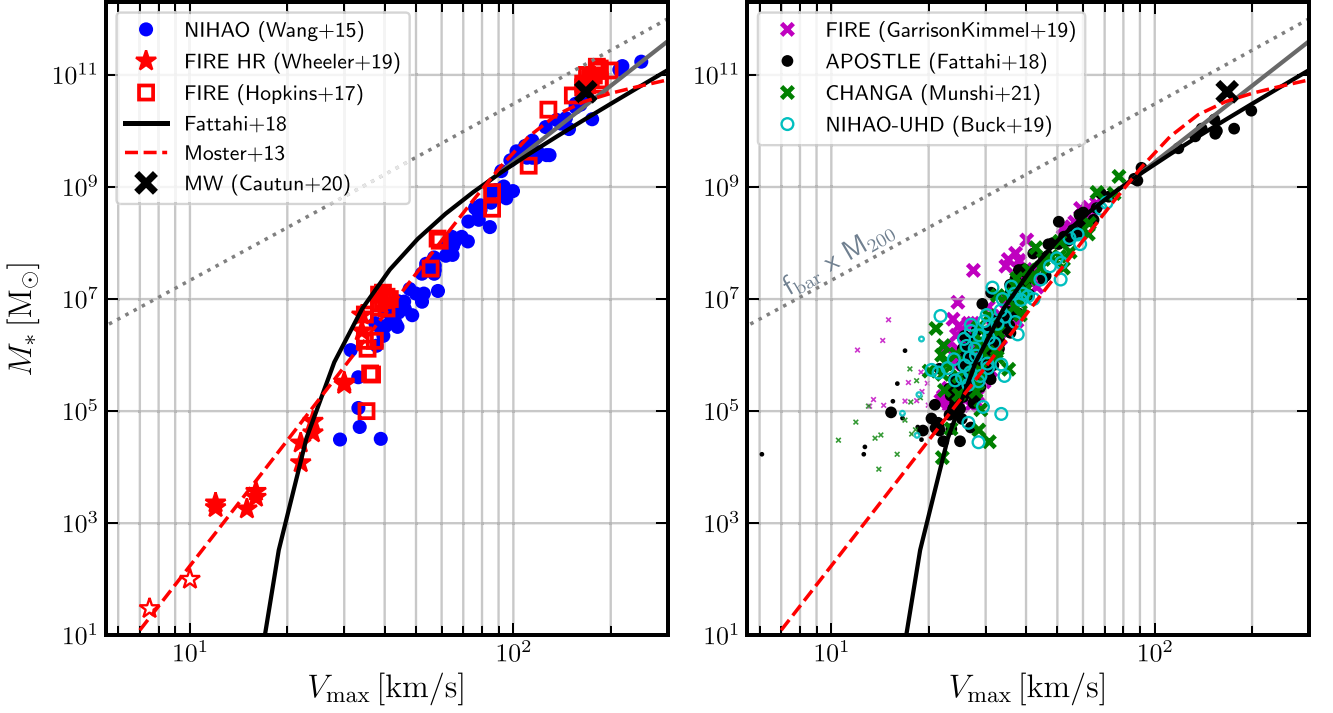


Figure 3. M_* – V_{\max} relations for simulated ‘central’ galaxies at $z = 0$ from recent cosmological hydrodynamical simulations (see legends). These have been divided into two groups, depending on whether they either roughly follow a ‘power-law’-like relation (left-hand panel), or a relation with a sharp ‘cut-off’ in M_* at low V_{\max} (right-hand panel) like the ones assumed in this work. The power-law results are well described by the extrapolated abundance-matching relation from Moster et al. (2013) (red dashed line); the cut-off results may be approximated by the fit to APOSTLE data from Fattahi et al. (2018) (black line). Open red star symbols indicate FIRE-HR galaxies with 1 and 15 stellar particles. Backsplash galaxies in APOSTLE are shown as smaller black points. In the right-hand panel, we mark all simulated galaxies from the other samples showing $V_{\max} < 20 \text{ km s}^{-1}$ with smaller symbols as they are likely to have been affected by tides as well. (Indeed, the NIHAO-UHD sample avoids most backsplash galaxies by selecting central dwarfs outside $2.5 \times r_{200}$ of a massive primary, and shows in general $V_{\max} > 20 \text{ km s}^{-1}$.) For reference, a dotted grey line marks the maximum total amount of baryons inside M_{200} as expected from the cosmic mix, where $f_{\text{bar}} = \Omega_{\text{bar}}/\Omega_{\text{m}}$.

In what follows, we shall express the stellar mass–halo mass relation in terms of V_{peak} , defined as the maximum circular velocity of a satellite before infall, or, for centrals, as V_{\max} at $z = 0$.

3.2.1 Scatter

As is clear from Fig. 3, the M_* – V_{peak} relation has substantial scatter. We account for this assuming that M_* follows, at given V_{peak} , a lognormal distribution with a dispersion, σ_{M_*} , that increases towards decreasing halo masses. Following Garrison-Kimmel et al. (2017) and Munshi et al. (2021), we parametrize σ_{M_*} as a broken power law of V_{peak} , as illustrated by the solid black line in the right-hand panel of Fig. 4,

$$\sigma_{M_*} = \begin{cases} \sigma_0 & V_{\text{peak}} > 57 \text{ km s}^{-1} \\ \kappa \log_{10}(V_{\text{peak}}/V_0) & V_{\text{peak}} < 57 \text{ km s}^{-1} \end{cases} \quad (5)$$

with $\sigma_0 = 0.24 \text{ dex}$, $\kappa = -1.26$, and $V_0 = 88.6 \text{ km s}^{-1}$. These parameters have been chosen arbitrarily but loosely guided by the measured scatter in APOSTLE galaxies (see open black circles in the right-hand panel of Fig. 4) and by the scatter parametrization to CHANGA galaxies in Munshi et al. (2021) (see cyan line). Note that Munshi et al. (2021) parametrize the scatter in terms of M_{peak} instead of V_{peak} , but indicate that in the latter case they find a scatter floor of 0.17 dex and an increasing scatter that reaches $\sim 1 \text{ dex}$ at their lowest V_{peak} s. As an approximation, here we assume it reaches the same

scatter at low V_{peak} as it does with their M_{peak} based model (dashed line).

For simplicity, we assume that both the ‘cut-off’ and the ‘power-law’ model have the same M_* scatter dependence on V_{peak} given by equation (5). The shaded bands in the left-hand and middle panels of Fig. 4 indicate the resulting 10–90 percentiles in the M_* distribution at a given V_{peak} assuming equation (5) in each model. The assumed cut-off relation with scatter reproduces the APOSTLE results well (see left-hand panel of Fig. 4). The middle panel of Fig. 4 shows that our choice, albeit arbitrary, also accommodates well other simulations too, such as FIRE, which we take as further validation of our assumed scatter model. Note that, while at fixed V_{peak} the scatter in M_* is identical in both models, the ‘cut-off’ model is steeper than the ‘power-law’ model at low V_{peak} , and therefore the shaded area is approximately constant in contrast with the visible increase in dispersion seen in the middle panel for the ‘power law’.

We have checked that changing the details of this scatter model (i.e. scatter floor, steepness of slope) makes negligible difference to results with the ‘cut-off’ model, because changes apply to the V_{peak} regime below the intrinsic threshold. Although changes do affect somewhat results with the ‘power-law’ model, the final relative differences between the satellite mass functions obtained with the ‘cut-off’ and ‘power-law’ models remain robust.




Figure 4. M_* – V_{peak} relations for the two models studied in this paper. Left: ‘cut-off’; Middle: ‘power law’. The ‘cut-off’ relation has been calibrated to match the APOSTLE simulation data for subhalo centrals at $z = 0$ and for subhalo satellites of MW/M31 analogue primaries before infall, when their circular velocities peak (black circles). We show our ‘power law’ model in comparison to FIRE simulation data (black symbols in middle panel; Hopkins et al. 2018; Wheeler et al. 2019). The shaded bands show the 10–90 percentile range in M_* at given V_{peak} . The scatter is assumed to be a function of V_{peak} , following equation (5), which is described in the rightmost panel. Throughout this paper, we assume the same scatter model for both the ‘cut-off’ and ‘power-law’ models. The black line in the rightmost panel compares our assumed scatter (solid black line) to the scatter in M_* measured in bins of V_{peak} for APOSTLE simulated galaxies (open circles) and to the model proposed by Munshi et al. (2021) after transforming their M_{peak} parametrization to V_{peak} (see text for more details).




Figure 5. Stellar mass of APOSTLE satellite galaxies measured at $z = 0$ versus that measured at the time when the maximum circular velocity of subhalo peaks, z_{peak} . The 1:1 correspondence is marked with a dashed grey line.

3.3 Stellar mass-loss

As mentioned above, our simple modelling shall neglect tidally induced stellar mass-loss in subhaloes. This is clearly a simplification, but finds support in the results for APOSTLE satellites, which show that the effects of stellar mass-loss are quite modest. This is shown in Fig. 5, where we plot the stellar mass of APOSTLE satellite galaxies at $z = 0$ versus that at z_{peak} , the redshift when its maximum circular velocity peaked. Unlike V_{max} , M_* changes, on average, very little

after infall into the main halo. Half of APOSTLE satellites have lost less than ~ 22 per cent of their peak mass, and only 10 per cent have lost more than ~ 58 per cent since infall. In the interest of simplicity, we have decided not to include any corrections for stellar mass-loss, but have checked that none of our main conclusions are altered if a correction of the magnitude suggested by Fig. 5, is implemented.

3.4 The cut-off and ‘power-law’ models

The assumptions discussed above allow us to compute the expected satellite stellar mass function for a system of arbitrary virial mass. To summarize, for a halo of given V_{200} we first draw a realization of the subhalo V_{peak} function consistent with the tide-corrected equation (1), assuming Poisson scatter. For each subhalo, we then draw a stellar mass using either the ‘cut-off’ or the ‘power-law’ models described in Section 3.2, with scatter as given by equation (5). Unless otherwise specified, we shall always show median results obtained by combining at least ~ 100 independent realizations of each primary, together with the 10–90 percentile range. We have confirmed that this number of realizations yields converged results by running our model with up to ~ 5 times more iterations with which we find no significant differences.

4 RESULTS

4.1 The cut-off model and APOSTLE

We start by comparing the results of the ‘cut-off’ model with satellite mass functions from the APOSTLE simulations. We do this to check that our ‘cut-off’ model is able to roughly reproduce the APOSTLE satellite stellar mass functions down to the resolution allowed by the simulations. Indeed, while we have chosen an average ‘cut-off’ M_* – V_{peak} relation based on APOSTLE, it is not obvious a priori that our simple model can yield satellite mass functions overall consistent with APOSTLE results.




Figure 6. Satellite stellar mass functions of APOSTLE primaries (grey) compared to those predicted by the ‘cut-off’ model (blue). Each panel corresponds to results for different bins in virial velocity of the primaries, V_{200} , as given in the legend. The number of APOSTLE primaries in each bin, N_{pri} , is also indicated in the legend. The blue lines show results from the ‘cut-off’ model as described in Section 3, applied to primaries with the same V_{200} as APOSTLE ones in each bin. Solid lines show the median results and shaded areas the 10–90 percentile range. For APOSTLE, results below $M_* < 10^5 M_\odot$ are shown with a dashed linestyle to indicate that objects below this mass are likely not well resolved.

For example, our analytical ‘cut-off’ model includes a fully independent sampling of the subhalo mass function directly taken from Λ CDM simulations and corrected statistically by tidal stripping, and does not use the subhalo mass functions from APOSTLE. A good agreement between our analytical model and the APOSTLE results is a necessary benchmark for our analytical models.

This comparison is shown in Fig. 6, where each panel shows, in grey, the APOSTLE satellite mass functions for central galaxies, binned by halo virial velocity. The average V_{200} and standard deviation in each bin is given in the legend of each panel. Solid lines show the median satellite mass function in the bin, while the shaded area represents the 10–90 percentile distribution.

Although we show mass functions down to stellar masses as low as $M_* > 10^2 M_\odot$ we note that objects with $M_* < 10^5 M_\odot$ in APOSTLE are resolved with fewer than 10 star particles. Therefore, below that mass APOSTLE results are best regarded as lower limits rather than actual simulation predictions.

By construction, the first bin (leftmost panel) includes the 10 primaries that are considered MW and M31 analogues in the APOSTLE volumes. For these APOSTLE primaries ($V_{200} \approx 150 \text{ km s}^{-1}$, or, equivalently, $M_{200} \sim 1.2 \times 10^{12} M_\odot$), the median number of satellites with $M_* > 10^5 M_\odot$ is $\sim 24.1^{+18.5}_{-6.7}$, where the uncertainties represent the 10–90 percentile range.

For dwarf primaries with $V_{200} \approx 62 \text{ km s}^{-1}$ (third panel from the left) the number of satellites in APOSTLE is drastically reduced by the cut-off in the M_* – V_{peak} relation, with a median of only $\sim 2.0^{+1.7}_{-1.0}$ satellites with $M_* > 10^5 M_\odot$. Finally, the last panel shows that no luminous satellites are found in APOSTLE around primaries with $V_{200} \lesssim 35 \text{ km s}^{-1}$.

The blue bands in Fig. 6 show the results of the ‘cut-off’ model, applied to a sample of primaries whose number and V_{200} distribution matches that in each APOSTLE bin. We use 10 independent realizations of the satellite mass function of each primary to obtain robust results.

There is in general good agreement between the analytical ‘cut-off’ model and the simulation results, especially for satellites with $M_* > 10^5 M_\odot$. Even the number of massive ($M_* > 10^8 M_\odot$) satellites is well reproduced, with a median of ~ 1 – 2 LMC or Small Magellanic

Cloud (SMC)-mass satellites expected around MW-mass primaries (leftmost panel).

This is not unexpected, given that we have motivated the model on APOSTLE results, but it provides validation for our approach. It also allows us to predict the population of dwarfs fainter than currently resolved by APOSTLE and other simulations. Importantly, the ‘cut-off’ model predicts a steady decline in the number of satellites surrounding dwarfs of decreasing mass, approaching zero as the mass of the primary approaches the threshold mass (rightmost bin in Fig. 6).

4.2 Cut-off versus power-law model satellite mass functions

We now compare the satellite stellar mass functions predicted by each model, as a function of the stellar mass of the primary. This is shown in Fig. 7, where each panel corresponds to a different M_*^{pri} , given in the panel legends (cut-off in blue, power-law in red). The most obvious difference is the large difference in the number of faint satellites predicted by each model. Hundreds of ultrafaints with $M_* > 10^2 M_\odot$ are expected in the ‘power-law’ model, even for primaries as faint as the Magellanic Clouds, whereas ultrafaint numbers are much less numerous in the case of the ‘cut-off’ model.

The difference between models is more clearly appreciated when comparing the normalized satellite mass functions; i.e. the satellite mass function expressed in terms of $M_*^{\text{sat}}/M_*^{\text{pri}}$. This is shown in Fig. 8 for all primary stellar mass bins in the ‘cut-off’ model (left) and ‘power-law’ model (right). For the ‘power-law’ model the normalized satellite mass function changes little with primary stellar mass. In particular, primaries with $M_*^{\text{pri}} < 10^9 M_\odot$ would be expected to share the same normalized satellite mass function, as shown by the overlap of the red dotted and long-dash-dotted lines in the right-hand panel of Fig. 8.

As discussed by Sales et al. (2013), this near ‘self-similarity’ arises because the subhalo mass function and the stellar mass–halo mass relation in this model are both close to power laws, and thus scale free. This is particularly true at $M_* < 10^9 M_\odot$ (see middle panel Fig. 4), which explains why lower mass bins overlap in their normalized satellite mass function. On the other hand, if the stellar–halo mass relation is not scale free, as is the case for the ‘cut-




Figure 7. Comparison of the satellite stellar mass functions predicted by the ‘cut-off’ (blue) and ‘power-law’ (red) models. Each panel shows results assuming a different fixed value of the stellar mass of the primary, M_*^{pri} (see legend). Solid lines show median results and shaded bands the 10–90 percentile range.




Figure 8. Scaled satellite stellar mass functions (i.e. $M_*^{\text{sat}}/M_*^{\text{pri}}$) predicted by the ‘cut-off’ (blue) and ‘power-law’ (red) models. Different linestyles show results assuming a different fixed value of the stellar mass of the primary, M_*^{pri} (see legend). Lines show median results and shaded bands the 10–90 percentile range. Note that for $M_*^{\text{pri}} \lesssim 5 \times 10^9 M_\odot$ the ‘power-law’ mass function becomes independent of primary mass, as discussed by Sales et al. (2013).

off’ model, the normalized satellite mass function declines with decreasing primary mass (see blue curves on the left-hand panel of Fig. 8). The large differences between models suggest that the satellite mass function around isolated primaries spanning a wide range of mass (and, in particular, including dwarfs) may be used to infer the shape of the stellar mass–halo mass relation at the faint end.

Another, perhaps more intuitive contrast between models may be obtained by comparing the expected total number of satellites more massive than a given stellar mass as a function of primary halo virial mass. We show this in Fig. 9, where different linestyles indicate the cumulative number of satellites above a given M_*^{sat} , as labelled on the left-hand panel, as a function of either the virial velocity of the host (a proxy for the primary halo mass; lower x -axis) or the corresponding stellar mass of the primary according to each of the two models (upper x -axis).

For massive satellites (i.e. $M_*^{\text{sat}} > 10^6 M_\odot$, solid line), the predictions of the two models are rather similar, with ~ 2.5 satellites on average in hosts with $V_{200} = 75 \text{ km s}^{-1}$ and 10–30 satellites for hosts in our most massive bin, $V_{200} \geq 150 \text{ km s}^{-1}$. However, the predictions of the two models differ appreciably when considering fainter satellites and, in particular, in the regime of ultrafaint dwarfs.

For example, in the ‘power-law’ model, a dwarf primary with $M_*^{\text{pri}} \sim 10^9 M_\odot$ (like the LMC) is expected to host ~ 70 satellites with $M_*^{\text{sat}} \geq 10^3 M_\odot$, ~ 50 of which would be ultrafaint ($M_*^{\text{sat}} < 10^5 M_\odot$). On the other hand, in the ‘cut-off’ model only eight satellites are expected with $M_*^{\text{sat}} \geq 10^3 M_\odot$ for the same primary. As we have seen before, the population of ultrafaint satellite dwarfs is heavily suppressed in models with a sharp cut-off in the stellar mass–halo mass relation like the one explored here. Deep imaging and spectroscopic surveys of the surroundings of isolated dwarfs designed to constrain the satellite population within their virial radius

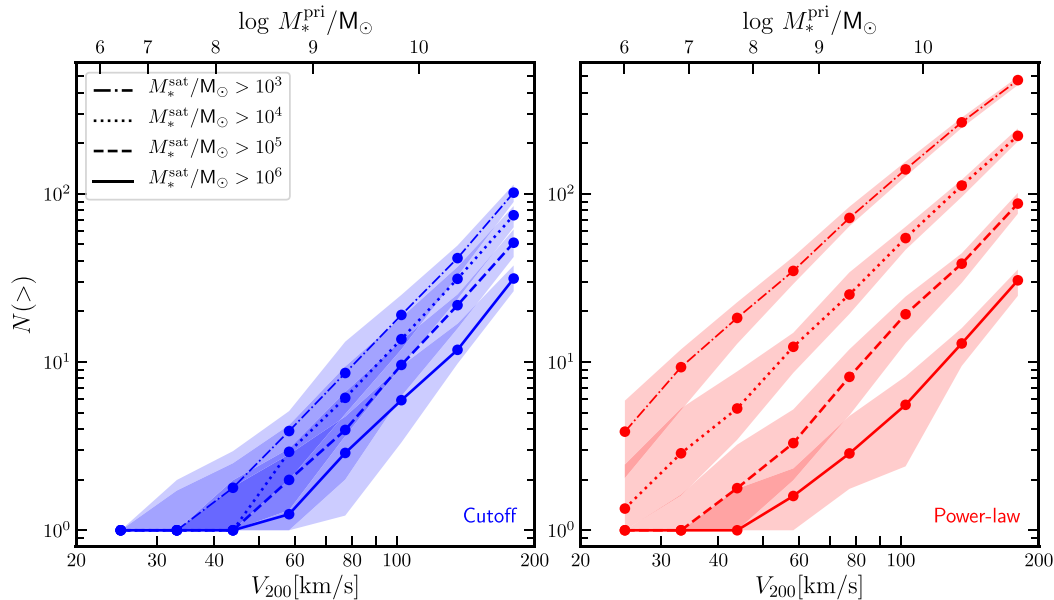


Figure 9. The cumulative number of satellites with M_*^{sat} above certain values (see legend) predicted by the ‘cut-off’ and ‘power-law’ models for primaries with fixed values of V_{200} , as indicated in the x -axis. The corresponding M_*^{pri} values, according to each model, can be read from the top x -axis. Lines show median results and shaded bands show the 10–90 percentile range.

should thus yield key insights into the stellar mass–halo mass relation at the faint end.

5 MODELS VERSUS OBSERVATIONS

5.1 Milky Way and M31 satellites

The most complete available census of faint satellites is in the Local Group, which provides therefore a good testbed for the ideas explored above. We compare in Fig. 10 the predictions of the theoretical models with data for MW and M31 satellites (left-hand and middle panels).² Black symbols connected by a solid curve show the observational data, taken from McConnachie (2012)’s updated compilation of Local Group dwarfs where objects within 300 kpc of the MW/M31 are considered satellites.

For the models, we choose a virial velocity of $V_{200} = 150 \text{ km s}^{-1}$ for the MW, and a somewhat larger $V_{200} = 165 \text{ km s}^{-1}$ for M31, in agreement with current available mass constraints (see e.g. Fardal et al. 2013; Sofue 2015; Cautun et al. 2020). Each virial velocity is sampled 100 times; the resulting median and 10th–90th percentiles are shown in blue for the ‘cut-off’ model and in red for the ‘power-law’ model.

The number of MW satellites with $M_*^{\text{sat}} \geq 10^6 M_\odot$ is in reasonable agreement with both models (see left-hand panel in Fig. 10), as well as with data from the SAGA survey, which targeted the bright end of the satellite population within 300 kpc of MW-like primaries (Mao et al. 2021). We note that our models refer to satellites within the virial radius of the assumed halo ($r_{200} \sim 215 \text{ kpc}$ for our choice of $V_{200} = 150 \text{ km s}^{-1}$) rather than the 300 kpc used in the observational data. The thin black line in the left-hand panel of Fig. 10 shows the

known MW satellites inside that smaller radius; the difference is quite small.

Both the ‘power-law’ and the ‘cut-off’ model predict the same number of satellites with $M_* > 10^6 M_\odot$ (roughly ~ 18), interestingly well in excess of the known number of such systems orbiting the MW. The discrepancy worsens between 10^4 and $10^6 M_\odot$, where the MW satellite mass function appears to have a sizable ‘gap’. It is unclear what the significance of such gap may be, but it is tempting to associate it with increasing incompleteness in observational detections (see e.g. the discussion in Fattahi, Navarro & Frenk 2020). The numbers climb rapidly in the 10^2 – $10^4 M_\odot$ range, to almost match the predictions of the ‘cut-off’ model.

As discussed in Section 4.2, it is in the ultrafaint regime where the ‘power-law’ and ‘cut-off’ models can be best differentiated. For faint dwarfs with $M_*^{\text{sat}} < 10^5 M_\odot$, the ‘cut-off’ model predicts substantially fewer ultrafaints than the ‘power-law’ model. Interestingly, this comparison suggests that if the stellar mass–halo mass relation does indeed have a low-mass cut-off, the majority of ultrafaint dwarfs in the MW might have already been discovered, leaving little room to accommodate a large missing population of ultrafaints. On the other hand, the ‘power-law’ model suggests the presence of a numerous, yet undetected population of ultrafaints in the MW halo. Upcoming surveys of the MW satellite population, especially those which account for satellites hidden behind the disc, or missing due to the incomplete spatial and surface brightness coverage of existing surveys, should be able to distinguish clearly between the two models proposed here.

The middle panel of Fig. 10 compares model predictions with current estimates of the M31 satellite population. Although the surveyed population in M31 does not go as deep as in the MW, the total number of satellites with $M_* \geq 10^5 M_\odot$ seems to fall below the ‘power-law’ model predictions, at least for the virial mass explored here. There is a hint that the observed satellite mass function compares more favourably with the ‘cut-off’ model, which predicts

²Stellar masses for observed Local Group satellites have been estimated using luminosities from McConnachie (2012) and assuming appropriate mass-to-light ratios according to Woo, Courteau & Dekel (2008).




Figure 10. Predictions of the ‘cut-off’ (blue) and ‘power-law’ (red) models for the satellite stellar mass functions of observed host galaxies. Left: MW; Middle: M31; Right: LMC (note the different y-axis limits). For each case, we assume a fixed value of V_{200} of the primary motivated by literature estimates (see text for details). Solid lines show median results and shaded bands the 10–90 percentile range. Black points show the observed satellite stellar mass function of each host. In the case of MW and M31, thicker lines show results for satellites within 300 kpc, while thinner lines correspond to satellites within r_{200} (214 kpc for $V_{200} = 150 \text{ km s}^{-1}$, and 236 kpc for $V_{200} = 165 \text{ km s}^{-1}$, assuming $h = 0.7$). The first panel additionally shows the satellite stellar mass function of MW-mass analogues observed as part of the SAGA survey (Mao et al. 2021). For the case of the LMC, we show model results assuming $V_{200} = 50$ (coloured, solid) as well as 100 (coloured, long-dashed) km s^{-1} , compared to likely Magellanic satellites according to Santos-Santos et al. (2021).

roughly half as many satellites in that mass range as the ‘power-law’ model.

The ‘cut-off’ model predicts at least ~ 50 new ultrafaint M31 satellites in the range $10^2 < M_*/M_\odot < 10^{4.6}$ (the mass of And XX, the least massive M31 satellite known), bringing the total population to ~ 90 – 110 total dwarfs above a stellar mass $100 M_\odot$. By contrast, the ‘power-law’ model predicts a total of ~ 680 – 740 satellites with $M_* > 10^2 M_\odot$. We note that these numbers are quite sensitive to the choice of virial mass for the M31 halo; doubling the mass (i.e. increasing V_{200} to 208 km s^{-1}) would yield roughly twice as many satellites for either model, although the relative differences in mass function shape would be preserved.

5.2 Satellites of isolated LMC-like dwarfs

Finally, the right-hand panel of Fig. 10 shows the predictions for the satellite population of isolated dwarf galaxies with stellar mass comparable to that of the LMC ($M_* \sim 3 \times 10^9 M_\odot$), or, more precisely, dwarf primaries inhabiting haloes with V_{200} in the range 50 to 100 km s^{-1} . This is consistent with the virial mass range ($2.5 < M_{200}/10^{10} M_\odot < 45$) of galaxies with comparable stellar mass in the APOSTLE simulations (see e.g. Santos-Santos et al. 2021). These authors use kinematic information to identify LMC-associated dwarfs; their list of most likely LMC satellites include seven satellites: the SMC, Hydrus 1, Horologium 1, Carina 3, Tucana 4, Reticulum 2, and Phoenix 2 (dashed black line). A less likely, but still plausible association is also ascribed to Carina, Horologium 2, Grus 2, and Fornax, bringing the total to 11 (solid black line).

In the context of the ‘cut-off’ model, these numbers seem to rule out a virial velocity as low as 50 km s^{-1} (bottom blue curve), and suggest a virial velocity a little below 100 km s^{-1} (top blue curve). In contrast, a virial velocity near the lower bound would be favoured in the case of the ‘power-law’ model. An LMC halo as massive as 100 km s^{-1} seems quite inconsistent with the data in this case. Note

that the predictions of the two different models differ substantially even for satellites with $M_* \sim 10^4 M_\odot$. This limit seems within reach of what may be achievable in future surveys of LMC-like primaries, turning them into strong constraints of the stellar mass–halo mass relation of faint galaxies, a subject we address in more detail below.

5.3 Predictions for future surveys

Beyond the Local Group, several ongoing (and future) observational efforts have the potential to measure the satellite population of isolated LMC-like galaxies, and thus deliver strong constraints on the stellar mass–halo mass relation at the faint end. To reduce fluctuations due to object-to-object scatter, it is desirable to survey several primaries of similar stellar mass while simultaneously reaching the ultrafaint satellite regime. This is why dwarf galaxies are the most promising primaries: within the Local Volume (i.e. within 10 Mpc from the MW) there are only eight MW-like galaxies ($M_* \geq 10^{10.5} M_\odot$) outside the Local Group but there are 112 known dwarfs with $10^8 < M_*/M_\odot < 10^{9.5}$ (Tully et al. 2009; Tully, Courtois & Sorce 2016).

As an example, we provide in Fig. 11 expectations from the ‘cut-off’ and ‘power-law’ models for the satellite mass function of four LMC-like dwarfs expected to be surveyed as part of the DES DELVE campaign (Drlica-Wagner et al. 2021). This includes NGC 300, NGC 55, IC 5152, and Sextans B, which span a stellar mass range $M_* = [0.5, 3] \times 10^9 M_\odot$. The model predictions are based on 100 realizations of dwarfs with fixed virial velocity, $V_{200} = 75 \text{ km s}^{-1}$, and are shown by the top red curves and the bottom blue curves. (Thin lines correspond to four individual realizations, to illustrate the expected object-to-object scatter.)

As in our earlier discussion, this figure makes clear that reaching satellites with $M_* \sim 10^4 M_\odot$ should be enough to differentiate between models, since the ‘power-law’ model predicts almost three times more such satellites than the ‘cut-off’ model. The differ-




Figure 11. Satellite stellar mass functions predicted by the ‘cut-off’ (blue) and ‘power-law’ (red) models, assuming a fixed primary halo virial velocity of $V_{200} = 75 \text{ km s}^{-1}$. A solid line shows the median result, while the shaded band shows the 10–90 percentile range. In addition, the dashed lines show the results obtained assuming *no scatter* in the M_* – V_{peak} relation. This has a strong effect on the power-law results (in orange), but affects very little the cut-off ones. These results summarize the predictions of our two models for the satellite mass functions of the four DELVE Magellanic analogues NGC 300, NGC 55, IC5152, and SextansB, with stellar masses in the $[0.5, 3] \times 10^9 M_{\odot}$ range (Drlica-Wagner et al. 2021). Thin individual random realizations of the model (with scatter), to illustrate the expected object-to-object variation.

ence is most striking when reaching ultrafaints with $M_* \sim 10^2 M_{\odot}$, where only 10 satellite dwarfs are expected around LMC analogues in the case of a cut-off whereas more than ~ 160 are predicted for the ‘power-law’ model. Should future surveys fail to discover a large number of ultrafaint dwarfs around isolated LMC analogues, this would be strong evidence in favour of some kind of cut-off in the stellar mass–halo mass relation.

5.4 Comparison with previous work on satellites of LMC-like hosts

It is interesting to compare our results with previous work in the literature on the satellite population of LMC-like hosts. For instance, assuming a power-law relation between stellar and halo mass, Nadler et al. (2020) predict 48 ± 8 satellites³ with $M_* \geq 10^2 M_{\odot}$, about a factor of 3 lower than the ~ 160 dwarf satellites predicted by our ‘power-law’ model. This is not due to differences in our assumptions about the primary virial mass nor about the subhalo abundance: we have explicitly checked that the number of *subhaloes* in our LMC-like primaries is consistent with Nadler et al. (2020). Indeed, we find 48–58 subhaloes (10th–90th percentiles) with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ within the virial radius of primaries with $V_{200} =$

75 km s^{-1} , in good agreement with the 52 ± 8 quoted by those authors. The difference must therefore be due to the way each model populates those subhaloes with galaxies.

The slope in the low-mass end of the M_* – V_{peak} relation inferred by Nadler et al. (2020) is somewhat shallower than the one adopted here (~ 5.1 compared to ~ 7.4 in our ‘power-law’ model), but we have identified two main factors contributing to the smaller number of dwarf satellites predicted by Nadler et al. (2020) compared to our work. One is that we model the scatter in the M_* – V_{peak} relation as velocity-dependent, increasing from ~ 0.22 dex for MW-like objects to ~ 1 dex in haloes with $V_{\text{peak}} \sim 10 \text{ km s}^{-1}$. On the other hand, fig. 6 (in combination with their table 1) in Nadler et al. (2020) suggests that their model infers a roughly constant upper limit of ~ 0.2 dex scatter in the dwarf regime in order to reproduce the completeness-corrected number of observed MW satellites.

The effect of the larger assumed scatter in our model is appreciable. Indeed, assuming zero scatter in the stellar mass–velocity relation, the ‘power-law’ model would decrease the predicted numbers from ~ 160 to ~ 67 satellites with $M_* \geq 10^2 M_{\odot}$ (see middle dashed orange curve in Fig. 11), in better agreement with the 48 ± 8 predicted in Nadler et al. (2020). This is also in agreement with the ~ 70 satellites with $-7 < M_V < -1$ predicted by Jethwa, Erkal & Belokurov (2016) via dynamical modelling of the Magellanic Cloud satellite population. In summary, these results show qualitative consistency with Munshi et al. (2021) who find that a scatter that grows with halo mass or V_{peak} steepens the slope of the faint end of the resultant satellite mass function. We refer the reader to Garrison-Kimmel et al. (2017) for a detailed discussion of the degeneracies in the slope/scatter of abundance-matching models and the expected number of dwarfs and to Munshi et al. (2019) for an example of how different sub-grid physics and resolution might impact the slope/scatter of the stellar–halo mass relation.

A second factor affecting the number of ultrafaints in Nadler et al. (2020) is that their model infers an occupation fraction such that below $V_{\text{peak}} \sim 9 \text{ km s}^{-1}$ an increasing fraction of haloes with decreasing V_{peak} remain dark and never host a galaxy (modelled according to their equation 3), while our ‘power-law’ model assumes an occupation fraction equal to 1 at all V_{peak} . We note that adding an occupation fraction to a power-law M_* – V_{peak} relation effectively makes it steeper and more comparable to the ‘cut-off’ model, lowering the total number of predicted faint satellites.

Our predictions may also be compared with the work of Dooley et al. (2017), who explored the satellite population of LMC-like hosts using (power-law) extrapolations of several abundance-matching models, including that of Moster et al. (2013). The main difference with our own ‘power-law’ model is that they also include an occupation fraction to model the effects of reionization. As such, their predictions are more similar to our ‘cut-off’ model, with ~ 10 – 15 (median, depending on which particular abundance-matching relation) dwarf satellites with $M_* \geq 10^2 M_{\odot}$ within a 50 kpc radius of their hosts. These results are bracketed by the predictions of our ‘cut-off’ model, with 6–13 (10th–90th percentiles), and our ‘power-law’ model, with 145–178 satellites, although our numbers are within a larger volume of $r_{200} \sim 107 \text{ kpc}$ (corresponding to $V_{200} = 75 \text{ km s}^{-1}$).

Our predictions in Fig. 11 might also inform other satellite searches around LMC analogues in the field such as the LBT-SONG survey or MADCASH (Carlin et al. 2021). At least two faint satellites have been identified around the Magellanic dwarf NGC 628 (Davis et al. 2021) surveying only a fraction of its inferred virial extension with the Large Binocular Telescope as well as the confirmation of DDO113 as (interacting) satellite of NGC 4214. Additionally, two

³Nadler et al. (2020) quote numbers above an absolute V magnitude $M_V = 0$, corresponding to an $M_* \sim 90 M_{\odot}$ assuming a mass-to-light ratio $M/L = 1$.

dwarfs have been confirmed as satellites of the Magellanic analogues NGC 2403 and NGC 4214 with *HST* data for the Hyper Suprime Cam survey MADCASH. Müller & Jerjen (2020) report, in addition, two candidate faint dwarfs possibly associated with NGC 24 in the Sculptor group using the Dark Energy Camera. As data continue to accumulate around dwarf primaries, the census of their satellite population is starting to emerge as the most promising and effective way to constrain the galaxy–halo connection at the low-mass end.

6 SUMMARY AND CONCLUSIONS

We have studied the effects of different stellar mass–halo mass relations on the predicted population of faint and ultrafaint dwarf satellites around primaries spanning a wide range of stellar mass. The models are motivated by results of recent state-of-the-art cosmological hydrodynamical simulations, extrapolated to the ultrafaint regime, down to $M_* \sim 10^2 M_\odot$.

Two faint-end stellar mass–halo mass model relations are explored: one is a ‘power law’ motivated by recent semi-analytical results about the abundance of satellites in the Local Group, and by recent high-resolution simulations from the FIRE project, which follow closely a power-law extrapolation to the faint regime of the abundance-matching results from Moster et al. (2013).

A second is a ‘cut-off’ model where the stellar mass–halo mass relation gradually steepens towards decreasing mass so that no luminous dwarf exists beyond a minimum threshold halo virial mass of the order of $M_{200} \sim 10^9 M_\odot$. The ‘cut-off’ model is motivated by results from the APOSTLE simulations, and by analytical considerations that disfavour the formation of galaxies in haloes below the ‘hydrogen-cooling’ limit (see e.g. Benitez-Llambay & Frenk 2020, and references therein). We assume the same subhalo mass function and the same mass-dependent scatter in the M_*-V_{peak} relation in both models.

Our main finding is that satellite mass functions of primary galaxies spanning a wide range of mass are an excellent probe of the shape of the stellar mass–halo mass relation at the faint end. Satellite mass functions are particularly constraining if they reach deep into the ultrafaint regime. For example, the ‘cut-off’ model predicts ~ 9 (19) times fewer dwarfs with $M_* \geq 10^2 M_\odot$ than the ‘power-law’ model for primaries with $M_* \sim 3 \times 10^{10}$ (3×10^7) M_\odot . The difference becomes more marked as the stellar mass of the primary decreases, implying that the satellites of dwarf primaries, in particular, provide particularly strong constraints on the stellar mass–halo mass relation at the faintest end.

The models also predict different *normalized* satellite mass functions, i.e. the number of satellites expressed as a function of $M_*^{\text{sat}}/M_*^{\text{pri}}$ rather than M_*^{sat} . While the normalized function declines with decreasing primary stellar mass in the ‘cut-off’ model, it is nearly independent of primary mass in the ‘power-law’ model. This self-similar behaviour results because the subhalo mass function is also a power law in LCDM, as discussed by Sales et al. (2013).

Our findings have important implications when applied to nearby galaxies, where the surveying of ultrafaint dwarfs is or will become feasible in the near future. For an MW-mass primary (i.e. $V_{200} \sim 150 \text{ km s}^{-1}$), the ‘power-law’ and ‘cut-off’ models predict $\sim 612_{-25}^{+35}$ versus $\sim 77_{-11}^{+10}$ satellites above the nominal $M_* = 10^2 M_\odot$ mass cut, respectively. Interestingly, in the MW itself the number of already discovered satellite dwarfs is quite close to the ‘cut-off’ model prediction, leaving only little room for the detection of large numbers of new ultrafaint dwarfs. This is a prediction that should also be testable by ongoing and future surveys of the satellite population around nearby galaxies.

LCDM predicts that dwarf galaxies should also host a number of fainter companions. The models described above may be used to compute the number of dwarf satellites expected around LMC-mass systems in the field. We find, on average, that the ‘cut-off’ model predicts ~ 3 – 22 satellites with $M_* > 100 M_\odot$; the number, on the other hand, grows to ~ 65 – 300 for the power-law case, where the range corresponds to assuming a virial velocity range between 50 and 100 km s^{-1} for the LMC halo. This highlights the potential for ultrafaint discoveries in regions surrounding Magellanic-like dwarfs in the field, a particularly exciting prospect in light of ongoing efforts such as DELVE (Drlica-Wagner et al. 2021), MADCASH (Carlin et al. 2021), or LBT-SONG (Davis et al. 2021), which target the surroundings of isolated LMC-like dwarfs.

We conclude that the satellite population of dwarf galaxies in the field offers a powerful way to constrain the faint end of the stellar mass–halo mass relation (see also Wheeler et al. 2015). Only if the relation extends well below the HCL (as envisioned in the ‘power-law’ model), then one would expect dwarfs to have numerous ultrafaint companions. In the cut-off case, as the mass of the primary approaches the cut-off, the number of satellites should decline rapidly. For the particular cut-off we explore in this paper, dwarfs with $M_* \sim 3 \times 10^8 M_\odot$ or less should have virtually no luminous satellites, regardless of luminosity.

Dwarf primaries are also good probes because the galaxy mass is, in relative terms, much smaller than that of their surrounding halo. The galaxy’s effect on the subhalo population is therefore proportionally reduced compared to galaxies like the MW, where the disc is massive enough to affect noticeably the evolution and survival of subhaloes in its vicinity (Jahn et al. 2019). Finally, dwarf primaries are more abundant in the Local Volume than giant spirals like the MW or M31, so surveying a statistically meaningful sample becomes, in principle, more feasible.

We conclude that the satellite mass function of dwarf galaxies in the field represents an efficient and attractive approach for constraining the mapping between stars and dark matter haloes in the low-mass end with deliverables expected in the foreseeable future.

ACKNOWLEDGEMENTS

ISS was supported by the Arthur B. McDonald Canadian Astroparticle Physics Research Institute. LVS is grateful for financial support from the NSF-CAREER-1945310, AST-2107993, and NASA-ATP-80NSSC20K0566 grants. JFN is a Fellow of the Canadian Institute for Advanced Research. AF was supported by a UKRI Future Leaders Fellowship (grant no. MR/T042362/1). This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University, and STFC operations’ grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The APOSTLE simulation data and model results underlying this article can be shared on reasonable request to the corresponding author. For data from other simulations, we refer the interested reader to the corresponding references cited in each case. The observational data for Local Group satellites used in this article come from McConnachie (2012; see http://www.astro.uvic.ca/~alan/Nearby_LDwarf_Database_files/NearbyGalaxies.dat, and references therein).

REFERENCES

- Baldry I. K. et al., 2012, *MNRAS*, 421, 621
- Behroozi P. S., Wechsler R. H., Conroy C., 2013, *ApJ*, 770, 57
- Benítez-Llambay A., Frenk C., 2020, *MNRAS*, 498, 4887
- Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, *MNRAS*, 398, 1150
- Brook C. B., Stinson G., Gibson B. K., Shen S., Macciò A. V., Obreja A., Wadsley J., Quinn T., 2014, *MNRAS*, 443, 3809
- Buck T., Macciò A. V., Dutton A. A., Obreja A., Frings J., 2019, *MNRAS*, 483, 1314
- Bullock J. S., Boylan-Kolchin M., 2017, *ARA&A*, 55, 343
- Carlin J. L. et al., 2021, *ApJ*, 909, 211
- Cautun M. et al., 2020, *MNRAS*, 494, 4291
- Conroy C., Wechsler R. H., Kravtsov A. V., 2006, *ApJ*, 647, 201
- Crain R. A. et al., 2015, *MNRAS*, 450, 1937
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- Davis A. B. et al., 2021, *MNRAS*, 500, 3854
- Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
- Dooley G. A., Peter A. H. G., Carlin J. L., Frebel A., Bechtol K., Willman B., 2017, *MNRAS*, 472, 1060
- Drlica-Wagner A. et al., 2021, *ApJS*, 256, 2
- Errani R., Navarro J. F., 2021, *MNRAS*, 505, 18
- Fardal M. A. et al., 2013, *MNRAS*, 434, 2779
- Fattahi A. et al., 2016, *MNRAS*, 457, 844
- Fattahi A., Navarro J. F., Frenk C. S., Oman K. A., Sawala T., Schaller M., 2018, *MNRAS*, 476, 3816
- Fattahi A., Navarro J. F., Frenk C. S., 2020, *MNRAS*, 493, 2596
- Gao L., Navarro J. F., Frenk C. S., Jenkins A., Springel V., White S. D. M., 2012, *MNRAS*, 425, 2169
- Garrison-Kimmel S., Bullock J. S., Boylan-Kolchin M., Bardwell E., 2017, *MNRAS*, 464, 3108
- Garrison-Kimmel S. et al., 2019, *MNRAS*, 487, 1380
- Gnedin N. Y., 2000, *ApJ*, 542, 535
- Graus A. S., Bullock J. S., Kelley T., Boylan-Kolchin M., Garrison-Kimmel S., Qi Y., 2019, *MNRAS*, 488, 4585
- Guo Q., White S., Li C., Boylan-Kolchin M., 2010, *MNRAS*, 404, 1111
- Hopkins P. F. et al., 2018, *MNRAS*, 480, 800
- Jahn E. D., Sales L. V., Wetzel A., Boylan-Kolchin M., Chan T. K., El-Badry K., Lazar A., Bullock J. S., 2019, *MNRAS*, 489, 5348
- Jethwa P., Erkal D., Belokurov V., 2016, *MNRAS*, 461, 2212
- Kelley T., Bullock J. S., Garrison-Kimmel S., Boylan-Kolchin M., Pawlowski M. S., Graus A. S., 2019, *MNRAS*, 487, 4409
- Komatsu E. et al., 2011, *ApJS*, 192, 18
- Ludlow A. D., Bose S., Angulo R. E., Wang L., Hellwing W. A., Navarro J. F., Cole S., Frenk C. S., 2016, *MNRAS*, 460, 1214
- McConnachie A. W., 2012, *AJ*, 144, 4
- Mao Y.-Y., Geha M., Wechsler R. H., Weiner B., Tollerud E. J., Nadler E. O., Kallivayalil N., 2021, *ApJ*, 907, 85
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Müller O., Jerjen H., 2020, *A&A*, 644, A91
- Munshi F., Brooks A. M., Christensen C., Applebaum E., Holley-Bockelmann K., Quinn T. R., Wadsley J., 2019, *ApJ*, 874, 40
- Munshi F., Brooks A., Applebaum E., Christensen C., Sligh J. P., Quinn T., 2021, *ApJ*, 923, 35
- Nadler E. O. et al., 2020, *ApJ*, 893, 48
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
- Okamoto T., Gao L., Theuns T., 2008, *MNRAS*, 390, 920
- Sales L. V., Wang W., White S. D. M., Navarro J. F., 2013, *MNRAS*, 428, 573
- Santos-Santos I. M. E., Fattahi A., Sales L. V., Navarro J. F., 2021, *MNRAS*, 504, 4551
- Sawala T. et al., 2016, *MNRAS*, 457, 1931
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Simon J. D., 2019, *ARA&A*, 57, 375
- Sofue Y., 2015, *PASJ*, 67, 75
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Springel V. et al., 2008, *MNRAS*, 391, 1685
- Tully R. B., Rizzi L., Shaya E. J., Courtois H. M., Makarov D. I., Jacobs B. A., 2009, *AJ*, 138, 323
- Tully R. B., Courtois H. M., Sorce J. G., 2016, *AJ*, 152, 50
- van den Bosch F. C., Ogiya G., Hahn O., Burkert A., 2018, *MNRAS*, 474, 3043
- Wang J., Frenk C. S., Navarro J. F., Gao L., Sawala T., 2012, *MNRAS*, 424, 2715
- Wang L., Dutton A. A., Stinson G. S., Macciò A. V., Penzo C., Kang X., Keller B. W., Wadsley J., 2015, *MNRAS*, 454, 83
- Wetzel A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, *ApJ*, 827, L23
- Wheeler C., Oñorbe J., Bullock J. S., Boylan-Kolchin M., Elbert O. D., Garrison-Kimmel S., Hopkins P. F., Kereš D., 2015, *MNRAS*, 453, 1305
- Wheeler C. et al., 2019, *MNRAS*, 490, 4447
- Woo J., Courteau S., Dekel A., 2008, *MNRAS*, 390, 1453

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