UNIVERSITY OF CALIFORNIA,
IRVINE

The Evolution of Satellite Galaxies in Local Group-like Systems

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Physics

by

John Isaac Phillips

Dissertation Committee:
Assistant Professor Michael Cooper, Chair
Professor James S. Bullock
Professor Manoj Kaplinghat

2016

DEDICATION

To everyone who has ever inspired me, with the hope that I may inspire those who come after me.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

This dissertation is the result of the hard work of many more people than myself, whithout whom it nor I would be here today.

I would firstly like to thank my advisor and mentor, Michael Cooper. Without his encouraging me and pushing me to be my best, I would not be in the position I am today, and I can't thank him enough for that. I would also like to thank James Bullock, who has also been an incredible mentor to me these past few years. If I can be half the scientist you guys are, I will consider myself a success.

Thank you to my parents, without whose support, love and encouragement I would not have made it this far. Thanks to my wonderful siblings, Charles, Sarah, and Michael, whose lives have wonderfully enriched mine. I love you guys. And thank you to Jessie, who has filled my life with such joy and love! And of course, thank you to Cassie and Percy!

FInally to my coworkers, collaborators, and friends in Irvine, you guys have made the department an amazing community that I'm so glad to be a part of. In particular, I would like to thank Basilio Yniguez, Shea Garrison-Kimmel, Coral Wheeler, Oliver Elbert, Andrew Graus, Andrew Pace, Dan Carson, and Ye Pei, as well as Jose Onorbe, Mike Boylan-Kolchin, Sarah Miller, and Evan Kirby. You guys are all amazing, and I can't wait to hear about what amazing things you do with the rest of your careers.

I would like to thank specific collaborators for certain chapters in this dissertation, in particular Mike Boylan-Kolchin in Chapters 2,3 and 4, Rodrigo Ibata in Chapter 2, and Kevin Bache and Pierre Baldi in Chapter 5. Thank you so much for your help.

Support for this work was partially provided by NASA through Hubble Fellowship grant 51316.01 awarded by the Space Telescope Science Institute, NSF grants AST-1009973 and AST-10999 and NASA grant NNX09AD09G, and NASA through a Hubble Space Telescope theory grant (program AR-12836) from the Space Telescope Science Institute (STScI).

CURRICULUM VITAE

John Isaac Phillips

EDUCATION

Doctor of Philosophy in Physics and Astronomy **2016**
University of California, Irvine *Irvine, California*

Bachelor of Science in Physics **2010**
Duke University *Durham, North Carolina*

RESEARCH EXPERIENCE

Graduate Research Assistant **2010–2016**
University of California, Irvine *Irvine, California*

TEACHING EXPERIENCE

Teaching Assistant **2010–2012**
University of California, Irvine *Irvine, California*

REFEREED JOURNAL PUBLICATIONS

A dichotomy in satellite quenching around L* galaxies **2014**
Monthly Notices of the Royal Astronomical

The mass dependence of satellite quenching in Milky Way-like haloes **2015**
Monthly Notices of the Royal Astronomical

Are rotating planes of satellite galaxies ubiquitous? **2015**
Monthly Notices of the Royal Astronomical

ABSTRACT OF THE DISSERTATION

The Evolution of Satellite Galaxies in Local Group-like Systems

By

John Isaac Phillips

Doctor of Philosophy in Physics

University of California, Irvine, 2016

Assistant Professor Michael Cooper, Chair

The evolution of galaxies is profoundly impacted by the environments in which they live. Satellite galaxies or larger hosts are particularly susceptible to this effect, and this dissertation examines the ways in which satellite galaxies and host galaxies interact and co-evolve. The picture we present coherently indicates how satellite galaxy mass impacts how satellite galaxies quench over five orders of magnitude in mass, with low-mass objects efficiently quenched, high mass objects inefficiently quenched. We also provide evidence for a sharp delineation between the two regimes, and indicate that the mass and star-formation rate of the host is correlated with dramatic changes to our model.

Chapter 1

Introduction

One of the most fundamental realizations in the study of galaxy evolution is that galaxies form a cleanly separated dichotomy. Some galaxies are star forming and active; while others are passively evolving, no longer forming stars. Such galaxies are said to be quenched (Strateva et al., 2001; Baldry et al., 2004; Blanton et al., 2005a). The increase in the number of quenched galaxies over cosmic time, roughly an increase by a factor of ~ 2 since $z = 1$ paints the picture of galaxies evolving from actively forming stars to passive, "red and dead" states, in a process called "quenching" (Cooper et al., 2006; Willmer et al., 2006).

The exact cause of galaxy quenching remains unknown, but many possible explanations have been posited. The stellar mass of a galaxy seems to play a significant role in how susceptible that galaxy is to quenching; more massive galaxies (Milky Way-mass and above) are significantly more likely be quenched (e.g. Kauffmann et al., 2003; Bell et al., 2003; Baldry et al., 2004). On the other hand, field galaxies below $M_\star \sim 10^8 M_\odot$ are extremely unlikely to be quenched (Geha et al., 2012).

Environment, however, can play a significant role in galaxy quenching. The clearest ways in which this occurs is through direct interaction between a satellite galaxy and the cir-

cumgalactic material of its host. If this material is dense enough, or if the galactic velocity with respect to the material is fast enough, the cold gas from which stars are formed can be violently ripped directly from the galaxy in a process called "ram pressure stripping" (e.g. Gunn & Gott, 1972; Bekki, 2009). Even if the forces involved are not powerful enough to rip the star forming material from the galaxy directly, they can evaporate away the reservoir of warmer gas which would eventually cool onto the galaxy. In this scenario, the galaxy continues to form stars, but once it exhausts its supply of cold gas, star formation ceases in a process called "strangulation" (e.g. Larson, Tinsley & Caldwell, 1980; Kawata & Mulchaey, 2008).

Subtler environmental effects may also be at play. It has long been seen that passive galaxies tend to be more closely clustered (e.g. Davis & Geller, 1976; Dressler, 1980; Lewis et al., 2002; Balogh et al., 2004; Hogg et al., 2004; **?**; **?**; **?**; **?**; **?**), and the tendency for galaxies of similar star formation rate to be found closer in proximity to each other, dubbed "galactic conformity," has also been more recently observed (Weinmann et al., 2006). Indeed, still more recent observations suggest that this "conformity" effect may occur on Mpc scales (Kauffmann et al., 2013). The puzzle surrounding the extent to which these environmental effects are present and possibly entangled with each other remains one of the more significant problems in galaxy evolution, and it is primarily this issue that this dissertation addresses.

The structure of this dissertation is as follows: Chapter 2 presents a problem with how environment affects the phase space distributions of satellites, asking the question, "Are rotating planes of satellite galaxies ubiquitous?" Chapter 3 points to an important dichotomy in the nature of satellite quenching existing between star-forming and quenched hosts. Chapter 4 examines how satellite quenching is affected by the masses of the galaxies involved, both host mass and satellite mass. Finally, Chapter 5 discusses a novel technique to examine low mass satellites, marrying background-subtracted photometric observations with machine learning techniques to push the frontier of satellite galaxy observations further than it has ever been.

This dissertation contains previously published material in Chapters 2,3, and 4.

Chapter 2

Co-rotating Planes of Satellites

We compare the dynamics of satellite galaxies in the Sloan Digital Sky Survey to simple models in order to test the hypothesis that a large fraction of satellites co-rotate in coherent planes. We confirm the previously-reported excess of co-rotating satellite pairs located near diametric opposition with respect to their host, but show that this signal is unlikely to be due to rotating discs (or planes) of satellites. In particular, no overabundance of co-rotating satellites pairs is observed within $\sim 20^\circ - 50^\circ$ of direct opposition, as would be expected for planar distributions inclined relative to the line-of-sight. Instead, the excess co-rotation for satellite pairs within $\sim 10^\circ$ of opposition is consistent with random noise associated with undersampling of an underlying isotropic velocity distribution. Based upon the observed dynamics of the luminous satellite population, we conclude that at most 10% of isolated hosts harbor co-rotating satellite planes (as traced by bright satellites).

2.1 Introduction

Within the ΛCDM paradigm, the growth of cosmic structure proceeds as overdensities collapse into dark matter halos, which eventually serve as the sites for galaxy formation (White & Rees, 1978; Blumenthal et al., 1984; Davis et al., 1985). Over time, the hierarchical accretion and merging of halos drives the development of substructure, such that some halos reside within the bounds of larger parent halos (Moore et al., 1998). The galaxies hosted by these subhalos are typically referred to as satellites and are important probes of the evolution of substructure, as they serve as tracers of dark matter on small scales.

High-resolution N-Body and hydrodynamic simulations confirm this picture of hierarchical structure formation, while also making predictions regarding the properties of subhalos and the satellite galaxies they host. In particular, simulated subhalos are not isotropically distributed with respect to their parent dark matter halo. Instead, simulations across a broad range of mass scales predict that satellite galaxies should preferentially lie along orbits aligned with the major axis of the host halo (e.g. van den Bosch et al., 1999; Knebe et al., 2004; Libeskind et al., 2005; Kang et al., 2007; Lovell et al., 2011; Wang, Frenk & Cooper, 2013). Two possible physical drivers are often associated with this predicted alignment of substructure with the shape of the larger gravitational potential: [i] preferential destruction (or suppression) of satellites on orbits anti-aligned with the halo's major axis (Zaritsky & Gonzalez, 1999; Peñarrubia, Kroupa & Boily, 2002; Pawlowski et al., 2012) or [ii] accretion of satellites along preferred directions, perhaps associated with large-scale filaments (Zentner et al., 2005; Libeskind et al., 2011).

Observations of galaxies in nearby groups and clusters largely support the predicted anisotropies found in simulations, such that satellites in massive dark matter halos are preferentially aligned with the major axis of the central galaxy and with the larger-scale, filamentary structure (e.g. West & Blakeslee, 2000; Plionis et al., 2003; Faltenbacher et al., 2007; Hao

et al., 2011; Tempel et al., 2015). When pushing to lower-mass, more-isolated halos, studies based on large spectroscopic samples similarly find that satellites preferentially reside along the major axis of red (or early-type) hosts, while the distribution of satellites around blue (or late-type) hosts is consistent with being isotropic (Brainerd, 2005; Sales & Lambas, 2004, 2009; Yang et al., 2006; Azzaro et al., 2007; Bailin et al., 2008, but see also Zaritsky et al. 1997). This apparent lack of spatial anisotropy for satellites of late-type hosts is potentially driven by random misalignment between the major axis of the host's disc and the dark matter halo, such that the satellites may be aligned with the latter but not the former (Bailin et al., 2005; Libeskind et al., 2007; Deason et al., 2011).

In contrast to the satellites of comparable star-forming hosts in the local Universe, observations of the Local Group suggest that the spatial distribution of satellites around both the Milky Way and M31 are significantly anisotropic. In particular, the satellites of the Milky Way preferentially reside near the northern and southern Galactic poles (i.e. along the minor axis of the Milky Way disc, Holmberg, 1969), possibly following a planar arrangement (Lynden-Bell, 1976; Kunkel & Demers, 1976; Metz, Kroupa & Jerjen, 2007; Pawlowski, Pflamm-Altenburg & Kroupa, 2012). The satellites of M31 are similarly anisotropic in their distribution, with a large subset belonging to a thin disc or plane (Karachentsev, 1996; Koch & Grebel, 2006; McConnachie & Irwin, 2006; Metz, Kroupa & Jerjen, 2009; Conn et al., 2013).

When including velocity information, the anisotropy of the Local Group satellite distribution becomes even more pronounced, with many of the Milky Way satellites following polar obits, consistent with a vast, coherently-rotating plane (Metz, Kroupa & Libeskind, 2008; Pawlowski, Kroupa & Jerjen, 2013; Pawlowski & Kroupa, 2013). For M31, a yet more-striking planar structure is observed, such that a large number of satellites exhibit coherent rotation along the line-of-sight to the Milky Way, forming a vast plane with a diameter of $\sim$ 400 kpc and a thickness of less than $\sim$ 15 kpc (Ibata et al., 2013). While simulations

predict that satellites should preferentially align with the major axis of the host dark matter halo, the strong anisotropies observed for the Local Group satellites (especially those around M31) are inconsistent with the expectations of simulated subhalo populations (Kroupa, Theis & Boily, 2005; Kroupa et al., 2010; Pawlowski et al., 2012, 2014; Ibata et al., 2014c, but see also Buck, Macci'o & Dutton 2015, who argue that co-rotating planar arrangements of satellites are predicted by ΛCDM).

The striking nature of the M31 satellite disc has served as fuel for many recent studies investigating the possibility of similar, strongly-anisotropic satellite distributions around galaxies outside of the Local Group, such as the discovery of possible planar structure in the satellite distribution of the Centaurus A group (Tully et al., 2015; Libeskind et al., 2015). In particular, recent analysis of satellite pairs in the Sloan Digital Sky Survey (SDSS, York et al., 2000) points towards the possibility of co-rotating planar satellite structures around nearby massive galaxies (Ibata et al., 2014a, hereafter I14); for 20 out of 22 systems, with satellite pairs located on diametrically-opposed sides of the host galaxy, I14 detect co-rotation along the line-of-sight, suggesting that thin satellite planes – similar to that of M31 – may be relatively common. Specifically, this result, which is supported by an analysis of the spatial positions of photometrically-selected satellite samples, indicates that $\gtrsim 50\%$ of the satellite population may reside in thin co-rotating planes (Ibata et al., 2014b, although Cautun et al. 2015b argue that the evidence for the ubiquity of such planar structures is not robust). Given the scarcity of such structures in modern simulations (Ibata et al., 2014c; Pawlowski & McGaugh, 2014, see Cautun et al. 2015a for an argument that the diversity of properties of these structures accounts for their perceived rarity.), the analysis of I14 poses a strong test of the ΛCDM cosmology and thereby warrants further investigation.

In this paper, we re-examine the kinematic evidence for the existence of co-rotating planes of satellites around nearby massive hosts by comparing the coherence of line-of-sight velocities of observed satellite galaxies to simple models of satellite spatial distributions and kinematics.

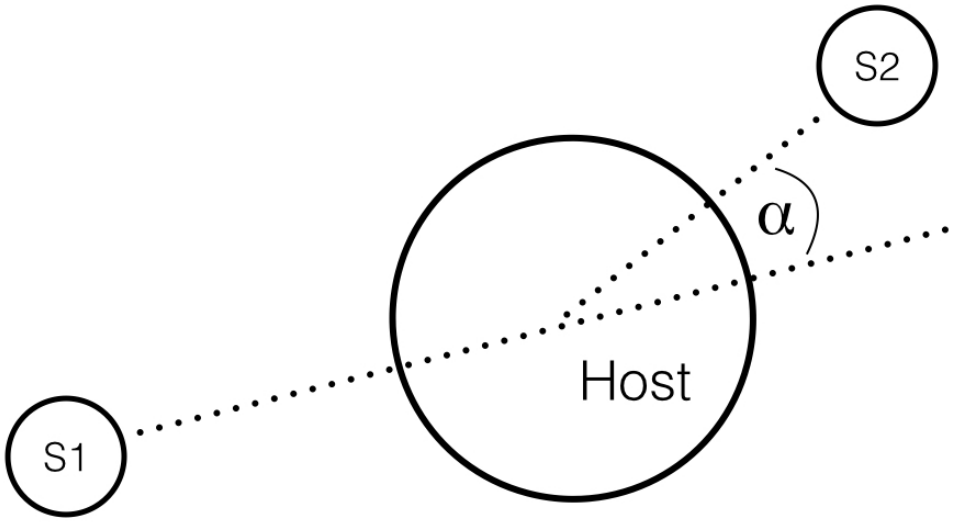

Figure 2.1: Definition of the opening angle (α) between satellite pairs (here, S1 and S2), as measured with respect to the host galaxy. Each satellite pair has a uniquely defined α, ranging from 0° to 180°, such that satellites on diametrically-opposed sides of a host correspond to $\alpha = 0°$.

The structure of the paper is as follows: in §2.2, we discuss the selection of the observational sample and the measured abundance of co-rotating satellite pairs. In §2.3, we introduce our numerical models and compare the mock observations derived from the models to the observational data. Finally, in §4.4, we discuss our results in the context of the search for M31-like planes elsewhere in the Universe. Throughout our analysis, we employ a Λ cold dark matter (ΛCDM) cosmology with WMAP7+BAO+H_0 parameters $\Omega_\Lambda = 0.73$, $\Omega_m = 0.27$, and $h = 0.70$ (Komatsu et al., 2011a), and unless otherwise noted all logarithms are base 10. Throughout the paper, we use the terms "[satellite] disc" and "[satellite] plane" interchangeably, referring to co-rotating planar satellite configurations.

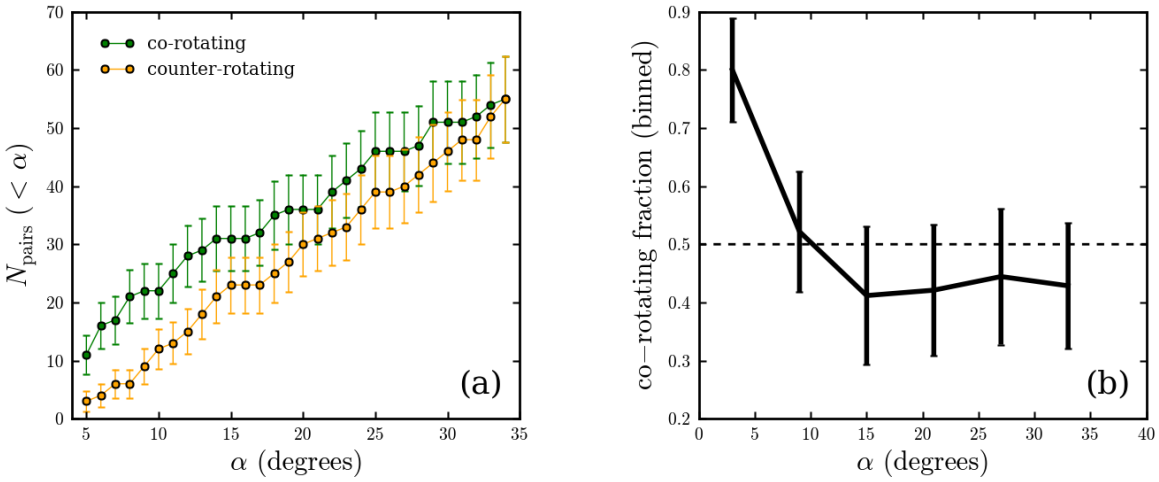

Figure 2.2: *Left*: the cumulative number of co-rotating (green line) and counter-rotating (orange line) satellite pairs as a function of opening angle (α), including uncertainties based on Poisson statistics. *Right*: the fraction of co-rotating satellite pairs as a function of opening angle, computed in distinct bins of α (bins have width 6°), with error bars derived according to the binomial theorem. There is a significant excess of co-rotating satellite pairs over counter-rotating pairs at small α (i.e. for satellites that are nearly diametrically opposed to each other). This overabundance of co-rotating pairs suggests the possible presence of coherently-rotating planar satellite structures around nearby massive host galaxies (however, see Fig. 2.3).

2.2 Observational data

2.2.1 Sample Selection

We draw our observational data from Data Release 7 (DR7, Abazajian et al., 2009a) of the SDSS, making use of the derived data products from the NYU Value-Added Galaxy Catalog (VAGC, Blanton et al., 2005a) including absolute r-band magnitudes (M_r) that are K-corrected to $z = 0.1$ using KCORRECT (Blanton & Roweis, 2007). Throughout this work, we restrict our analysis to regions of the SDSS where the spectroscopic completeness exceeds 70% (i.e. FGOTMAIN > 0.7). We also reject all galaxies with line-of-sight velocity errors greater than 25 km/s.

In selecting our galaxy sample, we adhere closely to the procedure of I14. We select a sample of hosts in the magnitude range $-23 < M_r < -20$ and within a redshift range of $0.002 < z < 0.05$. A host is considered isolated if there are no brighter objects within 500 kpc (in projection) on the sky and within 1500 km/s in velocity space. Only isolated systems are retained, reducing the number of hosts to 22,780 isolated galaxies. From this set of host systems, we identify galaxies as satellites of a given host if

(i) their magnitudes fall in the range
$M_{r,\mathrm{host}} + 1 < M_{r,\mathrm{sat}} < -16$,

(ii) they are located between 20 kpc and 150 kpc from their
host in projected distance (d_{proj}), and

(iii) their velocity offset from the host lies in the range

$$25 \text{ km/s} < |V_{\mathrm{sat}} - V_{\mathrm{host}}| < 300 \text{ km/s} \times e^{-(d_{\mathrm{proj}}/300 \text{ kpc})^{0.8}}.$$

This velocity bound is taken from I14, and is designed to reduce the contamination from interlopers in the satellite sample. Since our interest is in pairs of satellites, we retain only hosts with two or more satellites. Our final sample contains 427 such hosts, with 965 associated satellites. Note that individual hosts are allowed to harbor more than two satellites; on average, the SDSS hosts (as well as our model hosts, see §2.3) have 2.3 satellites.

2.2.2 Co-rotation signal

In this subsection, we investigate pairs of satellites for evidence of co-rotation with respect to their host. To facilitate this, we introduce the parameter α, defined as the angle between the line extending from one satellite through the host and the position vector of the second satellite relative to the host, as projected on the sky (see Figure 2.1). We define this opening angle α such that a satellite pair located on diametrically-opposed sides of a host will have an opening angle of 0°. For the duration of this work, we will refer to a satellite pair as "co-rotating" if the satellites have opposite-signed (i.e. one + and one -) line-of-sight velocity offsets relative to their host and their associated opening angle (α) is less than 90°, or if they have same-signed velocity offsets relative to their host and their associated α is greater than 90°. Otherwise, the satellite pair is deemed to be counter-rotating.

Figure 2.2a shows the cumulative number of co-rotating and counter-rotating satellite pairs in our sample as a function of opening angle at $\alpha < 35°$. At small opening angles there is a clear excess of co-rotating pairs, as first reported by I14. The overabundance of co-rotating pairs as a function of α is better illustrated in Figure 2.2b, which shows the fraction of satellite pairs that are co-rotating as a function of opening angle, computed in distinct bins of α. While the surplus of co-rotating pairs at small opening angle is readily apparent, at $10° < \alpha < 35°$ the signal is consistent with the sample being divided equally between co-rotating and counter-rotating (i.e. a co-rotation fraction of 0.5, as would be expected in

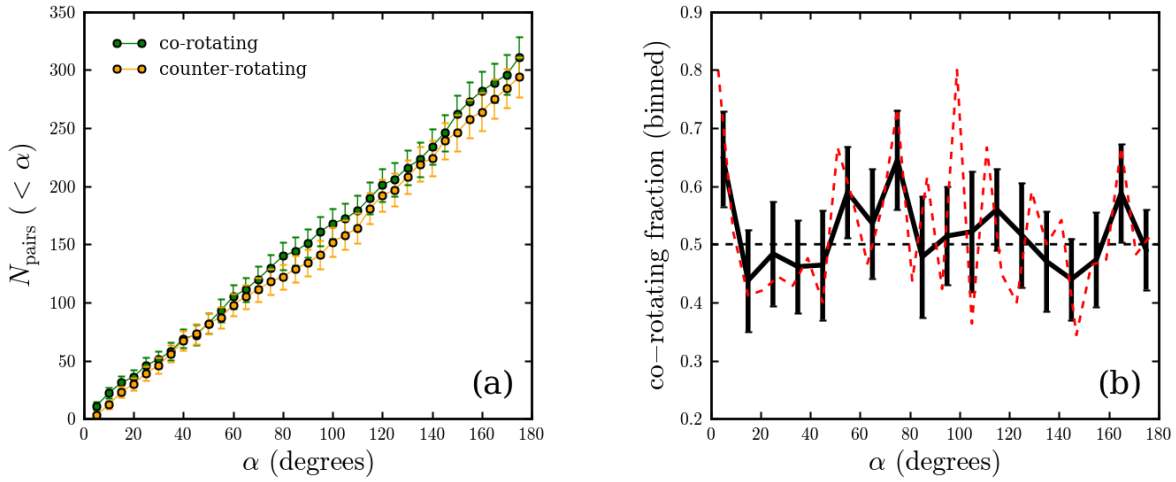

Figure 2.3: *Left*: the cumulative number of co-rotating (green line) and counter-rotating (orange line) satellite pairs as a function of α, including Poisson errors, over the full 180° domain. *Right*: the fraction of co-rotating satellite pairs as a function of opening angle, computed in distinct bins of α (bins have width 10°), with error bars derived according to the binomial theorem (solid black line). To facilitate visualization, a coarser binning is adopted than that employed in Fig. 2.2b; the dash dashed line, however, shows the corresponding dependence of co-rotating fraction on α utilizing this narrower binning. Over the full range of opening angles, the observations are largely consistent with no excess of co-rotating satellite pairs, suggesting an isotropic velocity distribution.

the absence of any co-rotating structure). At first glance, the data would seem to indicate the presence of coherently rotating structures that can only be detected at small values of α (i.e. co-rotating planes of satellites viewed close to edge-on).

When examining the co-rotating fraction of satellite pairs over the full range of opening angles (i.e. $0^\circ < \alpha < 180^\circ$), however, the evidence for planes of satellites is far less convincing. In Figure 2.3, we show (a) the cumulative counts of co-rotating and counter-rotating satellite pairs and (b) the co-rotating fraction, again in discrete bins of α, over the full range of opening angles.[1] In this light, the excess of co-rotating pairs at $\alpha < 10^\circ$ seems less likely to result from structured, coherent rotation associated with planar distributions of satellites. For example, repeatedly resampling 400 satellite pairs placed randomly in phase space (as in our "isotropic" model in §2.3) will frequently produce satellite samples with excess co-rotating fractions at random opening angles that are, by definition, not indicative of any underlying physics. Thus, care must be taken not to overinterpret the observed overadundance of co-rotating pairs at small angles, if indeed it is merely the result of random fluctuations associated with undersampling of an underlying isotropic distribution. The remainder of this paper will examine the argument that the excess of co-rotating pairs at small α is significant, and indicative of ubiquitous coherent co-rotation (similar to that observed for M31's satellite population) by comparing the SDSS data to statistical models of satellite kinematics.

2.3 Comparison to Toy Models

In order to gain insight as to whether the data presented in Figures 2.2 and 2.3 does indeed argue for the existence of coherently rotating satellite structures, we compare the SDSS data

[1]Note that the co-rotating fraction as a function of opening angle is computed using different binning procedures in Fig. 2.2b and Fig. 2.3b, so as to aid in visualization. Our results do not strongly depend on how the data are binned. The increase in co-rotating fraction at small opening angle appears less significant in Fig. 2.3b as a result of allowing satellite pairs with opening angles of $\sim 10^\circ$ in the innermost bin.

to mock observations of simple, idealized "toy" models of satellite systems. These models are not intended to give a detailed description of satellite phase-space distributions.[2] However, they do provide a meaningful basis for comparison to the observational data. We begin by detailing how each toy model is constructed:

(i) Isotropic model — For this model, each host is randomly assigned $2-5$ satellite galaxies.[3] The position and velocity of each satellite, with respect to the host, is randomly chosen to be between 0 and 200 kpc from the host (with random angular coordinates) and 0 and 200 $\mathrm{km\,s^{-1}}$, respectively. The model is randomly rotated and observed along the z direction — i.e. the z direction is taken to be the line-of-sight and the xy plane is taken to be the plane of the sky.

(ii) Disc model — In this model, $2-5$ satellites are placed randomly between 0 and 200 kpc from the origin on the xy plane (prior to any rotation taking place) and then randomly given a z coordinate between -10 and 10 kpc. All satellites are assigned a 3D velocity of 100 $\mathrm{km\,s^{-1}}$, such that each satellite is in circular motion about the host, initially rotating in the xy plane. The model is then randomly rotated and viewed along the z axis. Finally, to mimic observational error in the line-of-sight velocities, we add to the z component of each satellite's velocity a random offset drawn from a normal distribution with a standard deviation of $\sigma_V = 20$ $\mathrm{km\,s^{-1}}$. Our qualitative and quantitative results are not strongly dependent upon the assumed velocity structure or thickness of the model satellite discs; since we are only concerned with the sign (+ or -) of the z component (post-rotation) of the satellite's velocity vector, the magnitude of that vector – and any radial dependencies it

[2]In particular, the toy models do not account for potential velocity correlations due to group infall, which Cautun et al. (2015b) argues may be important, nor do they capture the complexities of observing against a background of interloper galaxies with potentially corellated velocities.

[3]In each case, where the model is permitted to have more than two satellites, we set the probability of a host having n satellites to be four times greater than the probability of having $n+1$ satellites.

might have – is largely unimportant.

(iii) M31 model — This model is based on the position and velocities of the 13 satellites belonging to the co-rotating plane identified around M31 by Ibata et al. (2013). The three-dimensional positions of the satellites are taken from McConnachie (2012) and the line-of-sight velocities are compiled from McConnachie (2012) and Collins et al. (2013). Note that we only consider the 13 satellites exhibiting coherent rotation; the two satellites aligned with the planar structure, but with counter-aligned line-of-sight velocities, are excluded. We assign each mock satellite a velocity, such that the radial component of its velocity is consistent with the observed line-of-sight velocities of the true M31 satellites, such that each satellite's total velocity puts it in circular motion around the host.[4] We then randomly select $2-5$ satellites to mock observe (independent of the luminosity of the true M31 satellites), and the system is randomly rotated and viewed along the z axis. We again add to the z component of the velocity a random offset drawn from a normal distribution with $\sigma_V = 30$ km s^{-1}, representative of measurement error in the line-of-sight velocity. While constructed using the positions and velocities of the co-rotating satellites in the M31 plane, it is useful to note that our model is not a true analog of the observed system as every member of the M31 planar structure is fainter (by $\sim 1-2$ magnitudes or more) than our satellite luminosity limit ($M_r < -16$).

(iv) Dumbbell model — In this model, each host is restricted to exactly two satellites. Once the first satellite is randomly placed on the xy plane (again, prior to any rotation), the placement of the second satellite is restricted, such that the opening angle between the two satellites is less than 10° when the system is viewed along the z axis. From there, each satellite is assigned a z coordinate between -10 and 10 kpc and the system is subject to random rotation, assigned of line-of-sight velocity errors, and finally viewed along the z axis. In essence, this model requires two satellites to be on opposite sides of their hosts

[4]The origin prior to rotation is taken to be the position of a Milky Way observer.

and orbiting in rigid-body rotation. As was the case for the disc model, our results do not strongly depend on the adopted thickness of the dumbbell.

For each realization of a model, we rotate the system to a random orientation before observing along the z axis. In our analysis, we simulate a variety of statistical samples, each consisting of $N = 10^6$ model realizations, where most samples include a mix of realizations drawn from the isotropic model along with one of the other three models. For example, the 50% disc + 50% isotropic model consists of 5×10^5 realizations of the disc model and 5×10^5 realizations of the isotropic model. Note that such a sample does *not* consist of 10^6 hosts whose satellites have a 50% probability to be placed in a disc and a 50% probability to be placed randomly. While we do not explicitly explore cases of this type, they are equivalent to the cases we explore that have a percentage disc composition of approximately $p_{\rm sat}^2$, where $p_{\rm sat}$ is the probability of an individual satellite being placed in a disc. Note that this is only an approximation, as some hosts in our models have three or more satellites. For example, a case where satellites of each host independently have a 50% chance of being placed in a disc, would be equivalent to our 25% disc + 75% isotropic model.

In Table 1, we present a summary of our analysis of the χ^2 values describing the goodness of fit for several statistical samples of modeled systems in comparison to the observed fraction of co-rotating satellite pairs as a function of opening angle in the SDSS (see Fig. 2.3b). The number of degrees of freedom that enter into each p-value calculation is based on how restrictive the model under consideration is: disc models are taken to have one fewer degree of freedom than the purely isotropic model or the model that is based on the observed positions of M31 satellites, since satellites are restricted to being placed in the disc. The dumbbell model essentially introduces an additional constraint on the disc model, so dumbbell models have yet one fewer degree of freedom.

Model	$p_{\rm sat}$	χ^2	$\tilde{\chi}^2$	p
100% Disc	1.0	460.07	28.76	< 0.001
50% Disc + 50% Isotropic	0.71	179.77	11.39	< 0.001
25% Disc + 75% Isotropic	0.50	71.20	4.45	< 0.001
10% Disc + 90% Isotropic	0.32	22.44	1.40	0.13
50% M31 + 50% Isotropic	0.71	221.79	13.05	< 0.001
50% Dumbbell + 50% Isotropic	0.71	8.70	0.58	0.89
10% Dumbbell + 90% Isotropic	0.32	21.62	1.44	0.12
100% Isotropic	0	11.40	0.67	0.83

Table 2.1: χ^2, reduced χ^2, and p values based on a comparison of the observed fraction of co-rotating satellite pairs in the SDSS versus that for various statistical samples of model satellite distributions as described in §2.3. Models in which a large fraction of hosts harbor discs of satellites (including the model based on the M31 plane) are disfavored relative to our dumbbell model or the simple isotropic case. Calculations are made taking satellite pairs over the full range of α (i.e. $0^\circ < \alpha < 180^\circ$), but the p-values are largely unchanged when we restrict the comparison to a narrower range of opening angles (e.g. $\alpha < 35^\circ$). Also shown is the probability that an *individual* satellite is found in a disc traced by bright satellites, $p_{\rm sat}$.

2.3.1 Disc Model and M31 model

In Figure 2.4, we show the observed fraction of SDSS satellite pairs that are co-rotating as a function of the opening angle α in comparison to mock observations of various disc models. The thick blue line corresponds to a statistical sample comprised purely of satellite discs, while the green, orange, and magenta lines correspond to samples composed of 50%, 25%, and 10% satellite discs, respectively, with the remainder of each sample consisting of isotropic satellites. In addition, the dashed green line corresponds to a sample with satellites for 50% of the simulated hosts following a M31 model and the other 50% distributed according to an isotropic satellite population. For comparison, the solid black line shows the co-rotating fraction for a purely isotropic sample.

We find strong disagreement between the fraction of co-rotating satellite pairs as a function of opening angle for models in which 25% or more of hosts harbor planes of satellites in comparison to that for the observed SDSS sample. In particular, the presence of inclined planes (relative to the line-of-sight) in the toy models results in a significant overabundance

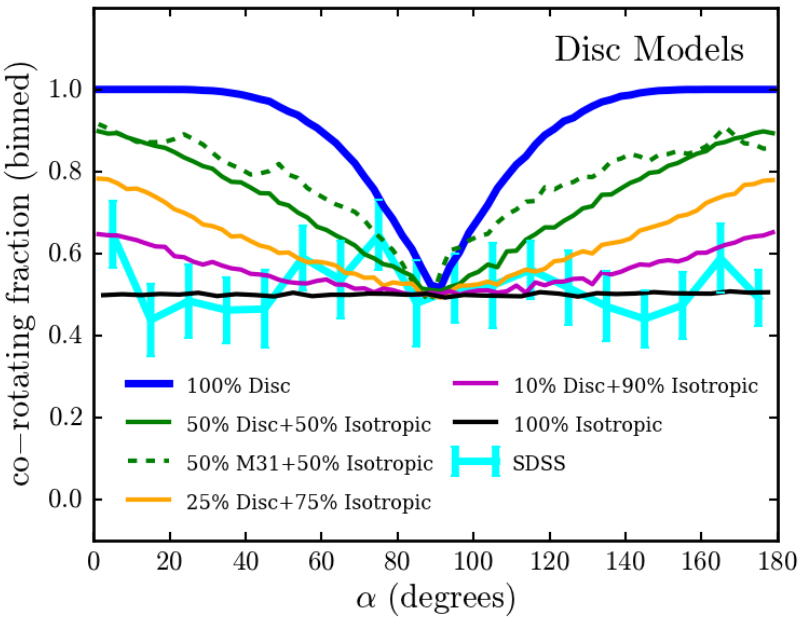

Figure 2.4: The fraction of co-rotating satellite pairs as a function of opening angle (α) for various versions of the disc model, including the model based on the M31 plane, in comparison to the observed co-rotating fraction measured in the SDSS (cyan line). The observed kinematics of bright satellites in the SDSS are consistent with at most 10% of hosts having disc-like or planar satellite distributions, with the remainder being distributed isotropically.

of co-rotating pairs at $20^\circ \lesssim \alpha \lesssim 60^\circ$ in comparison to the SDSS observations. Overall, the SDSS data agree reasonably well with models where at most $\sim 10\%$ of hosts have satellites residing in planes (or $\sim 90 - 100\%$ of the hosts have satellites distributed isotropically in phase space); however, the 100% isotropic model does fail to reproduce the overabundance of co-rotating pairs at very small opening angles ($\alpha \lesssim 10^\circ$), as measured in the SDSS sample. As shown in Fig. 2.4, the M31 sample follows the 50% disc + 50% isotropic sample very closely, as they fundamentally represent the same satellite arrangement (with the caveat that the positions and velocities of the satellites in the M31 model are tailored to match the observed positions of the true M31 satellites). While our results suggest that coherently rotating discs of bright satellites are not common, the objection could be raised that the velocity selection criteria used to select the SDSS systems systematically removes inclined satellite planes (i.e. those systems with satellite pairs at opening angles of $10^\circ \lesssim \alpha \lesssim 170^\circ$); we address this possible selection effect in §2.3.2.

2.3.2 Velocity Modeling and Cuts

As highlighted in §2.3.1, those toy models, in which a high fraction of host galaxies harbor satellite planes, are disfavored in part due to the lack of excess co-rotating satellite pairs at intermediate opening angles (i.e. $20^\circ \lesssim \alpha \lesssim 60^\circ$), corresponding to satellites in discs at non-zero inclination angles. If our sample selection criteria, in particular our velocity cut, are biasing us strongly against such systems, we could perhaps reconcile the apparent discrepancies between the disc models and the SDSS data. In a simplified test case, where each satellite orbits their host in a disc at a velocity of V_0, imposing a velocity cut of exactly V_0 on the host-satellite velocity offset would retain only perfectly edge-on discs, leading to a signal much like the one observed at small opening angles. Relaxing this velocity cut would permit progressively more face-on discs, and imposing no velocity cut would in principle permit any disc inclination angle. In constructing our model, we assigned a characteristic

velocity of 100 $\mathrm{km\,s^{-1}}$ to the satellites and only selected those satellites that have a 1D velocity offset (relative to their host) greater than $\sqrt{2} \times 25$ $\mathrm{km\,s^{-1}}$. Since the toy model is essentially scale-free, this is equivalent to removing satellites that have a 1D velocity offset less than $\sqrt{2} \times 25\%$ of the characteristic 3D velocity for satellites of the host — i.e. a velocity threshold of 0.35 V_0 is applied, where V_0 is the characteristic 3D velocity of the satellites. The higher this velocity cut, the more we would expect a contribution to the co-rotating fraction at intermediate α from inclined discs to be suppressed, as such systems would have a significant component of their velocity moving perpendicular to the line-of-sight.

In Figure 2.5, we demonstrate the impact of altering our velocity selection criterion on the measured co-rotation fraction versus opening angle for the 10% disc + 90% isotropic model. The various red lines range from selection limits of 30% of the 3D velocity to 90% of the 3D velocity, with the SDSS data included for comparison (cyan line). Reproducing the lack of excess co-rotation at $20^\circ \lesssim \alpha \lesssim 60^\circ$ — or similarly the sharpness of the increase in co-rotating satellite pairs at very small α — requires a very high velocity cut (i.e. $\sim 0.9 V_0$). Given that our adopted velocity limit is $\Delta V_{los} > \sqrt{2} \times 25$ $\mathrm{km\,s^{-1}}$, such a strong selection bias would require that the characteristic velocity (V_0) of satellites in planes would need to be ~ 40 $\mathrm{km\,s^{-1}}$. If the planar satellites have such low velocities, a velocity cut of $\sqrt{2} \times 25$ $\mathrm{km\,s^{-1}}$ would indeed correspond to 0.9 times the characteristic velocity of the plane members, and we could confidently state that we had removed all but edge-on planes.

The "toy model" adopted by I14 as a comparison to their measurements of the co-rotating satellite fraction in the SDSS utilizes exactly this characteristic velocity (40 $\mathrm{km\,s^{-1}}$); as a result, their toy model only includes edge-on (or nearly edge-one) discs, thereby reproducing the overabundance of co-rotating satellite pairs at small opening angles (see their Fig. 1b). Assuming a 3D characteristic velocity of 40 $\mathrm{km\,s^{-1}}$ for satellites in planes, however, is inconsistent with the broad expectations from subhalo kinematics in N-body simulations as well as the observed line-of-sight velocities of satellites in the M31 plane, which have a median

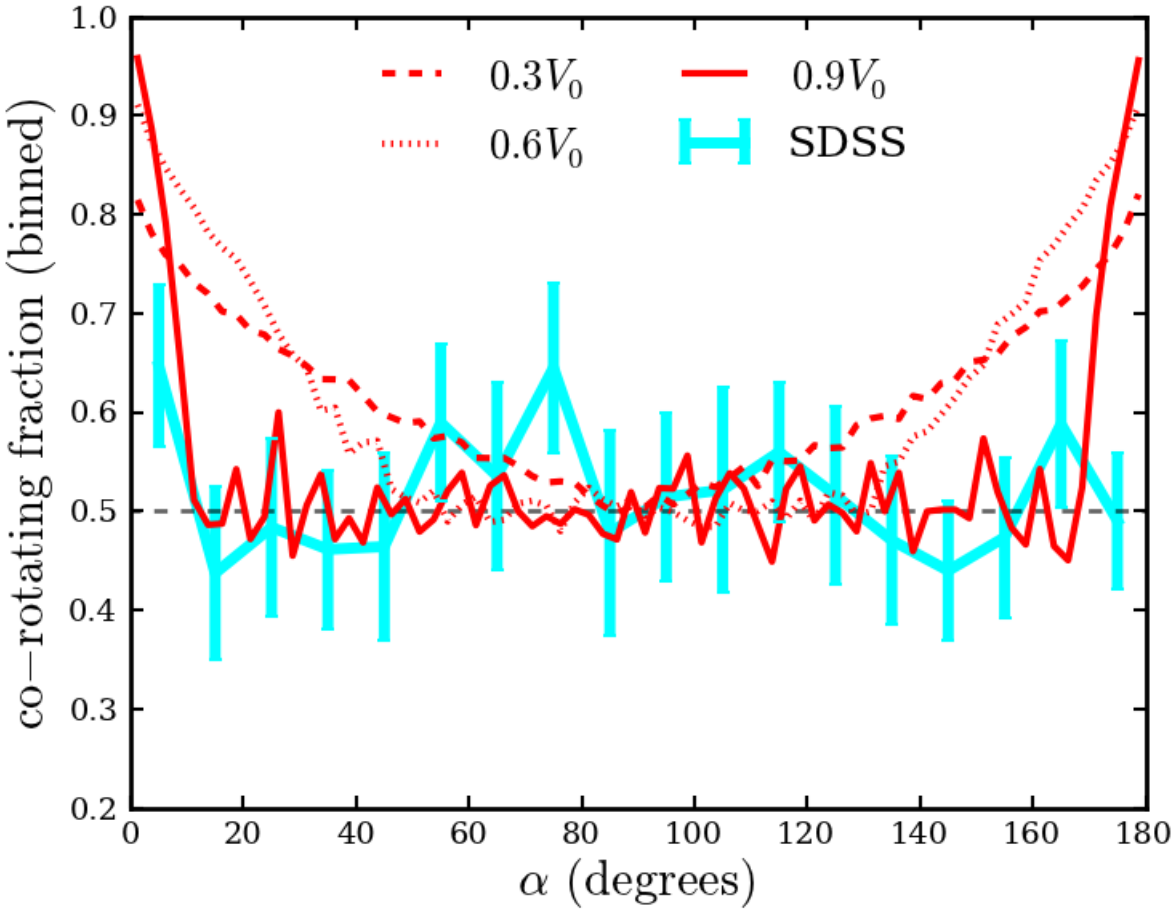

Figure 2.5: The fraction of co-rotating satellite pairs in the SDSS satellite sample (cyan line) and in the 25% disc + 75% isotropic model (red lines), varying the value of the velocity selection limit applied to the models. The dotted, dashed, and solid red lines correspond to models in which only systems with $\Delta V_{los} > 0.3$, 0.6, 0.9 V_0 are included, respectively, and where V_0 denotes the 3D characteristic velocity of satellites in the plane/disc. The dramatic overabundance of co-rotating pairs at small α in the SDSS data is only captured by models that contain nearly exclusively edge-on discs (i.e. applying the $\Delta V_{los} > 0.9$ V_0 cut).

value of $|\Delta V_{los}| \sim 92$ km s^{-1}. Moreover, such a low characteristic velocity directly conflicts with the measured line-of-sight velocity offsets of the co-rotating satellite pairs identified by I14 (see their Table 1) as well as those in our sample.

Figure 2.6 shows the host-satellite (line-of-sight) velocity offset plotted against host-satellite projected distance for each satellite in our SDSS sample, highlighting those systems belonging to pairs with $\alpha < 10^{\circ}$. The mean absolute value of the 1D (line-of-sight) velocities for the 44 satellites in co-rotating pairs with $\alpha < 10^{\circ}$ is 109.4 km s^{-1}, consistent with that for the overall sample (111.4 km s^{-1}). This indicates that a velocity cut of $\sqrt{2} \times 25$ km s^{-1} would be insufficient to select only edge-on planes, such that we should also detect co-rotation from planes at moderate inclination angles (i.e. yielding an elevated co-rotating fraction at $10^{\circ} < \alpha < 60^{\circ}$). As a result, we argue that the observational data are inconsistent with co-rotating planes being ubiquitous in the local Universe — at least with respect to satellites brighter than $M_r = -16$.

2.3.3 Dumbbell model

Recognizing that the observed variation in co-rotating satellite fraction with opening angle is strongly inconsistent with discs or planes of satellites being prevalent, we now discuss a potential alternative scenario: the dumbbell model (see §2.3). Figure 2.7 shows the co-rotating satellite fraction as a function of opening angle for the SDSS satellite sample in comparison to two formulations of the dumbbell model, one with 50% dumbbells (and 50% isotropic satellites) and one with 10% dumbbells (and 90% isotropic satellites). Both models are in relatively good agreement with the observational data. In particular, the dumbbell model is able to reproduce the observed overabundance of co-rotating pairs at small opening angles and the corresponding sharp decrease at $\alpha \sim 10^{\circ}$. In §4.4, we address the physical motivation (or lack thereof) for dumbbell-like satellite configurations. For now, we note that

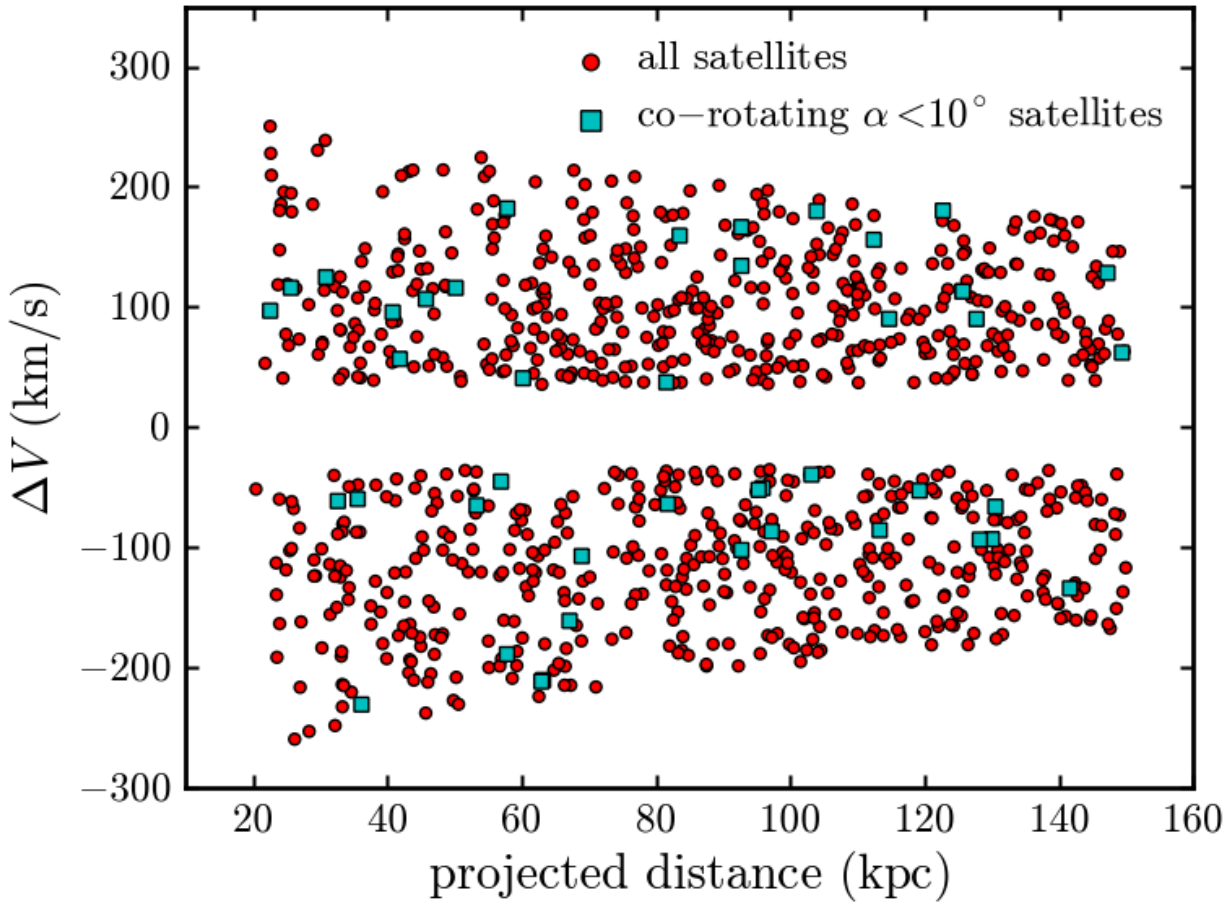

Figure 2.6: The line-of-sight velocity offset for each satellite in the SDSS sample relative to its host as a function of its projected distance from the host (red circles). Those satellites corresponding to co-rotating pairs with opening angles of $\alpha < 10°$ are highlighted as cyan squares. Given that most satellites in the sample (especially those that are co-rotating and at small opening angles) have velocity offsets greater than the lower selection limit, we conclude that our observational sample is not strongly biased against inclined planes (or discs).

the dumbbell model yields a better fit to the observed kinematic data than the disc model. The p-value for the toy model with 10% contribution from dumbbells is 0.948, meaning we fail to reject it at any confidence level. We also fail, at the 90% confidence level, to reject the 50% dumbbell model ($p = 0.152$). While we similarly fail to reject the 10% disc model at 90% confidence ($p = 0.223$), we reject all other disc models at $> 99\%$ confidence (see Table 2.1).

2.4 Discussion

Our analysis is motivated by the work of I14, which found an increase in the incidence of co-rotation in satellite pairs very near diametric opposition with respect to their host. If this excess co-rotation is the result of physical processes, it would tell us something significant about the co-evolution of satellite systems and the behavior of dark matter on small scales – perhaps indicating a serious problem with the ΛCDM model. On the other hand, we must consider the possibility that the co-rotation signal of I14 is not robust and is the product of random chance. In this section, we weight these competing possibilities.

In selecting our sample, we adhere closely to the selection criteria of I14; however, we did not reproduce the I14 satellite sample exactly. Of the host systems listed in Table 1 of I14, two galaxies fail our selection criteria. In one of these cases, the host magnitude fell just outside of our selection window, while the other host had a neighbor of nearly equal, but slightly brighter magnitude ($\Delta M_r = 0.01$). In addition, our selection includes several systems not identified by I14, though we do not expect these differences between our samples to bias our results in any meaningful way. Of the 22 satellite pairs with $\alpha < 8^\circ$ in I14, 20 (91%) appear in our sample. Moreover, we reproduce the principal result from I14, finding an excess of co-rotating satellite pairs located on diametrically-opposed sides of their host galaxy (i.e. at small opening angles). At $\alpha < 8^\circ$, we find 21 out of 27 satellite pairs to be co-rotating,

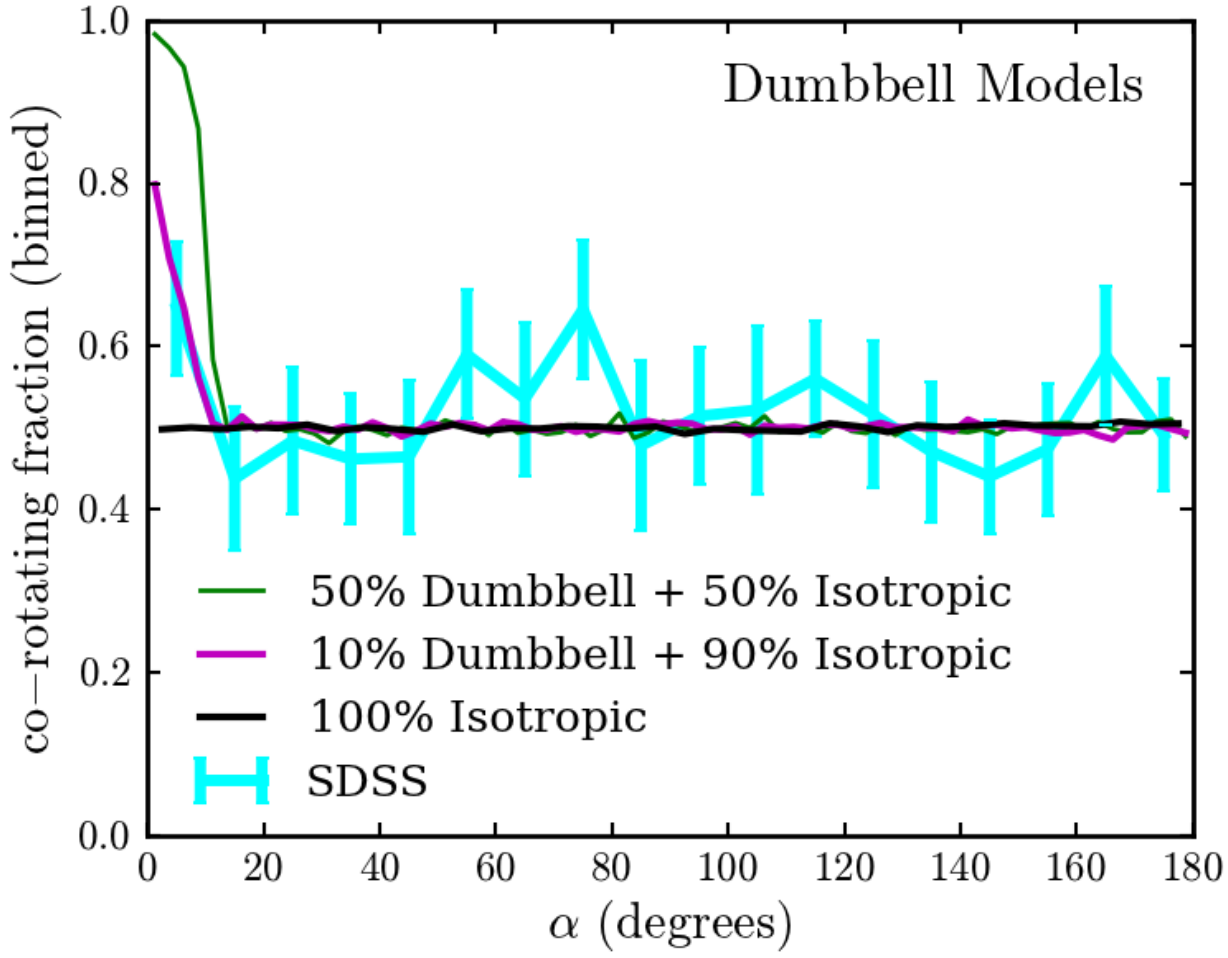

Figure 2.7: The fraction co-rotating satellite pairs as a function of opening angle α for the observed SDSS sample (cyan line) and various simulated dumbbell samples (green, magenta, and black lines). While unlikely to be physically accurate, a model in which a small fraction ($\sim 10\%$) of massive galaxies host satellites in dumbbell-like distributions is able to reproduce the overabundance of co-rotating satellites pairs at small opening angles.

corresponding to a co-rotating fraction of $0.78 \pm 0.08\%$.

However, when examining the kinematics of the entire population of satellite pairs (i.e. at all opening angles), the dependence of co-rotating fraction on α is inconsistent with the expectations from simple models of rotating discs (or planes) of satellites. For example, at large α, we would naively expect an overabundance of co-rotating satellites mirroring that detected at small opening angles. Ibata et al. (2014b) argue that caution must be taken when considering satellites on the same side of the host, as relative motion of the satellites with respect to each other (e.g. in orbital binary pairs) could dominate the motion of a satellite around the host, particularly when viewed in projection. As a result, when examining satellite pairs with large opening angles, Ibata et al. (2014b) require that the brighter of the satellite pair be at least two magnitudes fainter than the host to minimize the impact of infalling sub-groups; to mitigate deblending problems, identified pairs closer than 25 kpc in separation are also excluded. With these cuts in place, the I14 sample contains 15 pairs of satellites with $\alpha > 172^{\circ}$, 10 of which are co-rotating (Ibata et al., 2014b). Applying identical cuts to our sample, we find 13 out of 21 satellites co-rotating over the same range in α, while removing these additional selection criteria yields 44 pairs of satellites, 23 of which are co-rotating. In each case, the satellite sample at large α is divided nearly evenly between co-rotating and counter-rotating pairs, and the co-rotating excess at large α is significantly less substantive than that at small α. This agrees well with the work of Cautun et al. (2015b), which also finds no evidence of excess co-rotation in satellite pairs located on the same side of their host.

The differences between the expected kinematics of satellites belonging to discs (or planes) and that observed in the SDSS satellite sample extend far beyond the large α regime. If the observed overabundance of co-rotating satellite pairs at small α is associated with a large population of satellite planes, then we would expect that planes at non-zero inclination angles to contribute co-rotating pairs at intermediate opening angles to the observed sample,

resulting in a gradual decrease in the co-rotating fraction at $0^\circ < \alpha < 90^\circ$ (see Fig. 2.4). Instead, what is observed is a precipitous drop in the co-rotating fraction, with 21 out of 27 ($78 \pm 8.0\%$) satellite pairs found to be co-rotating at $\alpha < 8^\circ$, but only 27 out of 63 pairs ($44 \pm 6.2\%$) co-rotating over $8^\circ < \alpha < 28^\circ$. Moreover, the dependence of co-rotating fraction on α is quite noisy (see Fig. 2.3), such that excesses in the co-rotating fraction – comparable to that found at $\alpha < 10^\circ$ – are detected at other relatively narrow ranges of opening angle. For example, 20 out of the 31 satellite pairs with $70^\circ < \alpha < 80^\circ$ are co-rotating, which (if significant) is difficult to reconcile with the predicted kinematics of satellites in planes. The discrepancies between the data and the disc models are borne out in the statistical analysis, with models involving more than 10% contribution from discs (or planes) strongly disfavored (see Table 1). This is in good agreement with the predictions of Cautun et al. (2015b), which argue that only $\sim 15\%$ of bright satellites should share a coherent sense of rotation to within 25°, as would be required for planar co-rotation. Our analysis strongly disfavors the claim that planar configurations of satellites are ubiquitous over the magnitude range ($M_r < -16$) considered.

While the disc model does a relatively poor job of replicating the observational data, our toy model with satellites arranged in dumbbell configurations provides a better fit to the observed satellite kinematics. By construction, the satellites in the dumbbell model are located at opposition, such that they able to replicate the sharpness of the increase in co-rotating fraction at small α while being consistent with no excess co-rotation at larger opening angles. The physicality of the dumbbell model is questionable, however, as it is difficult to imagine a scenario that involves satellites co-rotating in such a highly-constrained configuration. Additionally, the satellites belonging to M31's plane are not in a dumbbell configuration. For these reason, we conclude that the data are not likely to be the result of dumbbell configurations either.

In the absence of viable alternative models, we argue that the excess of co-rotating satellite

pairs at small α is very likely the result of random noise. Among the disc models, only the model with a 10% contribution from satellite discs is a comparably good fit to the SDSS data as the purely isotropic model. Moreover, 25% of random realizations of an isotropic model comprised of 400 hosts (i.e. comparable in number to the observed SDSS sample) yield excess co- and counter-rotating pairs of satellites comparable in significance to that measured in the SDSS sample, although just $\sim 1-2\%$ show an overabundance of co-rotating pairs at small opening angles ($\alpha < 20^{\circ}$).

These results should not be taken to say that there are no planes of satellites in the Universe. Statistically, the observational data are roughly consistent with $\lesssim 10\%$ of isolated, massive galaxies playing host to planes, or more specifically planes with multiple luminous satellites. Our result does not exclude the possibility that planar structures preferentially populated by faint satellites could be ubiquitous; if true, this could provide powerful insight into the formation of planar satellite structures. With the caveat that all of the satellites under our consideration here are brighter than the satellites belonging to the M31 plane, we exclude the possibility of ubiquitous planes analogous to the M31 plane at a very high confidence level.

2.5 Conclusions

In this work, we investigate the ubiquity of co-rotating planes of satellites similar to that observed around M31. Using data drawn from the SDSS, we study the orientation and kinematics of bright ($M_r < -16$) satellite pairs around isolated galaxies, selected to be roughly analogous to the Milky Way and M31. By comparing the fraction of co-rotating satellites pairs as a function of opening angle to the predictions of simple models, we investigate the signatures of coherent rotation arising from satellites arranged in planar structures. Our findings are as follows:

- We confirm an excess of co-rotating pairs of satellites at opening angles of $\alpha < 10^{\circ}$ — i.e. satellite pairs located on diametrically-opposed sides of their host (see Fig. 2.2). This overabundance of co-rotating pairs at small opening angles, as first identified by I14, has been cited as evidence that 50% of satellites belong to coherently-rotating planes (Ibata et al., 2014b).

- We find that the excess of co-rotating pairs at small opening angles is unlikely to be due to ubiquitous co-rotating planes of satellites. While the behavior at small α is suggestive of planes aligned along the line-of-sight (i.e. with an orbital inclination of zero), the signal is strongly inconsistent with mock observations of satellite galaxy planes (or discs). In particular, we find no contribution to the co-rotating fraction from planes inclined relative to the line-of-sight (i.e. with satellites configured at opening angles of $10^{\circ} \lesssim \alpha \lesssim 60^{\circ}$). For our sample of isolated systems, we find that at most $\sim 10\%$ of hosts harbor satellites planes – as traced by the luminous (LMC-like) satellite population.

- The excess of co-rotating pairs of satellites at small α is better fit by a "dumbbell" model, where satellites have co-rotating partners located opposite the host galaxy (but not a true plane). However, this model is very likely unphysical.

- We do not rule out the possibility that $\sim 10\%$ of hosts having co-rotating planes of satellites, or correspondingly $\sim 30\%$ of satellites residing in such planes. This case is similar enough to the isotropic case that the observed SDSS co-rotating fraction as a function of opening angle is roughly consistent with $\lesssim 10\%$ of massive host galaxies harboring planes of satellites.

Chapter 3

A Dichotomy in Satellite Quenching

We examine the star formation properties of bright ($\sim 0.1\,L^*$) satellites around isolated $\sim L^*$ hosts in the local Universe using spectroscopically confirmed systems in the Sloan Digital Sky Survey DR7. Our selection method is carefully designed with the aid of N-body simulations to avoid groups and clusters. We find that satellites are significantly more likely to be quenched than a stellar mass-matched sample of isolated galaxies. Remarkably, this quenching occurs only for satellites of hosts that are themselves quenched: while star formation is unaffected in the satellites of star-forming hosts, satellites around quiescent hosts are more than twice as likely to be quenched than stellar-mass matched field samples. One implication of this is that whatever shuts down star formation in isolated, passive L^* galaxies also plays at least an indirect role in quenching star formation in their bright satellites. The previously-reported tendency for "galactic conformity" in colour/morphology may be a by-product of this host-specific quenching dichotomy. The Sérsic indices of quenched satellites are statistically identical to those of field galaxies with the same specific star formation rates, suggesting that environmental and secular quenching give rise to the same morphological structure. By studying the distribution of pairwise velocities between the hosts and satellites, we find dynamical evidence that passive host galaxies reside in dark matter haloes that are

$\sim 45\%$ more massive than those of star-forming host galaxies of the same stellar mass. We emphasize that even around passive hosts, the mere fact that galaxies become satellites does not typically result in star formation quenching: we find that only $\sim 30\%$ of $\sim 0.1L^*$ galaxies that fall in from the field are quenched around passive hosts, compared with $\sim 0\%$ around star forming hosts.

3.1 Introduction

In the now-favored dark energy plus cold dark matter (ΛCDM) cosmological model, the abundance of dark matter subhaloes as a function of parent dark matter halo mass is predicted with high precision in N-body simulations (e.g. Boylan-Kolchin et al., 2009; Klypin, Trujillo-Gomez & Primack, 2011). In contrast, the observed abundance and properties of satellite galaxies in the local Universe are poorly understood within the current cosmological framework. For example, modern semi-analytic (e.g. Somerville et al., 2008; Guo et al., 2011) and hydrodynamic (e.g. Davé, Oppenheimer & Finlator, 2011) models of galaxy formation are overly effective at halting star formation in satellite systems, such that current models dramatically overpredict the number density of quenched (or passive) satellites at $z \sim 0$ (Kimm et al., 2009; Weinmann et al., 2006, 2010, 2012).

Additionally, current models of galaxy formation fail to reproduce the relationship between host and satellite galaxy properties. In particular, recent studies of satellite and central galaxies in the Sloan Digital Sky Survey (SDSS, York et al., 2000) find that satellites of red (or passive) host galaxies are more likely to also be passive relative to their counterparts around star-forming hosts (Weinmann et al., 2006). This correlation between host and satellite properties, commonly dubbed "galactic conformity," is also poorly replicated by modern semi-analytic models (Kauffmann et al., 2013), potentially due to a lack of decoupling between the growth of satellite galaxies and the growth of their dark matter haloes in the

models (Weinmann et al., 2012). Ultimately, understanding this connection between the properties of host and satellite galaxies may be a powerful way to constrain the physical mechanisms responsible for halting star formation in passive systems.

To better understand the physics driving the relationship between the properties of host and satellite galaxies, it is particularly interesting to examine the satellite populations around L^* hosts. At this luminosity (or stellar mass), central galaxies are roughly equally likely to be star forming or quenched, while at higher (or lower) luminosities, the host population is increasingly dominated by passive (or star forming) systems (e.g. Kauffmann et al., 2003; Bell et al., 2003; Baldry et al., 2004). As such, this scale ($\sim L^*$) is well-suited for exploring connections between host and satellite properties. Moreover, our own Milky Way falls into this class of galaxies; understanding the statistical properties of L^* galaxies will therefore inform our understanding of the Milky Way and its satellites, the objects that can be studied in the most detail.

In this work, we investigate satellite quenching mechanisms by comparing satellites around isolated L^* galaxies to a statistically identical sample of field galaxies. In doing so, we build upon the work presented in Tollerud et al. (2011a), who used strict isolation criteria to explore satellite galaxy counts around galaxies in the field. By focusing on isolated hosts, we are able to eliminate the known effects of massive haloes (groups or clusters) and large-scale structure on satellites. We further split the samples based on the star forming activity of the host to investigate the origins and strength of correlations between star formation properties of satellites and hosts (i.e. the physical drivers of galactic conformity). In Section 2, we introduce our observational sample, use N-body simulations to develop criteria for selecting isolated L^* hosts, and quantify the expected level of interloper contamination. In Section 3.3, we describe the selection criteria as applied to the SDSS to create our galaxy samples. In Section 5.4, we present our findings on the relationship between star formation properties of satellites and their host galaxies. Finally, we discuss and summarize the implications

of the observed trends on galaxy formation models in Section 4.4 and Section 3.6. Unless otherwise noted, all logarithms are base 10. Halo virial masses are defined with respect to an overdensity of 94 relative to the critical density of the universe (Bryan & Norman, 1998). Throughout our analysis, we employ a Λ cold dark matter (ΛCDM) cosmology with WMAP7+BAO+H0 parameters $\Omega_\Lambda = 0.73$, $\Omega_m = 0.27$, and $h = 0.70$ (Komatsu et al., 2011b).

3.2 Using Simulations to Define the Observational Sample

Before we detail our approach to isolate host galaxies and their satellites, a brief description of our adopted terminology is in order.

We use the term "satellite" to refer to a galaxy that is within the virial volume of a central "host" galaxy's halo. Given that the observational data are subject to projection effects, we refer to our raw sample of observed galaxy pairs as "primaries" and "secondaries". We also construct theoretical mock catalogs, and use the terms host and satellite once the full three-dimensional information in the simulation is used to confirm whether an "observed" secondary is actually within the virial volume of the primary. Finally, based on these mock observations, our observationally-derived statistics will be corrected for contamination effects; accordingly, we will use the terms "host" and "satellite" in discussing contamination-corrected observed parameter distributions as well.

3.2.1 Observational Data

In selecting our observational sample, we employ data from Data Release 7 (DR7) of the Sloan Digital Sky Survey (SDSS York et al., 2000; Abazajian et al., 2009b). In particular, we utilize the MPA-JHU derived data products, including median total stellar masses, photometrically derived according to Kauffmann et al. (2003, see also Salim et al. 2007), and median total star formation rates, measured from the SDSS spectra as detailed by Brinchmann et al. (2004). Note that their procedure includes methods for estimating specific star formation rates when they are unobtainable from emission lines. Supplemental information is drawn from the NYU Value-Added Galaxy Catalog (NYU-VAGC Blanton et al., 2005b), such as Sérsic indices (Sérsic, 1968) that are derived from one-component fits according to Blanton et al. (2003, 2005a).

As a first step in identifying satellite and host galaxies, we compile a list of secondaries containing all galaxies with a stellar mass of $10^{9.5}\mathrm{M}_{\odot} < \mathrm{M}_{\star} < 10^{10.5}\mathrm{M}_{\odot}$ and a list of primaries with $\mathrm{M}_{\star} > 10^{10.5}\mathrm{M}_{\odot}$. As discussed in more detail in Section 3.2.3, we then apply isolation criteria to the potential primary sample and spectroscopically search for physically-associated secondaries. We also impose a spectroscopic completeness (*fgotmain*) cut of 0.7. This cut corresponds to a mean sample spectroscopic completness of 92%, making it extremely unlikely that an object for which no redshift was obtained would impact our isolation procedure. Furthermore, we do not expect the small remaining incompleteness to bias our satellite selection. Our final samples of primary and secondary galaxies have mean stellar masses of $\mathrm{M}_{\star} = 10^{10.80}\mathrm{M}_{\odot}$ and $\mathrm{M}_{\star} = 10^{9.97}\mathrm{M}_{\odot}$, respectively. We impose a limiting redshift of $z = 0.032$, within which we are complete to a stellar mass of $10^{9.5}$ $\mathrm{M}_{\odot}$. The mean redshift of our sample is 0.027. Our final sample contains 457 primary/secondary systems. These parameters are similar to the selection criteria described in Tollerud et al. (2011a).

We further divide the primaries into active and passive categories, with the dividing line

between the the two star formation rate classes given by

$$\log(\mathrm{SSFR_{host}}) = -0.6\ \log(M_{\star,\mathrm{host}}) - 5.2, \tag{3.1}$$

where $\mathrm{SSFR_{host}}$ denotes the specific star formation rate of the host galaxy. This equation is motivated by the established blue cloud/red sequence dichotomy of galaxies in the SDSS (Strateva et al., 2001; Baldry et al., 2004; Blanton et al., 2005a). The slope was selected to match the slope of the red sequence in our sample.

3.2.2 Numerical Simulation

Our goal of studying satellites of isolated L^* galaxies requires a rigorous and accurate identification of truly isolated L^* primaries, systems that are largely excluded from residing within group and cluster environments. Accordingly, we use the Millennium-II Simulation (hereafter, MS-II; Boylan-Kolchin et al. 2009) to test our selection and isolation procedures. The MS-II has a box size of 137 Mpc, well-matched to our observational volume at $z < 0.032$. Its high mass resolution – $m_p = 9.4 \times 10^6\,\mathrm{M_\odot}$ – ensures that it is complete for halo masses in excess of $\sim 2 \times 10^{10}\,\mathrm{M_\odot}$, corresponding to stellar masses of $\sim 10^{7.5-7.75}\,\mathrm{M_\odot}$ (Guo et al., 2010; Behroozi, Conroy & Wechsler, 2010; Behroozi, Wechsler & Conroy, 2012; Leauthaud et al., 2012; Moster, Naab & White, 2013). Using the MS-II, Tollerud et al. (2011a) demonstrate the effectiveness of this technique in obtaining clean samples of isolated hosts with halo masses of roughly a few $\times 10^{12}\,\mathrm{M_\odot}$.

Following the abundance matching prescription of Guo et al. (2010), we find that our lower limit for the stellar mass of secondaries ($\mathrm{M_\star} = 10^{9.5}\mathrm{M_\odot}$) corresponds to $v_{\mathrm{max}} = 94.8\,\mathrm{km\,s^{-1}}$. The cutoff between what we consider to be a primary and a secondary, $v_{\mathrm{max,cut}}$, corresponds to a value of $v_{\mathrm{max,cut}} = 166.5\,\mathrm{km\,s^{-1}}$. We use these values to construct mock catalogs below.

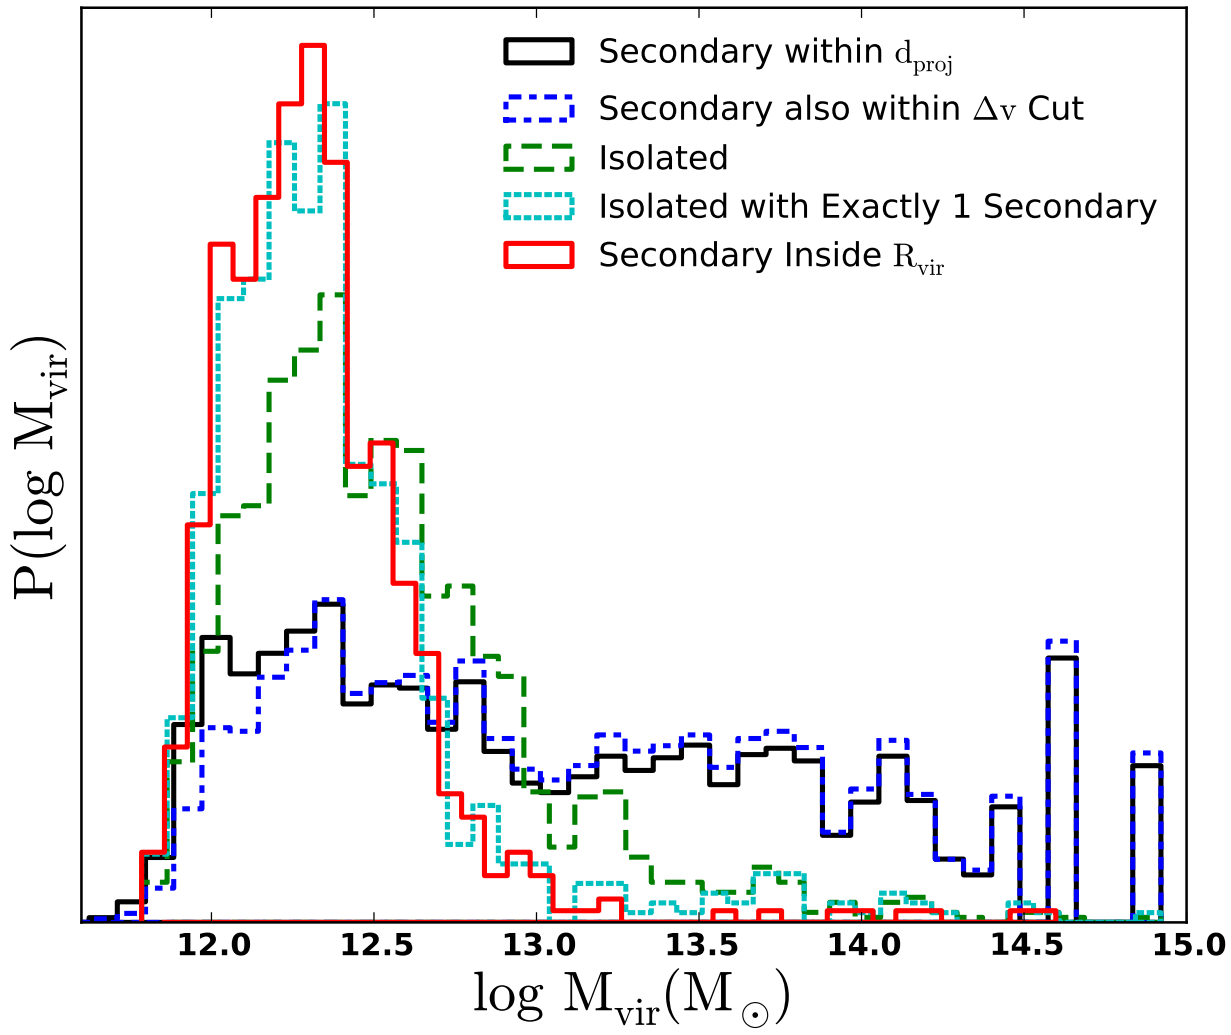

Figure 3.1: Probability distribution of virial masses of primaries associated with each secondary within a projected distance of $d_{\rm proj} = 355\,{\rm kpc}$ for various selection cuts in the MS-II mock catalogs. Using only this $d_{\rm proj}$ criterion (black solid line), many of the selected secondaries are associated with very massive dark matter host primaries. Adding a restriction of $\Delta v_{\rm sel} = 500\,{\rm km\,s^{-1}}$ does not help in isolating MW-size hosts, as cluster and group haloes typically contain galaxy pairs that satisfy this criterion. Further restricting this subset of primaries to those that that are isolated (according to our fiducial set of isolation criteria as defined in Section 3.2.4) *significantly* reduces the high-mass tail (green dashed line). Adding the constraint that primaries have *exactly* one secondary (cyan dotted line) further decreases the number of group-mass ($\sim 10^{13}\,{\rm M}_\odot$) haloes selected. This line represents the halo mass distribution for our full sample. The red line demonstrates that in most group-scale systems in our sample, the secondary is outside the virial radius of the group's central halo. Thus, even for the small number of group-sized haloes in our sample ($< 7\%$ of secondaries), we do not expect physical mechanisms specific to the group environment to significantly bias our results.

3.2.3 Obtaining Isolated L^* Galaxies

Figure 4.1 illustrates the difficulty in selecting satellite galaxies around isolated L^* host haloes without significant contamination by cluster or group members. We use the MS-II simulation to perform a series of mock observations to identify secondaries and then determine the host virial mass of the associated subhaloes (horizontal axis) for various selection criteria. The black solid histogram shows the host halo virial mass distribution for all secondaries within a projected distance of $d_{\rm proj} = 355\,{\rm kpc}$ of an identified primary. We count one host halo for each secondary identified, implying that that same host halo can appear more than once if it hosts multiple secondary candidates. It is evident that this selection picks out a large percentage of satellites that reside within group and cluster-mass haloes ($M_{\rm vir} > 10^{13}\,{\rm M}_\odot$).

A standard additional criterion for spectroscopic secondary selection is a cut on the velocity difference between the primary and secondary, Δv. As we might expect satellites of group sized haloes to be kinematically warm, this cut could prove useful for excluding such massive hosts. However, restricting pairs to lie within a projected velocity limit $\Delta v_{\rm sel} = 500\,{\rm km\,s^{-1}}$ only weakly eliminates haloes within the cluster and group mass regime (blue dashed histogram), while preferentially removing galaxy-sized ($< 10^{12.5}\,{\rm M}_\odot$) haloes (see blue dashed histogram in Fig. 4.1). Such a cut on velocity difference is ineffective at removing higher-mass haloes since group environments contain many objects – with a wide range of velocities – in the secondary range. As such, it is very likely that one of these objects will meet both our projected distance and projected velocity criteria, even if many or even most will not.

In order to filter out satellites associated with massive host haloes, we adopt an approach pioneered by Barton et al. (2007), which relies on strict isolation cuts (see also Tollerud et al., 2011a; Edman, Barton & Bullock, 2012). Specifically, we impose limitations on $n_{\rm iso}$, the maximum number of other haloes (galaxies) in the primary $v_{\rm max}$ ($M_\star$) range that we allow within an annulus bounded on the interior by $d_{\rm sel}$ and on the exterior by $d_{\rm iso}$, and bounded

in velocity space by $\pm\Delta v_{iso}$. We never allow any other primaries within $d_{\rm sel}$ in projected distance space and $\pm\Delta v_{iso}$ in velocity space. The green dashed histogram in Figure 4.1 shows the host halo mass distribution when we restrict our selection to only those haloes that (1) have no other primaries within $d_{\rm sel} = 355\,{\rm kpc}$ and $\Delta v_{\rm iso} = 1000\,{\rm km\,s^{-1}}$ and (2) have no more than one other primary within $d_{\rm iso} = 1\,{\rm Mpc}$ and $\Delta v_{iso} = 1000\,{\rm km\,s^{-1}}$. We will show below that this particular set of criteria are optimal for isolating MW-sized haloes and their LMC-sized companions, and also for maximizing the purity of our pair sample. With the isolation criteria imposed, we greatly reduce the number of haloes in our sample that inhabit clusters and somewhat reduce those in the group mass regime as well.

As a final cut, we select only those haloes that have *exactly* one secondary within 355 kpc and $|\Delta v| < 500, {\rm km\,s^{-1}}$. Applying this additional selection criterion (see the cyan dotted histogram in Fig. 4.1) removes nearly all pairs that fall within groups and cluster, yielding a population dominated by Milky Way-like haloes; only 6.7% of objects satisfying our isolation and satellite number cuts occupy haloes with $M_{\rm vir} > 10^{13}\,{\rm M_\odot}$ within the MS-II. Moreover, for the small tail of secondaries associated with massive primaries, only 38% of them actually lie inside the virial radius of the identified primary – the rest sit at the outskirts of their associated friends-of-friends group in the MS-II catalog. The red solid histogram in Figure 4.1 shows the mass distribution for secondary hosts that lie within the virial radius of their respective host haloes. The resultant histogram clearly picks out the galaxy-mass scale for hosts. The median host mass for this distribution is ${\rm M_{vir}} = 10^{12.29}\,{\rm M_\odot}$ with a 68% spread of $10^{12.17} - 10^{12.40}\,{\rm M_\odot}$. While it would be impossible to employ this in our observational sample, we do not expect satellites that lie outside of their host's virial radius to contribute significantly to a quenching signal (Wetzel et al., 2013b).

The selection cuts illustrated by the cyan histogram in Figure 4.1 represent our fiducial choices for five parameters that we have tuned to identify exclusively bright satellites around isolated $\sim L^*$ hosts. These choices are based on tests against mock observations in the MS-II

simulation:

- $d_{\rm sel} = 355\,{\rm kpc}$: The maximum projected separation between primary and secondary.
- $d_{\rm iso} = 1\,{\rm Mpc}$: The maximum projected distance within which we check for neighbouring primaries.
- $n_{\rm iso} = 1$: The maximum number of neighbouring primaries allowed between $d_{\rm proj}$ and $d_{\rm iso}$ within $\Delta v_{\rm iso}$ in velocity. We always require that no neighbouring primary can fall within $d_{\rm sel}$ in projected distance and $\Delta v_{\rm iso}$ in velocity, such that the primary in our sample have no other primary within a distance of $d_{\rm sel}$.
- $\Delta v_{\rm sel} = 500\,{\rm km\,s^{-1}}$: The maximum line-of-sight velocity difference between primary and secondary.
- $\Delta v_{\rm iso} = 1000\,{\rm km\,s^{-1}}$: The minimum line-of-sight velocity difference required between the primary and any other neighbouring primary within a projected distance of $d_{\rm iso}$ to not count towards $n_{\rm iso}$. No neighbouring primary is allowed within $d_{\rm sel}$ and $\Delta v_{\rm iso}$.

In the next subsection we illustrate how these fiducial choices were motivated to balance purity and sample size.

3.2.4 Balancing Purity and Sample Size

A crucial characteristic of any observational sample is its purity, $f_{\rm purity}$, which we define within the MS-II to be the fraction of identified host-satellite pairs having (1) primaries that are actually host haloes (i.e. not satellites of a larger system); (2) secondaries that are actually satellites (i.e. not lower-mass primaries); and (3) satellites that are actually satellites of the host (and not satellites of a different host that are included because of projection). All of the parameters listed in Section 3.2.3 could be tuned to increase the purity of the sample,

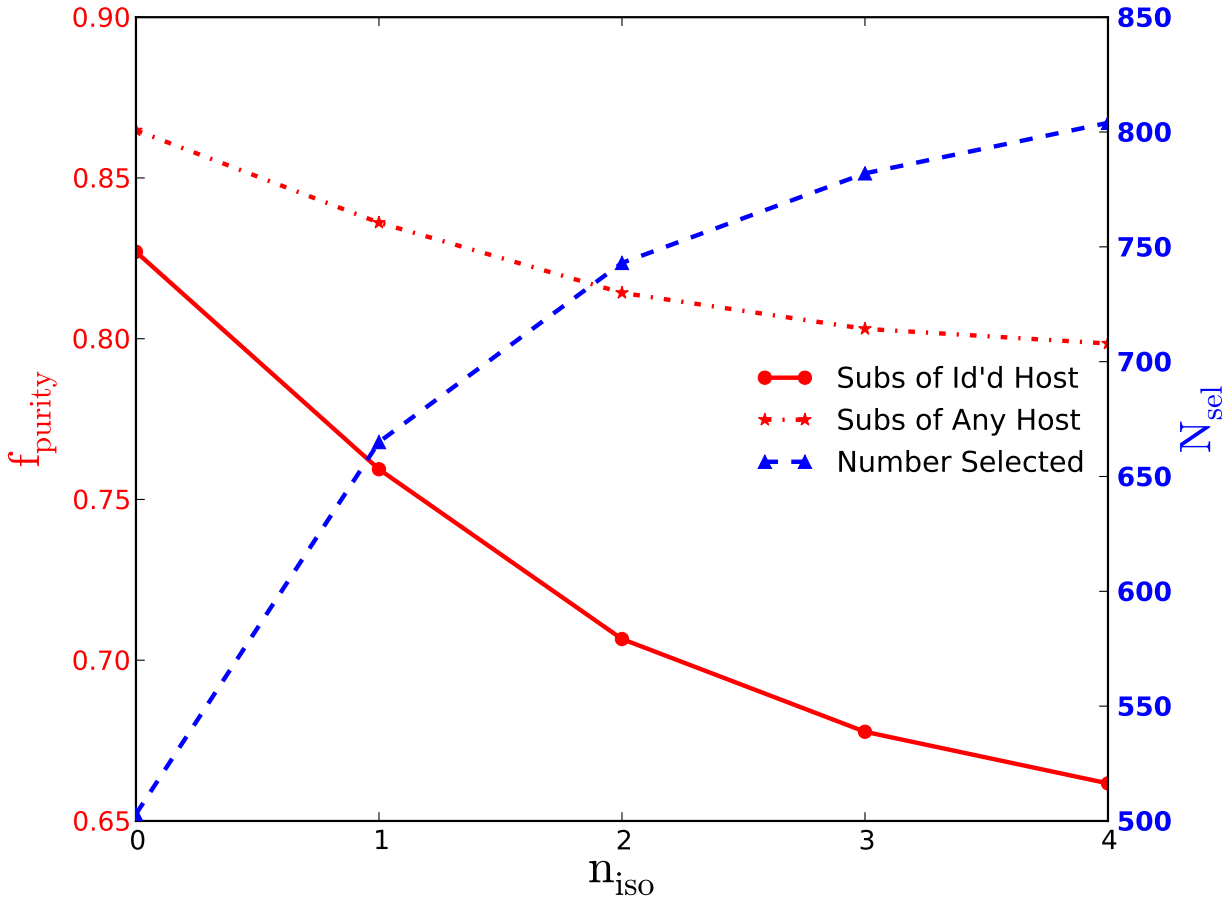

Figure 3.2: Fraction of the pair sample (in the MS-II) for which the secondary is a subhalo of the primary (red solid line) and for which the secondary is a subhalo of any host (red dash-dotted line) alongside the number of pairs selected (blue dashed line) as a function of one of our isolation criteria, $n_{\rm iso}$. While purity rises sharply as $n_{\rm iso}$ approaches 0, the number of pairs selected decreases precipitously. As a compromise between sample size and purity, we define our selection limit to be $n_{\rm iso} = 1$. At this value of $n_{\rm iso}$, the secondary is a subhalo of the primary in 75% of pairs while the secondary is a subhalo of any host in just over 85% of the sample.

however this usually involves a trade-off with regard to sample size. As an illustration, we present one such example of this tuning: Figure 3.2 illustrates the dependence of sample purity on our isolation criterion, $n_{\rm iso}$, holding all other parameters fixed at their fiducial values. As expected, increasing $n_{\rm iso}$ increases the number of interlopers and decreases $f_{\rm purity}$ (solid red line). While the purity of our sample increases with the strictness of our isolation criterion, $n_{\rm iso}$, the sample size simultaneously decreases (blue dashed line). As a compromise between sample size and purity, we define our selection limit to be $n_{\rm iso} \leq 1$ throughout the remainder of this work, yielding a sample purity of 75%. While using $n_{\rm iso} = 0$ would increase the purity to 82%, it would also reduce the sample size by more than 25%.

Note that if we relax the definition of a true pair to require only that the secondary be a subhalo, rather than being a subhalo of the chosen host, our purity reaches approximately 85% and depends little on $n_{\rm iso}$ (dotted red line). In subsequent sections, we perform corrections to account for the presence of interlopers in the observational data using this modified definition of purity. We make the assumption that the interloper population is comprised only of field galaxies having the same specific star formation rate distribution as the control sample (see Section 5.4).

Purity could also vary within our sample as a function of separation between the primary and secondary galaxy, either in physical space or velocity space. While sample size and purity exhibit little dependence on Δv, we find significant variation in purity with the projected separation between the primary and secondary. As shown in Figure 3.3, sample purity increases with decreasing $d_{\rm proj}$ for both our standard and less-restrictive definition of purity (blue solid and red dashed lines, respectively). For close separations ($\sim 60\,$kpc), approximately 88% of secondaries are subhaloes of their respective primaries. At larger distances, however, the problem of contamination increases noticeably. At $\sim 200\,$kpc, the purity drops to $\sim 77\%$, while it falls to just over 50% at $300\,$kpc. Even when considering whether the secondary is a subhalo of any host, the purity decreases from $\sim 96\%$ in the innermost bin to

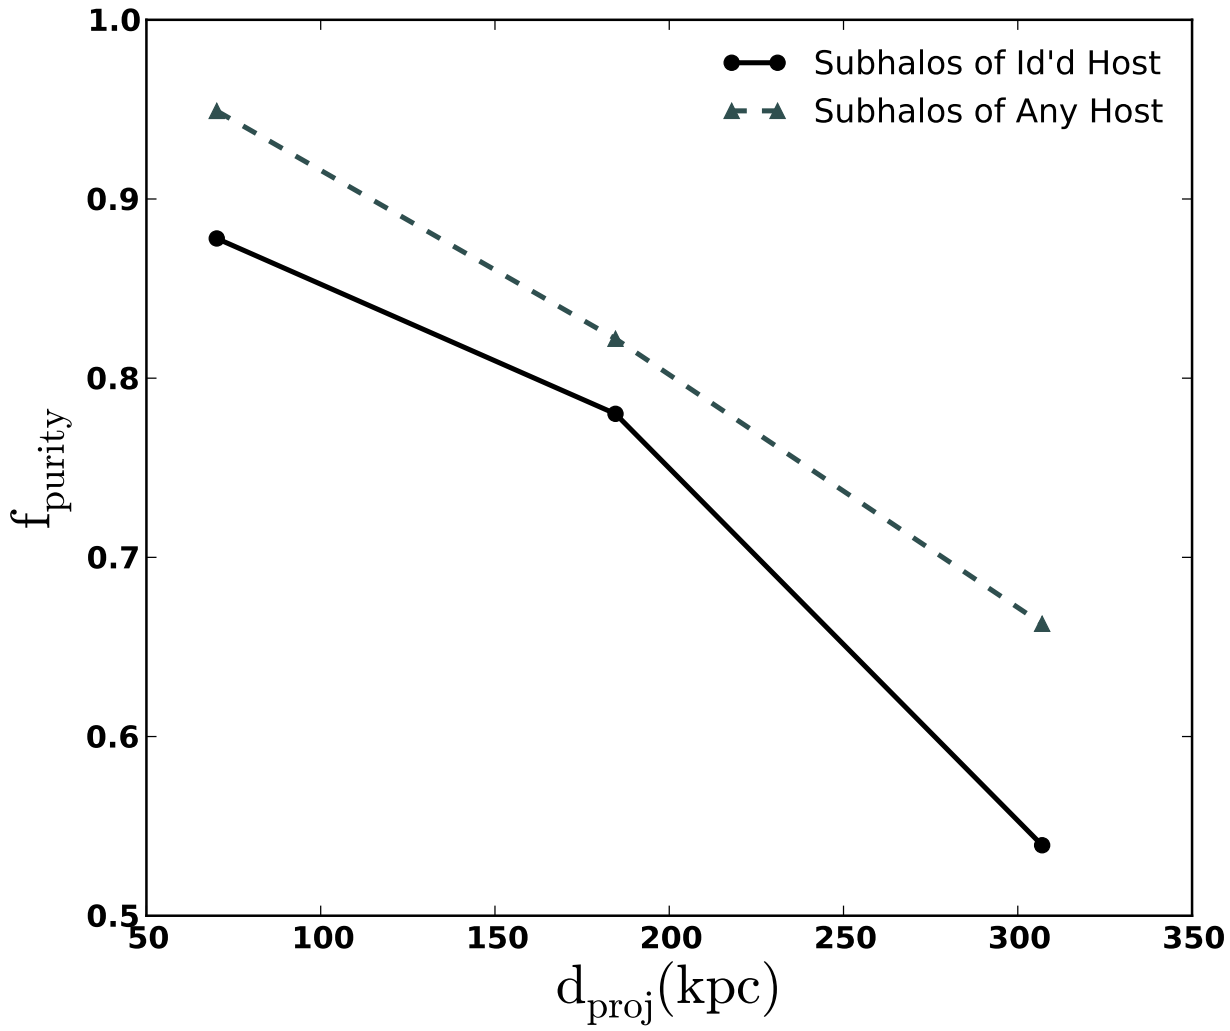

Figure 3.3: Fraction of the pair sample for which the secondary is a subhalo of the primary (solid line) and for which the secondary is a subhalo of any host (dashed line) as a function of projected distance between the primary and secondary in the MS-II. Due to an increased chance of interlopers at larger d_{proj}, the purity decreases with projected distance from the primary. The fraction of secondaries that are true subhaloes of the primary decreases from $\sim$88% in the innermost bin to just above 50% at 355 kpc, while the fraction of secondaries that are subhaloes of any host decreases from $\sim$96% in the innermost bin to 70% at 355 kpc.

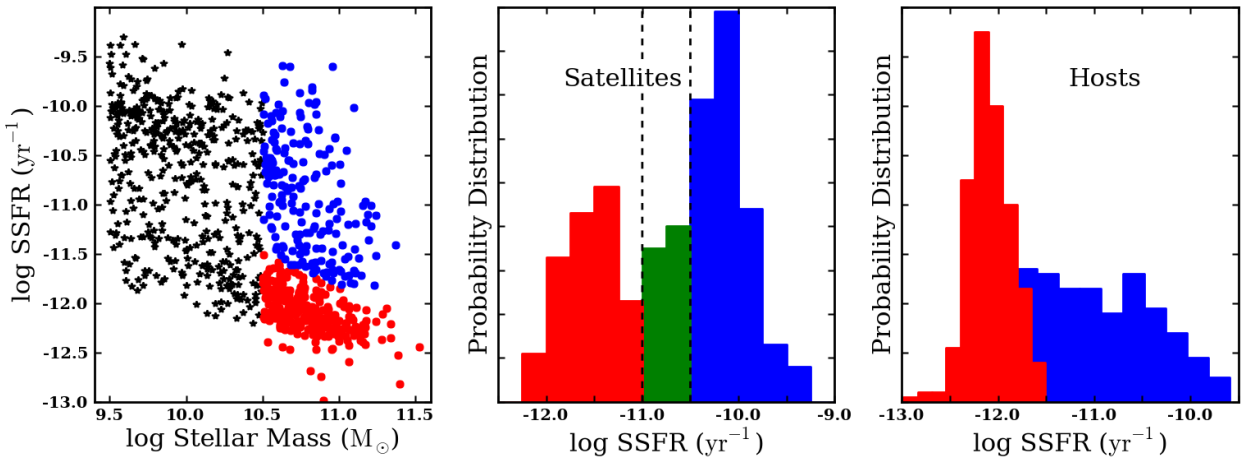

Figure 3.4: (*Left*): The distribution of specific star formation rate (SSFR) as a function of stellar mass for the galaxy sample employed in this study. Galaxies with $M_\star > 10^{10.5} M_\odot$ are categorized as hosts (red and blue points), with galaxies at lower mass categorized as satellites (black stars). Hosts are divided into two samples, where those coloured blue are star forming and those coloured red are passive. Each host has exactly one satellite. (*Centre*): The SSFR distribution for the satellite population. For the purposes of quantifying quenching effects, we subdivide the satellite population into three categories of star formation activity: satellites with SSFRs greater than $10^{-10.5}$ yr^{-1} are defined as vigorously star forming (blue histogram), satellites with SSFRs between $10^{-10.5}$ yr^{-1} and 10^{-11} yr^{-1} are defined as moderately star forming (green histogram), and satellites with SSFRs less than 10^{-11} yr^{-1} are deemed passive (red histogram). (*Right*): Distribution of SSFRs for host galaxies. Hosts are divided into passive (red histogram) and star forming (blue histogram) subsamples according to blue cloud/red sequence membership as detailed in Equation 1.

86% at 200 kpc and 70% at 300 kpc. When we consider trends in quenching with projected radius, we will take this radial dependence on purity into consideration.

3.3 Isolated L^* hosts in SDSS

3.3.1 Isolation Procedure Applied to Data

Having identified an optimal set of selection criteria, we turn to applying these criteria to our observational data. To identify a set of potential primary galaxies, we begin by selecting

all galaxies in the SDSS DR7 catalog with stellar mass greater than $10^{10.5}\mathrm{M}_{\odot}$. The isolation criteria described in Section 3.2.3 are then applied. In particular, we discard (1) all primaries with one (or more) neighbour of stellar mass $\mathrm{M}_{\star} > 10^{10.5}\mathrm{M}_{\odot}$ within a projected distance of 355 kpc and within a velocity difference of $|\Delta v| \equiv c|z_1 - z_2| < 1000\,\mathrm{km\,s^{-1}}$ along the line-of-sight, and (2) all primaries with two (or more) neighbours of stellar mass $10^{10.5}\mathrm{M}_{\odot}$ having $355\,\mathrm{kpc} < d_{\mathrm{proj}} < 1\,\mathrm{Mpc}$ and $|\Delta v| < 1000\,\mathrm{km\,s^{-1}}$. All galaxies that pass the isolation criteria are deemed "isolated host-mass primaries." We then compile a catalog of host/satellite pairs, our "pairs catalog," by selecting those isolated host mass primaries with exactly one neighbour in the stellar mass range $10^{9.5}\mathrm{M}_{\odot} < \mathrm{M}_{\star} < 10^{10.5}\mathrm{M}_{\odot}$ with $d_{\mathrm{proj}} < 355\,\mathrm{kpc}$ and $|\Delta v| < 500\,\mathrm{km\,s^{-1}}$. As shown in the previous section, the requirement of exactly one satellite allows us to better select hosts that have MW-sized haloes (i.e. $\mathrm{M}_{\star} \sim 10^{12}\mathrm{M}_{\odot}$). After the full isolation procedure, our pairs catalog consists of 483 host-satellite partners. Based on the results of Section 3.2.3, we expect that the hosts have a median virial mass of $1.9 \times 10^{12}\,\mathrm{M}_{\odot}$.

3.3.2 Observational Sample

The left panel of Figure 3.4 shows our hosts and satellites in the specific star formation rate-stellar mass plane. Black stars denote the satellite population, while our host population is coloured by star formation properties (see Eqn. 4.1): red and blue circles correspond to the passive and star forming hosts, respectively. Hosts with multiple bright satellites and satellites near two host-like objects were removed from the sample, such that there is a one-to-one correspondance between the hosts and satellites in the pairs sample.

The distribution of specific star formation rates for the satellite population is given in the central panel of Figure 3.4. We place vertical lines at 10^{-11} yr^{-1} and $10^{-10.5}$ yr^{-1} in this and all subsequent plots as the borders of our defined regimes of star formation activity for the satellite population: passive (red histogram), moderately star forming (green histogram),

and vigorously star forming (blue histogram). Finally, the SSFR distribution for the hosts is shown in the right panel of Figure 3.4. Hosts are divided into passive and star forming categories according to Equation 4.1.

3.3.3 Satellite-host velocity distributions

In Figure 3.5, we show the distributions of Δv for our full primary/secondary sample (dash-dotted grey) along with those for our star forming primary sample (blue curve) and a stellar mass-matched subsample of our passive primaries (red curve). These histograms represent "stacked" satellite velocity distributions for hosts. By bootstrapping our distributions, we find evidence (at approximately two standard deviations) that passive hosts reside in slightly more massive haloes *at fixed stellar mass*, with measured velocity dispersions $\sigma_{\rm passive} = 165.1 \pm 8.3$ km/s and $\sigma_{\rm starforming} = 145.4 \pm 9.4$ km/s. This suggests that the haloes of passive L^* galaxies are $\sim 45\%$ more massive at fixed stellar mass than those of star forming galaxies, assuming a mass-scaling proportional to σ^3, where the exponent is derived from the assumption of a linear relation between halo maximum circular velocity and the spread of stacked pairwise host/satellite velocity distributions at fixed halo mass. [1] Also plotted is the distribution of velocity differences for the objects in MS-II chosen by mock observations that mirror our selection criteria (black histogram). Remarkably, the mock catalog falls almost perfectly between the red and blue histograms, providing an independent dynamical verification that we are indeed selecting the haloes we have set out to study using abundance matching.

The evidence we see for a halo mass trend with SSFR at fixed stellar mass is qualitatively consistent with the results of Mandelbaum et al. (2006), who find that early-type galaxies

[1] The qualitative result that passive hosts reside in kinematically warmer and thus more massive haloes is robust to the exact nature of the scaling between stacked satellite velocities and halo mass. We have used simulated observations described in Section 3.2.2 to compare the observed velocity dispersions of satellite subhalos within hosts stacked in virial mass bins and found that the naive scaling adopted here does hold for the halo mass range of interest.

with stellar masses of $\sim 10^{11}$ $M_\odot$ tend to live in haloes that are ~ 3 times more massive than their star forming counterparts at the same stellar mass. The Mandelbaum et al. (2006) work samples all environments and thus the mass trend may be a result of the colour-density relation where red galaxies prefer overdense regions – i.e. more massive haloes (Hogg et al., 2004; **?**; **?**; **?**). Our sample focuses on lower-mass systems, and specifically avoids group and cluster haloes.

3.3.4 Control Sample

As a reference sample by which to compare our satellite population, we define a set of galaxies in the field with stellar masses similar to that of our satellite galaxies. As with our pairs sample, we tune our selection criteria by maximizing purity, where purity here is defined to be the fraction of objects that are actually hosts (i.e. centrals) and not satellites within the MS-II. Since we have no shortage of objects (SDSS contains 12447 such objects in our stellar mass range), conserving sample size is less important, and we can tune our selection parameters to create a maximally pure control sample. The final selection parameters for the control sample are $n_{\rm iso,control} = 0$, $d_{\rm iso,control} = 2.9\,{\rm Mpc}$ and $\Delta v_{\rm iso,control} < 400\,{\rm km\,s^{-1}}$. With these values for the selection and isolation criteria, we obtain $f_{\rm p,control} = 0.97$. The impurity in the sample is approximately equal to the Poisson noise from the number of objects, $1 - f_{\rm purity,control} = 1/\sqrt{N}$.

Implementing these selection criteria, we construct a catalog of isolated galaxies in the mass range $10^{9.5} < M_\star < 10^{10.5}$ to use as our control sample. Our isolation criteria are similar to, but more stringent than, those used in selecting the pairs catalog. We reject any galaxy with a neighbour of stellar mass $\log(M_\star/M_\odot) > 9.5$ with $d_{\rm proj} < 2.9$ Mpc and with $\Delta v_{\rm iso} < 400\,{\rm km\,s^{-1}}$. This catalog, which we will call the "field dwarf catalog" contains 581 galaxies, with a mean stellar mass of $10^{9.95}$ $M_\odot$ and a mean redshift of 0.024.

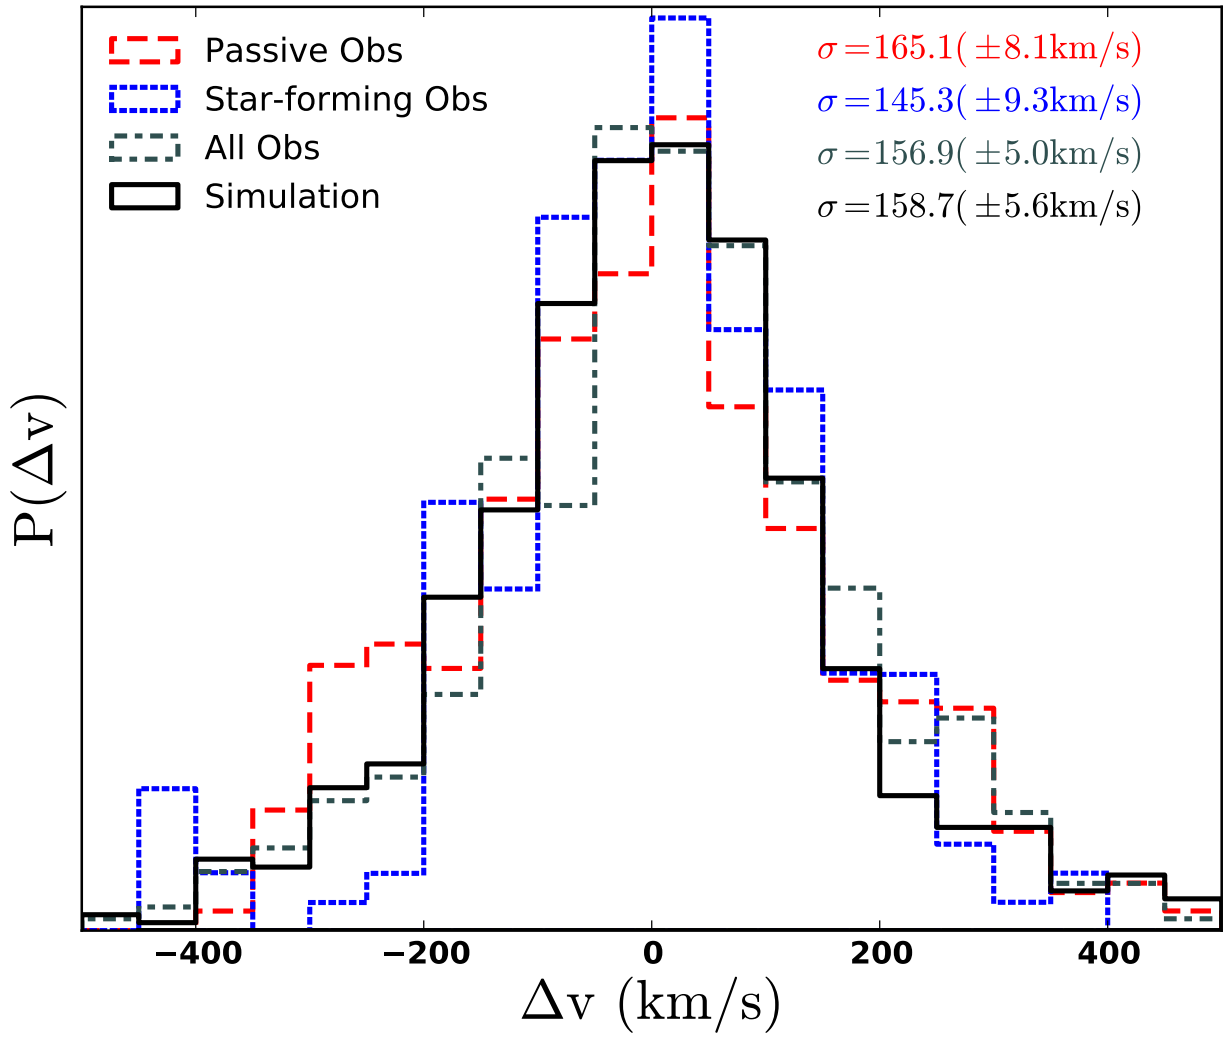

Figure 3.5: Probability distribution of velocity offsets Δv for all observed primary-secondary pairs in our sample (dash-dotted grey) along with pairs chosen from stellar mass-matched samples of star forming (blue) and passive (red) hosts. The root-mean-square of each distribution is given in the figure. The passive host subsample appears kinematically warmer than the star forming host subsample, suggesting that these passive galaxies reside in more massive dark matter haloes at fixed stellar mass. Also shown are the host-satellite velocity distributions for our mock catalogs with the same selection criteria applied to MS-II (black line), which matches well the observed sample. This remarkably consistency shows that our selection procedure is identifying host haloes as expected, with a virial mass distribution consistent with that shown by the red histogram in Figure 1.

3.3.5 Parameter-Matching Procedure

It will frequently be useful in our analysis to explore the variation in a given galaxy property (e.g. specific star formation rate) across two related samples of galaxies (e.g. isolated galaxies versus satellites). When performing such a comparison, we would like to control for correlations between the varied parameter and other properties of the chosen samples — for example, the observed decrease in SSFR with increasing stellar mass (Brinchmann et al., 2004; Noeske et al., 2007; Elbaz et al., 2007). To do so, we match the stellar mass distribution of a subsample of the field dwarf catalog to the subsample of satellite galaxies under study; this ensures that the two samples have statistically identical distributions of stellar mass. We divide both samples into bins (typically 10, although we allow this to vary based on sample size) in the relevant parameter, then randomly select galaxies with replacement from the sample with the larger number in that bin, until both subsamples have an equal number in the bin. This procedure is repeated 100 times, and the results given below are mean values.

3.4 Satellite Quenching as a Function of Central Galaxy Properties

The primary goal of this paper is to explore the degree to which satellite galaxies have suppressed or quenched star formation relative to similar galaxies in the field. In order to numerically interpret our results, we introduce a conversion fraction ($f^{\rm X}_{\rm convert}$), designed to indicate the fraction of galaxies that are converted from star forming to suppressed/quenched *because* they are satellites.

Mathematically, the conversion fraction is defined as follows. Let X = log(SSFR) indicate a variable value for the specific star formation rate. Define the unquenched fraction (with

$\log \mathrm{SSFR} \geq X$) of satellites and control galaxies to be u_{sat} and u_{control}, respectively. The associated quenched fractions with $\log \mathrm{SSFR} < \mathrm{X}$ are $q_{\{\mathrm{sat,control}\}} = 1 - u_{\{\mathrm{sat,control}\}}$. The conversion fraction, $f^{\mathrm{X}}_{\mathrm{convert}}$, is then given by

$$f^{\mathrm{X}}_{\mathrm{convert}} = \left. \frac{q_{\mathrm{sat}} - q_{\mathrm{control}}}{u_{\mathrm{control}}} \right|_{\mathrm{SSFR}=10^{\mathrm{X}}}. \tag{3.2}$$

For example, if 100% of the galaxies in the control sample are star forming (at the given SSFR threshold) but only 80% of the satellite galaxies are star forming (at the same threshold), then $f_{\mathrm{convert}} = 0.2$ and we may conceivably argue that $\sim 20\%$ of star forming galaxies were converted to quenched galaxies after becoming satellites. Likewise, for star forming fractions of 60% and 40% for the control and satellite samples (respectively), the conversion fraction would be $f_{\mathrm{convert}} = 0.33$.

This definition of conversion fraction is similar to the similar quenching efficiency specified by e.g. van den Bosch et al. (2008a); Peng et al. (2012), although $f^{\mathrm{X}}_{\mathrm{convert}}$ is generalized to be a function of SSFR. The two conversion fractions we will primarily consider are $\mathrm{X} = -11$ and $\mathrm{X} = -10.5$. The conversion fraction evaluated at the passive SSFR threshold for satellites, $f^{-11}_{\mathrm{convert}}$, can be thought of as the fraction of satellites that were star forming prior to infall and have subsequently been converted to passive satellite systems, making it an effective method of quantifying quenching. The conversion fraction evaluated at the border between vigorous and moderate star formation, $f^{-10.5}_{\mathrm{convert}}$, is representative of the fraction of vigorous star forming galaxies that upon infall had their star formation either reduced to a moderate level or halted entirely. In general, we will refer to a satellite's SSFR being lowered below $10^{-11}\, yr^{-1}$ as "quenching," and the general case of a satellite's SSFR being lowered as "suppression" (e.g. we will often discuss $f^{-10.5}_{\mathrm{convert}}$ as suppression of vigorous star forming satellites). We will make use of these parameters throughout the remainder of this work.

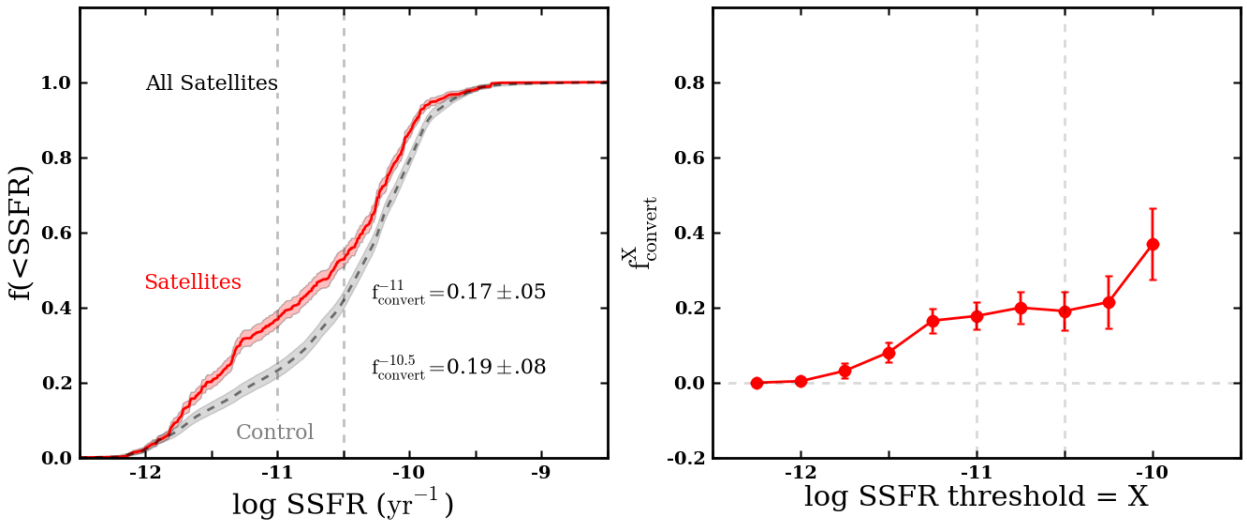

Figure 3.6: (*Left*): Cumulative distributions of specific star formation rate for the observed satellite sample (red dashed line), the purity-corrected satellite sample (red solid line), and the control sample (black dashed line). After a modest statistical correction for interloper contamination, we find that L^* hosts quench $\sim$20% of their infalling star forming satellites. Black and red shaded regions correspond to 1σ binomial errors for the control and purity-corrected satellite samples, respectively. (*Right*): Conversion fraction as a function of specific star formation threshold X, which provides an estimate for the fraction of satellites that have had their star formation suppressed below a threshold value X as a result of becoming satellites. This quantity is defined in equation 4.3. The conversion fraction, $f_{convert}$, is quite flat over a wide range in satellite star formation rates. Vertical dashed lines correspond to our definitions of "quenched" and "vigorously star forming" satellites.

3.4.1 All Satellites

Figure 3.6 shows the cumulative distributions of SSFRs for all satellite galaxies in our sample. The dashed red line is the observed distribution SSFR for the satellite population, while the solid red line is the corrected satellite distribution, where the observed distribution is adjusted to take into account the presence of interlopers. This is done by subtracting the control distribution scaled by the probability that a randomly selected galaxy is an interloper, (1-$f_{\rm purity}$), from the satellite distribution. The adjusted distribution is then re-normalized.[2] This correction for purity is made in each plot of satellite SSFR that follows, and the uncorrected result is omitted. The black line shows the distribution for the mass-matched field control sample — i.e. for a sample of isolated (central) galaxies with the same stellar mass distribution as the satellites. For both the corrected satellite and control samples, binomial errors are shown by the corresponding shaded region according to the formula

$$\sigma^2 = \frac{(q_{\rm sample})(u_{\rm sample})}{N_{\rm sample}} \tag{3.3}$$

where $q_{\rm sample}$ is the sample's quenched fraction, $u_{\rm sample}$ is the sample's unquenched fraction, and $N_{\rm sample}$ is the total number of objects in the sample.

Globally, $f_{\rm convert}^{-11} = 0.18$ for our L^* host sample, with a similar value for $f_{\rm convert}^{-10.5}$. This implies that approximately 20% of satellites are quenched as a result of falling into their host's virial radius. The right panel of Figure 3.6 generalizes $f_{\rm convert}$ to an arbitrary SSFR threshold. Bright satellite galaxies residing in haloes similar to that of the Milky Way show an overall flat increase in conversion fraction over a dex in satellite SSFR ranges. We use this style

[2]This assumes that the interlopers follow the same specific star formation distribution as the control sample.

of presentation in the next three plots; the left panels show the cumulative distributions of SSFR for the satellite and control samples and the rightmost panel shows $f_{\rm convert}$ plotted as a function of satellite threshold SSFR (as defined by Equation 4.3).

3.4.2 Dependence on Host Star Formation Rate

In order to investigate the correlation of quenching efficiency with host star formation activity, we define a mass-dependent passive threshold according to SSFR in our host sample, given as Equation 1 above. Hosts that fall below the relation are considered passive ($N = 267$, $\langle {\rm M}_\star \rangle = 10^{10.81}$ ${\rm M}_\odot$), while those above it are considered star forming ($N = 190$, $\langle {\rm M}_\star \rangle = 10^{10.78}$ ${\rm M}_\odot$). Our goal of determining the correlation between star formation activity in satellites and hosts is complicated by the fact that passive hosts tend to have slightly higher stellar masses. To avoid biasing our results, we mass-match the host samples — i.e. we ensure that the distribution of stellar masses is the same between our star forming and passive host samples.

Comparisons of the cumulative distributions of satellite SSFRs for the two mass-matched host samples are shown in Figure 3.7. The satellite population around quenched hosts (middle panel) is itself quenched relative to the mass-matched field sample, with $f_{\rm convert}^{-11} = 0.28$ and $f_{\rm convert}^{-10.5} = 0.35$. This result is qualitatively consistent with the observed correlation between host and satellite properties otherwise known as "galactic conformity" (e.g. Weinmann et al., 2006; Kauffmann et al., 2013; Robotham et al., 2013).

Remarkably, the satellite population around star forming hosts (left panel) is markedly different, with a SSFR distribution that is *indistinguishable* from the mass-matched sample of field galaxies. The right panel of Figure 3.7 emphasizes this stark difference — at all choices of SSFR threshold, the conversion fraction for satellites of star forming hosts (blue line) is consistent with zero. Meanwhile, the conversion fraction for passive hosts increases

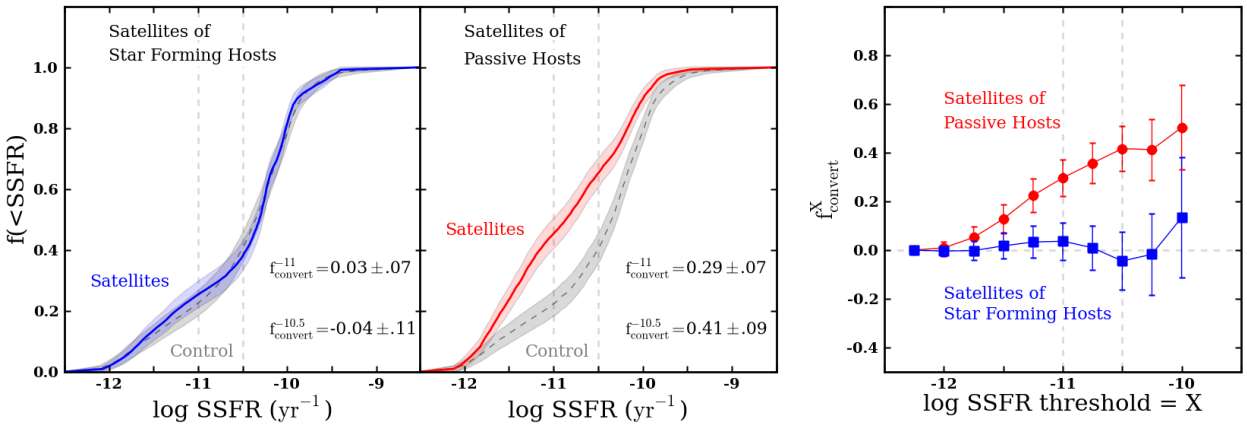

Figure 3.7: (*Left*): Cumulative distributions of specific star formation rate for satellites surrounding star forming hosts (blue line) and passive hosts (red line), with relevant mass-matched control samples (black dashed lines). Passive hosts quench $\sim 30\%$ of their infalling star-forming satellites, while star forming hosts have satellite populations that very closely match the field. Shaded regions are binomial errors. (*Right*): Conversion fraction vs. threshold SSFR (see equation 4.3) for satellites of star forming and passive hosts. At all values of X, the satellites of star forming hosts are indistinguishable from the field population at the same stellar mass. In stark contrast, satellites of passive hosts show an increasing difference from the field at higher values of SSFR.

monotonically as the SSFR threshold decreases, suggesting that passive hosts are more likely to suppress vigorous star formation in their satellites than strongly quench it to passivity.

Figure 3.7 provides a compelling demonstration that satellite quenching and central quenching are strongly related: among bright $\sim 0.1L^*$ satellites, only those that inhabit quenched hosts are themselves quenched (relative to the field population). This is somewhat distinct from the view usually discussed as galactic conformity, in which the colour or star formation rate of a satellite is correlated with the colour or star formation rate of its host. Rather, it appears that bright satellites of isolated, star forming L^* galaxies are essentially unaffected by their host, whereas approximately 25-35% of satellites of quenched hosts are quenched as a result of being a satellite. The implications of this finding will be discussed in Section 4.4.

3.4.3 Dependence on Satellite Morphology

The previous subsection demonstrated that the quenching of $\sim 0.1L^*$ satellites only occurs around galaxy-size hosts that are themselves quenched. A natural extension of this result is to ask whether the morphologies of the satellites are affected along with their star formation rates, where in this work we will use Sérsic index as a proxy for morphology. In other words, is the quenching mechanism connected to a morphological change?

Figure 3.8 explores this question by comparing the Sérsic index distributions for passive and star forming satellites of quenched hosts to mass-matched control samples of the same star formation category (passive or star forming). Neither satellite sample shows substantial deviation from the field. One might expect to observe no morphological difference between the star forming satellite sample. However, it is less obvious why the passive field dwarf sample resembles the passive satellite sample when the field galaxies are manifestly not being quenched by an environmental process, and a substantial fraction of the passive satellites of passive galaxies are.

It is clear from Figure 3.8 that whatever is causing satellite galaxies to be quenched at a higher rate than field galaxies, this process results in morphological properties that are indistinguishable from field galaxies that are presumably quenched via some distinct secular process. One potential implication is that the morphological differences that divide quenched and star forming galaxies are a *result* of a galaxy being quenched. This is distinct from the idea of "morphological quenching" that suggests that morphological changes themselves give rise to suppressed star formation (Martig et al., 2009). We return to a discussion of these issues in Section 5.

3.4.4 Dependence on Projected Separation

We expect that the efficiency of some potential quenching processes may vary with the separation of the host and satellite, possibly reflecting radial variation in the circumgalactic medium (CGM) surrounding the host, the strength of the local tidal field, or the time since first infall for the satellite. In order to investigate trends with host/satellite distance, we place our satellite sample into three linearly spaced bins in projected separation. Conversion fractions are shown as a function of projected distance from the host, with samples divided according to SSFR in Figure 3.9. Subsamples in each bin of projected separation are matched on host mass and compared to control samples matched to the corresponding satellite mass distributions. When examining the dependence of $f^{11}_{\rm convert}$ on satellite/host separation, we find a gradient in quenching efficiency. Around passive hosts, quenching persists only out to the intermediate projected radius bin, which has its outer edge at 236 kpc, less than the characteristic virial radius of our primary sample in MS-II ($\sim$400 kpc). Around star forming hosts, however, we find results consistent with no quenching at all radii.

The right panel of Figure 3.9 shows $f^{10.5}_{\rm convert}$, which is sensitive to the *damping* of vigorous star formation, if not full quenching. The trends are highly similar in shape; the largest difference is the central value in the innermost bin of the star forming host sample. Again, we see no evidence of quenching around star forming hosts, and strong quenching around passive hosts.

3.5 Discussion

The main result from our study of central and satellite galaxies in the SDSS is that bright satellites of $\sim L^*$ galaxies are only quenched beyond what is seen in the field around passive hosts, indicating a strong connection between the quenching of satellites and centrals

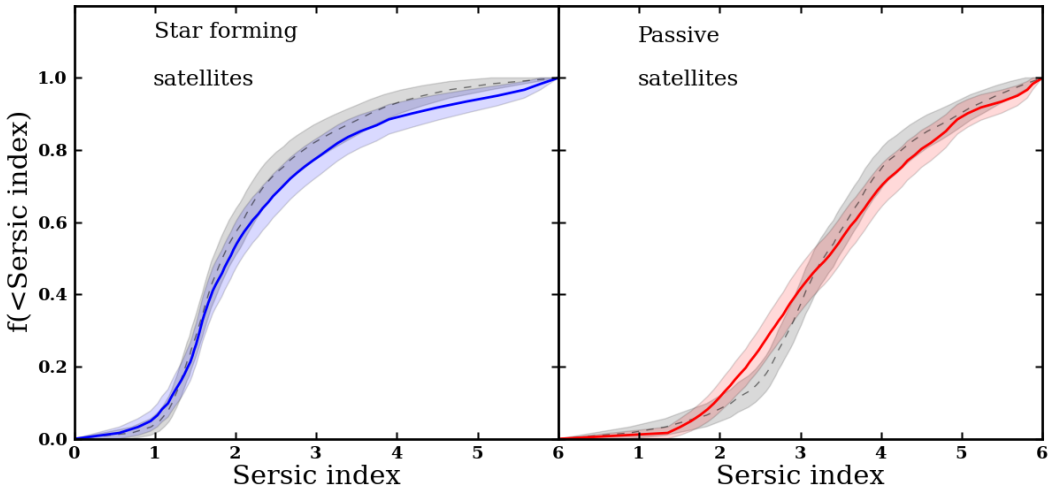

Figure 3.8: Cumulative distributions of Sérsic indexes for star forming (left panel) and passive (right panel) satellites around passive hosts and mass matched control samples (black dashed lines). Satellites show no significant difference in morphology relative to field galaxies at fixed stellar mass and star formation activity. Shaded regions are binomial errors.

in Milky Way-like haloes. This host-specific dichotomy in satellite quenching is likely (or at least partially) responsible for the observed correlation between the properties of central and satellite galaxies in the local Universe, a phenomenon commonly dubbed "galactic conformity." As first shown by Weinmann et al. (2006), studying more massive haloes than those isolated in our work, groups dominated by a "late-type" central host a higher fraction of "late-type" satellites relative to groups populated by an "early-type" central.[3] More recently, results from Wang & White (2012) highlight a similar conformity in the rest-frame $g-r$ colour of satellite galaxies in the SDSS, examining systems across a broad range of host stellar mass. Lastly, studying low-mass galaxy groups identified by the Galaxy And Mass Assembly project (GAMA, Driver et al., 2011), Robotham et al. (2013) find that conformity is also evident in the morphological properties of satellites relative to their hosts, such that bulge-dominated central galaxies tend to host bulge-dominated satellites in contrast to disk-dominated centrals. Previous works have demonstrated that semi-ananlytic models do show evidence of galactic conformity, however the nature of the conformity does not match

[3]Note that the terms late- and early-type as used by Weinmann et al. (2006) correspond to the colour and/or SSFR of the galaxy and do not refer to a morphological classification.

observations. For example, Wang & White (2012) argue that the semi-analytic model of Guo et al. (2011) overpredicts the number of red satellites in g-r (colour). Similarly, we have analyzed the Guo et al. (2011) model results independently and found that they predict a conversion fraction of 0.6 around passive hosts and 0.3 around star-forming hosts, where we apply our same isolation criteria to the SAM. This is in stark disagreement with the dichotomy we see in the data.

In general, our results are in agreement with each of the aforementioned observational analyses. However, unlike these previous studies that compare the composition of $\gtrsim 10^{12}$ $M_{\odot}$ haloes separated according to the properties of the central galaxy (i.e. excluding the field population), our analysis extends this work to focus on the comparison of satellite samples (again divided according to the properties of the central) to field galaxies of like stellar mass. This is a critical step towards understanding the physical driver of satellite quenching in the local Universe, as it directly compares satellite galaxies to their parent population (i.e. field systems of like mass).

One possible explanation for the observed dichotomy in satellite quenching evident in our sample could be a dependence of quenching efficiency on halo mass. That is, passive hosts could live in more massive dark matter haloes relative to their star forming counterparts at fixed stellar mass, with quenching operating more efficiently in more massive haloes. As detailed in Section 3.3, we tailored our host selection criteria to isolate a sample that spans a relatively small spread in halo mass (FWHW ~ 0.5 dex) and peaks at a few $\times 10^{12}$ $M_{\odot}$. Studies of gas accretion onto dark matter haloes indicate that there is a transition in the dominant accretion mode at roughly this halo mass, with infalling gas shock-heated at the virial radius in haloes above a threshold mass of $\sim$ a few $\times 10^{12}$ $M_{\odot}$, while cold gas reaches a smaller radius, possibly falling all the way to the galaxy in less-massive haloes (e.g. Binney, 1977; Rees & Ostriker, 1977; Birnboim & Dekel, 2003; Kereš et al., 2005, 2009; Stewart et al., 2011a). Such a variation in accretion mode could manifest itself as a correlation between

the star forming properties of central (and also satellite) galaxies and the mass of their host dark matter haloes, such that passive central galaxies reside in more massive dark matter haloes relative to star forming central galaxies of comparable stellar mass.

Recall that in our analysis, we tune our isolation criteria to identify haloes comparable in mass to that of the Milky Way, but select our sample of star forming and passive hosts for comparison by matching them in stellar mass (i.e. potentially allowing a weak correlation between halo mass and host properties). In Figure 3.5, we investigate this possibility by plotting separately the stacked velocity distribution of stellar mass-matched subsamples of satellites around our passive and star forming host galaxies. While we find that the red hosts are biased towards slightly more massive haloes (approximately 45% more massive), the difference is only significant at the $\sim 1-2\sigma$ level. Furthermore, previous work examining galactic conformity by Weinmann et al. (2006) found that the correlation between the colour of central and satellite galaxies exists at fixed halo mass, looking at groups in the Yang et al. (2007) catalog, where the halo masses computed in the group catalog are based on the total luminosity or stellar mass of all member galaxies above a given luminosity threshold. This suggests that a potential dependence of quenching efficiency on halo mass is unable to explain our results in concert with those of Weinmann et al. (2006); while we find evidence for conformity at fixed *stellar* mass (and similar, though possibly not identical halo mass), Weinmann et al. (2006) claims conformity at fixed *halo* mass, although that result is consistent with ours.

The dichotomy in satellite quenching also could be driven by differences in the circumgalactic medium (CGM) of passive and star forming hosts. The CGM of the host halo could potentially affect the star formation of a satellite galaxy via two physical processes: (*i*) "ram pressure stripping," where the cold interstellar gas from which stars could be actively formed is removed from the satellite, rapidly truncating its star formation (e.g. Gunn & Gott, 1972; Bekki, 2009) and (*ii*) "strangulation," where the satellite retains its cold gas, allowing active

star formation, but loses its reservoir of warm gas that would otherwise replenish the cold gas as it is exhausted (e.g. Larson, Tinsley & Caldwell, 1980; Kawata & Mulchaey, 2008). In either scenario, our results could be explained by the presence of a hot gaseous corona preferentially surrounding passive hosts, such that it prevents the infall of cold gas onto the central galaxy, thereby quenching it, while similarly halting the star formation of associated satellites. A CGM dichotomy could also arise in association with the the transition between cold and hot mode accretion. While realistically we expect there to be some scatter in the halo mass scale where the transition occurs (Kereš et al., 2009), there might naturally be a correlation between star formation in the central galaxy and ongoing cold-mode accretion. In hot mode haloes, star formation may be naturally suppressed by a lack of fuel, and the associated build-up of a hot corona could then provide a CGM-related suppression mechanism.

In our discussion of the hot CGM interpretation, we must adress the results of Tumlinson et al. (2011): quasar line-of-sight probes of the CGM of nearby systems show an enhancement in ionized oxygen around star forming galaxies relative to their quenched counterparts, indicating that star forming galaxies host more warm gas in their haloes. The lack of O VI around passive hosts could be explained as a temperature effect, with passive galaxies surrounded by a typically hotter CGM. Observations of X-ray coronae around local ellipticals support this conclusion (O'Sullivan, Forbes & Ponman, 2001; Sun et al., 2007), including recent work to study the hot gas surrounding lower-mass, isolated ellipticals (Mulchaey & Jeltema, 2010; Humphrey et al., 2011, 2012).

In Section 3.4.3, we also examined the morphology of the satellites in our sample, finding consistent Sérsic indices at fixed stellar mass between the quenched satellites and quenched galaxies in the field. The same is true for star forming satellites. This result could be compared to similar results in van den Bosch et al. (2008a), who found consistent concentrations between satellites and centrals at fixed stellar mass and colour, and Bamford et al. (2009)

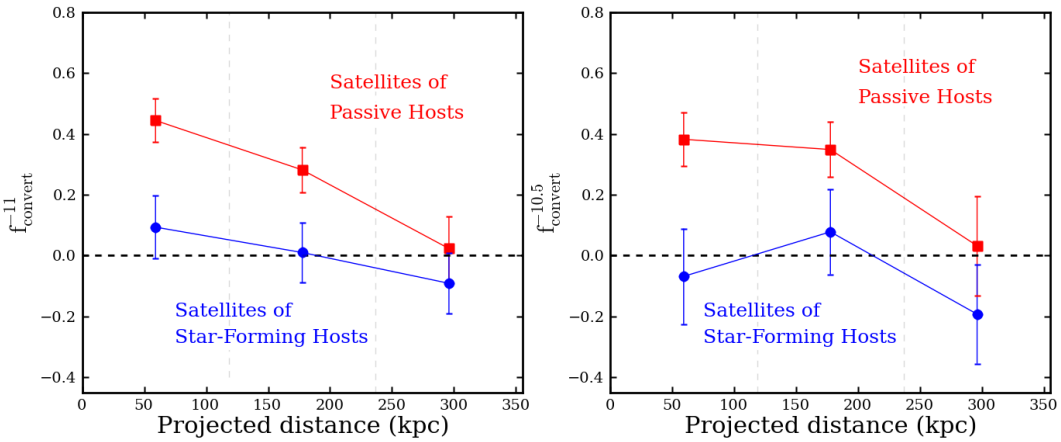

Figure 3.9: Conversion fractions of satellites around passive (red line) and star forming hosts (blue line) as a function of projected distance from the central galaxy. The dashed vertical grey lines denote the edges of the radial bins. (*Left*): $f_{\rm convert}^{-11}$, representative of the fraction of star-forming satellites that have their star formation quenched upon infall. Around star-forming hosts, quenching is only seen at small projected radii, while around passive hosts, quenching is observed out to our radial limit. (*Right*): $f_{\rm convert}^{-10.5}$, representative of the fraction of vigorously star forming satellites that have their star formation suppressed upon infall. For star forming hosts, the quenching observed in the left panel is no longer apparent; the frequency of vigorous star forming satellites suppressed around star forming hosts is consistent with the field at all projected radii.

who found no trend in the early-type fraction of galaxies with environment in the sparse regime (and a weak trend, as compared to the trend in red fraction, over all environments). The lack of any difference between satellite and field morphology could provide a clue to the origin of the well-established link between structure and star formation in galaxies (e.g. Kauffmann et al., 2003; Blanton et al., 2005a; Wuyts et al., 2011; Bell et al., 2012; Cheung et al., 2012), with quenched galaxies typically having more bulge-dominated morphologies than star forming ones.

One idea to explain the correlation between morphology and star formation activity posits that the morphological transformation itself is the driving agent. Specifically, in the "morphological quenching" scenario (Martig et al., 2009), the presence of a bulge component stabilizes a galaxy against fragmentation to bound clumps, thereby halting star formation. In such a picture, a bulge-dominated galaxy may remain quenched even if it accretes gas (i.e. regardless of the assumed halo mass-dependent accretion mode). Our results may be at odds with such a scenario: satellites that are apparently quenched via environmental mechanisms never experienced by quenched galaxies in the field nevertheless have exactly the same distribution of morphological indexes. If morphology were really the driver of quenching, it seems unlikely that this would be the case, unless somehow the passive-halo environment causes exactly the same type of bulge formation as do processes in the field. One might argue that tidal forces could incite secular instabilities, but this would also produce quenching for satellites around non-star forming hosts. No such signal is seen.

Often, galaxy mergers are relied upon as a means for driving the correlation between star formation activity and morphology (Springel, Di Matteo & Hernquist, 2005; Hopkins et al., 2006). However, our results suggest that mergers, which occur preferentially in group-like environments or the field (McIntosh et al., 2008; Fakhouri & Ma, 2008; Darg et al., 2010; Lin et al., 2010), are unlikely to be responsible for the properties of quenched satellites in passive galaxy haloes. Instead, whatever mechanism(s) is responsible for shaping the

morphology of our quenched satellite population needs to operate as effectively in the field (i.e. lower-mass haloes) as it does for subhaloes within Milky Way-like systems. While many models suggest that bulge formation is merger-driven (e.g. Toomre, 1977; Springel, Di Matteo & Hernquist, 2005; Cox et al., 2006), our results would favor a different origin for bulge-dominated morphology (e.g. Kormendy & Kennicutt, 2004; Jogee, Scoville & Kenney, 2005; MacArthur, Courteau & Holtzman, 2003), since we find no variation in Sérsic index with environment. Collectively, our results may support a picture where the cessation of star formation (e.g. via gas removal from ram pressure/strangulation or blow-out) triggers an associated morphological transformation.

Finally, we also examined the dependence of quenching efficiency on projected host/satellite separation, finding a radial gradient in the conversion fraction, $f^{\rm X}_{\rm convert}$, around passive galaxies. Recent work by Kauffmann et al. (2013) reports a correlation in the specific SFR of satellites extending out to comparable spatial scales around massive central galaxies; for hosts with stellar masses of $>10^{10}$ $M_{\odot}$, Kauffmann et al. (2013) argue that this correlation is confined to spatial scales less than 1-2 Mpc. Analyzing satellite galaxies around massive ($\sim 10^{11}$ $M_{\odot}$) central galaxies selected from the GAMA project, Prescott et al. (2011) also find a significant radial dependence for the satellite red fraction out to separations of roughly 500 Mpc.

The observed radial trend measured in our sample could reflect a correlation between infall time and the cessation of star formation, such that more-recently accreted satellites are more likely to be star forming (see also Hearin & Watson, 2013). However, the connection between infall time and projected distance is particularly poor inside the virial radius (Oman, Hudson & Behroozi, 2013). Instead, if the CGM drives the quenching of star formation via ram-pressure stripping or strangulation, the radial dependence of quenching efficiency could be the result of radial variation in the properties of the CGM — a denser CGM would yield more effective gas stripping, as the ram pressure force scales with the density of the stripping

medium (Gunn & Gott, 1972).

Altogether, our study has uncovered two remarkable results with regard to the properties of relatively massive satellite galaxies around $\sim L^*$ hosts in the local Universe: (i) falling into a $\sim 10^{12.3}$M$_\odot$ halo hosted by a star forming central galaxy has a negligible impact on the star formation rate of $\sim$ 0.1 L^* galaxies, such that massive satellites of star forming hosts are indistinguishable from a field population of like stellar mass, and (ii) $\sim$30% of star forming galaxies accreted into a similar halo hosting a passive central galaxy are quenched, but with no corresponding change in morphology. When combined, these results support a picture where the mechanism(s) by which bright satellites of $\sim L^*$ galaxies are quenched only occur in haloes where the central galaxy is itself quenched. This could either imply that the same quenching mechanism is operating on host and satellite, or it could point to a secondary effect, perhaps associated with differences in the CGM of quenched and star forming galaxies. Whatever the quenching mechanism, it impacts the morphologies of satellites in a manner that mirrors distinct quenching mechanisms in the field.

It is worth emphasizing one particularly interesting implication of our findings: even for passive hosts, the quenching mechanism(s) must operate(s) inefficiently (possibly with a long timescale, De Lucia et al., 2012; Wetzel et al., 2013a; Trinh et al., 2013), such that only $\sim$30% of satellites that have fallen in from the field are quenched. This relatively low quenching efficiency is mirrored in studies of lower-mass satellites (Geha et al., 2012). It is not clear why quenching occurs around passive hosts (and not star forming hosts) but the fact that it is apparently fairly inefficient may be a clue to its origin. Moreover, while our results focus on fairly massive satellites (M$_\star > 10^{9.5}$ M$_\odot$), we know from studies of dwarf galaxies in the Local Group that quenching must become very efficient for smaller satellites. The vast majority of low-mass satellites (M$_\star < 10^8$ M$_\odot$) within the virial radii of the Milky Way and M31 are gas-poor and quenched, in contrast to the field population (McConnachie, 2012). Our ongoing work involves investigating this potential mass dependence along with the exploration of

simple quenching models that can replicate the relatively inefficient quenching of massive satellites (Wheeler et al., in prep).

3.6 Conclusions

In this work, we have investigated quenching and suppression of star formation of satellite galaxies in the mass range $10^{9.5}M_\odot < M_\star < 10^{10.5}M_\odot$ orbiting hosts that are selected to have $M_\star > 10^{10.5}$ $M_\odot$. Our hosts are chosen to be central galaxies of $\sim 10^{12.3}M_\odot$ dark matter haloes, selected via careful isolation criteria to the group and cluster environment. In order to evaluate the degree of quenching among satellite systems, we make comparisons to stellar-mass matched control samples of isolated, low-mass field galaxies.

Our key result (Figure 7) is that there is a dichotomy in satellite quenching based on host star formation strength: passive hosts quench roughly 30% of their infalling satellites, while satellites of star forming hosts *of the same stellar mass* show a SSFR distribution consistent with field galaxies. Around passive galaxies, quenching is only present in the inner ~ 200 kpc, significantly less than the virial radius of these hosts, and increases in efficiency with decreasing host/satellite separation. We also show that passive satellites and passive field galaxies show no morphological distinction at fixed stellar mass.

This paper presents a clear dichotomy in massive satellite quenching; however, we anticipate that other trends will be present. For example, smaller ($M_\star < 10^8M_\odot$) satellites in the Local Group are universally quenched, suggesting that quenching might be related to satellite stellar mass. We will explore trends with host and satellite stellar mass in a companion paper.

Chapter 4

Mass Trends in Satellite Quenching

Using the Sloan Digital Sky Survey, we examine the quenching of satellite galaxies around isolated Milky Way-like hosts in the local Universe. We find that the efficiency of satellite quenching around isolated galaxies is low and roughly constant over two orders of magnitude in satellite stellar mass ($M_\star = 10^{8.5} - 10^{10.5}\,M_\odot$), with only $\sim$ 20% of systems quenched as a result of environmental processes. While largely independent of satellite stellar mass, satellite quenching does exhibit clear dependence on the properties of the host. We show that satellites of passive hosts are substantially more likely to be quenched than those of star-forming hosts, and we present evidence that more massive halos quench their satellites more efficiently. These results extend trends seen previously in more massive host halos and for higher satellite masses. Taken together, it appears that galaxies with stellar masses larger than about 10^8 $M_\odot$ are uniformly resistant to environmental quenching, with the relative harshness of the host environment likely serving as the primary driver of satellite quenching. At lower stellar masses ($< 10^8$ $M_\odot$), however, observations of the Local Group suggest that the vast majority of satellite galaxies are quenched, potentially pointing towards a characteristic satellite mass scale below which quenching efficiency increases dramatically.

4.1 Introduction

It is well documented that galaxy properties, such as morphology and star-formation rate, depend upon the local galaxy density, often referred to as the "environment" in which a galaxy is located. For instance, passive systems are systematically overrepresented in high-density environments at both low and intermediate redshift (e.g. Davis & Geller, 1976; Dressler, 1980; Lewis et al., 2002; Balogh et al., 2004; Hogg et al., 2004; **?**; **?**; **?**; **?**; **?**). This observed dependence of galaxy properties on local environment is most apparent at lower stellar masses, with satellite galaxies in groups and clusters exhibiting redder rest-frame colors, more bulge-dominated morphologies, as well as older and more metal-rich stellar populations than their counterparts of equal stellar mass in the field (e.g. Baldry et al., 2006; Yang et al., 2007; van den Bosch et al., 2008b; Peng et al., 2010; Pasquali et al., 2010; Woo et al., 2013).

A variety of physical mechanisms are potentially responsible for the generally lower star-formation rates and higher incidence of bulges for satellite galaxies relative to their field counterparts. In particular, processes such as strangulation (Larson, Tinsley & Caldwell, 1980; Balogh, Navarro & Morris, 2000; Balogh & Morris, 2000), ram-pressure stripping (Gunn & Gott, 1972; Quilis, Moore & Bower, 2000), and harassment (Farouki & Shapiro, 1981; Moore et al., 1996) may preferentially suppress star formation and transform the structure of satellite galaxies in more massive halos (i.e. higher-density environments). At present, our understanding of which mechanism(s) dominate(s) the evolution of low-mass satellite galaxies is woefully incomplete. For example, modern semi-analytic models of galaxy evolution dramatically overpredict the number of passive satellite galaxies in the local Universe (Kimm et al., 2009; Weinmann et al., 2010, 2012) and many hydrodynamical models fail to reproduce the low-mass end of the stellar mass function at $z \sim 0$ (Crain et al., 2009; Davé, Oppenheimer & Finlator, 2011).

Understanding the mass dependence of satellite quenching represents a critical step towards identifying the specific physical processes at play in driving the evolution of low-mass galaxies as a function of environment. Simple models of ram-pressure stripping and strangulation predict an increased efficiency at low mass, such that less-massive satellites should be more readily quenched by these processes. In addition, understanding the dependence of environmental effects on host mass is similarly fundamental, as more massive host halos are likely to harbor a hotter and denser circumgalactic medium that may more efficiently strip or strangulate infalling satellite systems. In a previous paper (Phillips et al., 2014a, hereafter P14), we pointed out a significant trend in satellite galaxy quenching that may result from this dependence on host mass: passive Milky Way-mass galaxies quench their satellites whereas star forming Milky Way-mass galaxies do not. Specifically, for massive satellites (with $M_\star \sim 10^{10}$ $M_\odot$), *passive* $\sim L^*$ hosts quench roughly 30% of their infalling satellites, while the satellites of *star-forming* $\sim L^*$ hosts exhibit the same star-forming activity as a field sample with the same stellar mass distribution. Using stacked line-of-sight satellite kinematics to estimate the host halo mass for these systems suggests that this dichotomy in quenching may be partially related to larger dark matter halo masses for passive L^* hosts relative to their star-forming counterparts at $z \sim 0$.

In this work, we expand upon the analysis of P14 by examining trends in satellite quenching beyond the dichotomy in host star-formation activity, focusing on the mass dependence of satellite quenching for Milky Way-like hosts. In particular, we compare satellites with stellar masses of $10^{8.5} - 10^{10.5}$ $M_\odot$ around a carefully-selected sample of isolated L^* galaxies to field galaxies with equivalent stellar mass and/or specific star-formation rate (SSFR) that are the central galaxy in their respective dark matter halo. By focusing on isolated hosts, we are able to probe systems residing in dark matter halos comparable to that of our Milky Way (a few $\times$ 10^{12} $M_\odot$, Deason et al., 2012; Boylan-Kolchin et al., 2013; van der Marel et al., 2012), while eliminating the known effects of more massive halos and large-scale structure on satellites. As motivated by the now well-established bimodality of galaxies in color-versus-

luminosity space (i.e. the red sequence and blue cloud) and the dichotomy of quenching for which we argue in P14, we will often separately consider trends where the central host galaxy is passive from those where the central galaxy is star-forming. By characterizing the mass dependence of satellite quenching, we aim to constrain the physical mechanisms dominating the evolution of satellite systems in Milky Way-like halos.

The paper is structured as follows: In Section 5.3, we describe our methodology for identifying Milky Way-like halos and describe the specific criteria applied to create our central/host and satellite samples as well as our control samples. In Section 5.4, we present our primary findings on the dependence of satellite quenching efficiency on the properties of the satellite and host galaxies. Finally, in Section 4.4, we discuss the implication of the observed trends on galaxy formation models. Throughout our analysis, we employ a Λ cold dark matter (ΛCDM) cosmology with WMAP7+BAO+H0 parameters $\Omega_\Lambda = 0.73$, $\Omega_m = 0.27$, and $h = 0.70$ (Komatsu et al., 2011b). Unless otherwise noted, all logarithms are base 10, and all quoted virial masses are derived from a spherical top-hat model according to the stellar mass/virial mass relation given in Guo et al. (2011).

4.2 Sample Selection

In selecting our observational sample, we employ data from Data Release 7 (DR7) of the Sloan Digital Sky Survey (SDSS, York et al., 2000; Abazajian et al., 2009b). In particular, we utilize the MPA-JHU derived data products, including median total stellar masses, photometrically derived according to Kauffmann et al. (2003, see also Salim et al. 2007), and median total star formation rates, measured from the SDSS spectra as detailed by Brinchmann et al. (2004). Spectroscopic completeness information as a function of position on the sky (i.e. *fgotmain*) is drawn from the NYU Value-Added Galaxy Catalog (NYU-VAGC, Blanton et al., 2005b). Our selection criteria follow very closely those of P14, motivated by careful analysis of the

Millennium II simulation (MS-II, Boylan-Kolchin et al., 2009). We will summarize them here; for a full discussion we refer the reader to that paper. Throughout this work, we refer to objects associated on the sky as "primaries" and "secondaries," while populations that have undergone correction for the presence of interlopers will be referred to as "hosts" and "satellites," as the corrected data more accurately reflect the properties of true host-satellite systems.

4.2.1 Sample Selection

For our sample of primaries, we select all objects in the SDSS spectroscopic sample with a stellar mass of $M_\star > 10^{10.5}$ $M_\odot$ and with $z < 0.032$, restricting to SDSS fiber plates with a spectroscopic completeness of >0.7 for the main galaxy sample. We then apply the following isolation criteria to preferentially select galaxies residing in halos of mass comparable to that of the Milky Way (a few $\times$ 10^{12} $M_\odot$), where the halo mass of host systems in our sample is calibrated by applying our selection criteria to the MS-II simulation. First, we allow no other galaxies with a stellar mass of $M_\star > 10^{10.5}$ $M_\odot$ within a cylinder defined by a radius of 350 kpc in projection and a length in velocity space of $\pm\,1000$ km s^{-1} along the line-of-sight centered on the primary. In addition, we define an annulus with an inner (outer) radius of 350 kpc (1 Mpc), wherein we allow no more than one galaxy of stellar mass $M_\star > 10^{10.5}$ $M_\odot$. Galaxies that pass these criteria are deemed isolated primaries.

To select spectroscopically-confirmed satellites or secondaries around our isolated primaries, we define a search region with a radius of 350 kpc on the sky and $\pm$ 500 km s^{-1} in velocity space. We search for secondaries in two mass ranges: "massive secondaries" with stellar mass of $10^{9.5}$ $M_\odot < M_\star < 10^{10.5}$ $M_\odot$ and a maximum redshift of $z = 0.032$, and "dwarf secondaries" with stellar mass of $10^{8.5}$ $M_\odot < M_\star < 10^{9.5}$ $M_\odot$ and a maximum redshift of $z = 0.024$. These limiting redshifts are designed to ensure that we are complete at all stellar

masses and SSFRs under consideration.

According to the number of massive secondaries identified about each primary, we divide our sample of isolated primaries, such that our main sample consists of primaries with exactly one secondary. As discussed by P14, restricting to systems with exactly one massive secondary in the SDSS identifies halos with a mass distribution sharply peaked at a few $\times$ 10^{12} $M_{\odot}$. We also consider primaries with exactly zero and exactly two massive secondaries. However, primaries with three or more massive secondaries are excluded, since such systems are strongly biased towards the group/cluster regime. Figure 4.1 shows the virial mass distributions for primaries in our mass range from comparison to the MS-II simulation.[1] Applying our isolation criteria removes most primaries residing in clusters, while restricting the sample to primaries with at most two massive satellites strongly selects against systems with $M_{\rm vir} > 10^{13}$ $M_{\odot}$. Primaries with two massive satellites represent a subsample of isolated objects with greater virial masses, such that these systems are more likely to have $M_{\rm vir} \gtrsim 10^{12.4}$ $M_{\odot}$ than primaries with a single massive satellite. We give no consideration to the number of dwarf secondaries orbiting a primary.

For both of the secondary samples, we construct corresponding control samples of isolated field galaxies over the same mass and redshift ranges, which we will refer to as the "massive control" and "dwarf control" samples, respectively. These samples are subjected to an isolation procedure more rigorous than that of the primaries; for both samples, we require that no galaxy more massive than the lower mass limit of the respective sample be within 3 Mpc on the sky and $\pm$ 400 km s^{-1} in velocity. This ensures that the control samples are almost completely comprised of objects that are themselves the primary galaxy in their dark matter halo ($f_{\rm purity} \sim 97\%$, see P14 for a full discussion of purity considerations) — i.e. the samples are free of satellite galaxies. The number of objects in each of our samples, including control samples, is given in Table 1.

[1]For further details regarding the analysis of the MS-II simulation, we refer the reader to the detailed discussion in P14 as well as Section 4.2.2 herein.

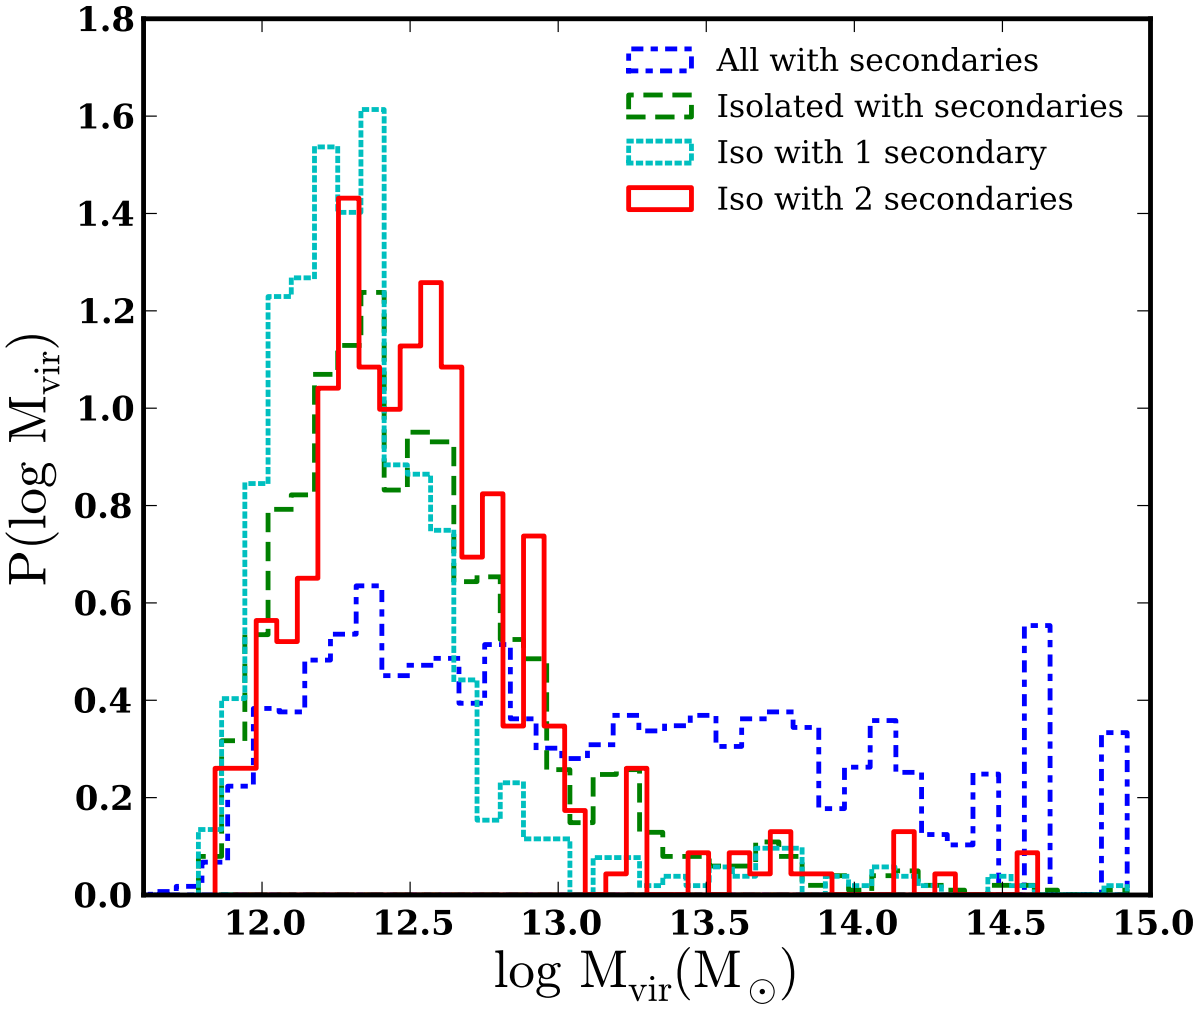

Figure 4.1: Virial mass distributions illustrating our host selection criteria. Shown are the virial mass distributions for all primaries with secondaries (blue short-dashed line), primaries that pass our isolation criteria and have at least one secondary (dashed green line), primaries that pass our isolation criteria and have exactly two secondaries (solid red line), and primaries that pass our isolation criteria and have exactly one secondary (dotted cyan line). Applying our selection criteria and restricting the number of massive satellites effectively removes massive halos from the sample. Primaries with fewer massive satellites tend to have lower virial masses. As shown by P14, restricting the sample to primaries with only one massive secondary efficiently selects Milky Way-like systems.

Mass Range	Sample	N	$< \log M_\star >$
Massive	One Satellite ("Main Sample")	457	9.98
	Two Satellites	306	9.98
$10^{9.5} M_\odot < M_\star < 10^{10.5} M_\odot$	Control	581	9.94
Dwarf	Satellite	665	8.95
$10^{8.5} M_\odot < M_\star < 10^{9.5} M_\odot$	Control	302	9.16

Table 4.1: Number of galaxies and mean log $M_\star$ in the satellite and control samples used in this study. Massive and dwarf satellites (along with their respective control samples) are restricted to stellar masses of $10^{9.5}$ $M_\odot$ $< M_\star <$ $10^{10.5}$ $M_\odot$ and $10^{8.5}$ $M_\odot$ $< M_\star <$ $10^{9.5}$ $M_\odot$, respectively.

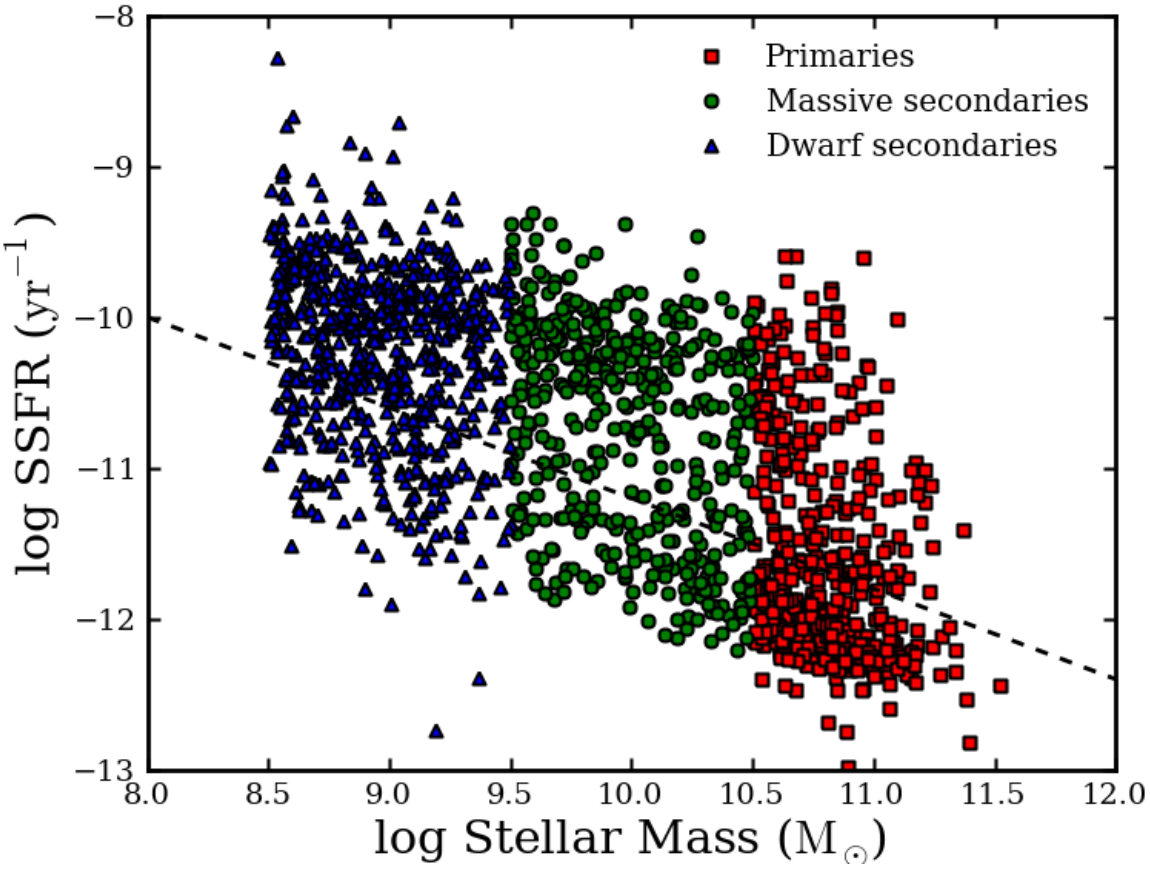

Figure 4.2: SSFR vs. stellar mass for our main sample. Plotted are all primaries at $z < 0.032$ and their massive ($10^{9.5}$ $M_\odot$ $< M_\star <$ $10^{10.5}$ $M_\odot$) secondaries. In our main sample, only primaries with exactly one massive secondary are considered. Also plotted is our dwarf ($10^{8.5}$ $M_\odot$ $< M_\star <$ $10^{9.5}$ $M_\odot$) secondary sample, which consists only of low-mass secondaries at $z < 0.024$. The dashed line separates objects into passive (below line) and star-forming (above line) categories.

As shown in P14, there are dramatic differences in the effectiveness of quenching around passive and star-forming Milky Way-size hosts, such that only massive ($\sim$ 10^{10} $M_\odot$) satellites of passive Milky Way analogs are quenched relative to a stellar mass-matched sample of isolated field galaxies. Recognizing this dichotomy in satellite quenching, we divide our sample, including both primaries and secondaries, according to SSFR, with the division between passive and star-forming galaxies set to be

$$\log(\mathrm{SSFR_{host}}) = -0.6\ \log(\mathrm{M_\star}) - 5.2. \tag{4.1}$$

This equation is motivated by the established blue cloud/red sequence bimodality of galaxies in the SDSS (Strateva et al., 2001; Baldry et al., 2004). This passivity threshold matches the slope of the stellar mass-SSFR relation for star-forming galaxies in a self-consistent way. Figure 4.2 shows where our main and dwarf satellite samples reside in SSFR-stellar mass space, with the dividing line between passive and star-forming galaxies overplotted. We note that our qualitative results are not particularly dependent on our chosen division between star-forming and passive/quenched, such that other reasonable choices of quenching definition (e.g. a quenching threshold of SSFR $< 10^{-10.5}$ yr^{-1}) give similar results.

4.2.2 Interloper Corrections and Parameter Matching

In this subsection, we discuss two numerical procedures employed to approximate the distributions of "true satellite" properties (rather than simply secondary properties) and to compare samples to corresponding control samples in a self-consistent way.

In P14, we describe in detail how we connect observations to the dark-matter only Millennium II Simulation (Boylan-Kolchin et al., 2009). In short, we wish to use the simulations to disentangle objects that are truly bound to their hosts from objects that, despite having

velocities that would suggest their being part of the system, are not. This requires two critical pieces of information: perfect phase-space information about the objects in question, which the simulations provide, and a functional model that links dark matter halos to galaxies in the real Universe. For the latter, we adopt the subhalo abundance matching (SHAM) prescription of Guo et al. (2011). While abundance matching has difficulty predicting the dark matter halo masses of the local dwarf galaxy population (Boylan-Kolchin, Bullock & Kaplinghat, 2011, 2012; Garrison-Kimmel et al., 2014), it is very successful in reproducing the clustering of more massive galaxies, including the mass range probed by our work (Berrier et al., 2006; Conroy, Wechsler & Kravtsov, 2006).

In order to ascertain the distributions and values of parameters for the true satellite population, we correct for the presence of interlopers statistically. This is done by taking the cumulative distribution of a given parameter in the secondary population and subtracting the distribution of the control sample multiplied by the probability that a randomly selected galaxy is an interloper (i.e. $1-\mathrm{f_{purity}}$). The result is an unnormalized cumulative distribution of the parameter in the *satellite* population. We renormalize so as to produce a well behaved cumulative distribution function. The equation for the satellite distribution of a parameter, here F, is given by

$$\mathrm{F_{satellite}} = \frac{\mathrm{F_{secondary}} - (1 - \mathrm{f_{purity}})\mathrm{F_{control}}}{\mathrm{f_{purity}}}. \tag{4.2}$$

Note that this requires the assumption that all "impurities" are isolated interlopers. To ease this assumption, we make the above calculation using a modified definition of $\mathrm{f_{purity}}$, where $\mathrm{f_{purity}}$ measures the fraction of secondaries in our sample that are satellites of any host, not necessarily the identified host. We calculate $\mathrm{f_{purity}}$ for our main sample and adopt this value for other samples; small variations in $\mathrm{f_{purity}}$ across samples will affect our results minimally.

Often, as in the case of comparisons between satellite and control samples, we wish to control

for distributions of parameters correlated with those under investigation. For example, to control for mass dependencies, we match the stellar mass distributions of a given satellite subsample to a control sample by dividing both subsamples into bins in the relevant parameter (e.g. stellar mass). For each bin, we then randomly select, with replacement, an object from the second sample for each object in the first sample, yielding two samples of equal number that are matched on a given parameter. Whenever this is done, the process is repeated 100 times.

Finally, connecting our observational results to the simulation allows us to investigate the extent to which we may be biasing our results by subsampling our systems. In particular, we might be concerned that satellites with, e.g., lower stellar mass might be biased towards hosts with lower virial mass. Figure 4.3 shows the stellar mass of the secondaries in our dwarf and main observational samples plotted against the stellar mass of their primaries. The data are consistent with trend in secondary stellar mass with primary stellar mass. Similarly, in the simulation the distribution of host virial masses for low virial mass satellites match the corresponding distribution for high virial mass satellites. Additionally, the simulation allows us to identify with certainty which object within a system is the host and which object is the satellite; in the observational data, this is not possible. Cases where the host and satellite are very near in mass may lead to misidentification of the host, potentially affecting our results. We find, however, that removing the borderline cases, cases where $\log(\mathrm{M}_{\star,host}) - \log(\mathrm{M}_{\star,sat}) < 0.3$, from our main sample does not affect our qualitative results.

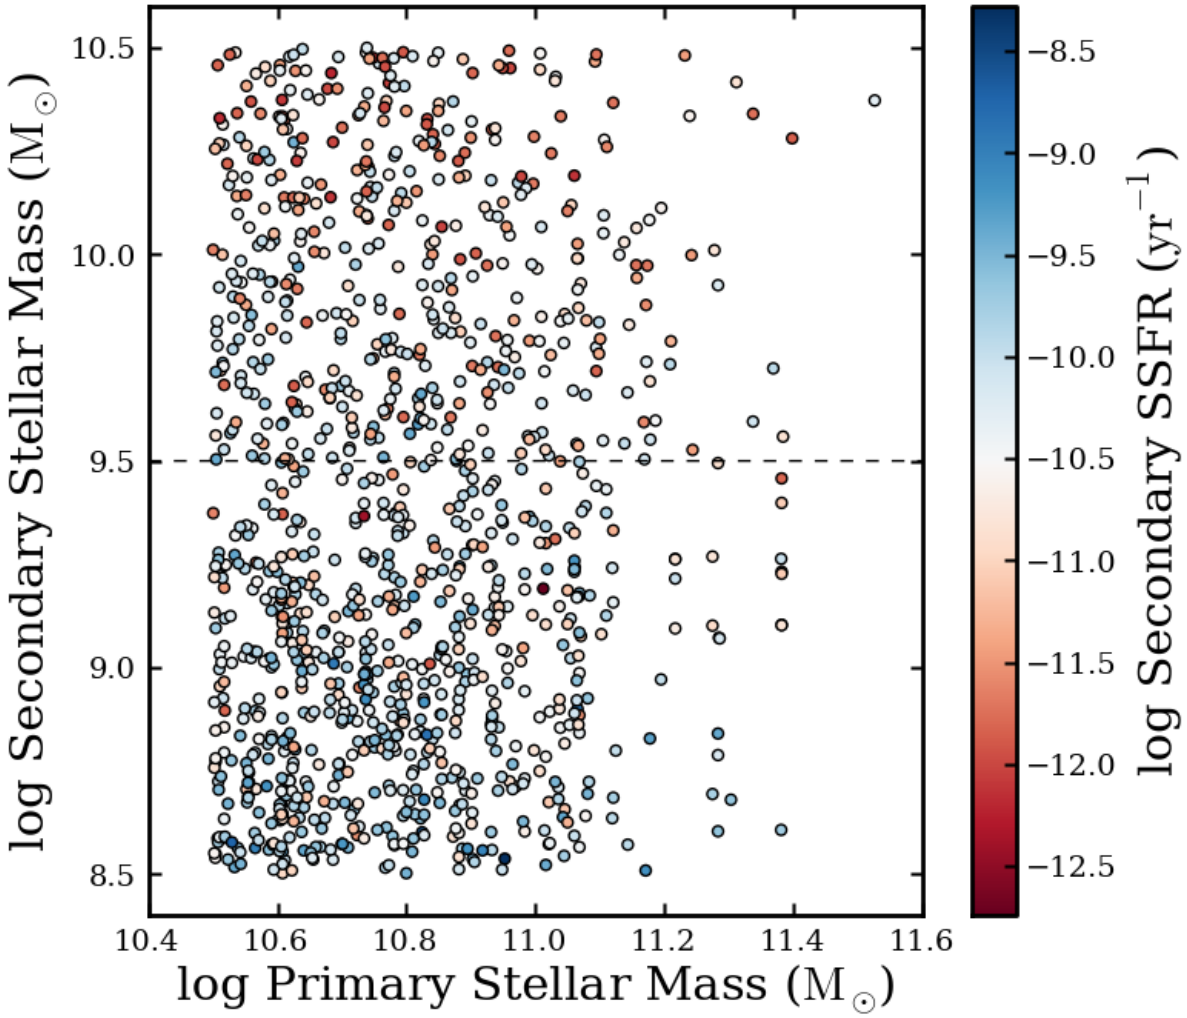

Figure 4.3: Secondary stellar mass vs. primary stellar mass for our main sample and dwarf sample. Data points are colored by the specific star formation rate of the secondary. Points that lie above the dashed line belong to the main sample, while points below it belong to the dwarf sample.

4.3 Satellite Quenching as a Function of System Properties

In P14, we introduce a parameter to quantify quenching efficiency: the conversion fraction ($f_{convert}$), which is defined as the difference in the quenched fraction between the satellite and control samples relative to the unquenched fraction of the control sample. In other words, let the unquenched fraction $u_{\{sat,control\}}(\log \mathrm{SSFR} = X)$ equal the fraction of satellites or control galaxies with $\log \mathrm{SSFR} > X$, and the quenched fraction $q_{\{sat,control\}}(\log \mathrm{SSFR} = X) = 1 - u_{\{sat,control\}}$ be the fraction of satellites or control galaxies with $\log \mathrm{SSFR} < X$. The conversion fraction, $f_{convert}$, is then given by

$$f_{convert} = \frac{q_{sat} - q_{control}}{u_{control}}. \tag{4.3}$$

In short, $f_{convert}$ corresponds to the fraction of star-forming galaxies that have been quenched upon infall onto the halos of a given set of hosts. This relies on the assumption that the properties of the control sample are adequately representative of the properties of the progenitors of the satellite samples, which is not a perfect assumption. Most satellites are star-forming, such that they continue to increase their stellar masses over time, even though they fell onto a host some time in the past. Other works have used a similar statistic to our conversion fraction, including van den Bosch et al. (2008b) and Peng et al. (2012).

Throughout this analysis, errors on the quenched fractions are computed according to binomial statistics as

$$\sigma_{q,sample} = \sqrt{\frac{(q_{sample})(1 - q_{sample})}{N_{sample}}}, \tag{4.4}$$

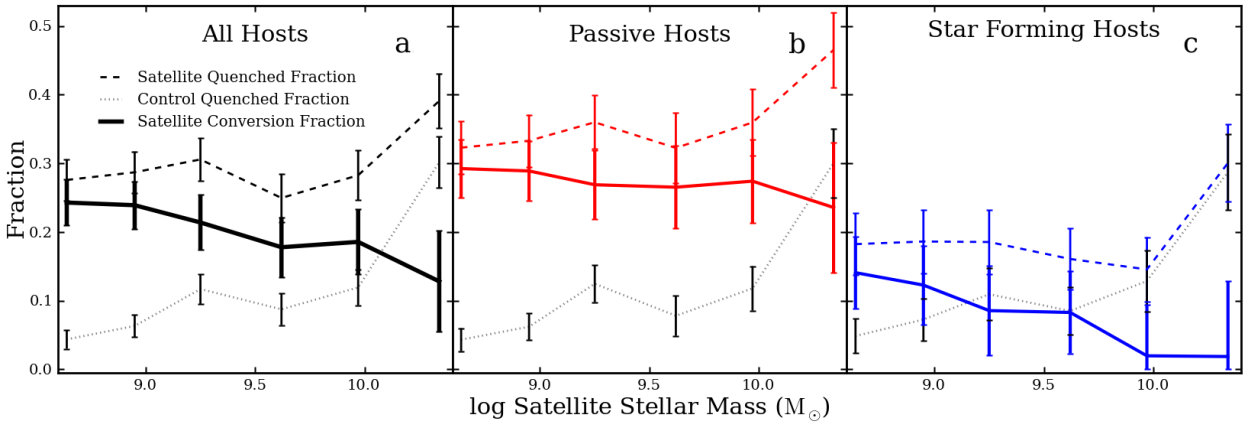

Figure 4.4: The conversion fraction for satellites of all hosts (*left panel*), passive hosts (*center panel*), and star-forming hosts (*right panel*) as a function of satellite stellar mass in our main sample (*solid lines*). For reference, the quenched fraction for the respective satellite samples (*dashed lines*) and mass-matched control samples (*dotted lines*) are also shown. The conversion fraction is independent of stellar mass in the passive host sample, while exhibiting a negative correlation with satellite mass in the star-forming host sample, such that less-massive satellites of star-forming hosts are more likely to be quenched upon infall.

where $\sigma_{\mathrm{q,sample}}$ is the error on the quenched fraction of the sample, $\mathrm{q_{sample}}$ is the quenched fraction of the sample, and $\mathrm{N_{sample}}$ is the number of objects in the sample. These errors are then propagated according to Equation 4.3 to determine the associated errors on the conversion fractions.

4.3.1 Dependence on Satellite Mass

Considering both our main and dwarf subsamples, our full sample of satellite galaxies spans two orders of magnitude in stellar mass, $10^{8.5} < \mathrm{M_\star/M_\odot} < 10^{10.5}$. To explore the dependence of quenching on stellar mass, we separate our satellite sample into six distinct bins, by independently dividing the main and dwarf subsamples according to the 33^{rd} and 67^{th} percentiles in stellar mass. The resulting six independent bins are bounded in stellar mass space by $10^{8.5}$, $10^{8.80}$, $10^{9.11}$, $10^{9.5}$, $10^{9.79}$, $10^{10.14}$, and $10^{10.5}$ $\mathrm{M_\odot}$. In Figure 4.4a, we show the passive fraction for the satellite samples in each stellar mass bin (dashed line) along side that of the control samples of isolated galaxies (dotted line) with the same stellar mass

distribution (i.e. "mass-matched"). As highlighted by previous studies of satellite galaxies in the local Universe, predominantly in more-massive halos, we find that satellites are preferentially passive relative to the field population (e.g. Weinmann et al., 2006; Tollerud et al., 2011b; Geha et al., 2012; Wang & White, 2012).

Following Equation 4.3, the measured passive fractions for the satellite and control samples yield a conversion fraction as a function of satellite mass that weakly increases with decreasing stellar mass, such that less-massive satellites are slightly more likely to be quenched upon infall to a Milky Way-like halo (see solid line in Fig. 4.4a). Within the errors, however, the measured conversion fractions are largely consistent with no dependence on satellite stellar mass at $10^{8.5}$ $M_{\odot} < M_{\star} < 10^{10.5}$ $M_{\odot}$. Using repeated Monte Carlo resampling and assuming the stated errors are normally distributed,[2] we estimate the slope of the relation between conversion fraction and $\log M_{\star}$ to be -0.069 ± 0.036, marginally inconsistant with no correlation and consistent with a slight anti-correlation.

As shown by P14, the conversion fraction (and thus efficiency of satellite quenching) varies significantly between passive and star-forming Milky Way-like hosts, such that massive ($\sim 10^{10}$ $M_{\odot}$) satellites are only quenched around passive hosts. Given this observed dichotomy of massive satellite quenching, we compute the passive and conversion fractions as a function of satellite stellar mass for passive and star-forming hosts separately (see Figure 4.4b,c). At all satellite stellar masses probed, passive hosts are more effective at quenching than star-forming hosts. Moreover, we again find that at high satellite masses, passive hosts are the sole drivers of satellite quenching, with a conversion fraction for satellites of passive hosts of roughly 30% relative to nearly 0% around star-forming systems. Remarkably, across the entire range of satellite stellar masses studied, this moderate quenching efficiency (of $\sim$30%) is relatively independent of satellite mass for passive hosts. Using the Monte Carlo method described above, we find a slope in the conversion fraction-satellite stellar mass relation of

[2]Our samples are large enough that the assumption of normally-distributed errors is valid, typically on the order of n = 60 - 100.

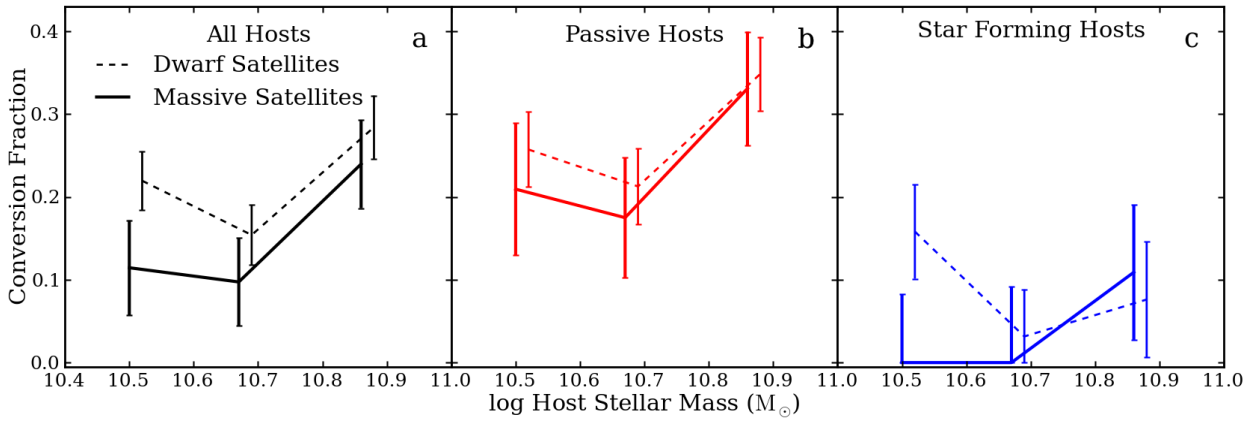

Figure 4.5: Conversion fractions for satellites in our main sample (*solid lines*) and dwarf sample (*dashed lines*) as a function of the stellar mass of their host, divided according to the star-forming properties of the host. Left panel: all hosts, center panel: passive hosts, right panel: star-forming hosts. We find little evidence for a dependence of satellite quenching effiency on host stellar mass.

-0.027 ± 0.048, consistent with no dependence of quenching efficiency on satellite stellar mass for passive Milky Way-like hosts. Conversely, we find weak evidence for a modest increase in the conversion fraction at lower stellar masses around star-forming systems, with a slope of -0.089 ± 0.058. This suggests that the weak negative slope seen with passive and star forming hosts taken together is driven by the anti-correlation in quenching efficiency and satellite stellar mass found around star-forming hosts.

4.3.2 Dependence on Host Mass

It is natural to expect that the efficiency of satellite quenching correlates with the stellar mass of the host. For example, galaxies of higher stellar mass tend to live in more massive dark matter halos (Mandelbaum et al., 2006; Conroy et al., 2007; Behroozi, Wechsler & Conroy, 2012; McGaugh, 2012; Moster, Naab & White, 2013; Miller et al., 2011, 2014), which may quench satellites more effectively. This trend is observed on large scales, with clusters hosting significantly higher fractions of passive galaxies relative to comparable field samples (e.g. Dressler, 1980; Postman et al., 2005). With this thought in mind, we study the dependence

of quenching efficiency on host stellar mass within our sample of less-extreme, lower-mass halos. Hosts are divided into three mass bins bounded by the $33^{\rm rd}$ and $67^{\rm th}$ percentiles of the host stellar mass distribution, which correspond to stellar masses of $10^{10.69}$ $M_\odot$ and $10^{10.86}$ $M_\odot$, respectively.

To examine how quenching efficiency depends on the stellar mass of the host, we examine the satellites in the main and dwarf sample separately, as shown in Figure 4.5, and divide the hosts according to their status as passive or star-forming. This gives us four independent subcategories that we can use to examine the host stellar mass dependence of quenching efficiency: massive satellites of passive hosts, massive satellites of star-forming hosts, dwarf satellites of passive hosts, and dwarf satellites of star-forming hosts. Of the four subsamples, all but the dwarf satellites of star-forming hosts exhibit a similar correlation between $f_{\rm convert}$ and host stellar mass: constant quenching efficiency in the two lower host mass bins, with a slight increase in quenching efficiency in the highest mass bin. We again use Monte Carlo resampling to test whether the slope of the conversion fraction-host stellar mass relation is consistent with zero in each subsample — that is, wether or not the increase in quenching efficiency at high host stellar mass is statistically significant. In all cases, we find that the observed quenched efficiencies are largely consistent with no dependence on host mass, such that the measured slopes of the conversion fraction versus $\log(M_\star)$ relation are consistent with zero at $< 2\sigma$.

The subsample that differs significantly from the other three categories, dwarf satellites of star-forming hosts, sees an increase in quenching efficiency around the lowest mass hosts. We showed in the previous subsection that lower mass satellites of star-forming hosts are more likely to be quenched; however, the increase in the efficiency with which dwarf satellites are quenched in low-mass star-forming hosts is likely not driven by an decrease in the characteristic stellar mass of satellites of such hosts. A Kolmogorov-Smirnov (KS) test fails to show that the stellar masses of the objects in the three dwarf satellite/star-forming host bins

are drawn from different underlying distributions ($p_{lower,middle} = 0.71$, $p_{lower,upper} = 0.37$). While the behavior of dwarf satellites of star-forming galaxies does seem odd, the data are consistent with no dependance in conversion fraction on host stellar mass.

4.3.3 Dependence on Satellite Number

As discussed in §5.3, we have to this point restricted our analysis to hosts with exactly one massive satellite, as this preferentially selects Milky Way-like systems in lieu of more massive dark matter halos (see Fig. 4.1). Here, we relax this restriction on the number of massive satellites and consider separately hosts that have exactly one massive satellite and hosts with exactly two massive satellites. These categories represent hosts with different characteristic virial masses. As shown in Fig. 4.1, we find that hosts with two massive satellites live in systematically more massive dark matter halos based on comparison to the MS-II simulation.[3] For ease of comparison with P14, we adopt a simplified threshold for quenching, by which satellites will be considered quenched if their SSFR is below 10^{-11} yr^{-1}. In Figure 4.6, we plot the cumulative distribution of specific star formation rates for massive satellites around passive and star-forming hosts separately, matching the stellar mass distributions of all four host samples. For hosts with only one massive satellite, the dichotomy of quenching discovered in P14 is readily apparent, with satellites of passive hosts more than twice as likely to be quenched than a stellar mass-matched field sample and satellites of star-forming hosts indistinguishable from their field counterparts. For systems with multiple massive satellites, however, the behavior is different. Star-forming hosts now exhibit non-zero quenching efficiency, in fact quenching their satellites with approximately the same efficiency as possive hosts, possibly pointing to a trend in quenching efficiency with host halo mass. Cutting the two satellite sample on host and satellite stellar mass yields results similar to the one satellite case

[3]The median halo mass for isolated hosts with two massive secondaries is approximately 60% larger than that for isolated systems with only one massive secondary in the MS-II simulation.

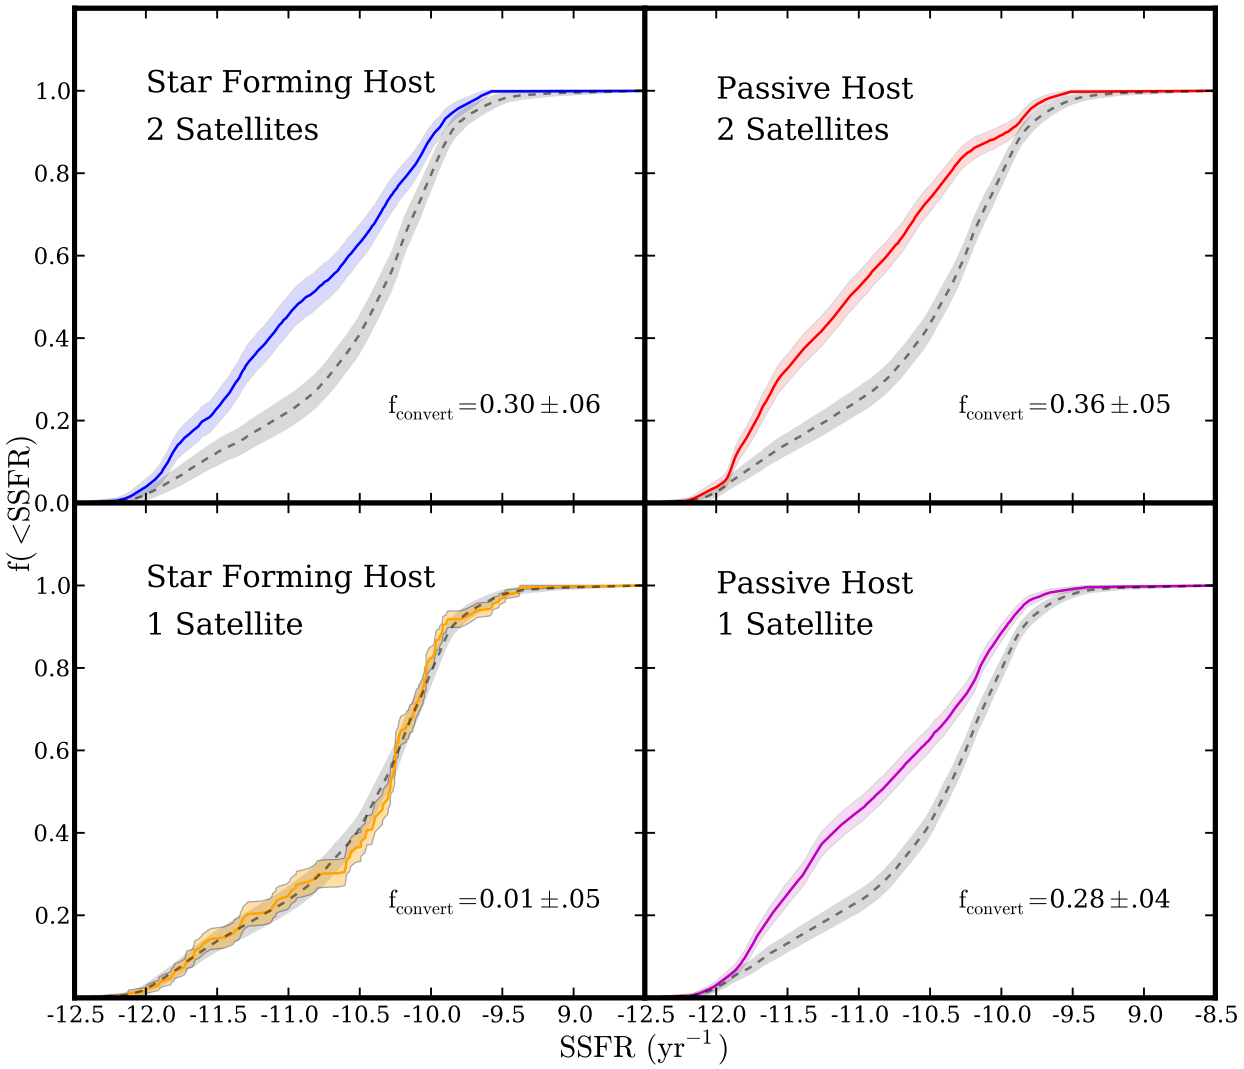

Figure 4.6: Cumulative distributions of specific star formation rate for massive satellites of passive (right column) and star-forming (left column) hosts with exactly one (bottom row) and exactly two (top row) satellites. The black dashed lines show the corresponding distributions for stellar mass-matched samples of isolated field galaxies. Colored and grey shaded regions correspond to 1σ binomial errors for the satellite and control samples, respectively. Hosts with two massive satellites are more effective at quenching those satellites than corresponding hosts with only one satellite, and star-forming hosts with two satellites have a non-zero quenching efficiency, breaking the dichotomy of quenching presented in P14.

In light of the evidence that more massive halos quench their satellites more effectively, we might hypothesize that there exists a critical halo mass threshold above which a host will quench its satellites extremely efficiently. This could be considered a potential limiting case of quenching scenarios, whereby in a two-satellite system the conditional probability of finding a satellite quenched given that its partner is quenched is near unity. An alternate limiting case might be that the probabilities of finding either satellite quenched are independent of each other. How well the data conform to either limiting case can inform quenching models, and by examining our sample host-by-host, we can examine how the satellites are distributed among the hosts.

Considering each two-satellite system individually, we ask whether the satellites are matched or unmatched in their star-forming properties. Figure 4.7 shows the frequency of passive, star-forming, and mis-matched satellite pairs around passive and star-forming hosts. For comparison, we also plot the binomial distributions for both samples — i.e. the expectation from randomly drawing satellites from the population of all satellites. In the observational sample, mis-matched pairs are better represented than the matched pairs of passive or star-forming satellites. While the incidence of mis-matched pairs is lower than the binomial distribution and the frequencies of matched passive and star-forming pairs are higher than the binomial distribution, the disagreement is fairly minor. For a system with two massive satellites, having one satellite be passive corresponds to a slightly higher probability of its partner being passive and vice versa for star-forming satellites. We note that the distributions are very similar between the two host types, and despite statistically signifiant deviation from the binomial distribution, the distribution of observed pairs is much better modeled by a binomial distribution than the limiting case of all satellites being found in matched pairs.

Satellite galactocentric distance can provide information about the environment in which it lives, such as what type of CGM it is embedded in or how long it has been interacting with its host. With that in mind, we examine the projected distance distribution of satellite

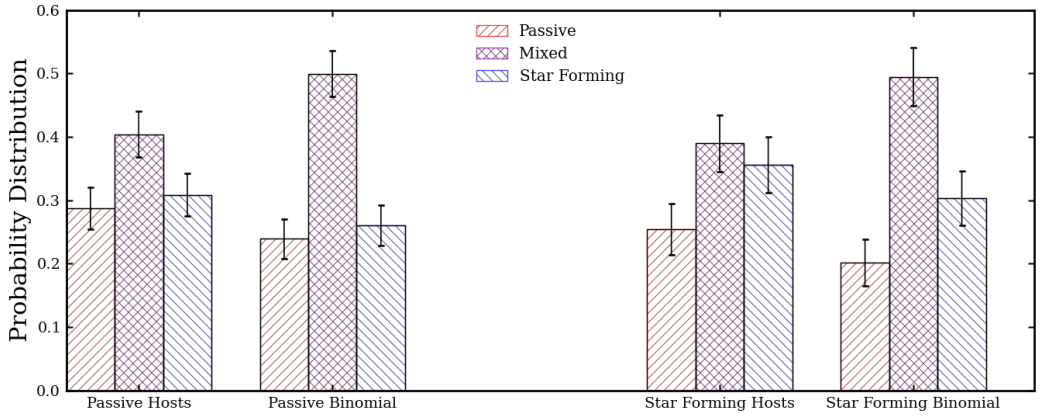

Figure 4.7: The relative frequency of matched or mis-matched satellite properties in two satellite systems for passive (*left*) and star-forming (*right*) hosts. Shown separately are the frequency of matched passive pairs (red bars), matched star-forming pairs (blue bars) and mixed pairs (i.e. one passive satellite and one star-forming satellite, purple bars). Also plotted are binomial distributions calculated from the red fraction of satellites in the parent passive and star-forming host categories, representative of the limiting case where satellites are randomly assigned to hosts. In systems with two satellites, having one passive satellite corresponds to a slightlyhigher likelihood of the second satellite also being passive and vice versa for star-forming satellites.

galaxies in systems with exactly two massive satellites. In Figure 4.8, we examine the radial distribution of satellites in systems with two massive satellite, grouping the satellites according to their star-forming properties as well as that of the host. The top and bottom sets of three panels show the distributions for satellites of passive and star-forming hosts, respectively.

Of particular interest is the difference between the radial distributions of passive satellites compared to their star-forming counterparts. We find that passive satellites are more likely to be found at small projected radii, while there is a corresponding overabundance of star-forming satellites in the outer regions of the halo. This trend is apparent in the mixed satellite cases, and is readily seen upon comparing passive and star-forming satellites across the various host subsamples. A Kolmogorov-Smirnov (KS) test rejects the null hypothesis that the distributions of projected distances are identical between star-forming and passive

satellites ($p = 0.004$). This trend is particularly apparent in the cases where one satellite is star-forming and one satellite is passive (middle panels of Fig. 4.8); for both passive and star-forming hosts, the star-forming satellites are more likely to be found in the outer regions of the host halo relative to the passive satellites.

The one category that seems not to follow the trend is the category of star-forming hosts with paired passive satellites. Here we find, despite low statistics, that satellites are mostly found at large projected distances. One possible interpretation of this result is that a relatively high portion of these objects were not environmentally quenched by their host, but rather quenched in the field. We might expect that if the satellites had been environmentally quenched at a similar rate to the other categories, they would be found at projected distances more in line with those seen in the other categories. Alternately, the apparent inversion of the radial trend seen in the other subsamples could be the result of host misidentification. In situations where the hosts and satellites have a nearly one-to-one mass ratio, if the true host of the system is misidentified as a distant satellite, the unexpected case of a host having a distant passive satellite would indeed be the expected case of a host having a distant star-forming satellite. Of the 30 objects identified as satellites in this subsample, 12 have projected distances greater than 250 kpc. Of these, 5 are within a factor of three of their host in stellar mass, comprising 4 of the 15 systems. If these four systems are disregarded, the radial distribution of passive satellites becomes flat, more in line with the other subsamples of passive galaxies.

4.4 Discussion

In this paper, we have examined the dependence of satellite quenching on the mass of the satellite galaxy, the mass of the host galaxy, and the multiplicity of satellites in the system. For our host and satellite sample, we show that satellite quenching is not strongly correlated

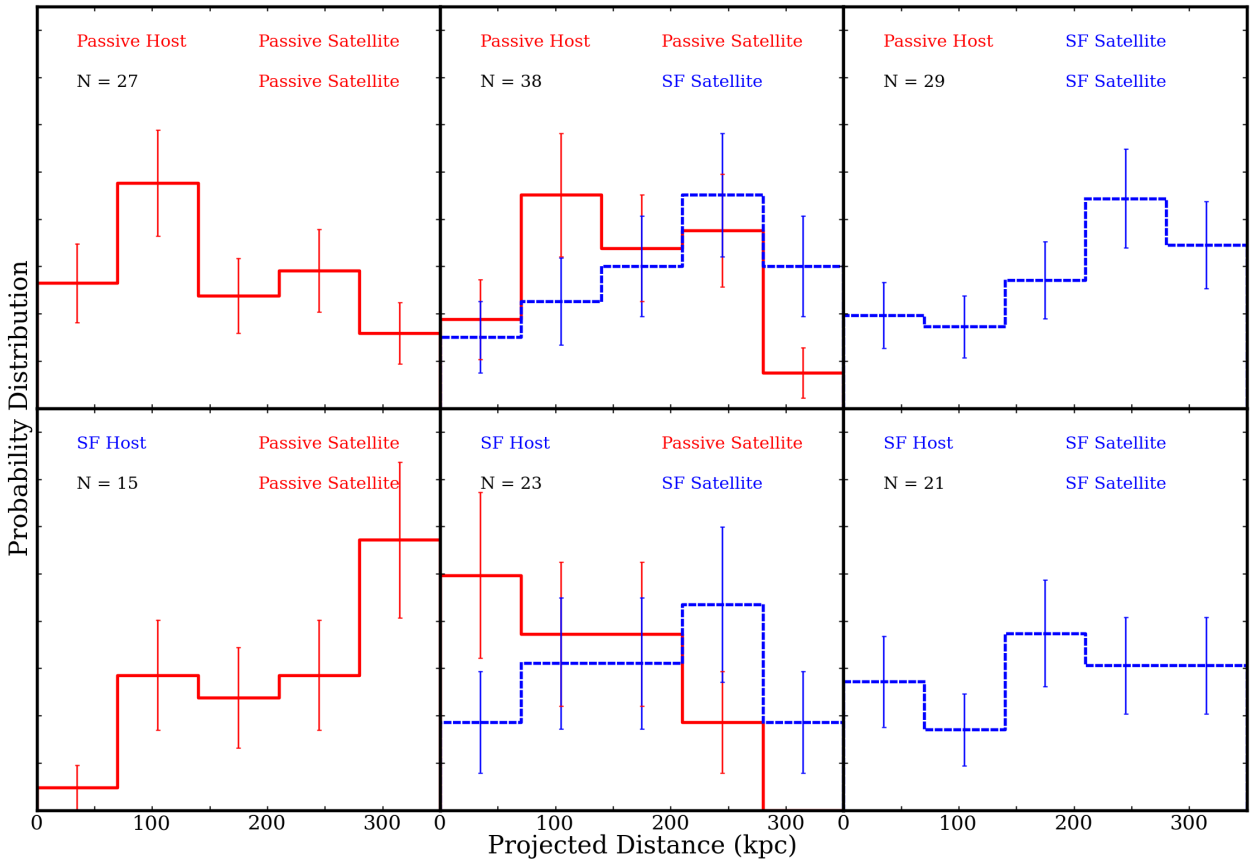

Figure 4.8: Radial distributions of massive satellites for passive (*top row*) and star-forming (*bottom row*) hosts with exactly two massive satellites. In each row, we divide the subsamples into passive satellite pairs (*left column*), mixed satellite pairs (*center column*), and star-forming satellite pairs (*right column*). In each case, passive satellite distributions are represented by a solid red line and star-forming satellite distributions by a dashed blue line. There is an overall trend for passive satellites to be more centrally concentrated than their star-forming counterparts. The overabundance of passive satellites at large projected distances in the bottom left panel may be a result of mis-identifying the host galaxy.

with satellite mass and potentially increases with increasing host virial mass, as best shown by an increase in conversion fraction in systems with two massive satellites. We also demonstrate that within such systems, passive and star-forming satellites are found around hosts roughly with frequencies described by a binomial distribution, implying little correlation between the properties of satellite pairs. In the following section, we discuss these results in comparison to previous studies in the local Universe and highlight the implications of these findings on models of galaxy evolution.

4.4.1 The Mass Dependence of Satellite Quenching

Across a broad range in stellar mass, $10^{8.5}$ $M_\odot < M_\star < 10^{10.5}$ $M_\odot$, we find little evidence for a correlation between conversion fraction (i.e. the efficiency of satellite quenching) and satellite stellar mass, with Milky Way-like hosts quenching $\sim$ 20% of infalling satellites on average. This relatively inefficient quenching is strongly correlated with host properties, such that the satellites of passive hosts, at all masses probed, are more likely to be quenched than their counterparts around star-forming hosts. Moreover, for the most massive satellites ($\sim$ 10^{10} $M_\odot$), quenching is entirely driven by the halos of passive hosts, with massive satellites of star-forming hosts indistinguishable from the field population — a confirmation of the dichotomy of satellite quenching shown in P14. At lower satellite stellar masses ($\sim$ $10^{8.5}$ $M_\odot$), however, we measure a conversion fraction of $\sim$ 15% for satellites of star-forming hosts. This breaks the strong dichotomy of satellite quenching observed at higher masses, though the mass dependance of satellite quenching efficiency for satellites of star-forming hosts is mild.

The lack of significant correlation between satellite stellar mass and conversion fraction is surprising. In accordance with subhalo abundance matching, lower stellar mass galaxies tend to occupy lower virial mass (and thus less dense) halos. Given their shallower potential

wells, these galaxies would therefore be expected to lose their gas more easily and become quenched, thereby yielding a higher conversion fraction at lower satellite masses. However, this is not observed.

One possibility is that the expected trend for low-mass galaxies to lose their gas (and thus quench) more easily may be balanced out by a tendency for these systems to possess larger gas reservoirs. Under a scenario where hot gas is stripped from a satellite upon infall but cold gas is retained (i.e. "strangulation"), galaxies with higher cold gas fractions would thereby take longer to use up their gas and ultimately quench. According to observations of atomic hydrogen in local star-forming galaxies, lower stellar mass systems are generally found to have higher atomic gas fractions and longer atomic depletion timescales (SFR/M_{HI}) than their more massive counterparts (Skillman, Côté & Miller, 2003; Geha et al., 2006; Leroy et al., 2008; Schiminovich et al., 2010). Furthermore, recent studies of molecular gas, which more closely traces ongoing star formation (Wong & Blitz, 2002; Kennicutt et al., 2007; Bigiel et al., 2008), also show increasing cold gas fractions at lower stellar masses (Saintonge et al., 2011; Boselli et al., 2014). Altogether, the overall trend for the progenitors of low-mass satellite to have high cold gas fractions and correspondingly long depletion times may serve to counteract the tendency for low-mass satellites to lose gas easily, resulting in an overall non-dependence of conversion fraction on satellite stellar mass.

While the efficiency of satellite quenching is observed to be relatively independent of satellite stellar mass, we do find evidence for a significant increase in the conversion fraction, $f_{convert}$, with increasing host halo mass. In particular, we find a greater satellite quenching efficiency for host systems with two massive satellites relative to those with only one massive satellite (see Fig. 4.6), where comparison to the MS-II simulation shows that systems with more massive satellites are preferentially biased towards greater dark matter virial masses (see Fig. 4.1). Our measurements of stacked satellite velocity dispersions confirm that hosts with two massive satellites preferentially reside in more massive halos. As in P14, we stack the

line-of-sight velocity distributions for satellites of passive and star-forming hosts at fixed host stellar mass, so as to measure the velocity dispersion as a proxy for the characteristic virial mass of the respective host populations. Passive hosts with one satellite have a satellite velocity dispersion of 165 ± 11 km s^{-1}, while passive hosts with two satellites have a satellite velocity dispersion of 197 ± 13 km s^{-1}, such that passive hosts with two satellites live in more massive dark matter halos than passive hosts with one satellite. Likewise, star-forming hosts with a single massive satellite yield a stacked satellite velocity dispersion of 148 ± 14 km s^{-1}, whereas star-forming hosts with two massive satellites have a satellite velocity dispersion of 189 ± 16 km s^{-1}, implying the same about star-forming hosts.

The apparent correlation between satellite quenching efficiency and host halo mass may also explain the increased prevalence of quenched satellites around passive hosts relative to their star-forming counterparts. As shown in Fig. 4.6, the conversion fraction for satellites of passive Milky Way-like hosts, in both the one- and two-satellite cases, exceeds that measured for satellites of star-forming hosts (see also Fig. 4.4). The comparison of stacked satellite velocity dispersions for these samples shows that even at fixed satellite number, as well as fixed stellar mass, passive hosts preferentially live in more massive dark matter halos. Thus, as discussed by P14, a quenching efficiency that depends on host halo mass in concert with the preference for passive hosts — at a given stellar mass — to reside in more massive halos directly explains the increased prevalence of quenched satellites around passive hosts.

Beyond dynamical tracers of halo mass, host stellar mass is also expected to track halo mass on average, with more massive host galaxies typically residing in more massive halos (e.g. Moster et al., 2010; Miller et al., 2014). Yet, while we do find mild evidence for an increase in quenching efficiency at high host stellar mass, our results are largely consistent with no correlation between conversion fraction and the stellar mass of the host. One possible explanation for this lack of observed correlation between satellite quenching and host stellar mass is that our sample spans a relatively narrow range in host stellar mass, such that

we cannot resolve a clear stellar mass-halo mass relation — i.e. the stellar mass-halo mass relation may be largely dominated by scatter over the range in halo masses probed by our hosts. To test this possibility, we again stack the satellites of the hosts in each host stellar mass bin and measure the velocity dispersions of the massive satellite populations. We find no trend in satellite velocity dispersion with increasing host stellar mass: the lowest bin in host stellar mass has a velocity dispersion of 163 ± 10 km s^{-1}, the middle bin has a velocity dispersion of 162 ± 11 km s^{-1} and the upper bin has a velocity dispersion of 155 ± 9 km s^{-1}. This suggests that the evidence we find for a strong dependence of quenching efficiency on host halo mass and for weak to no dependence of quenching efficiency on host stellar mass do not necessarily contradict each other and are consistent with a picture where higher virial mass hosts are more likely to quench their satellites.

As discussed in P14, the dependence of quenching efficiency on host halo mass may reflect the preference for more massive dark matter halos to harbor hot gas coronas, which are then able to quench infalling satellite galaxies via ram-pressure stripping. Studies of gas accretion onto dark matter halos indicate that there is a transition in the dominant accretion mode at a halo mass of roughly a few $\times 10^{12}$ M$_\odot$. Infalling gas is shock-heated at the virial radius in halos above this threshold, while cold gas reaches a smaller radius, possibly falling all the way to the galaxy, in less-massive halos (e.g. Binney, 1977; Rees & Ostriker, 1977; Birnboim & Dekel, 2003; Kereš et al., 2005, 2009; Stewart et al., 2011b). Models of galaxy formation in which quenching only occurs above this critical halo mass show significant promise in reproducing the observed dependence of the galaxy quenched fraction on stellar mass and environment at $> 10^{9.5}$ M$_\odot$ (Gabor & Davé, 2014). Our results, as presented in Fig. 4.6, are largely consistent with this picture of gas accretion and stripping, such that star-forming hosts with one massive satellite preferentially reside in halos below the critical halo mass and passive hosts with one massive satellite as well as all hosts with two massive satellites inhabit more massive halos with established hot coronas. In the less massive halos, the satellite galaxies largely mirror the field population, while roughly 30% of infalling satellites

are quenched in more massive systems, potentially due to the presence of a hot halo.

It should be noted, however, that models such as these, which depend on interactions with host circumgalactic media to drive satellite quenching, may struggle to reproduce the observed correlation between the star-forming properties of massive galaxies and that of lower-mass systems located at distances of several virial radii (e.g. Kauffmann et al., 2013; Wetzel et al., 2014). Alternatively, work by Hearin, Watson & van den Bosch (2014) presents a picture where the importance of intra-halo quenching mechanisms is overstated, and large scale conformity, i.e. two-halo effects, is important.

While clear evidence is found for a correlation between quenching efficiency and halo mass, the observations of two-satellite systems are not – at first glance – entirely consistent with this picture. As shown in Fig. 4.7, the star-forming properties of satellites are consistent with being randomly drawn from the parent population, such that if one massive satellite is quenched in a system then the second satellite is only marginally more likely to also be quenched. In a scenario where quenching is driven entirely by host halo mass, one might expect that quenched satellites would be preferentially found within a particular subset of halos (e.g. within those more massive halos with a CGM capable of stripping an infalling satellite). This reasoning, however, assumes that satellites quench on a reasonably short timescale, such that there is a relatively small chance of observing a star-forming satellite within a halo capable of quenching it. Recent studies of quenching timescales for low-mass satellites suggest that this assumption is flawed, with satellites with a stellar mass of $\sim 10^{10}$ $M_{\odot}$ estimated to quench ~ 6 Gyr after infall (De Lucia et al., 2012; Wetzel et al., 2013a; Wheeler et al., 2014). Adopting this quenching timescale within the MS-II simulation, where infall times are known, and assuming that all systems with two massive satellites are above the critical halo mass, we are able to precisely reproduce the distributions of matched and mis-matched satellite pairs shown in Figure 4.7, under the assumption that only subhalos accreted more than 5 Gyr ago are quenched. Additional models could potentially

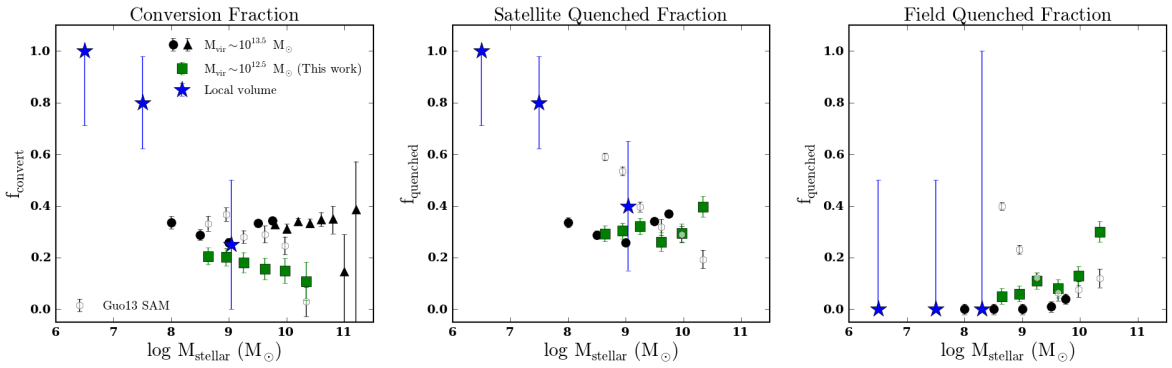

Figure 4.9: Conversion fractions (left panel), satellite quenched fractions (center panel), and field quenched fractions (right panel) for our sample of Milky Way-like systems ($M_{vir} \sim 10^{12.5}$ $M_\odot$; green squares). For comparison, we show various samples of local galaxies including group-scale systems ($M_{vir} \sim 10^{13.5}$ $M_\odot$, black circles: Geha et al. 2012; Wheeler et al. 2014; black triangles: Wetzel et al. 2013a), the Local Group/Local Volume (blue stars) and simulated galaxies from the Guo et al. (2013) SAM (open grey hexagons). At stellar masses greater than 10^8 $M_\odot$, conversion fractions show little to no dependance on satellite stellar mass at fixed host halo mass. At fixed satellite stellar mass, hosts with greater virial masses quench their satellites more effectively. A significantly greater fraction of Local Group satellites are quenched than higher mass satellites, indicating a potential critical mass scale for satellite quenching at $\sim 10^8$ $M_\odot$. The differences between the SAM and the observational data is likely driven by overpredicting the field quenched fraction at these masses. Note that we exclude likely "backsplash galaxies" (i.e. galaxies believed to have previously interacted with their host, namely Cetus) from the Local Field sample.

be constrained with quenching timescales that vary with host mass and a critical halo mass above which a host can quench and below which it can not.

4.4.2 Comparison to Previous Studies

Using data drawn from the SDSS, several recent studies similarly conclude that satellite quenching efficiency is largely independent of satellite stellar mass at $10^8 \lesssim M_\star/M_\odot \lesssim 10^{11}$ (Geha et al., 2012; Wetzel et al., 2013a; Wheeler et al., 2014). In particular, Wetzel et al. (2013a) employ the group catalog of Yang et al. (2005, 2007) to study the frequency of quenched satellites at $> 10^{9.7}$ $M_\odot$, showing that $\sim 30\%$ of massive satellites are quenched upon infall to a host halo. Using the data of Geha et al. (2012) to push to lower satellite

masses, Wheeler et al. (2014) extend this work by studying the quenching efficiency around local hosts with stellar masses of $> 10^{10.4}$ $M_\odot$, finding a comparable quenching efficiency. Figure 4.9 shows our measurements of (*i*) satellite conversion fraction, (*ii*) satellite quenched fraction, and (*iii*) field quenched fraction as function of satellite stellar mass for Milky Way-like hosts (i.e. our "main sample"). For comparison, we include complementary results from Wheeler et al. (2014) and Wetzel et al. (2013a) along side corresponding measurements for the Local Group and the Guo et al. (2013) semi-analytic model (SAM). Both the Wheeler et al. (2014) and Wetzel et al. (2013a) studies find higher conversion fractions at all stellar masses. This relative increase in quenching efficiency is likely driven by variations in the sample selection that lead to significant differences in the typical halo masses probed. As shown by Wheeler et al. (2014), the typical host in the Geha et al. (2012) sample has a halo mass of $\sim 10^{13.5}$ $M_\odot$, comparable to that of the group catalog employed by Wetzel et al. (2013a). Our methodology differs in that we apply isolation criteria so as to restrict our analysis to 10^{12} $M_\odot$ halos. This places our hosts at systematically lower virial masses than the above works, which in turn biases our results towards environments of lower quenching efficiency.

Our work, taken together with the complementary results described above, paints a picture of relatively low satellite quenching efficiency ($\sim 30\%$) that is independent of satellite stellar mass across an impressively broad range in mass ($10^8 < M_\star/M_\odot < 10^{11}$). At stellar masses of $> 10^8$ $M_\odot$, observations of the Local Group are remarkably consistent with these results. Among the massive satellites of the Milky Way and M31, the LMC, SMC, and M33 are star-forming while NGC 205 and M32 are quenched, yielding a quenched fraction of 40% and conversion fraction of roughly 25%. Here, we define objects in the Local Group and surrounding field population as quenched according to their observed atomic gas fractions as reported by McConnachie (2012), such that objects with $M_{HI}/M_\star < 0.1$ are classified as quenched. For the field population surrounding the Local Group, we find that nearly all galaxies are star-forming, where the rare exceptions (e.g. Cetus and Tucana) are likely

systems that previously interacted with either the Milky Way or M31 (i.e. "backsplash" or "super-virial" galaxies, Mamon et al. 2004, Wetzel et al. 2014, Garrison-Kimmel et al. in prep). This broadly agrees with our observations of low-mass isolated systems as well as the results of Geha et al. (2012).

While at stellar masses of $> 10^8$ $M_\odot$ the satellite population in the Local Group shows broad agreement with the observed properties of satellites around our sample of Milky Way-like systems, at lower stellar masses Local Group satellites exhibit a dramatically greater quenched fraction and quenching efficiency (see Fig. 4.9). With one exception, IC10, every satellite galaxy of both the Milky Way and M31 below 10^8 $M_\odot$ is quenched. This dramatic increase in satellite quenching efficiency at low stellar masses ($< 10^8$ $M_\odot$) points towards a potential critical mass scale for satellite quenching, such that environmental quenching is highly efficient for very low-mass satellites.

Observations of systems in the local Universe comparable to the Local Group generally confirm the increased satellite quenching efficiency at low masses. For example, the low-mass ($\sim$ a few $\times 10^6$ $M_\star$) satellites of M81 are universally quenched, with observed star formation rates of $< 10^{-5}$ $M_\odot$ yr^{-1} (Kaisin & Karachentsev, 2013). In addition, recent observations of the NGC 4258 group find an overall satellite quenched fraction of $\sim 50\%$ at $< 10^8$ $M_\odot$, including some objects with blue rest-frame colors likely to be satellites at stellar masses equal to Fornax and Leo I (Spencer, Loebman & Yoachim, 2014). Such observations, however, must be interpreted under the caveat that there is significantly greater uncertainty in determining which objects are satellites and which objects belong to the field in systems beyond the Local Group, where line-of-sight distances are more poorly constrained. Given that the field is highly dominated by star-forming systems at these masses, sample contamination will strongly bias results to systematically lower quenched fractions.

Altogether, the observational results from our work and others, which constrain the relative impact of self-quenching and environmental quenching at low masses, present a challenge

to models of galaxy evolution. Previous studies at higher stellar mass ($> 10^{9.5}$ $M_\odot$) find that modern semi-analytic models overpredict the fraction of quenched satellites (Weinmann et al., 2010; Wang & White, 2012). In particular, Kimm et al. (2009) find that while semi-analytic models overestimate the number density of quenched satellites at high masses, the same models are generally able to reproduce the observed quenched fraction in the field, suggesting that conflicts with observations are largely driven by overprediction of the the impact of environmental quenching (i.e. overestimation of the satellite quenching efficiency) at $> 10^{9.5}$ $M_\odot$.

Applying our sample selection criteria to the Guo et al. (2013) semi-analytic model yields a satellite quenched fraction significantly elevated with respect to observations at $< 10^{9.5}$ $M_\odot$.[4] However, the increased number of quenched satellites in the model is largely driven by a corresponding overprediction of quenched field systems. At stellar masses $> 10^{9.5}$ $M_\odot$, where our data overlap with that of Kimm et al. (2009), we similarly find that the model reproduces the field quenched fraction, while at lower masses the Guo et al. (2013) SAM yields an excessive number of quenched centrals. Thus, while the model does overpredict the efficiency of satellite quenching (i.e. $f_{convert}$) at low masses, the effect is subdominant. Below 10^9 $M_\odot$, the observed conversion fractions agree with those in the semi-analytic model to within 50%, while the model overpredicts the quenched fraction of field objects by roughly a factor of 5. Combined with the results of Kimm et al. (2009), we find that modern models primarily fail to accurately describe the physics of feedback (i.e. self-quenching) within low-mass galaxies and the environmental quenching of high-mass galaxies.

[4]In the SAM, we consider an object "quenched" if it has a SSFR less than 10^{-11} yr^{-1}.

4.5 Conclusions

In this work, we study the quenching of satellite galaxies in Milky Way-like systems, primarily focusing on the dependence of quenching efficiency on both satellite and host mass. Our principal results are as follows:

- The efficiency of satellite quenching is largely independent of satellite mass over roughly three orders of magnitude in stellar mass, 10^8 $M_\odot < M_\star < 10^{11}$ $M_\odot$. Comparison to the Local Group suggests that satellite quenching efficiency may be significantly greater at yet lower stellar masses ($< 10^8$ $M_\odot$), perhaps indicating a critical mass for satellite quenching.

- Satellite quenching efficiency is well correlated with host halo mass, such that satellites of more massive halos are more likely to be quenched. A model in which satellite quenching only occurs in halos above a given critical halo mass is consistent with (*i*) the observed increase in satellite quenching in systems with two massive satellites (i.e. more massive halos) relative to those with one massive satellite in addition to (*ii*) the higher incidence of quenched satellites around passive hosts relative to their star-forming counterparts with one massive satellite.

- Discrepancies between the observed quenched fractions of low-mass ($< 10^{9.5}$ $M_\odot$) field and satellite galaxies and the predictions of the Guo et al. (2013) semi-analytic model are primarily driven by overly-effective internal processes (i.e. feedback or self-quenching mechanisms) that yield an overabundance of quenched field systems in the models. While the SAM overpredicts the efficiency of satellite quenching at low masses, the excess number of quenched satellites in the model is largely a product of the over-abundance of quenched field systems. In contrast, at higher masses ($> 10^{9.5}$ $M_\odot$), SAMs are generally able to reproduce the star-forming properties of field (or central) galaxies,

while they instead fail to accurately model the environmental quenching mechanisms, thereby overpredicting the number of quenched satellites at $> 10^{9.5}$ $M_{\odot}$.

Chapter 5

Using Machine Learning to Probe Low Mass Satellites

We use a novel machine learning method to analyze faint photometric satellite galaxies in the Canada-France-Hawaii Legacy Survey (CFHTLS). We use both a neural network-based and a decision tree-based algorithm, using training data from the Sloan Digital Sky Survey (SDSS) and the Galaxy and Mass Assembly survey (GAMA). Our method allows us to explore star formation in satellites down to $M_\star = 10^7 M_\odot$. We find that satellite quenching increases markedly with decreasing stellar mass, from $\sim 35\%$ at high masses, to $\sim 80\%$ at low masses. Our high mass result agrees with previous studies, while our low mass result agrees with observations of Local Group satellites, suggesting that the increased in quenched fraction at low stellar mass seen in local satellites is a generic feature of satellite galaxy evolution.

5.1 Introduction

Local Group (LG) observations serve, in some sense, as a Cosmic Rosetta Stone. Observations of nearby galaxies provide great detail into their formation histories (Hodge, 1989; Mateo, 1998; Tolstoy, Hill & Tosi, 2009; McConnachie, 2012; Ibata et al., 2013; Weisz et al., 2014a,b), providing a lens through which cosmology and galactic evolution in the early Universe can be understood (Bullock, Kravtsov & Weinberg, 2000; Freeman & Bland-Hawthorn, 2002; Ricotti & Gnedin, 2005; Bovill & Ricotti, 2011; Brown et al., 2012). However, while Boylan-Kolchin et al. (2016) argues that LG observations represent a reasonable volume of the Universe at high redshift, it remains to be shown that the local population of dwarf is suitably representative of similar objects elsewhere in the Universe.

Studies of the satellite systems of Milky Way analogs have revealed much about how these hosts and their satellites co-evolve (Weinmann et al., 2006; De Lucia et al., 2012; Kauffmann et al., 2013; Wetzel et al., 2013a; Wheeler et al., 2014; Phillips et al., 2014b), frequently showing that such galaxies quench their most massive satellites inefficiently over long timescales, in concert with what is seem in the brightest LG satellites. However, large scale spectroscopic surveys can not probe satellites below $10^{8.5}M_{\odot}$ in stellar mass (Phillips et al., 2015).

Photometric studies, on the other hand, can provide some comparisons to low-mass LG systems, particularly in their mass and phase-space distributions (Nierenberg et al., 2011; Strigari & Wechsler, 2012), but do not replicate the detailed star formation information obtained in spectroscopic studies. This is important, as studies of LG objects indicate quenching is substantially more efficient at low masses than at high masses (Fillingham et al., 2015, Fillingham16 in prep.). However, this strong dependence in quenching efficiency with stellar mass around $M_{\star} \sim 10^{8.5}M_{\odot}$ is based only on local observations. Confirming that this feature of LG satellites is a general property of satellite galaxy evolution is a crucial step in understanding how the LG fits into the larger picture of galaxy evolution.

In this work, we present a novel methodology for extending spectroscopic studies to lower masses than they are capable of directly probing through machine learning. The paper is organized as follows: in Section 5.2, we describe the observational data we use, in Section 5.3, we describe our procedure for identifying and classifying satellites, in Section 5.4, we give our results, and in Section 5.5, we discuss the implications of our findings.

5.2 Observations

The observational data used in this study is compiled from several different surveys:

- The *training set* was drawn from the Sloan Digital Sky Survey (SDSS), using data from value-added galactic catalogs of Blanton et al. (2005b) and **?**. These data were used to assign each point of four dimensional color-color space a value that maps to the probability that a galaxy residing at that point in parameter space is quenched.

  From the SDSS, we select for the training sample galaxies of stellar mass $7 < \log M_\star$. Galaxies are labelled star-forming or quenched based on where they lie in SFR-stellar mass space. If they lie above the line

$$\mathrm{SFR} = -0.7 \times \log M_\star - 7.7 \tag{5.1}$$

  they are considered star forming, otherwise they are taken to be quenched. Star formation rates are taken from the catalog, and are derived from emission lines and the D4000 spectral index. For the purposes of this study, it is critical that the star formation rates used in this sample are derived from galactic spectra, as we will be seeking to link spectroscopic properties to ones derived from photometric observations of faint objects.

Additional training set data was drawn from the Galaxy and Mass Assembly (GAMA) survey, a spectroscopic survey of ~ 300000 galaxies covering $\sim 286\text{deg}^2$ of equatorial sky. Spectra were taken AAOmega multi-object spectrograph on the Anglo-Australian telescope. The survey is complete down to $r < 19.8$, making it 2 magnitudes deeper than SDSS in spectroscopy. Since we are primarily interested in faint objects, and our algorithm weighs low-mass objects more heavily than high-mass objects, this added depth is beneficial.

- The *test set* is drawn from the Canada France Hawaii Telescope Legacy Survey (CFHTLS), a photometric survey comprised of a "wide" component and a "deep" component. For the purposes of this study, we consider only three our of the four"wide" fields, corresponding to those with spectroscopic overlap, fields W1 (72 square degrees), W3 (49 square degrees), and W4 (25 square degrees). The chosen fields overlap with the NASA-Sloan Atlas, which combines Galaxy Evolution Explorer (GALEX) photometry with SLOAN data; this overlap allows us to select photometric satellites of spectroscopically confirmed hosts.

  In selecting our test set, we apply a number of quality cuts to CFHTLS objects, most importantly requiring they be flagged as galaxies by SEXTRACTOR. At this point, we apply no cut on distance to nearby object; below we will discuss how our test set is divided into photometric satellites and background objects. The photometric satellites form our primary scientific sample; for this reason it is important to understand our completeness, both in an absolute sense and the extent to which completeness depends on galaxy properties, particularly color.

  To evaluate how complete our test set is, we first examine how complete our training set is. Figure 5.2 illustrates this; the red and blue points represent the passive and star-forming galaxies in our training sample, respectively. The black dashed line represents the region of parameter space where the training set is complete. As a fiducial measure,

we transpose the completeness curve to fainter magnitudes to bring it into agreement with the Next Generation Virgo Cluster Survey (NGVS, see Ferrarese et al. (2012)), a study that used the MegaCam instrument to observe Virgo cluster objects (i.e. objects with known redshift). We then adjusted our completeness limit to account for the difference in depth between the NGVS and the CFHTLS, and adjusted again to account for the differential completeness between MegaCam bands. This ensures that a galaxy observed in, e.g., the r band would not be unobserved in the i band.

These adjustments give us a final completeness magnitude of $r \sim 19.7$, two magnitudes fainter than SDSS. Using our assumed relationship between magnitude and mass, we can compute the redshifts our study is complete to as a function of mass; we are able to observe galaxies of mass $10^9 M_\star$ out to z = 0.05, of which there are 75 hosts meeting our criteria in our survey area; galaxies of mass $10^8 M_\star$ out to z = 0.027, giving us 26 hosts; and galaxies of mass $10^7 M_\star$ out to z = 0.013, giving us just 4 hosts. We would be able to track galaxies down to satellites of mass $10^6 M_\star$ to z = 0.006, however there are no such nearby hosts in our sample. These critical redshifts are noted in Figure 5.2.

5.3 Satellite Characterization

The procedure of assigning dwarf galaxies to satellite systems can be broken down into two parts: identifying the families of satellite galaxies, and characterizing their star formation. In this section, we discuss each of these in turn.

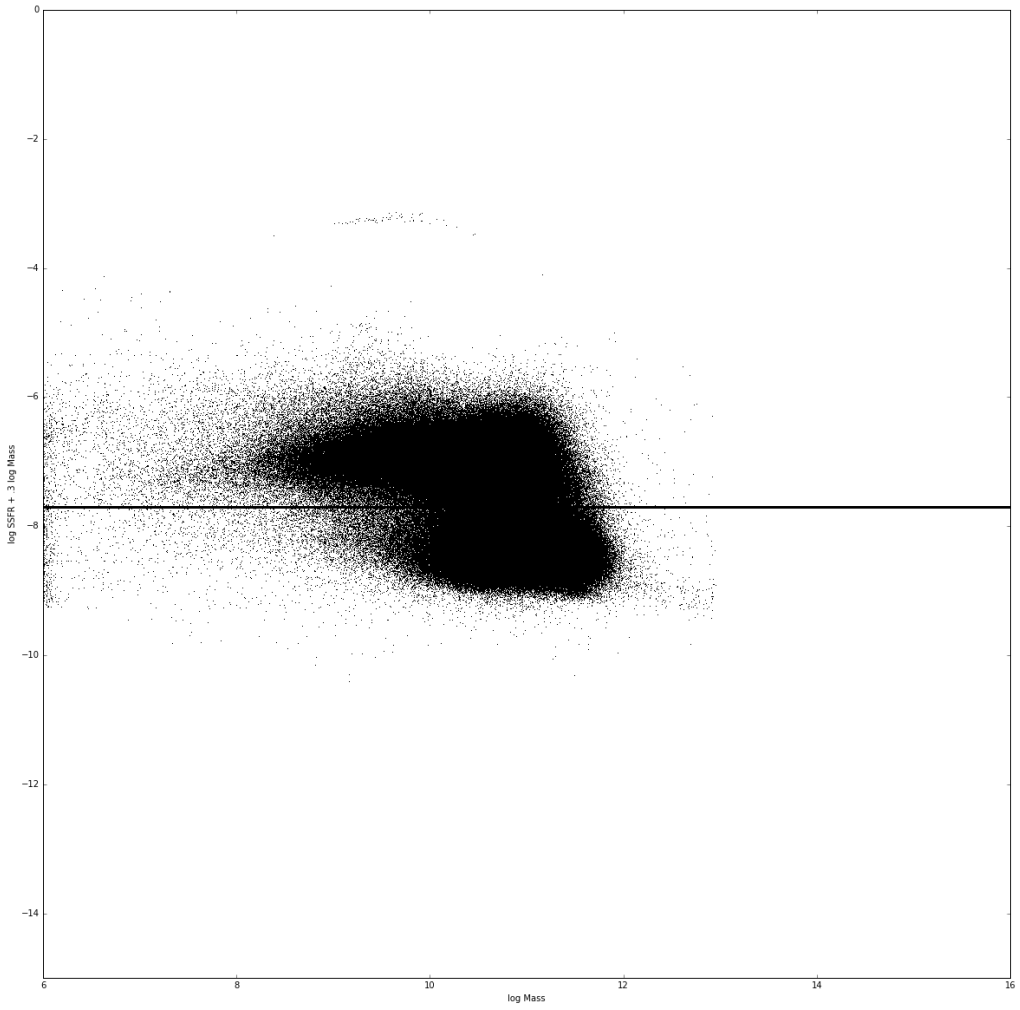

Figure 5.1: Modified specific star formation rate (log SFR - .7 log Stellar Mass) plotted against stellar mass for objects in the SDSS training sample. The horizontal line is the dividing line between star forming and quenched galaxies.

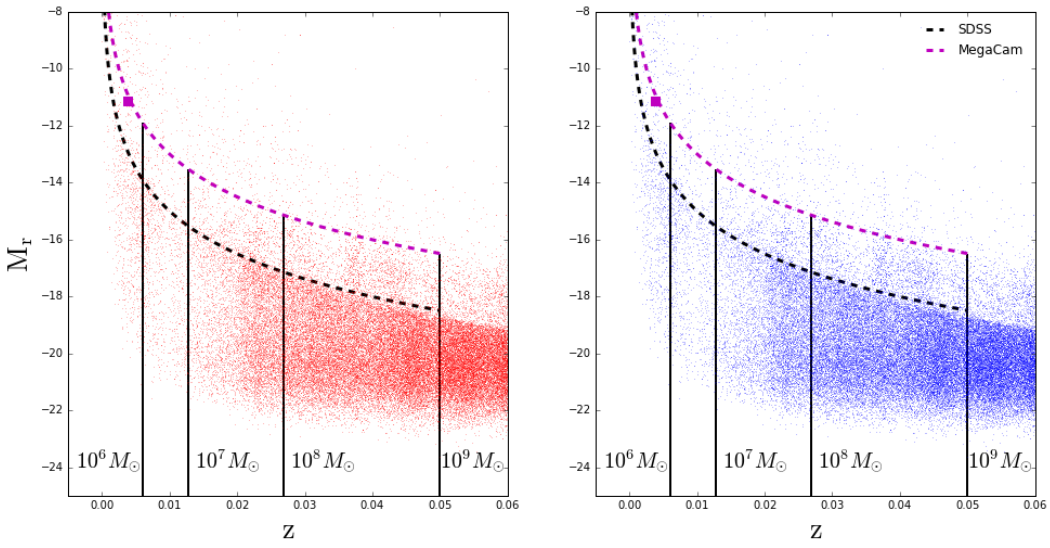

Figure 5.2: Absolute r band magnitude plotted against redshift for galaxies in the training set. Completeness curves are shown for both the training set (black curve) and the test set (magenta curve). The magenta square shows the test set completeness derived from Ferrarese et al. (2012) (see text for discussion). Vertical black lines indicate the redshifts at which the test set is complete to the labeled mass.

5.3.1 Satellite Finding

The first step in obtaining satellite galaxy samples is host selection. As our goal is to select hosts similar to the Milky Way, we select hosts from the NASA-Sloan Atlas between $10^{10.4}$M$_\star$ and 10^{11}M$_\star$ in stellar mass. Hosts are retained only if they fall on our three CFHTLS fields. A redshift cut is then applied to the hosts so that only hosts closer than z=.04 end up in the final sample (see previous section). We then select all objects within 150 projected kpc to be photometric satellites. To account for edge effects, we draw a circle of radius 150 kpc around each galaxy and calculate the ratio of the area of a circle that lies within the field to the full area of the circle, calling this ratio $F_{\rm area}$. Hosts with $F_{\rm area}$ lower than 0.75 are discarded; for the remainder, $1/F_{\rm area}$ serves a correction factor for the missing area lying outside of the field.

Once satellites are identified, we assign a mass to them by assuming that the mass of each

satellite is a power law function of its apparent r band magnitude, where the log-space intercept of the function is set by redshift. Our investigation confirms that, in the training set, mass is indeed well-described by a power law in r band magnitude. By interpolating the redshift-intercept relationship, we can define a power law relation between stellar mass and r magnitude for satellites of any redshift, where we assume the redshift of each satellite is approximately equal to that of its host.

The fundamental challenge of assigning satellite group membership to photometric objects is that it is not known where they lie in velocity space, and thus whether they may be associated with a given host. There are two ways of addressing this issue: making use of photometric redshifts derived from SED fitting, or by systematically subtracting out an assumed background population based on observations of "blank" fields. It is this second method that we make use of. For each of the three CFHTLS fields we use in this study, we take use 500 random "pencil beam" pointings within the field to compute the on-sky background density. We use the 40th percentile two-dimensional density to account for the skewedness in the distribution of densities due to random pointings in the direction of low-z, i.e. foreground, overdensities, as the background density for each.

5.3.2 Star Formation Modeling

Our modeling of star formation is motivated by a desire to classify photometric satellites into star formation categories based on their position in color-color parameter space in a data-driven way. We take a machine-learning approach to the problem using two algorithms, one built around decision trees and one which uses a neural network. Both of these are calibrated on the training set, then applied to the test set. We discuss our pair of algorithms further in the following subsections.

Decision Tree Algorithm

Our decision tree (DT) algorithm uses boolean logic to divide color-color space into regions that are labeled as either star forming or quenched based on how the training set populates the space. The code used to run the algorithm is built around the SKLEARN package DecisionTreeClassifier. Our procedure involves first taking the galaxies in the SDSS sample below a mass of $10^9 M_\odot$. We divide this subsample in half, thereby ensuring we are not testing the algorithm on the data points used to train it, eliminating redundancy in evaluating the algorithm. We take one of the halves and train 100 different random decision trees on it. Each iteration of the decision tree selects, with replacement, 15000 galaxies on which to train. The ensemble of these trees creates a "random forest" algorithm which takes as input a set of four features and gives as output a number between 0 and 1, which corresponds to the fraction of trees that return a "quenched" label for the set of features. We then evaluate the effectiveness of the algorithm using the data not used in the training procedure.

In evaluating the algorithm, the two values we are concern with are the "precision" (often called purity) and "recall" (often called completeness). In the case of our algorithm, precision is the probability of a galaxy identified as quenched by the algorithm is truly quenched. Recall is the probability that a truly quenched galaxy is identified as quenched. Since the SDSS gives us spectroscopic information on the galaxies in our training set, we may compute both precision and recall directly. Furthermore, in the case where objects have binary labels (i.e. "quenched" and "star-forming"), it can be shown that

$$F_q = f_q \times (P/R) \tag{5.2}$$

where F_q is the true quenched fraction, f_q is the quenched fraction determined by the algorithm, P is the precision, and R is the recall. There are two things to note from this

equation: Regardless of the values of P and R, if $P = R$ then the measured quenched fraction is the same as the true quenched fraction. More importantly, regardless of the values of P and R, if those values are known, F_q may be determined from f_q using the value of P/R, which we will refer to as the "correction factor," or C. Note that this can be easily generalized to a scenario where F_q, f_q, P, and R and all functions of mass, enabling us to define a mass-dependent correction factor $C(\mathrm{M}_\star)$.

Neural Network Algorithm

For the NN method, we trained a series of 500 distinct neural network models on the training set from the 4 color bands[1]. All the neural network models were built in Keras Chollet (2015) with a Theano Theano Development Team (2016) backend and trained on GPUs.

We used Spearmint, a Bayesian optimization frameworkSnoek, Larochelle & Adams (2012), to conduct an extensive hyperparameter search over neural network architectures and optimization hyperparameters. The parameters considered were: the number of hidden layers in the feed-forward neural network (range: 2 to 10), number of hidden units per layer (range: 4 to 400), the dropout probability for each layer Srivastava et al. (2014) (range: 0.0 to 0.8), the batch size for each stochastic update (range: 16 to 512), the base-10 log of the stochastic gradient learning rate (range: -5.0 to 0.0), the per-epoch learning rate decay factor (range: 0.90 to 1.0), the momentum coefficient for the stochastic gradient update rule Sutskever et al. (2013) (range: 0.0 - 1.0), whether to use Nesterov Nesterov (1983) or classical momentum updates, and whether to replace the stochastic momentum update rule the with AdaDelta update ruleZeiler (2012).

The best of the 500 models—as judged by the area under the receiver operating characteristic curve (AUC) on the validation set—was a model with 5 hidden layers, 388 neurons in each

[1]Mass information was not used to train the neural network.

layer, dropout probability of 0.1232 at each hidden layer, batch size of 86, and AdaDelta updates with a learning rate multiplier of 1.0. The model was trained on the training data until reaching a maximum AUC of 0.902 on the validation data after 6 epochs. This model achieved an AUC value of 0.900 on the hitherto-unseen test data[2].

Since the mass distribution for the objects in the training set may differ significantly from the mass distribution of the true objects of interest, we next conducted a sensitivity analysis to see how well this model would perform on data drawn from a different underlying mass distribution. For this purpose, we partitioned the test set into objects with (base 10) log masses in the ranges [6, 8), [8, 10), and [10, 12). We then sampled a series of datapoints (with replacement) from each of these bins in the test set to create new test sets with different mass distributions. If the number of datapoints sampled from the bins [6, 8), [8, 10), and [10, 12) are labeled n_6, n_8, and n_{10} respectively, we sampled datapoints such that $n_6/n_8 = n_8/n_{10} \equiv \gamma$ for γ values of 0.5, 1.0, 1.5, 2.0, 2.5, and 3.0. Finally, we used the model which had been trained on the un-resampled training data to make predictions for each of the resampled test sets.The test set AUC is slightly lower for the resampled datasets, but remains above 0.87 even for datasets with grossly different mass distribution than the training data.

5.3.3 Algorithm Efficacy

Previously, we have discussed the need to define a correction factor, which is equal to the precision of an algorithm divided by its recall. In Figure 5.3 we show the correction factor derived from our two algorithms plotted against stellar mass. The red squares represent points derived the DT algorithm, while the blue circles represent the same for the NN algorithm. For each algorithm, we assign a fit to serve as our modelled $C(\mathrm{M}_\star)$. The NN algorithm is well described by a linear fit; the DT algorithm, on the other hand, shows

[2]AUC on either random or modal predictions (i.e.: predicting 0 for all datapoints) would be 0.5, and perfect predictions would earn an AUC of 1.0.

inflection points arising from our choice of cut on the training set. To capture this behavior, we use a spline fit. We emphasize that, in principal, $C(\mathrm{M}_\star)$ could take any value and the true quenched fraction would be recoverable from the data; however, we tuned our algorithm such that $C(\mathrm{M}_\star)$ was near unity over the dynamic range we consider. In practice, choosing $C(\mathrm{M}_\star)$ negates the extent to which systematic errors will be amplified in applying the correction to the data.

5.4 Results

In Figure 5.4 we show the fraction of satellite galaxies that are quenched, as determined by both of our algorithms as a function of satellite mass. The DT algorithm is shown as black squares, while the NN algorithm is shown as blue circles. Both algorithms show a trend of decreasing quenched fraction with increasing stellar mass. The errors reported in Figure 5.4 are binomial errors on the satellite quenched fraction; we can estimate the errors associated with the algorithm itself to be $\sigma =\sim 0.1$, which would bring the data from both algorithms into relative agreement. With these errors in place, the trend is still seen.

In comparing Figures 5.3 and 5.4, there are two items of note: The trend noted above seems to be driven by the functional form of $\mathrm{C_m}$, and in particular, the difference between the DT and NN points at high mass is driven by the difference in C_m at high and low mass. The second point is easily explained by the difference in accounting for the mass dependence of the training sets for the respective algorithms: The DT algorithm used a "top hat" mass selection function, where only low mass objects were represented in the training set, whereas the NN method used a power law selection function, where low-mass objects were oversampled, but high-mass objects were still accounted for. Thus it follows that the NN algorithm out-performed the DT algorithm at high mass.

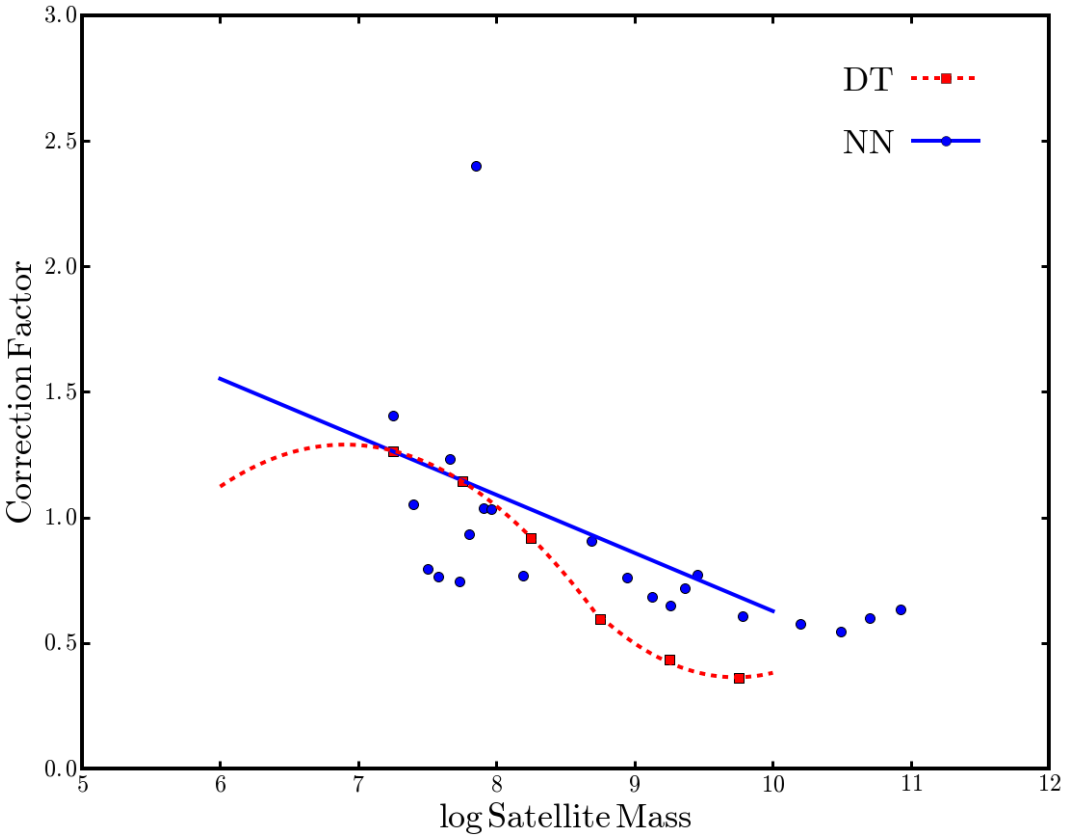

Figure 5.3: The correction factor, defined as the ratio of precision to recall, plotted against stellar mass for each of our two algorithms. The red squares and blue circles represent the correction factor derived from narrowly binning the data in the DT and NN algorithms, respectively. From these data, an analytic correction function is estimated as a function of stellar mass; a spline fit was used for the DT algorithm while a simple linear fit describes the NN data well. Note that the algorithms were tuned such that correction factors near unity were produced.

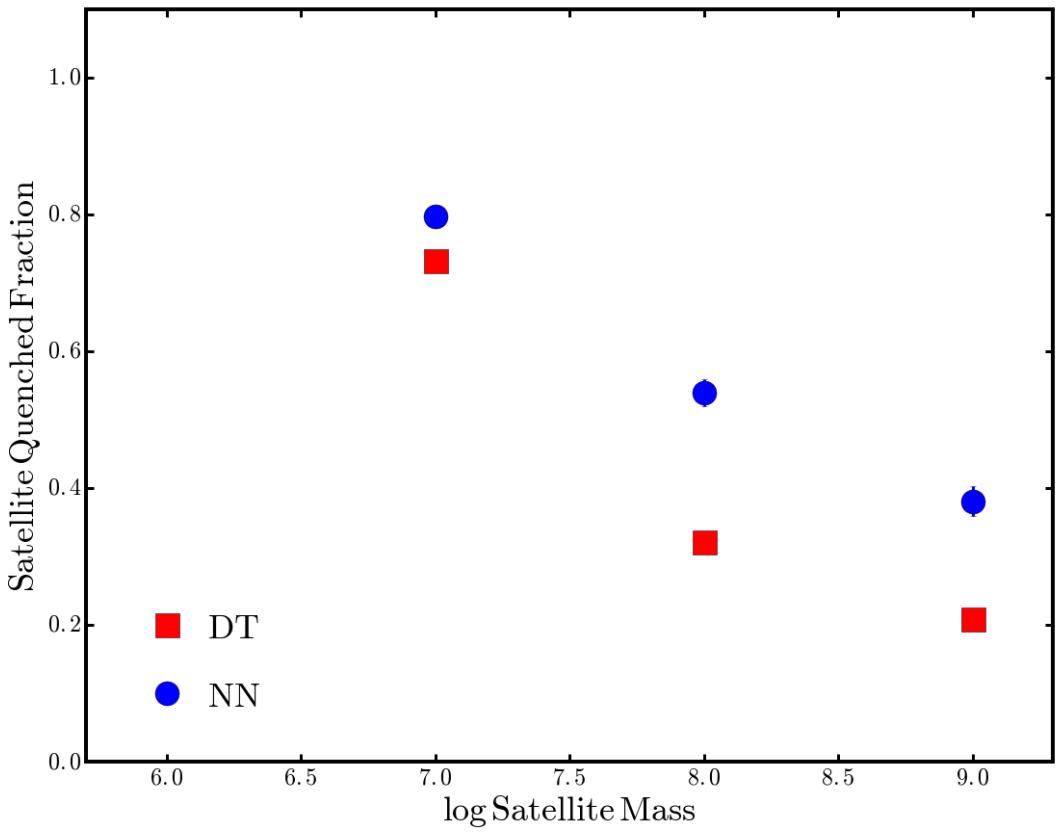

Figure 5.4: The quenched fraction of satellite galaxies as a function of satellite mass, derived from our two algorithms. Red squares denote quenched fractions measured using the *decision tree* algorithm, while blue circles denote those using the *neural network* algorithm (see §5.3). Our results indicate an elevated quenched fraction at low magnitudes, in broad agreement between the two algorithms.

The first point is more subtle; if our methodology failed entirely, the algorithms would, in essence, randomly assign "quenched" or "star-forming" labels to galaxies, and the dependence of quenched fraction on stellar mass would be entirely driven by the form of C_m. While at first this may appear to be a concern, there are reasons that this is not the case. Firstly, our high-mass data, in particular from the NN algorithm, is in good agreement with previous results (see 5.5); this serves as confirmation that our methodology "matches up" with known results at high mass. Furthermore, the functional form of the algorithm applied to the evaluation portion of the training set irrespective of environment shows a strong dependence on stellar mass, with low mass galaxies being quenched at a rate of about 20%, highly different than our test set result, suggesting we are indeed capturing an environmental effect.

5.5 Discussion

Our work extends previous results examining the quenched fraction of satellites in the $z \sim 0$ Universe. In Figure 5.5 we compare our results to this work, also including HI observations of Local Group satellites. Taking all the data together, the data paints a coherent picture of satellite quenching over a substantial dynamical mass range: high mass galaxies ($M_\star >\sim 10^9 M_\odot$) are quenched fairly inefficiently ($f_{quench} \sim .4$). This has been seen in previous studies, but is confirmed by our work. At intermediate masses ($M_\star \sim 10^8 M_\odot$), galaxies undergo a transition, where they begin to become quenched with high efficiency. At low masses ($M_\star <\sim 10^8 M_\odot$), galaxies are nearly uniformly quenched.

The mass range around $10^8 M_\star$ is particularly interesting, as it marks the transition between the low-quenched-fraction regime at high mass, and the high-quenched-fraction regime at low mass. Indeed, this is where environmental quenching begins to play a markedly more significant role, as the positive trend in quenched fraction with mass at high masses can be explained by non-environmental means. Our study identifies this mass regime is a near-

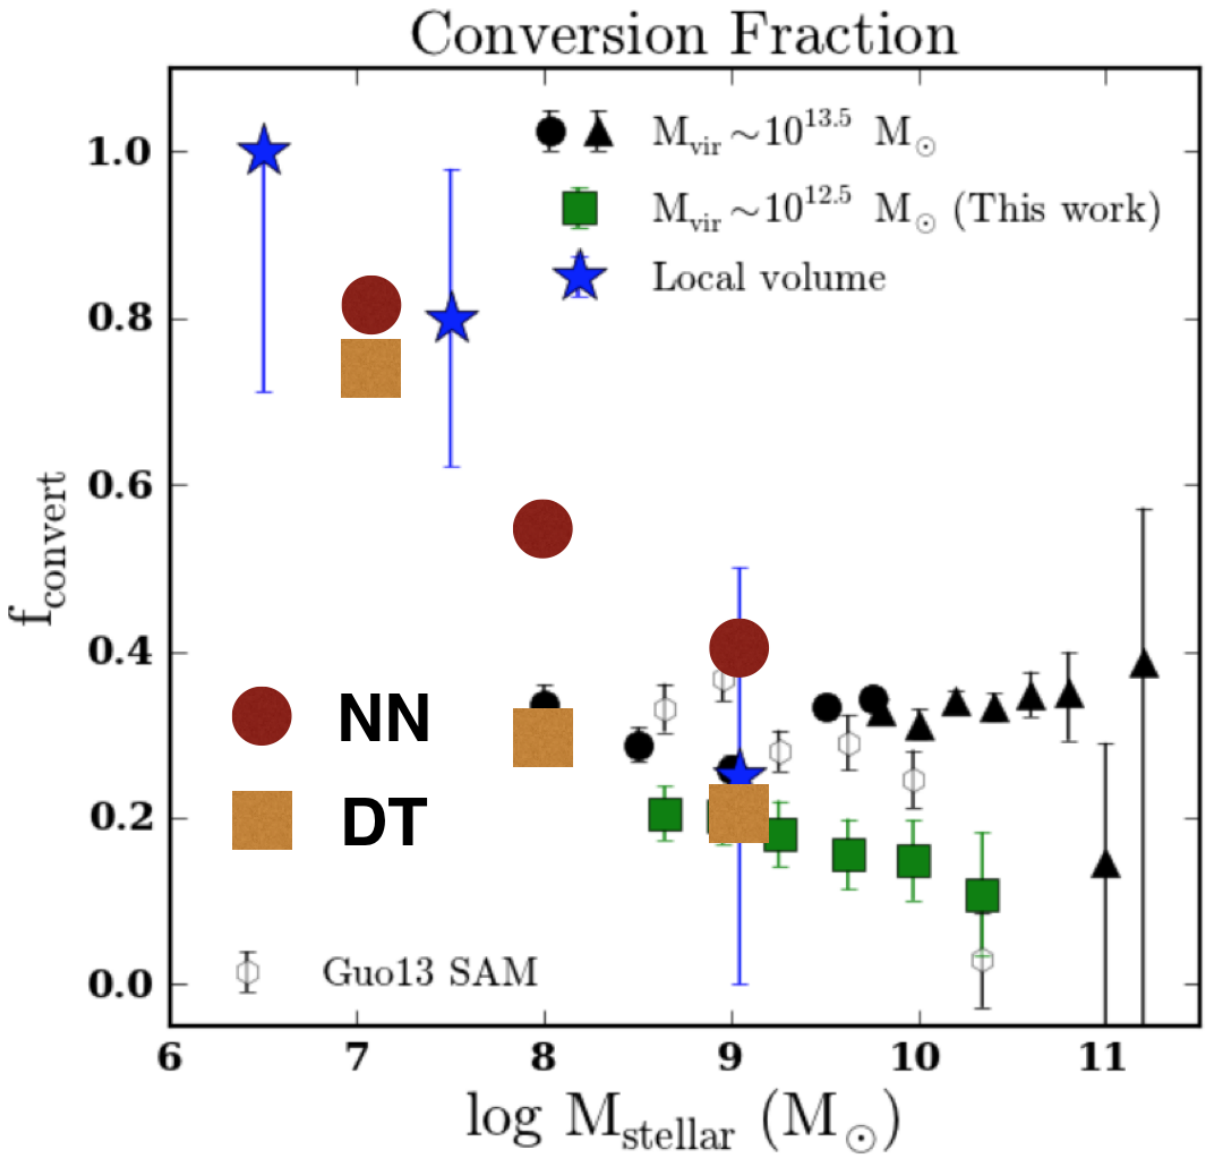

Figure 5.5: The environmentally quenched fraction of satellite galaxies studied here as a function of satellite mass, plotted along with previous work.

universal feature of satellite galaxy evolution, but lacks the resolution to identify exactly where the upturn in quenched fraction. Making this identification, and investigating how it depends on properties of the host, in particular properties that would affect the CGM, such as host mass, would provide valuable insight into the physical mechanisms at play.

On the other hand, our result serves as a clear indication of the validity of the assumption that the satellite systems in the Local Group are not strongly biased in their star formation histories. This is an important point, as local observations can provide strong leverage on addressing questions of high-Z galaxy formation and evolution, provided that they are free of bias. Our study provides a novel to addressing this issue, which already provides tantalizing hints at the universality of quenching in low-mass satellites, and which paves the way for future studies using deeper observations, such as the Large Synoptic Survey Telescope (LSST) survey. We estimate that an identical study carried out using LSST would provide data from more than 500 hosts reaching satellites of mass 10^6 and below, and tens of thousands of systems at which the critical mass window ($10^8 M_\star$) may be investigated.

Bibliography

Abazajian K. N. et al., 2009a, ApJS, 182, 543

Abazajian K. N. et al., 2009b, ApJS, 182, 543

Azzaro M., Patiri S. G., Prada F., Zentner A. R., 2007, MNRAS, 376, L43

Bailin J. et al., 2005, ApJ, 627, L17

Bailin J., Power C., Norberg P., Zaritsky D., Gibson B. K., 2008, MNRAS, 390, 1133

Baldry I. K., Balogh M. L., Bower R. G., Glazebrook K., Nichol R. C., Bamford S. P., Budavari T., 2006, MNRAS, 373, 469

Baldry I. K., Glazebrook K., Brinkmann J., Ivezić Ž., Lupton R. H., Nichol R. C., Szalay A. S., 2004, ApJ, 600, 681

Balogh M. L., Baldry I. K., Nichol R., Miller C., Bower R., Glazebrook K., 2004, ApJ, 615, L101

Balogh M. L., Morris S. L., 2000, MNRAS, 318, 703

Balogh M. L., Navarro J. F., Morris S. L., 2000, ApJ, 540, 113

Bamford S. P. et al., 2009, MNRAS, 393, 1324

Barton E. J., Arnold J. A., Zentner A. R., Bullock J. S., Wechsler R. H., 2007, ApJ, 671, 1538

Behroozi P. S., Conroy C., Wechsler R. H., 2010, ApJ, 717, 379

Behroozi P. S., Wechsler R. H., Conroy C., 2012, ArXiv e-prints

Bekki K., 2009, MNRAS, 399, 2221

Bell E. F., McIntosh D. H., Katz N., Weinberg M. D., 2003, ApJS, 149, 289

Bell E. F. et al., 2012, ApJ, 753, 167

Berrier J. C., Bullock J. S., Barton E. J., Guenther H. D., Zentner A. R., Wechsler R. H., 2006, ApJ, 652, 56

Bigiel F., Leroy A., Walter F., Brinks E., de Blok W. J. G., Madore B., Thornley M. D., 2008, AJ, 136, 2846

Binney J., 1977, ApJ, 215, 483

Birnboim Y., Dekel A., 2003, MNRAS, 345, 349

Blanton M. R. et al., 2003, AJ, 125, 2348

Blanton M. R., Roweis S., 2007, AJ, 133, 734

Blanton M. R. et al., 2005a, AJ, 129, 2562

Blanton M. R. et al., 2005b, AJ, 129, 2562

Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J., 1984, Nature, 311, 517

Boselli A., Cortese L., Boquien M., Boissier S., Catinella B., Lagos C., Saintonge A., 2014, A&A, 564, A66

Bovill M. S., Ricotti M., 2011, ApJ, 741, 17

Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2011, MNRAS, 415, L40

Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2012, MNRAS, 422, 1203

Boylan-Kolchin M., Bullock J. S., Sohn S. T., Besla G., van der Marel R. P., 2013, ApJ, 768, 140

Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, MNRAS, 398, 1150

Boylan-Kolchin M., Weisz D. R., Bullock J. S., Cooper M. C., 2016, ArXiv e-prints

Brainerd T. G., 2005, ApJ, 628, L101

Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, MNRAS, 351, 1151

Brown T. M. et al., 2012, ApJ, 753, L21

Bryan G. L., Norman M. L., 1998, ApJ, 495, 80

Buck T., Macci'o A. V., Dutton A. A., 2015, ArXiv e-prints

Bullock J. S., Kravtsov A. V., Weinberg D. H., 2000, ApJ, 539, 517

Cautun M., Bose S., Frenk C. S., Guo Q., Han J., Hellwing W. A., Sawala T., Wang W., 2015a, MNRAS, 452, 3838

Cautun M., Wang W., Frenk C. S., Sawala T., 2015b, MNRAS, 449, 2576

Cheung E. et al., 2012, ApJ, 760, 131

Chollet F., 2015, keras. https://github.com/fchollet/keras

Collins M. L. M. et al., 2013, ApJ, 768, 172

Conn A. R. et al., 2013, ApJ, 766, 120

Conroy C. et al., 2007, ApJ, 654, 153

Conroy C., Wechsler R. H., Kravtsov A. V., 2006, ApJ, 647, 201

Cooper M. C. et al., 2006, MNRAS, 370, 198

Cox T. J., Dutta S. N., Di Matteo T., Hernquist L., Hopkins P. F., Robertson B., Springel V., 2006, ApJ, 650, 791

Crain R. A. et al., 2009, MNRAS, 399, 1773

Darg D. W. et al., 2010, MNRAS, 401, 1552

Davé R., Oppenheimer B. D., Finlator K., 2011, MNRAS, 415, 11

Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371

Davis M., Geller M. J., 1976, ApJ, 208, 13

De Lucia G., Weinmann S., Poggianti B. M., Aragón-Salamanca A., Zaritsky D., 2012, MNRAS, 423, 1277

Deason A. J. et al., 2012, MNRAS, 425, 2840

Deason A. J. et al., 2011, MNRAS, 415, 2607

Dressler A., 1980, ApJ, 236, 351

Driver S. P. et al., 2011, MNRAS, 413, 971

Edman J. P., Barton E. J., Bullock J. S., 2012, MNRAS, 424, 1454

Elbaz D. et al., 2007, A&A, 468, 33

Fakhouri O., Ma C.-P., 2008, MNRAS, 386, 577

Faltenbacher A., Li C., Mao S., van den Bosch F. C., Yang X., Jing Y. P., Pasquali A., Mo H. J., 2007, ApJ, 662, L71

Farouki R., Shapiro S. L., 1981, ApJ, 243, 32

Ferrarese L. et al., 2012, ApJS, 200, 4

Fillingham S. P., Cooper M. C., Wheeler C., Garrison-Kimmel S., Boylan-Kolchin M., Bullock J. S., 2015, MNRAS, 454, 2039

Freeman K., Bland-Hawthorn J., 2002, ARA&A, 40, 487

Gabor J. M., Davé R., 2014, ArXiv e-prints

Garrison-Kimmel S., Boylan-Kolchin M., Bullock J. S., Kirby E. N., 2014, ArXiv e-prints

Geha M., Blanton M. R., Masjedi M., West A. A., 2006, ApJ, 653, 240

Geha M., Blanton M. R., Yan R., Tinker J. L., 2012, ApJ, 757, 85

Gunn J. E., Gott, III J. R., 1972, ApJ, 176, 1

Guo Q., White S., Angulo R. E., Henriques B., Lemson G., Boylan-Kolchin M., Thomas P., Short C., 2013, MNRAS, 428, 1351

Guo Q. et al., 2011, MNRAS, 413, 101

Guo Q., White S., Li C., Boylan-Kolchin M., 2010, MNRAS, 404, 1111

Hao J., Kubo J. M., Feldmann R., Annis J., Johnston D. E., Lin H., McKay T. A., 2011, ApJ, 740, 39

Hearin A. P., Watson D. F., 2013, ArXiv e-prints

Hearin A. P., Watson D. F., van den Bosch F. C., 2014, ArXiv e-prints

Hodge P., 1989, ARA&A, 27, 139

Hogg D. W. et al., 2004, ApJ, 601, L29

Holmberg E., 1969, Arkiv for Astronomi, 5, 305

Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Robertson B., Springel V., 2006, ApJS, 163, 1

Humphrey P. J., Buote D. A., Canizares C. R., Fabian A. C., Miller J. M., 2011, ApJ, 729, 53

Humphrey P. J., Buote D. A., O'Sullivan E., Ponman T. J., 2012, ApJ, 755, 166

Ibata N. G., Ibata R. A., Famaey B., Lewis G. F., 2014a, Nature, 511, 563

Ibata R. A., Famaey B., Lewis G. F., Ibata N. G., Martin N., 2014b, ArXiv e-prints

Ibata R. A., Ibata N. G., Lewis G. F., Martin N. F., Conn A., Elahi P., Arias V., Fernando N., 2014c, ApJ, 784, L6

Ibata R. A. et al., 2013, Nature, 493, 62

Jogee S., Scoville N., Kenney J. D. P., 2005, ApJ, 630, 837

Kaisin S. S., Karachentsev I. D., 2013, Astrophysics, 56, 305

Kang X., van den Bosch F. C., Yang X., Mao S., Mo H. J., Li C., Jing Y. P., 2007, MNRAS, 378, 1531

Karachentsev I., 1996, A&A, 305, 33

Kauffmann G. et al., 2003, MNRAS, 341, 54

Kauffmann G., Li C., Zhang W., Weinmann S., 2013, MNRAS, 430, 1447

Kawata D., Mulchaey J. S., 2008, ApJ, 672, L103

Kennicutt, Jr. R. C. et al., 2007, ApJ, 671, 333

Kereš D., Katz N., Fardal M., Davé R., Weinberg D. H., 2009, MNRAS, 395, 160

Kereš D., Katz N., Weinberg D. H., Davé R., 2005, MNRAS, 363, 2

Kimm T. et al., 2009, MNRAS, 394, 1131

Klypin A. A., Trujillo-Gomez S., Primack J., 2011, ApJ, 740, 102

Knebe A., Gill S. P. D., Gibson B. K., Lewis G. F., Ibata R. A., Dopita M. A., 2004, ApJ, 603, 7

Koch A., Grebel E. K., 2006, AJ, 131, 1405

Komatsu E. et al., 2011a, ApJS, 192, 18

Komatsu E. et al., 2011b, ApJS, 192, 18

Kormendy J., Kennicutt, Jr. R. C., 2004, ARA&A, 42, 603

Kroupa P. et al., 2010, A&A, 523, A32

Kroupa P., Theis C., Boily C. M., 2005, A&A, 431, 517

Kunkel W. E., Demers S., 1976, in Royal Greenwich Observatory Bulletins, Vol. 182, The Galaxy and the Local Group, Dickens R. J., Perry J. E., Smith F. G., King I. R., eds., p. 241

Larson R. B., Tinsley B. M., Caldwell C. N., 1980, ApJ, 237, 692

Leauthaud A. et al., 2012, ApJ, 744, 159

Leroy A. K., Walter F., Brinks E., Bigiel F., de Blok W. J. G., Madore B., Thornley M. D., 2008, AJ, 136, 2782

Lewis I. et al., 2002, MNRAS, 334, 673

Libeskind N. I., Cole S., Frenk C. S., Okamoto T., Jenkins A., 2007, MNRAS, 374, 16

Libeskind N. I., Frenk C. S., Cole S., Helly J. C., Jenkins A., Navarro J. F., Power C., 2005, MNRAS, 363, 146

Libeskind N. I., Hoffman Y., Tully R. B., Courtois H. M., Pomarede D., Gottloeber S., Steinmetz M., 2015, ArXiv e-prints

Libeskind N. I., Knebe A., Hoffman Y., Gottlöber S., Yepes G., Steinmetz M., 2011, MNRAS, 411, 1525

Lin L. et al., 2010, ApJ, 718, 1158

Lovell M. R., Eke V. R., Frenk C. S., Jenkins A., 2011, MNRAS, 413, 3013

Lynden-Bell D., 1976, MNRAS, 174, 695

MacArthur L. A., Courteau S., Holtzman J. A., 2003, ApJ, 582, 689

Mamon G. A., Sanchis T., Salvador-Solé E., Solanes J. M., 2004, A&A, 414, 445

Mandelbaum R., Seljak U., Kauffmann G., Hirata C. M., Brinkmann J., 2006, MNRAS, 368, 715

Martig M., Bournaud F., Teyssier R., Dekel A., 2009, ApJ, 707, 250

Mateo M. L., 1998, ARA&A, 36, 435

McConnachie A. W., 2012, AJ, 144, 4

McConnachie A. W., Irwin M. J., 2006, MNRAS, 365, 902

McGaugh S. S., 2012, AJ, 143, 40

McIntosh D. H., Guo Y., Hertzberg J., Katz N., Mo H. J., van den Bosch F. C., Yang X., 2008, MNRAS, 388, 1537

Metz M., Kroupa P., Jerjen H., 2007, MNRAS, 374, 1125

Metz M., Kroupa P., Jerjen H., 2009, MNRAS, 394, 2223

Metz M., Kroupa P., Libeskind N. I., 2008, ApJ, 680, 287

Miller S. H., Bundy K., Sullivan M., Ellis R. S., Treu T., 2011, ApJ, 741, 115

Miller S. H., Ellis R. S., Newman A. B., Benson A., 2014, ApJ, 782, 115

Moore B., Governato F., Quinn T., Stadel J., Lake G., 1998, ApJ, 499, L5

Moore B., Katz N., Lake G., Dressler A., Oemler A., 1996, Nature, 379, 613

Moster B. P., Naab T., White S. D. M., 2013, MNRAS, 428, 3121

Moster B. P., Somerville R. S., Maulbetsch C., van den Bosch F. C., Macciò A. V., Naab T., Oser L., 2010, ApJ, 710, 903

Mulchaey J. S., Jeltema T. E., 2010, ApJ, 715, L1

Nesterov Y., 1983in , pp. 372–376

Nierenberg A. M., Auger M. W., Treu T., Marshall P. J., Fassnacht C. D., 2011, ApJ, 731, 44

Noeske K. G. et al., 2007, ApJ, 660, L43

Oman K. A., Hudson M. J., Behroozi P. S., 2013, MNRAS, 431, 2307

O'Sullivan E., Forbes D. A., Ponman T. J., 2001, MNRAS, 328, 461

Pasquali A., Gallazzi A., Fontanot F., van den Bosch F. C., De Lucia G., Mo H. J., Yang X., 2010, MNRAS, 407, 937

Pawlowski M. S. et al., 2014, MNRAS, 442, 2362

Pawlowski M. S., Kroupa P., 2013, MNRAS, 435, 2116

Pawlowski M. S., Kroupa P., Angus G., de Boer K. S., Famaey B., Hensler G., 2012, MNRAS, 424, 80

Pawlowski M. S., Kroupa P., Jerjen H., 2013, MNRAS, 435, 1928

Pawlowski M. S., McGaugh S. S., 2014, ApJ, 789, L24

Pawlowski M. S., Pflamm-Altenburg J., Kroupa P., 2012, MNRAS, 423, 1109

Peñarrubia J., Kroupa P., Boily C. M., 2002, MNRAS, 333, 779

Peng Y.-j. et al., 2010, ApJ, 721, 193

Peng Y.-j., Lilly S. J., Renzini A., Carollo M., 2012, ApJ, 757, 4

Phillips J. I., Wheeler C., Boylan-Kolchin M., Bullock J. S., Cooper M. C., Tollerud E. J., 2014a, MNRAS, 437, 1930

Phillips J. I., Wheeler C., Boylan-Kolchin M., Bullock J. S., Cooper M. C., Tollerud E. J., 2014b, MNRAS, 437, 1930

Phillips J. I., Wheeler C., Cooper M. C., Boylan-Kolchin M., Bullock J. S., Tollerud E., 2015, MNRAS, 447, 698

Plionis M., Benoist C., Maurogordato S., Ferrari C., Basilakos S., 2003, ApJ, 594, 144

Postman M. et al., 2005, ApJ, 623, 721

Prescott M. et al., 2011, MNRAS, 417, 1374

Quilis V., Moore B., Bower R., 2000, Science, 288, 1617

Rees M. J., Ostriker J. P., 1977, MNRAS, 179, 541

Ricotti M., Gnedin N. Y., 2005, ApJ, 629, 259

Robotham A. S. G. et al., 2013, MNRAS, 431, 167

Saintonge A. et al., 2011, MNRAS, 415, 32

Sales L., Lambas D. G., 2004, MNRAS, 348, 1236

Sales L., Lambas D. G., 2009, MNRAS, 395, 1184

Salim S. et al., 2007, ApJS, 173, 267

Schiminovich D. et al., 2010, MNRAS, 408, 919

Sérsic J. L., 1968, Atlas de galaxias australes

Skillman E. D., Côté S., Miller B. W., 2003, AJ, 125, 593

Snoek J., Larochelle H., Adams R. P., 2012, in Advances in neural information processing systems, pp. 2951–2959

Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, MNRAS, 391, 481

Spencer M., Loebman S., Yoachim P., 2014, ArXiv e-prints

Springel V., Di Matteo T., Hernquist L., 2005, ApJ, 620, L79

Srivastava N., Hinton G., Krizhevsky A., Sutskever I., Salakhutdinov R., 2014, The Journal of Machine Learning Research, 15, 1929

Stewart K. R., Kaufmann T., Bullock J. S., Barton E. J., Maller A. H., Diemand J., Wadsley J., 2011a, ApJ, 735, L1

Stewart K. R., Kaufmann T., Bullock J. S., Barton E. J., Maller A. H., Diemand J., Wadsley J., 2011b, ApJ, 735, L1

Strateva I. et al., 2001, AJ, 122, 1861

Strigari L. E., Wechsler R. H., 2012, ApJ, 749, 75

Sun M., Jones C., Forman W., Vikhlinin A., Donahue M., Voit M., 2007, ApJ, 657, 197

Sutskever I., Martens J., Dahl G., Hinton G., 2013, in Proceedings of the 30th international conference on machine learning (ICML-13), pp. 1139–1147

Tempel E., Guo Q., Kipper R., Libeskind N. I., 2015, ArXiv e-prints

Theano Development Team, 2016, arXiv e-prints, abs/1605.02688

Tollerud E. J., Boylan-Kolchin M., Barton E. J., Bullock J. S., Trinh C. Q., 2011a, ApJ, 738, 102

Tollerud E. J., Boylan-Kolchin M., Barton E. J., Bullock J. S., Trinh C. Q., 2011b, ApJ, 738, 102

Tolstoy E., Hill V., Tosi M., 2009, ARA&A, 47, 371

Toomre A., 1977, in Evolution of Galaxies and Stellar Populations, Tinsley B. M., Larson D. Campbell R. B. G., eds., p. 401

Trinh C. Q., J. B. E., Bullock J. S., Zentner A. R., Wechsler R. H., 2013, ArXiv e-prints

Tully R. B., Libeskind N. I., Karachentsev I. D., Karachentseva V. E., Rizzi L., Shaya E. J., 2015, ArXiv e-prints

Tumlinson J. et al., 2011, Science, 334, 948

van den Bosch F. C., Aquino D., Yang X., Mo H. J., Pasquali A., McIntosh D. H., Weinmann S. M., Kang X., 2008a, MNRAS, 387, 79

van den Bosch F. C., Aquino D., Yang X., Mo H. J., Pasquali A., McIntosh D. H., Weinmann S. M., Kang X., 2008b, MNRAS, 387, 79

van den Bosch F. C., Lewis G. F., Lake G., Stadel J., 1999, ApJ, 515, 50

van der Marel R. P., Fardal M., Besla G., Beaton R. L., Sohn S. T., Anderson J., Brown T., Guhathakurta P., 2012, ApJ, 753, 8

Wang J., Frenk C. S., Cooper A. P., 2013, MNRAS, 429, 1502

Wang W., White S. D. M., 2012, MNRAS, 424, 2574

Weinmann S. M., Kauffmann G., von der Linden A., De Lucia G., 2010, MNRAS, 406, 2249

Weinmann S. M., Pasquali A., Oppenheimer B. D., Finlator K., Mendel J. T., Crain R. A., Macciò A. V., 2012, MNRAS, 426, 2797

Weinmann S. M., van den Bosch F. C., Yang X., Mo H. J., 2006, MNRAS, 366, 2

Weisz D. R., Dolphin A. E., Skillman E. D., Holtzman J., Gilbert K. M., Dalcanton J. J., Williams B. F., 2014a, ArXiv e-prints

Weisz D. R. et al., 2014b, ArXiv e-prints

West M. J., Blakeslee J. P., 2000, ApJ, 543, L27

Wetzel A. R., Tinker J. L., Conroy C., Bosch F. C. v. d., 2014, MNRAS, 439, 2687

Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2013a, MNRAS, 432, 336

Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2013b, ArXiv e-prints

Wheeler C., Phillips J. I., Cooper M. C., Boylan-Kolchin M., Bullock J. S., 2014, ArXiv e-prints

White S. D. M., Rees M. J., 1978, MNRAS, 183, 341

Willmer C. N. A. et al., 2006, ApJ, 647, 853

Wong T., Blitz L., 2002, ApJ, 569, 157

Woo J. et al., 2013, MNRAS, 428, 3306

Wuyts S. et al., 2011, ApJ, 742, 96

Yang X., Mo H. J., van den Bosch F. C., Jing Y. P., 2005, MNRAS, 356, 1293

Yang X., Mo H. J., van den Bosch F. C., Pasquali A., Li C., Barden M., 2007, ApJ, 671, 153

Yang X., van den Bosch F. C., Mo H. J., Mao S., Kang X., Weinmann S. M., Guo Y., Jing Y. P., 2006, MNRAS, 369, 1293

York D. G. et al., 2000, AJ, 120, 1579

Zaritsky D., Gonzalez A. H., 1999, PASP, 111, 1508

Zaritsky D., Smith R., Frenk C. S., White S. D. M., 1997, ApJ, 478, L53

Zeiler M. D., 2012, arXiv preprint arXiv:1212.5701

Zentner A. R., Kravtsov A. V., Gnedin O. Y., Klypin A. A., 2005, ApJ, 629, 219