Dwarves and giants: a tale of multi-scale galaxy assembly

Joseph Allen O'Leary



München 2021

Dwarves and giants: a tale of multi-scale galaxy assembly

Joseph Allen O'Leary

Dissertation at the faculty of physics of the Ludwig–Maximilians–University Munich

presesnted by Joseph Allen O'Leary from West Chester, PA, U.S.A.

Munich, 13. September 2021

Erstgutachter: Dr. Benjamin P. Moster Zweitgutachter: Dr. Klaus Dolag Tag der mündlichen Prüfung: 2. November 2021

Zwerge und Giganten: eine Geschichte der Multiskalen-Galaxienbildung

Joseph Allen O'Leary

Dissertation an der Fakultät für Physik der Ludwig–Maximilians–Universität München

vorgelegt von Joseph Allen O'Leary aus West Chester, PA, U.S.A.

München, 13. September 2021

Contents

Zusammenfassung

Abstract				xvii	
1	Intr	oducti	ion	1	
	1.1	Observ	vational facts and data	2	
		1.1.1	Expansion	2	
		1.1.2	Redshift	3	
		1.1.3	The cosmic microwave background	4	
		1.1.4	Dark matter	5	
		1.1.5	Galaxy demographics	6	
2 The galaxy-halo connection			y-halo connection	11	
	2.1	Cosmo	blogy and structure formation	11	
		2.1.1	Cosmology	11	
		2.1.2	Dark matter haloes	13	
	2.2	Galaxy	$y \text{ formation } \ldots $	13	
		2.2.1	Star formation	14	
		2.2.2	Feedback	15	
		2.2.3	Merging	17	
3	Numerical simulations				
	3.1	Initial	conditions	19	
	3.2	Gravit	ational force computation	20	
	3.3	Halo f	inders	22	
	3.4	Merge	r trees	23	
4	An	empiri	ical model for galaxy formation	27	
	4.1	Halo a	bundance matching	28	
		4.1.1	EMERGE	30	
		4.1.2	Satellite galaxies	33	
		4.1.3	Fitting and Model selection	38	

 $\mathbf{x}\mathbf{v}$

5	Gal	axy merger rates	47
	5.1	Dark Matter Simulations	48
		5.1.1 Obtaining halo merger trees	48
		5.1.2 Halo-Halo mergers	49
	5.2	The Galaxy-Galaxy merger rate	49
		5.2.1 Intrinsic merger rate	51
		5.2.2 Comparison to other theoretical predictions	61
	5.3	Which galaxies grow through mergers?	69
		5.3.1 Stellar mass fraction accreted through different merger types	69
		5.3.2 How frequent are different kinds of mergers?	73
	5.4	Discussions & Conclusions	74
6	Gal	axy merging timescales	77
Ū	6.1	Introduction	77
	6.2	The simulation and relevant emerge details	78
	0	6.2.1 The relation between merging and clustering	. e 78
	6.3	Observed merger rates	79
	0.0	6.3.1 Close galaxy pairs	80
		6.3.2 The merger rate from close pairs	83
	6.4	From simulation to observation	88
		6.4.1 Light-cone construction	88
		6.4.2 Identifying close pairs	89
	6.5	Merging probabilities and timescales	91
		6.5.1 Merging Probability	92
		6.5.2 Merging Timescales	93
	6.6	Reconstructing merger rates	97
	6.7	Discussion	100
		6.7.1 Fitting at low and high redshift	100
		6.7.2 Completeness	102
		6.7.3 Sources of uncertainty	102
	6.8	Conclusions	103
7	Dw	arf SHMB	107
•	7.1	Introduction	107
	7.2	Methods Observations Simulations etc	109
		7.2.1 N-body simulations	109
		7.2.2 Updates to EMERGE	109
		7.2.3 Observational data	109
	7.3	High redshift quenching	111
	1.0	7.3.1 Model variations	112
		7.3.2 Fitting	114
	7.4	Results	114
	=	7.4.1 Fits and model selection	114

		7.4.2	The stellar-to-halo mass relation	118
		7.4.3	Satellite populations	121
		7.4.4	Star formation history	125
	7.5	Discus	ssion and Conclusions	128
		7.5.1	Halo mass threshold for galaxy formation	128
		7.5.2	Is the Local Group representative?	128
		7.5.3	Other model options	129
		7.5.4	Further analysis	131
		7.5.5	Conclusions	132
8	3 Summary and Outlook			135
Α	A Observational data			139
В	B Displacing particles on a grid			
C high-z quenching model fits			147	
	C.1	Paran	neter correlations	147
	C.2	Comp	parison to observed data	150
A	Acknowledgements			168

List of Figures

1.1 1.2	Planck 2018 CMB thermal map and power spectrum	5 7
1.3 1.4 1.5 1.6	The observed galaxy projected two point correlation function in three mass bins	$7 \\ 8 \\ 8 \\ 9$
$3.1 \\ 3.2 \\ 3.3$	Barnes and Hut oct-tree	21 22 24
$4.1 \\ 4.2 \\ 4.3$	Comparison of halo mass function to Stellar mass function	29 31 37
$5.1 \\ 5.2 \\ 5.3$	Halo-halo merger rates	50 52
$5.4 \\ 5.5$	progenitor mass ratio Major mergers and the stellar-to-halo mass relation Cumulative galaxy merger rates as a function of descendant stellar mass and	55 56
5.6	progenitor mass ratio	57 60
5.7 5.8	Merger rates with EMERGE compared to other theoretical predictions The impact of model assumptions on major merger rates	62 63 67
5.9 5.10 5.11 5.12	Ex-situ mass contribution by merger type $\dots \dots \dots$	07 70 72 73
6.1 6.2 6.3	The redshift evolution of galaxy close pair fractions	81 84 90

6.4	The distribution of galaxies pairs as a function of projected separation and line of sight velocity difference	91
6.5	Merging fraction of close pairs as a function of projected radius and line of	00
66	Sight velocity difference	92
0.0	line of sight velocity difference	96
6.7	Galaxy merger rates reconstructed from mock pair catalogs	97
6.8	Ratio of predicted merger rates to simulated intrinsic merger rates	98
6.9	Merger rates derived from mock light cone catalogues	101
7.1	Quenching mass	115
7.2	Impact of high redshift quenching on the global stellar mass function \ldots	117
7.3	The stellar-to-halo mass relation in the dwarf regime	119
7.4	The distribution of halo masses at fixed stellar mass in the dwarf regime .	120
7.5	Cumulative satellite stellar mass function around Milky Way and Andromeda	
	like hosts	123
7.6	Normalised cumulative radial distribution of satellites around Milky Way	104
	and Andromeda like hosts	124
7.7	Star formation histories in the dwarf regime	125
(.8 7.0	The 90 per cent formation timescale in the dwarf regime	120
7.9 7.10	Lecally measured mass function along and effect from the average around	127
1.10	simulated Milky Way like bests	130
	Simulated winky way like hosts	130
C.1	Parameter correlations for 'instantaneous' quenching model	147
C.2	Parameter correlations for 'lin-max' quenching model	148
C.3	Parameter correlations for 'logistic' quenching model	149
C.4	Model SMF	150
C.5	Model CSFRD	150
C.6	Model SSFR	151
C.7	Model FQ	152

List of Tables

4.1	The best fit model parameters used for this work.	45
6.1	Summary of selection criteria for observed pair fractions	86
6.2	Best-fit parameters for pair fractions and intrinsic merger rates following	
	the functional form of eq. 6.2	87
6.3	Best fit parameters for merging probability eq. 6.6	94
6.4	Best fit parameters for merging timescale eq. 6.7.	95
7.1	Properties of Local Group dwarf galaxies	110
7.2	Best fit model parameters for each high redshift quenching model variant .	116
7.3	High redshift quenching model statistics	116
7.4	Typical halo peak mass values for fixed stellar mass intervals	122
7.5	Observed Milky Way and Andromeda properties	122
A.1	Observed quenched fractions	139
A.2	Observed clustering	139
A.3	Observed stellar mass functions	140
A.4	Observed cosmic star formation rate densities	141
A.5	Observed specific star formation rates	142

Zusammenfassung

In dieser Arbeit untersuchen wir drei aktive Forschungsbereiche zur Galaxienentstehung: Galaxienverschmelzungsraten, Beobachtungszeitskalen von Galaxienpaaren und die 'stellarto-halo mass relation' (SHMR) für Zwerggalaxien. Wir nähern uns diesen Themen durch das empirische Modell für Galaxienentstehung EMERGE, das Galaxien in simulierten Halos aus dunkler Materie so platziert, dass eine Reihe von Beobachtungen reproduziert werden.

Theoretische Modelle unterscheiden sich in den Vorhersagen der Verschmelzungsrate von Galaxien um eine Größenordnung. Wir treiben dieses Feld voran, indem wir die Verschmelzungsraten von Galaxien als Folge eines Modells untersuchen, das von Natur aus mit Beobachtungsdaten übereinstimmt. Diese Studie zeigt, dass Verschmelzungen für das Wachstum von Galaxien mit $\log_{10}(m_*/M_{\odot}) \gtrsim 10.75$ am bedeutendsten und bei geringeren Massen weitgehend irrelevant sind. Bei den massereichsten Galaxien dominieren große Verschmelzungen den akkretierten Massenanteil und tragen bis zu 90 Prozent der gesamten akkreierten Sternmasse bei. Unsere Ergebnisse bekräftigen, dass die Massenverhältnis-Abhängigkeit der Fusionsraten von der SHMR angetrieben wird. Wir zeigen jedoch auch, dass empirische Modelle, die die relevanten Observablen erfüllen, nicht garantiert die äquivalenten Verschmelzungsraten produzieren, was auf eine Diskrepanz in unserem Verständnis der Verbindung von Halo-Halo-Verschmelzungen mit Galaxie-Galaxie-Verschmelzungen hindeutet.

Die Analyse von Verschmelzungsraten aus beobachtbaren Daten bleibt eine große Herausforderung. Unterschiede in den Modellierungstechniken, kombiniert mit begrenzten Beobachtungsdaten, führen zu widersprüchlichen Schlussfolgerungen über die Zeitskalen der Verschmelzung von engen Paaren. Unser Modell stimmt besser mit den jüngsten Beobachtungen überein und weist eine Form auf, die am besten durch eine Potenzgesetz-Exponentialfunktion beschrieben wird. Wir beobachten, dass simulierte 'pair fraction' gut auf die intrinsische Verschmelzungsrate abgebildet werden können, indem eine Beobachtungszeitskala verwendet wird, die linear mit der Rotverschiebung abnimmt. Dies beruht auf der Annahme, dass alle beobachteten Paare bei z = 0 verschmelzen. In einer detaillierteren Analyse zeigen wir, dass die Verschmelzungswahrscheinlichkeiten einzelner Paare am besten durch eine logistische Funktion beschrieben werden und dass mittlere Verschmelzungszeitskalen durch eine lineare Beziehung in der projizierten Trennung und Sichtliniengeschwindigkeitsdifferenz jedes Paares gut angenähert werden können. Zusammen können unsere Anpassungsformeln Verschmelzungsraten von Galaxienpaaren bis mindestens $z \sim 4$ unter einer Vielzahl von Paarauswahlkriterien genau vorhersagen. Darüber hinaus zeigen wir, dass einige häufig verwendete Paarauswahlkriterien unter Stichprobenunvollständigkeit leiden können und stellen eine Korrekturformel für diese Fälle bereit. Schließlich schließen wir aus unserer Analyse, dass Beobachtungszeitskalen hauptsächlich von der Dynamik bestimmt werden und nicht stark von den Sternentstehungseigenschaften der einzelnen Galaxien beeinflusst werden.

Eines der Hauptziele bei der Untersuchung der Galaxienentstehung ist es zu verstehen, wie die leuchtende Komponente des Universums, Galaxien, mit dem Strukturwachstum zusammenhängt, das durch den Gravitationskollaps von Halos aus dunkler Materie dominiert wird. Die Massenbeziehung von Stern zu Halo untersucht, wie Galaxien Halos aus dunkler Materie besetzen und was dies für ihre Sternentstehungsgeschichte bedeutet. Wir liefern das erste selbstkonsistente empirische Modell, das die SHMR bis auf $\log_{10}(m_*/M_{\odot}) \leq 5.0$ einschränken kann, indem wir unser Modell direkt an die Daten der Zwerge der lokalen Gruppe anpassen. Dies wird erreicht, indem das Galaxienwachstum in spät entstehenden Halos mit geringer Masse unterdrückt wird und die Auswirkungen der Reionisation nachgeahmt werden. Dieser Prozess dient dazu, die Anzahldichte von Galaxien zu regulieren, indem die Streuung in M_h^{peak} bei fester Sternmasse verändert wird. Dies sorgt für eine engere Streuung, als sie für eine Beschreibung ohne unterdrückenden Mechanismus bei hoher Rotverschiebung der Fall wäre. Unsere Ergebnisse zeigen, dass die zuvor aufgestellte Doppelpotenz-Massenbeziehung von Stern zu Halo auf Galaxien mit $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M}_{\odot}) \gtrsim 10.0$ erweitert werden kann. Außerdem zeigen wir, dass Halos mit $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M}_{\odot}) \lesssim 9.3$ von z = 4 wahrscheinlich keine Galaxie mit $\log_{10}(m_*/{\rm M}_{\odot}) > 5.0$ beherbergen.

Abstract

In this thesis we investigate three active areas of research concerning galaxy formation: galaxy merger rates, galaxy pair observation timescales, and the stellar-to-halo mass relation for dwarf galaxies. We approach these topics through the empirical model for galaxy formation EMERGE, which populates galaxies in simulated dark matter haloes such that a suite of observations are reproduced.

Theoretical models differ in predictions of galaxy merger rates by as much as an order of magnitude. We advance this field by studying galaxy merger rates as a consequence of a model that agrees with observational data by design. This study shows mergers are most significant for the growth of galaxies with $\log_{10}(m_*/M_{\odot}) \gtrsim 10.75$ and largely irrelevant at lower masses. For the most massive galaxies major mergers dominate the accreted mass fraction, contributing up to 90 per cent of the total accreted stellar mass. Our results reinforce that the mass-ratio dependence of merger rates is driven by the stellar-to-halo mass relation. However, we also show that empirical models that meet the relevant observables are not guaranteed to produce the equivalent merger rates, suggesting a remaining disconnect in our understanding of how halo-halo mergers relate to galaxygalaxy mergers.

Extracting merger rates from observable data remains an outstanding challenge. Differences in modelling techniques, combined with limited observational data, drive conflicting conclusions on the merging time-scales of close pairs. Our model produces a galaxy pair fraction consistent with recent observations, exhibiting a form best described by a powerlaw exponential function. We find that simulated pair fractions can be well mapped to the intrinsic merger rate by adopting an observation timescale that decreases linearly with redshift, assuming all observed pairs merge by z = 0. In a more detailed analysis we demonstrate that the merging probabilities of individual pairs are best described by a logistic function and that mean merging timescales can be well approximated by a linear relation in the projected separation and line of sight velocity difference of each pair. Together, our fitting formulae can accurately predict merger rates from galaxy pairs to at least $z \sim 4$ under a wide variety of pair selection criteria. Additionally, we show that some commonly used pair selection criteria may suffer from sample incompleteness and provide a corrective formula for these cases. Finally, we conclude from our analysis that observation timescales are primarily driven by dynamics and are not strongly impacted by the star formation properties of the component galaxies.

One of the primary goals when studying galaxy formation is to understand how the

luminous component of the Universe, galaxies, relates to the growth of structure which is dominated by the gravitational collapse of dark matter haloes. The stellar-to-halo mass relation probes how galaxies occupy dark matter haloes and what that entails for their star formation history. We deliver the first self-consistent empirical model that can place constraints on the stellar-to-halo mass relation down to $\log_{10}(m_*/M_{\odot}) \leq 5.0$ by fitting our model directly to Local Group dwarf data. This is accomplished by penalising galaxy growth in late-forming, low-mass haloes by mimicking the effects of reionization. This process serves to regulate the number density of galaxies by altering the scatter in M_h^{peak} at fixed stellar mass, creating a tighter scatter than would otherwise exist without a high-z quenching mechanism. Our results indicate that the previously established double-power law stellar-to-halo mass relation can be extended to include galaxies with $\log_{10}(M_h^{\text{peak}}/M_{\odot}) \gtrsim 10.0$. Furthermore, we show that haloes with $\log_{10}(M_h^{\text{peak}}/M_{\odot}) \lesssim 9.3$ by z = 4 are unlikely to host a galaxy with $\log_{10}(m_*/M_{\odot}) > 5.0$.

Chapter 1 Introduction

The field of galaxy formation has advanced rapidly since their first observational confirmation by Hubble (1925), less than 100 years ago. Despite our broad advancement in the sub-fields relevant to galaxy formation (e.g., star formation, interstellar medium physics, black hole formation, etc.) there remain several open questions in the field of galaxy formation, not least among them is: How do galaxies form? While our understanding has advanced in nearly all aspects of the galaxy formation cycle, the disparate physical scales involved in galaxy formation make it difficult to simultaneously capture all relevant processes. Although galaxies can be simulated and studied on an individual basis, reproducing the large scale statistics of galaxy populations remains a challenge.

The contents of this thesis will focus on answering specific questions regarding statistical properties of galaxies in large volumes. Chiefly, where do galaxies acquire their mass? Here we are interested in the rate of galaxy-galaxy mergers as well as the contribution of various merger types to the growth of galaxies. Additionally, we address methods by which these quantities could be obtained from observation. On the other end of the mass spectrum we explore how the smallest galaxies relate to the formation of structure. Specifically the relationship between dwarf galaxies and dark matter haloes.

To address these topics the thesis is be organised as follows:

- 1. In this introduction we will provide the basic background terminology required to understand the problem statements. We also introduce a suite of observables relevant to galaxy populations in large volumes.
- 2. Next we briefly cover the most salient background theory relevant for this work, including: cosmology §2.1.1, dark matter halo formation §2.1.2 and galaxy formation §2.2.
- 3. In chapter 3 we discuss the numerical foundations that are included in our galaxy formation model. In this chapter we discuss: generating initial conditions §3.1, running N-body simulations §3.2, and finally locating dark matter haloes §3.3 and constructing merger trees §3.4

- 4. We introduce the empirical model EMERGE in chapter 4. In this chapter we motivate the empirical approach §4.1, discuss the inner workings of EMERGE §4.1.1 and the details of the model fitting procedure §4.1.3.
- 5. Following the material introduction we move onto the first science result of this thesis, determining the galaxy-galaxy merger rate. We approach this through a direct measure of mergers as computed by the code EMERGE.
- 6. The results of chapter 6 place our theoretical merger rates directly into an observational context. Here we propose formulations for both the merger probability of galaxy pairs, along with average merging timescales.
- 7. We break from the discussion of merger rates to discuss the **stellar-halo mass** relation (SHMR) in the dwarf galaxy regime in chapter 7. In this chapter we present the implementation of an EMERGE model extension that introduces high redshift quenching mechanisms and we discuss the implications for observed populations of dwarf galaxies.
- 8. Finally, in chapter 8 we summarise the results of this thesis and reiterate opportunities for future work.

1.1 Observational facts and data

Before approaching the details of this works it's important to layout some grounding observations. Together, the observations discussed here provide the motivation for the underlying cosmological model as well as layout some galaxy characteristics that ought to be reproduced to have confidence in our model.

1.1.1 Expansion

The first theoretical framework for an expanding universe were provided by Friedmann (1922) through his solutions of the Einstein's field equations for general relativity (see §2.1.1). Observational evidence for expansion were delivered first by Lemaître (1927) and independently confirmed by Hubble (1929) with their discovery that extra-galactic objects are receding from the observer at a rate proportional to the objects distance. This rate is colloquially described with **Hubble Parameter**, H. While initially described as a constant, we now know the Hubble parameter can evolve with time in a manner which depends on both the cosmological model, and the model parameters (see §2.1.1). However, it is commonplace to discuss the Hubble parameter in terms of its present day quantity H_0 . Initial results placed the value of the Hubble parameter at $H_0 \sim 500 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Modern measurements place H_0 at a considerably lower value in the neighbourhood of $H_0 \sim 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Komatsu et al., 2011; Planck Collaboration, 2016). In either case, the recession velocity of an object due to expansion can thus be calculated as $v_{cos} = H \cdot D$ where D is the distance to the object.

The implications of this observation are two fold: The first is that the Universe is not static but instead in a state of expansion; Second, all extra-galactic objects are on average moving away from one another which supports a conclusion that the universe has no preferred centre. These assumptions play a central role is constructing our cosmological framework (see §2.1.1).

1.1.2 Redshift

The recession velocity due to expansion becomes manifest in the observed **redshift** in the spectra of extra-galactic objects¹. Redshift describes the change in wavelength experienced by a photon from the time it was emitted to the time it arrives at the observer and is given by

$$z \equiv \frac{\lambda_0 - \lambda_e}{\lambda_e} \sim \frac{v}{c} \tag{1.1}$$

where z is redshift, λ_e is the wavelength at the time of emission and λ_0 is the observed wavelength. In the non-relativistic limit this can be expressed as the line of sight velocity v divided by the speed of light c. From this construction we can see that the cosmological redshift imparts not only a spatial coordinate with respect to the observer, but as we will soon see also a time coordinate. In this way, galaxies become a tool to probe the expansion history of the universe, and trace its matter content.

However, redshift measurements are complicated due to the individual peculiar velocities of the objects, which also contribute to the observed redshift. The redshift due to expansion and peculiar velocity can be written as:

$$1 + z_{\rm obs} = (1 + z_{\rm cos})(1 + z_{\rm pec}) . \tag{1.2}$$

If we then assume peculiar velocities are non-relativistic² then the peculiar redshift can be reduced to $z_{\text{pec}} = v_{\text{pec/c}}$, and the observed redshift can be written as:

$$z_{\rm obs} = z_{\rm cos} + \frac{v_{\rm pec}}{c} (1 + z_{\rm cos}) .$$
 (1.3)

How these coordinates then translate into physical dimensions depends on the expansion model we adopt. Throughout this work we use the cosmological redshift as a proxy for time when discussing the evolution of galaxy properties, and as a coordinate when discussing the location of galaxies with respect to the the local Universe. However, it's also clear from eq. 1.3 that the relative motions of galaxies blur this picture. In §6 we discuss how the resulting uncertainty in redshift measurements impact our ability to track the merger rate of galaxies through measurements of close pairs.

 $^{^{1}}$ The relation between cosmological redshift and peculiar velocities follows the description in Mo et al. (2010)

 $^{^{2}}$ A reasonable assumption as the fastest galaxies have peculiar velocities of only a few 1000 km s⁻¹(Karachentsev et al., 2006)

1.1.3 The cosmic microwave background

Our current understanding of cosmology is rooted in the Big Bang scenario, where our Universe as we know it was synthesised in the finite past from a point of singularity. Using arguments based on nucleosynthesis, the framework for a hot Big Bang was first espoused by Gamow (1946). He posited that if the Universe is expanding there must have been a time in the past where the Universe was much more dense and hot. In this line of reasoning the Universe likely possessed a non-negligible radiation density. During this epoch the Universe was opaque and too hot for neutral atoms to form. As the Universe expanded, it cooled. Once the Universe cooled sufficiently, electrons and protons could combine to form neutral hydrogen, which is known as recombination. At this point, the universe became transparent and photons could propagate freely. With these expectations he predicted the existence of a visible remnant of the hot Universe which should have been red shifted to a few degrees kelvin. Looking in all directions we can see the relic of this event, photons that originated from this point of last-scattering. Forming a nearly perfect black body, photons emitted from the surface of last-scattering have been heavily redshifted and are visible to us in the form of the **Cosmic Microwave Background** (CMB). His predictions were resoundingly confirmed with the discovery of the cosmic microwave background (CMB) (Penzias & Wilson, 1965; Dicke et al., 1965), with a measured local temperature of 3.5K. More recent observations have shown to the CMB to be astoundingly uniform with $T_0 =$ $2.72548 \pm 0.00057K$ (Fixsen, 2009).

Follow-up measurements of the CMB with the COBE satellite revealed both large and small scale temperature anisotropies (Smoot et al., 1992).Figure 1.1a shows a recent thermal map of the CMB where thermal fluctuations can be readily seen. These anisotropies were shown to be scale-invariant with a Gaussian distribution. These anisotropies can be described by a **power spectrum**, which characterises the amplitude of the anisotropies at different angular scales. Figure 1.1b shows the power spectrum measured from the same thermal map shown in figure 1.1a.

The study of CMB anisotropies provides a vital gateway into the early universe where the inhomogeneities prior to the CMB are encoded in the anisotropies of the CMB. Prior to recombination when the radiation density was non-negligible, photons were available to exert a force on the baryonic component of the Universe acting as a pressure to resist gravitational compression. Together these formed acoustic oscillations propagating in the primordial Universe altering the distribution of matter. At recombination acoustic modes at the extrema of their oscillation were frozen in place and appear as enhanced contrast regions on the CMB. The modes that are preserved are therefore sensitive to our cosmology and should be encoded in the power spectrum. While the details of this are beyond the scope of this thesis, we know that the location of the first peak is sensitive to the spatial curvature of the Universe, while the second peak is sensitive to the baryon density and the third peak places constraints on the physical dark matter density (see Hu et al., 2001, for details.). Research of the CMB is not only important to studies of cosmology, but are also essential to studies of galaxy formation as the CMB power spectrum details the distribution of matter in the early universe. Therefore, these observations can be used to determine



Figure 1.1: Planck 2018 CMB thermal map Planck Collaboration et al. (2020a) and power spectrum Planck Collaboration et al. (2020b)

the initial conditions necessary for large scale structure formation and subsequently galaxy formation. We discuss how this observed power spectrum informs numerical simulations in §3.1.

1.1.4 Dark matter

The first indication that the matter content of the universe could not be described with luminous material was put forward by Zwicky (1933). Through use of the virial theorem, he noticed the velocities of galaxies in the Coma cluster could not be explained solely by adding up the mass from visible sources. This prediction was validated by Ostriker et al. (1974) and Einasto et al. (1974) when they showed that the flat rotation profile of spiral galaxies could also not be explained by the visible content of a galaxy but instead required an extended massive halo.

Although the first clues of a dark Universe were determined through dynamical arguments, there have been many subsequent observations in support of dark matter. Despite the overwhelming observation data in support of dark matter little is currently known about what it actually *is*. While the precise nature of dark matter remains uncertain, modern measurements place the total matter density of the universe at around 30 per cent with only 5 per cent being attributed to normal baryonic matter. At these abundances dark matter dominates the formation of structure and drives dynamical interactions at large scales. The formation of structure from dark matter and how galaxies in turn form within that structure forms the back bone of the research in this thesis.

1.1.5 Galaxy demographics

Finally, we need to consider some bulk characteristics of galaxies. The substance of this work relies on constructing a model for galaxy formation that can reproduce a number of observed statistics as accurately as possible. In this section we will give a brief overview of the data we use to constrain our model.³

Galaxy luminosity and stellar mass functions

Arguably the most important statistic to consider is the distribution of galaxies according to their luminosity. The luminosity function describes the number density of galaxies of some type in a luminosity interval per volume and carries a few important properties. First, there is not a uniform distribution of luminosities. By numbers the least massive galaxies are the most abundant. The number density of galaxies at some luminosity can be approximated with a Schechter function (Schechter, 1976):

$$\Phi(L)dL = \Phi^* \left(\frac{L}{L^*}\right)^{\alpha} \exp\left(-\frac{L}{L^*}\right)$$
(1.4)

where Φ^* is a normalisation parameter and L^* is a characteristic luminosity that defines the transition between the low luminosity power-law distribution with slope α and the exponential cut off at high luminosity. Rather than use the luminosity function directly this work will adopt the 'equivalent'⁴ mass function, which conveys the number density of galaxies at a stellar mass derived from their observed luminosity. Figure 1.2 shows the observed evolution of the mass function at three redshift intervals. In this figure we can see the abundances increasing over time, with the most massive objects appearing later than less massive objects.

Throughout this work we will frequently discuss galaxy statistics in terms of their stellar mass, instead of luminosity. The reason for doing so, and why the **galaxy stellar mass function** (SMF hereafter) holds a special role in our model for galaxy formation is discussed in §4.1.

Clustering

The second important characteristic is clustering. Not only are galaxies not uniformly distributed in mass but their locations on the sky are not random! A common measure of galaxy clustering is the two point correlation function:

$$\xi(r) = \frac{DD(r)}{RR(r)} - 1,$$
(1.5)

³A complete list of the observational data used in this work is shown in §A

⁴The conversion of luminosity to mass is beyond the scope of this work. For our purposes we treat the mass function as observed data although galaxy mass is not the direct observable.



Figure 1.2: The observed stellar mass function at various redshift

where DD(r) is the number galaxy pairs with separation r within some bin Δr and RR(r) is the expected number of pairs with separation r within a random distribution. The twopoint correlation function therefore gives the excess probability of finding two galaxy at some preferred separation beyond random. Often we use the projected correlation function which integrates eq. 1.5 in redshift space between the interval $\Delta z \in [0, \pi_{max})$ to become

$$w_p(r_p) = \int_0^{\pi_{\max}} 2.0\xi(r_p, \pi) d\pi , \qquad (1.6)$$

where r_p is a projected radius. Figure 1.3 shows the observed projected correlation function in three mass bins. For this work the salient features of observed correlation functions are that clustering tends to increase toward smaller radii, and that more massive objects cluster more strongly. These observations not only play an active role in our model fitting (see §4.1.1), but also in our analysis of observed galaxy merger rates (see §6).



Figure 1.3: The observed galaxy projected two point correlation function in three mass bins.

Star formation characteristics

Finally, we need to consider the star formation properties of observed galaxies. We are interested in not only how star formation properties scale with stellar(halo) mass, but how those scaling relations evolve with redshift.



Figure 1.4: The observed CSFRD evolution in various frequency bands.

The first measure we consider is the **cosmic star formation rate density** (CSFRD). The CSFRD describes how the *comoving* volume averaged star formation rate evolves with redshift. Figure 1.4 shows the evolution of the CSRD with time. Each observable has been corrected to measure the star formation rate in the rest frame. The primary take away from this observation is that the star formation rate density increases with decreasing redshift, reaching a peak density near z = 2 before decreasing gradually toward z = 0. The increase in the CSFRD at early times is inline with the increase number density of galaxies as a whole shown in figure 1.2. Meanwhile, The drop in the CSFRD for $z \leq 2$ indicates that *in-situ* star formation may play less of a role in the continued growth of galaxies at late time.



Figure 1.5: The observed quenched fractions evolution in various frequency bands.

We can break down the CSFRD further to probe the redshift evolution of the **specific star formation rate** (sSFR). Rather than express the star formation rate in terms of M_{\odot} yr⁻¹the specific star formation rate defines the star formation rate per unit galaxy stellar mass. similar to the CSRD the sSFRD tracks the evolution of star formation with time. However, the sSFR more directly measures that evolution in discrete mass range intervals. Figure 1.5 shows how the sSFR scales with redshift in three mass bins. This observation shows us that generally low mass galaxies are more actively star forming than more massive galaxies at a fixed redshift. For the purposes of constructing a model for galaxy formation, this already hints that the processes involved in star formation are not equal across all galaxies types. Beyond that it raises the question of how massive galaxies came to be if their relative star formation is so low? We explore this topic in §5.



Figure 1.6: The observed quenched fraction scaling with mass in three redshift intervals

Finally, galaxies can be broadly defined in terms of their colour, which is a reflection of the stellar content of a galaxy. This follows the trend that massive hot stars are very luminous with spectra that peaks at shorter wave-lengths (blue), meanwhile low mass stars tend to peak at longer wave-lengths (red). Combined with the fact that the most massive stars are also short lived we can infer that if massive blue stars are present that a galaxy possess a younger stellar population that is actively forming stars. Under these categories, observations indicate galaxies fall into two populations. They are either blue (active, star forming) or red (passive, quenched)⁵. These categories motivate the observed quenched fractions. Figure 1.6 shows how the fraction of quenched galaxies depends on both mass and redshift. Analysis of galaxy star formation rates and quenched fractions indicates that the majority of massive galaxies are quenched by z = 0 while low mass galaxies remain star forming. In other words, the most massive galaxies are "red and dead" while lower mass galaxies are blue and star forming⁶.

⁵The distinction between active/passive is more quantitatively stated in §5.2.1

⁶This general trend may not hold for dwarf galaxies, see §7 for discussion

Chapter 2

The galaxy-halo connection

2.1 Cosmology and structure formation

2.1.1 The Λ CDM Universe

Modern cosmology is founded on the notion that on the largest scales the Universe is homogeneous and isotropic, this is what is known as the cosmological principle. Meaning that our place in the Universe is not unique, and has no preferred directionality. Under this powerful paradigm we are able to make remarkable predictions about the origin of structure and eventual evolution of the Universe. At cosmological scales gravity dominates, and General Relativity provides an unparalleled description of this physical phenomenon.

Einstein's field equations (Einstein, 1916) describe the curvature of space-time in the presence of matter and energy

$$\mathcal{R}_{\mu\nu} - \frac{1}{2} \mathcal{R} g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} - \Lambda g_{\mu\nu} , \qquad (2.1)$$

where $\mathcal{R}_{\mu\nu}$ is the Ricci tensor¹, \mathcal{R} is the Ricci scalar², $g_{\mu\nu}$ is the metric tensor, $T_{\mu\nu}$ is the stress-energy tensor, G is the gravitational constant, c is the speed of light and Λ is the cosmological constant with a purpose that will soon be made clear. Possible solutions to eq. 2.1 depend on the the chosen metric $g_{\mu\nu}$, which describes the structure of space-time. The **Friedmann-Lemaître-Robertson-Walker** (FLRW) metric (Friedmann, 1924; Lemaître, 1927; Robertson, 1935, 1936a,b; Walker, 1937) describes the geometry and causal structure of a Universe of variable size that fulfils the cosmological principle. A space-time interval under this metric takes the form

$$ds^{2} = \sum_{\mu,\nu} g_{\mu\nu} dx_{\mu} dx_{\nu} = -dt^{2} + a^{2}(t) \left[\frac{dr^{2}}{1 - kr^{2}} + r^{2}d\theta^{2} + r^{2}sin^{2}\theta d\phi^{2} \right]$$
(2.2)

where r, θ, ϕ spherical position coordinates and the value of k sets the curvature normalised such that $k \in \{+1, 0, -1\}$ corresponding to closed, flat or open geometry respectively.

¹The Ricci tensor is a contraction of the Riemman tensor $\mathcal{R}_{\mu\nu} = \mathcal{R}^{\lambda}_{\ \mu\lambda\nu}$

²The Ricci scalar is the trace of the Ricci tensor $\mathcal{R} = g^{\mu\nu} \mathcal{R}_{\mu\nu}$

Critically, this metric establishes a non-trivial coupling between the temporal and spatial components, where information regarding the time dependent physical size of the Universe is captured by the dimensionless scale factor a(t).

In a Universe endowed with the FLRW metric we can analytically solve eq. 2.1, deriving the equations of motion for the Universe itself. These are known as the Friedmann equations, and they describe the rate and acceleration of the scale factor.

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} + \frac{\Lambda}{3}$$
(2.3)

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3c^2}(c^2\rho + 3P) + \frac{\Lambda}{3}$$
(2.4)

Here, P is pressure and the density term ρ describes the energy density of matter (baryonic and dark matter), as well as radiation. Finally, the constant Λ describes dark energy, a 'force' driving the observed acceleration of the Universe (Riess et al., 1998).

In the study of cosmology and galaxy formation, it is common to adopt a dimensionless parameterisation of the first Friedmann equation (2.3). To do so, we define the *critical* density $\rho_c \equiv 3H_0^2/8\pi G$, which is the density required for a flat Universe. We can then introduce the dimensionless density parameter $\Omega_0 \equiv \rho_0/\rho_c$. With $\rho_0 = \rho_m + \rho_r + \rho_\Lambda$, where ρ_m is the matter density, ρ_r is the radiation density and $\rho_\Lambda = c^2\Lambda/8\pi G$ is the vacuum energy density. Thus, $\Omega_0 = \Omega_m + \Omega_r + \Omega_\Lambda$. Under these definitions, the curvature can be expressed in this parameterisations as $k = H_0^2 a_0^2 (\Omega_0 - 1)$. Taken together eq. 2.3 can be rewritten as

$$\left(\frac{\dot{a}}{a}\right)^2 = H^2(t) = H_0^2 \left(\frac{1-\Omega_0}{a^2} + \frac{\Omega_m}{a^3} + \frac{\Omega_r}{a^4} + \Omega_\Lambda\right) .$$
(2.5)

In this formulation, the scale factor is typically normalised such that $a_0 = 1$, the scale factor today. Similarly, this implies the density parameters are also represented in terms of their present day quantities.

The baryonic component of the Universe makes up only a small fraction of the matter content $\Omega_b = f_b \cdot \Omega_m$. However, because dark matter does not interact electromagnetically (hence the 'dark'), we are unable to observe it directly. Instead we rely on baryonic material cooling and collecting in the centre of massive virialized dark matter haloes, where it can condense and begin the process of star formation, eventually constructing the galaxies we see.

To summarise, we have described a Universe comprised of radiation, baryons, and dark matter with accelerated expansion driven by dark energy (Λ). The precise nature of dark energy and dark matter is the subject of debate. For this work we adopt the standard cosmological model where the dark matter component is assumed to be 'cold' (CDM) with velocities much smaller than c (Peebles, 1982). Based on recent measurements of the cosmological density parameters we further assume $\Omega_0 = 1$ corresponding to flat geometry (Planck Collaboration, 2016). Together these assumptions form the Λ CDM framework that we use as the basis for our galaxy formation analysis.

2.1.2 Dark matter haloes

Underpinning the study of galaxy formation is the study of dark matter halo formation. Dark matter haloes form from gravitational instability due to perturbations in the cosmic density field. we can characterize these perturbations in terms of the dimensionless overdensity field

$$\delta(\mathbf{x},t) \equiv \frac{\bar{\rho}(t) - \rho(\mathbf{x},t)}{\bar{\rho}(t)}$$
(2.6)

where $\delta(\mathbf{x}, t)$ describes the amplitude of a density fluctuation about the mean $\bar{\rho}(t)$. Physically, over-dense regions *can* collapse and form structure under their own self-gravity, while under-dense regions will expand, forming voids. Gunn & Gott (1972) proposed a model for spherical collapse in an expanding Universe. This model establishes a critical density δ_c , for which a perturbation *will* have collapsed to form a virialized object by time t if $\delta(\mathbf{x}, t) > \delta_c(t)$. Essentially, these regions where self-gravity can out-compete expansion will see a non-linear spike in the density contrast compared to the background.

Using the **Zel'dovich approximation**, an initial over-density field $\delta(\mathbf{x})_0$ can be linearly extrapolated to determine the over-density field at subsequent times. Taken in the context of the spherical collapse model, Press & Schechter (1974) proposed a means to determine the mass of structures that collapsed from $\delta(\mathbf{x}, \mathbf{t})$, a non-linear process. To segregate the density field into discrete regions, Press & Schechter (1974) smoothed the density field according to a spherical top-hat window function $W(\mathbf{x}; \mathbf{R})$, with characteristic radius R. It is then argued that the probability that $\delta_s > \delta_c$ at some instant in time is equivalent to the mass fraction of objects with mass greater than M. The final result is a co-moving number density of haloes of mass $M \to M + dM$:

$$n(M,t)dM = \sqrt{\frac{2}{\pi}} \frac{\bar{\rho}}{M^2} \frac{\delta_c}{\sigma} \exp\left(-\frac{\delta_c^2}{2\sigma^2}\right) \left|\frac{\mathrm{dln}(\sigma)}{\mathrm{dln}(M)}\right| dM , \qquad (2.7)$$

where σ is the mass variance of a smoothed density field δ_s and can be related to the initial density field through the power spectrum as

$$\sigma^2(M) = \langle \delta_s^2 \rangle = \frac{1}{2\pi^2} \int_0^\infty P(k) \tilde{W}^2(\mathbf{k}R) k^2 \mathrm{d}k \;. \tag{2.8}$$

Here P(k) is the power spectrum of the density field, and $W(\mathbf{k}R)$ is the Fourier transform of the window function. The halo mass function gives a statistical description of how many haloes we should expect for each mass interval at some epoch. As we will soon see, these haloes will play a critical role in our efforts to describe the assembly process of galaxies.

2.2 Galaxy formation

So far we have only eluded to the fact that galaxies and dark matter haloes are linked. Although both of these components have been known about for quite some time, its important to understand that a relationship between galaxies and dark matter haloes was not proposed until relatively late. Early models for galaxy formation only considered formation from massive collapsing clouds of gas (Eggen et al., 1962; Gott & Thuan, 1976), where various formulations for star formation and cooling were invoked to explain the observed populations of elliptical and spiral galaxies. It was not until White & Rees (1978) that an explicit connection between dark matter haloes and galaxies was proposed. While these two-stage models can successfully build galaxies in the context of a CDM cosmology, they remain incomplete. It was already shown quite early that additional properties need to be incorporated to produce not only the the correct galaxy type, but also their luminosity functions(Larson, 1974).

The intent of this section is not to provide a detailed description of all galaxy formation physics or the relative contribution from different physical processes. Instead this section should provide a general overview of what is and what *might* be relevant to galaxy formation. In §4.1.1 we will explain in detail how we can marginalise over the uncertain contributions of these processes to construct a complete view of galaxy formation at cosmological scales.

2.2.1 Star formation

Gravitational collapse

At the scale of dark matter haloes galaxy formation begins with gravitational collapse. As a rough order of magnitude we can evaluate the minimum halo halo mass that could collapse shortly after recombination. The conditions under which a gravitational perturbation could collapse and grow were first explore in Jeans (1902). This work showed that a region is only susceptible to collapse if gravity is able to overcome the internal pressure gradient of that region. This essentially specifies a minimum length scale required before collapse can occur, known as Jeans length $\lambda_{\rm J}$. Expressed in the context of a cosmological background this term can be defined as

$$\lambda_{\rm J} \equiv c_{\rm s} \sqrt{\frac{\pi}{G\bar{\rho}}} \tag{2.9}$$

where $c_{\rm s}$ is the sound speed, G is the gravitational constant and $\bar{\rho}$ is the mean cosmological matter density. Using this criterion and some basic assumptions we can establish a rough minimum halo mass through which we might be able to form a galaxy. If we define a sphere with radius $\lambda_{\rm J}/2$ the mass enclosed in that volume is

$$M_{\rm J} = \frac{\pi}{6} \rho_m \lambda_{\rm J}^3 \tag{2.10}$$

If we evaluate this at the time of recombination ($z \approx 1100$) where the mean temperature is $T \sim 3000k$, and assuming a monatomic gas we find that the smallest perturbations able to collapse are on the order of $\log_{10}(M_{\rm h}/{\rm M_{\odot}}) \sim 5.5$. In practice there are other processes that limited the lower bound of a halo at some given epoch. For our purposes we can use this to establish a rough lower limit on haloes that could host a galaxy.

This same processes can be further applied to gravitational instability in giant molecular clouds, the birth place of stars. The largest molecular clouds range between $10^5 - 10^6 M_{\odot}$

and themselves contain substructure in the form of clumps. There are several mechanisms which might prevent the collapse of clouds and clumps inhibiting star formation. For instance, turbulent motion can provide an additional source of pressure by increasing the effective sound speed, and the presence of magnetic fields may also provide further stability against gravitational collapse (but see Ibáñez-Mejía et al., 2021). So while molecular clouds can be quite massive, their ability collapse (at all scales) is impacted by numerous internal processes depending the local state of clouds/clumps. Ultimately these processes limit the fraction of a given cloud that is both available to form stars and the efficiency at which they can do so.

Stellar populations

As a cloud collapses it will begin fragmenting due to local instabilities. The largest clumps will collapse first with smaller clumps following. The stars formed in a molecular cloud will take on a range of mass that can generally be described by a power law at high masses and a modified exponential at low masses:

$$\xi(m) = \frac{dN}{dm} \propto \begin{cases} \exp\{-[\log_{10}(m/0.2M_{\odot})]^2/0.6\} & (0.1 \le m \le 1.0M_{\odot}) \\ m^{-1.35} & (m > 1.0M_{\odot}) \end{cases}$$
(2.11)

Where m is stellar mass. The distribution of stars formed from a molecular cloud is called the initial mass function (IMF), and eq. 2.11 specifically describes the Chabrier (2003) IMF. The shape of the IMF will ultimately determine the stellar content of a galaxy following successive generations of star formation. In particular the most massive stars will burn through their fuel supply more quickly than low mass stars. Ultimately the largest stars will explode in core collapse supernovae, seeding the interstellar medium with gas and metals. If the supply of cold gas is cutoff new stars will be unable to form, within a few 10Myr the massive stars will die off while less massive stars continue to burn. Over time this will result in a galaxy with an older, redder population of stars. Observations of stellar populations in galaxies can therefore be used to determine star formation history of a galaxy.

2.2.2 Feedback

Photoinoizing and stellar wind feedback

The first form of feedback is from the stars themselves. On the high mass end, O/B type stars can deposit significant energy into the ISM via ionizing radiation and stellar winds (Kudritzki & Puls, 2000; Puls et al., 2008; Haid et al., 2018). Stars more more massive than $m \gtrsim 10 M_{\odot}$ can deliver as much as 10^{39} erg/sec in ionizing radiation. Over the entire lifetime of the star (5-20 Myr) this can deposit as much as 10^{53} erg. For the most massive stars the photo-ionizing radiation dominates their energetic output, however winds can also impart sizeable energy deposition. While winds deposit energy at a much lower rate

 $\sim 10^{34}-10^{37}$ erg/sec, the total energy deposition can reach as high as 10^{51} erg over the star's life.

Both of these mechanisms can contribute to quenching through mechanical feedback which ejects gas from the local ISM, strangling star formation. Photo-ionization can also serve to prevent star formation by heating the surrounding gas preventing the cooling necessary for star formation. The lifetime energy deposition of these feedback mechanisms are as high and sometimes higher than that from supernovae.

Supernovae

Massive stars burn through their fuel on short timescales (3-60 Myr, Georgy et al., 2013) ending their lives in core-collapse supernova explosions that can also inject energy into the surrounding medium shutting off star formation. Core collapse supernovae deliver around 10^{51} erg into the ISM (Arnett, 1996; Janka, 2012). Additionally, Steinwandel et al. (2020) showed that supernovae occurring in low density environments, such as those caused by stellar winds, can drive significant outflows. This allows supernovae occurring in a region evacuated by the aforementioned stellar wind feedback to generate more efficient quenching outflows.

Active galactic nuclei

Observations suggest that most galaxies host a super massive black hole at their center with a mass proportional to that of the galaxy (Ferrarese & Merritt, 2000; Gebhardt et al., 2000)^{3,4}. Observations of high redshift galaxies further show that these black holes are accreting at astounding rates. The formation of an accretion disk around a growing black hole provides a substantial source of ionizing radiation that is suspected to reheat the halo shutting down star formation on galactic scales. The luminosity of an accretion disk around a black hole can be approximated as:

$$L_{\rm disk} = \eta \dot{M} c^2 \tag{2.12}$$

where η is the accretion efficiency, M is the mass accretion rate and c is the speed of light. If we assume a black hole accreting at the Eddington limit where the luminosity is given by

$$L_{\rm E} = \frac{4\pi c G M m_p}{\sigma_{\rm T}} \approx 3.3 \times 10^4 L_{\odot} \frac{M}{\rm M_{\odot}}$$
(2.13)

where m_p is the proton mass and σ_T is the Thompson-scattering cross section. We can see that a $10^9 M_{\odot}$ black hole could have $L \approx 3.3 \times 10^{13} L_{\odot}$, which is as bright or brighter than entire galaxies. If accreting with only $\eta = 0.1$ (Soltan, 1982; Davis & Laor, 2011) the black hole would consume ~ 23 M_{\odot} yr⁻¹. These massive luminous central black holes are termed active galactic nuclei (AGN).

³The examples here are drawn from Maoz (2016)

⁴More precisely the relation is between black hole mass and the velocity dispersion in the stellar bulge. Where the dispersion increases with mass

Recent estimates place the lifetime of AGN at $\sim 10^5$ yrs (Schawinski et al., 2015); as a rough measure the black hole in the example above would deposit 10^{59} erg if accreting at the Eddington limit over its lifetime. However, the process by which AGN actually couple to the ISM to shut down star formation is uncertain and remains one of the most active

2.2.3 Merging

In addition to the properties produced through in-situ star formation, Toomre & Toomre (1972) demonstrated through numerical simulations that the diverse morphological properties of observed galaxies could be produced through galaxies interacting with one another. In an extreme case, this proposition can be extended to say that all large elliptical galaxies were formed through merging of one or more disk galaxies. The contribution of mergers to galaxy growth is a central topic of this thesis. In §4.1.1 we will detail the physical merger process. The rate of galaxy mergers, implications for galaxy growth and the relation to dark matter haloes will be covered in §5.

areas of research in galaxy formation (King & Pounds, 2015; Harrison, 2017).

2. The galaxy-halo connection
Chapter 3

Numerical simulations

The topics studied in this work rely on an empirical model for galaxy formation which populates simulated dark matter halos with galaxies. Before we can employ the model we therefore must generate the necessary input data. This process can be summarized in four steps:

- 1. Generate initial conditions
- 2. Run N-body simulation
- 3. Identify dark matter halos
- 4. Link halos across time to form merger trees

In this chapter we briefly describe how each step is performed and the codes used.

3.1 Initial conditions

The first step in running an N-body simulation is obtaining the appropriate initial conditions. Initial conditions are simply a snapshot containing all properties of our simulation particles at some desired starting redshift. Imperatively, initial particle positions and velocities must be intrinsically linked to our cosmological density perturbations.

For the purposes of N-body simulations we do not start directly at the CMB, but instead at some later time to prevent the amplification of particle noise. Thus we need a power spectrum corresponding to the starting redshift of the simulation. If all of the density fluctuations are small at very early times, we can assume that only growing modes are significant at later times. The density field can then be written in the simplified form

$$\delta(\mathbf{x}, a) = D(a)\delta_0(\mathbf{x}) , \qquad (3.1)$$

where D(a) is the linear growth factor. This is known as the **Zel'dovich Approximation** (Zel'Dovich, 1970), and allows us to determine the density field at any time¹ if we have

¹A linear evolution is only applicable while $\delta \ll 1$. At this point our numerical work takes over.

information on the initial field $\delta_0(\mathbf{x})$. This correspondingly allows us to determine the power spectrum at later times as a function of linear growth:

$$P(k) \equiv \langle |\tilde{\delta}_{\mathbf{k}}|^2 \rangle = \alpha k^{n_s} \mathcal{T}^2(k) , \qquad (3.2)$$

where P(k) is the wave-space power spectrum, fully defined by the over-density field $\tilde{\delta}_{\mathbf{k}}$, which is simply the Fourier transform of the real-space over-density field $\delta_{\mathbf{x}}$ where k denotes wave number. The power spectrum spectral index is n_s , α is a normalisation constant and $\mathcal{T}(k)$ is the transfer function. The transfer function relates late time amplitudes to their initial values as

$$\mathcal{T}(k) \equiv \frac{\tilde{\delta}_k(z=0)}{\tilde{\delta}_k(z)D(z)} .$$
(3.3)

The transfer function is a catch all for physical effects and is specified for the detailed cosmology chosen. An appropriate transfer function is typically computed with software such as CAMB (Lewis et al., 2000; Howlett et al., 2012; Lewis, 2014). Next, the transfer function needs to be translated into particle displacements (see §B for a complete description). We utilise software, such as MUSIC (Hahn & Abel, 2011), to convert this Fourier space transfer function into a real space over-density field. Essentially, MUSIC samples a random Gaussian sample in k space for as many particles as we desire. This random Gaussian field is then applied to the power spectrum, such that the Fourier space over-density field becomes

$$\tilde{\delta}(\mathbf{k}) = \alpha k^{n_s/2} \mathcal{T}(k) \tilde{\mu}(\mathbf{k}) , \qquad (3.4)$$

where $\tilde{\mu}(\mathbf{k})$ is our Fourier space random Gaussian field. Finally, the Fourier space overdensity can be translated back into the real-space positions and velocities, and used to place the particles within our simulation volume. There are a multitude of software solutions available to generate initial conditions, and the one chosen depends on the specific applications and use cases of the resulting simulation.

3.2 Gravitational force computation

Now that we have the initial conditions, we can evolve these 'particles' forward to follow the non-linear growth of structure. At the base level, this involves solving the gravitational force equation for each particle in the simulation:

$$m_i \frac{d^2 \mathbf{x}_i}{dt^2} = \sum_{\substack{j=1\\j\neq i}}^{N} \frac{Gm_i m_j (\mathbf{x}_j - \mathbf{x}_i)}{||\mathbf{x}_j - \mathbf{x}_i||^3} .$$
(3.5)

A brute force approach is computationally straight forward, but suffers from practical scaling limitations. To properly model the Universe on cosmological scales we require $\sim 10^9$ particles, even for a massively parallelised code a brute force approach is too inefficient. To overcome the $\mathcal{O}(N^2)$ scaling, a more sophisticated algorithm is required.

There are a variety of N-body solvers available, each with their own advantages and disadvantages. We use simulations generated with GADGET-2 (Springel, 2005), which uses the so-called **Tree-PM** method. A hybrid method that splits the forces into short range and long range components, with short range forces calculated using the **Tree-algorithm** and long range forces with the **Particle-Mesh** (PM) method.

The tree algorithm works by subdividing the simulation volume into octant, also known as *nodes*, this division is shown for a 2D grid in figure 3.1. This process is carried out recursively until each node contains only 1 or 0 particles. Now one can walk through the tree, computing the force contributions from the other nodes. Where the force from a node is approximated by the multi-pole moment of the gravitational potential, in the case of GADGET-2 only the monopole is used. The tree walk continues down through each tree until the desired accuracy has been achieved, then the process is repeated for each node. The end result is a force calculation that scales as $\mathcal{O}(N \log N)$ (Barnes & Hut, 1986; Springel et al., 2001).



Figure 3.1: A Barnes and Hut oct-tree in two dimensions. As the box is split each node will contain only 0 or 1 particles, when walking through the tree, nodes containing no particles are skipped over.

In the Particle-Mesh method, the simulation volume is divided into a mesh with the particles sitting at the vertices. The long range forces can then be computed by solving the Poisson equation, in Fourier space, along the mesh:

$$\tilde{\Phi} = -4\pi G \frac{\tilde{\rho}}{k^2} \,. \tag{3.6}$$

Here, the tilde denotes the Fourier transformed components of the more familiar Poisson equation (B.3), and k is the wave-number. The advantage of this approach is in the simple form of the Fourier space Poisson equation, which can be easily transformed from real space with the Fast Fourier Transform algorithm (Brigham & Morrow, 1967).



Figure 3.2: These slices show the projected density field for 15 Mpc/h thick slices from the *Millennium Simulation* (Springel et al., 2001). The lighter, more yellow, regions are more dense than the dark regions.²

As the simulation is evolved through time, a series of **snapshots** are created at predefined redshifts. Each snapshot contains information about each particle in the simulation; an ID number, position, velocity, type, etc. It is from these snapshots that we can then generate catalogues of dark matter halos.

3.3 Halo finders

The next step is to take the snapshots and use them to construct a quantifiable description of the *bound* structure. There are a few approaches to do this, e.g. the Amiga halo finder (AHF), bound density maximum (BDM), and **Friend of Friends** (FoF) (see Knebe et al., 2011, for a complete comparison). In this work we focus on the advanced FoF algorithm employed by ROCKSTAR (Behroozi et al., 2013a).

The typical FoF approach use particle separation to define halo membership. The linking length between particles is set according to some desired density. Subsequently, a variety of lengths are used to identify structures of difference scales. In this way its possible to identify both halos, and their sub-halos. The deficiency of this approach is that only position space is considered to establish halo membership.

The more advanced algorithm implemented in ROCKSTAR builds upon these basic approaches by elevating from position space identification to phase-space identification, i.e. considering both the position and velocity of constituent particles. First the particles in the simulation are organised into the standard FoF groups for parallelisation. Then a *phase-space* linking length is adaptively chosen such that a fixed fraction of the particle group is contained within a subgroup. Here, the phase-space particle separation is defined as

$$d(p_1, p_2) = \left(\frac{|\mathbf{x}_1 - \mathbf{x}_2|^2}{\sigma_x^2} + \frac{|\mathbf{v}_1 - \mathbf{v}_2|^2}{\sigma_v^2}\right)^{1/2},$$
(3.7)

 $^{^{2}}$ These images are acquired from the *Millennium Simulation* media page, captions indicating redshift were added, but the aspect ratio and included scaling have not been altered.

where σ_x and σ_v are the position and velocity dispersion for the FoF group. i.e. the phasespace metric involves the normalization of sub-group member particles by the subgroup position and velocity dispersion. This operation is repeated until all successively deeper subgroups have been identified. Next, a *seed halo* is placed at the deepest levels of each group. Then the particles are hierarchically assigned to the nearest seed halo based on their phase-space separation from the halo centre:

$$d(h,p) = \left(\frac{|\mathbf{x}_h - \mathbf{x}_p|^2}{r_{dyn,vir}^2} + \frac{|\mathbf{v}_h - \mathbf{v}_p|^2}{\sigma_v^2}\right)^{1/2} , \qquad (3.8)$$

where

$$r_{dyn,vir} = v_{max} t_{dyn,vir} = \frac{v_{max}}{\sqrt{\frac{4}{3}\pi G\rho_{vir}}}, \qquad (3.9)$$

 v_{max} is the maximum circular velocity, and ρ_{vir} is the virial over-density as defined by Bryan & Norman (1998). Finally, once all particles have been assigned to a halo, the unbound particles can be removed, and the physical halo properties (position, velocity, mass, etc.) can be computed. The resulting product is a comprehensive halo list for each time step (snapshot) in the simulation.

3.4 Merger trees

Naturally, the physical properties of halos are not limited to their positions and velocities at a single epoch, but rather their temporal evolution must be considered for a complete description of structure growth. A **merger tree** is a list that describes the mass formation history of each halo in the final simulation snapshot. The final snapshot is at z = 0 throughout this work. Each tree accounts for growth from merging with progenitor systems, as well as from smooth accretion. To construct merger trees, identified halos must be meaningfully, and consistently identified across time steps, CONSISTENT TREES (Behroozi et al., 2013b) does just that.

A common approach to constructing merger trees is to connect a halo at one time step to a halo at the following time step, referred to as *progenitors* and *descendants* respectively, by means of particle membership. That is to say, a progenitor is linked with a descendant if most of the particles identified as belonging to the progenitor halo end up as member particles of the descendant halo in the next time step. However, a few problems exist for pure particle based merger trees:

- 1. Halos near the resolution limit may disappear and reappear between time steps. Leading to false mergers, or halos with no descendants.
- 2. Sub-halos near the centre of their host halo may be inconsistently identified across time steps, mistakenly identifying a merger, as well as creating a 'new' sub-halo in a later time step.

3. Particles from a sub-halo and host halo may be miss-allocated, resulting in false growth of the sub-halo.

Such issues not only cause problems with determining halo properties, but also in linking observed stellar quantities to the halo catalogue. CONSISTENT TREES seeks rectify these inconsistencies providing a more accurate halo merger history.

First, CONSISTENT TREES performs a standard particle association scheme to identify descendant halos. The descended halos are then evolved *backwards* in time to determine their most likely positions and velocities at the previous time step. This trace back is then used to cut links to halos that were falsely associated. After trimming the links to spurious halos, progenitors can be linked to likely descendants. For those halos *without* a progenitor, one will be instead created at the previous time step. If after several time steps no likely progenitor is found, the halo will be removed all together. Finally, halos at the previous time step with no descendant will be merged with the halo exerting the strongest tidal influence.



Figure 3.3: Example halo merger tree for a Milky Way-mass halo. Virial mass is indicated by the colour of each circle. The primary evolutionary tract (main branch) follows the lower most set of halos. In this view we can see not only halo growth through smooth accretion and merging, but mass loss through tidal stripping on subhalos.

After connecting the halos across time-steps we are able to realise a complete merger history for halos appearing at redshift zero. Figure 3.3 shows an example merger tree. Having formed a comprehensive understanding of how the dark matter halos have evolved through cosmic time, we are finally in a position to relate simulated dark matter halos to their stellar components.

Chapter 4

An empirical model for galaxy formation

In §2.1 we provided the background theory necessary to understand the growth of structure in ACDM context. We also discussed the basic ingredients necessary to form a galaxy in these conditions (§2.2). And while we have discussed how one can simulate the gravitational universe (§3) we have so far omitted how we might go about simulating the luminous universe. Although the fundamental physics involved in galaxy formation are well reasonably well understood, there are several computational challenges that hinder the studies of numerical galaxy formation in large volumes.

Many of the difficulties arising in numerical galaxy formation are a result of the scales involved. On the one hand, if we want to make robust predictions on statistical properties we will need a sufficiently large sample of galaxies to build such statistics; i.e. large volume. On the other hand, increasing volume at a fixed resolution can massively increase the computational expense of such simulations. As discussed §2.2 regulating star formation in galaxies requires information on feedback mechanism that occur on scales much smaller than that of a halo. Due to computational limitations, a single 'particle' in a cosmological simulation may represent some $10^6 M_{\odot}$ or more. Physical processes occurring below the mass and length scales inherent to a simulation thus require additional parameterisations, as physics below those scales cannot be modelled directly. Therefor, the need for large simulation volumes stands diametrically opposed to our need for small scale physics.

Theoretical models differ in their approach to linking dark matter haloes with galaxies and the numerical compromises that must be considered. *Ab-initio* methods provide a complete treatment of baryonic physics to build galaxies by directly computing the physical processes. These simulations are thus reliant on accurate treatments of the physics, such as gas cooling, star formation, and the relevant feedback mechanisms (Somerville & Davé, 2015; Naab & Ostriker, 2017). Due to their sophisticated nature they are time consuming and costly to run, which limits the resolution that can be achieved. As it is impossible to resolve the scales on which the fundamental forces act, most physical processes have to be combined into effective models – so-called subgrid models – with a number of free parameters which are tuned in order to reproduce a number of observational constraints. In this sense, there are currently no true *ab initio* methods in galaxy formation, as all simulations include free parameters in some form that need to be fitted or constrained by observational data, and thus rely on empirical evidence.

The two most commonly used methods that aim to model the baryonic physics are *hydrodynamical simulations* (Dubois et al., 2014; Hirschmann et al., 2014; Vogelsberger et al., 2014; Schaye et al., 2015; Pillepich et al., 2018; Hopkins et al., 2018) and *semi-analytical models* (Bower et al., 2006; Somerville et al., 2008; Benson, 2012; Henriques et al., 2015). Hydrodynamical simulations attempt to create galaxies from first principles through the calculation of gas physics along with the gravitational forces at the level of the resolution elements. These models rely heavily on subgrid approximations to address feedback from AGN, supernovae, star formation. without resolving the scales at which these processes occur it is difficult evaluate where they have be correctly coupled to the star formation model. Alternatively, semi-analytical models operate by post-processing DM-only simulations, populating dark matter haloes with galaxies using analytic prescriptions at the level of individual haloes. Both approaches have made vast progress in recent years, but still struggle to reproduce a large number of observations simultaneously, as it is very difficult to explore the parameter space of the subgrid models due to the computational cost.

An alternative approach known as *empirical models* of galaxy formation (Moster et al., 2013, 2018; Conroy & Wechsler, 2009; Behroozi et al., 2013c, 2019) presents a compelling alternative to ab-initio models. Instead of aiming to directly model the baryonic processes, these models use parameterised relations between the properties of observed galaxies and those of simulated DM haloes. The parameters of these relations are then constrained by requiring a number of statistical observations be reproduced. This approach has the advantage of accurately matching observations by construction, allowing us to analyse the evolution of galaxy properties with cosmic time, and investigate the different growth channels. Furthermore, as these models can very efficiently post-process DM-only simulations it is easy to probe large volumes to gather statistics across a large dynamic range. In this way we can reproduce the bulk statistical properties without necessarily being hampered by incomplete knowledge of sub-grid physics. The ability to reproduce observed properties of galaxies make these models the ideal vehicle to study the research topics in this thesis.

4.1 Halo abundance matching

There are several options available for constructing a model for galaxy formation directly constrained by observational data. So called HOD models (Benson et al., 2000; Peacock & Smith, 2000; Scoccimarro et al., 2001; Yang et al., 2003) for instance create a statistical description of galaxies by assuming that the number of satellites in a halo is based only on the mass of that halo alone. Another classical approach to empirical modelling is known as abundance matching, and serves as the predecessor to the more sophisticated model we use in this work.

The idea behind abundance matching is simple, if we count up all galaxies in some



Figure 4.1: This figure illustrates a basic relationship between galaxies and dark matter haloes via abundance matching. *left panel:* Blue lines show the number density of galaxies as a function of their stellar mass, orange lines shows the number density of dark matter haloes as a function of their virial mass. The x-axis indicates the mass for each line in their respective units. *centre panel*: The solid blue line illustrates the stellar mass-halo mass relationship based on abundance matching with no scatter. The black dashed line indicates a 1 : 1 relation between halo mass and stellar mass. *right panel*: The integrated baryon conversion efficiency. The blue line indicates the fraction of baryons that are contained in stars as a function of halo mass, assuming a constant universal baryon fraction.

volume and do the same for some predicted(simulated) population of dark matter haloes we can map galaxies into haloes monotonically. Here the most massive galaxy is assigned the most massive halo, the next most massive galaxy to the second most massive halo and so on. In practice this will ensure that galaxies will be mapped to haloes that appear at equivalent number density, hence the name.

This scheme is roughly visualised in the the first panel of figure 4.1 where blue lines indicate a stellar mass function and orange lines indicate a halo mass function in the same volume¹. Although, we included no physical prescriptions in this relation we can already draw some important conclusions regarding the relation between galaxies and haloes. First consider the second panel of Figure 4.1, here can clearly see that there is no 1 : 1 mapping of galaxies to haloes. If a constant relation was assumed, the abundances of galaxies at low mass and high masses would be over predicted compared with observation. The right panel of figure 4.1 reinforces this conclusion with the integrated baryon conversion efficiency. This parameter indicates the fraction of baryons that are contained within stars as a function of host halo mass. The most notable feature in this figure is the well defined peak near $M_{\rm h} \sim 10^{12} {\rm M}_{\odot}$ with sharp drop-offs to either side. From this we can infer that there is a characteristic halo mass at which galaxies can form most efficiently. To account for this behaviour Yang et al. (2003) and Moster et al. (2010b) propose a model with a parametrization which ascribes a stellar-to-halo-mass ratio that increases for low masses

¹The halo mass function shown here is constructed from an N-body simulation. The SMF shown is a simulated mass function generated using the EMERGE code applied to that same N-body simulation.

up to some peak, then decreases towards the high mass end (e.g. Yang et al., 2003; Moster et al., 2010b):

$$\frac{m(M)}{M} = 2\left(\frac{m}{M}\right)_0 \left[\left(\frac{M}{M_1}\right)^{-\beta} + \left(\frac{M}{M_1}\right)^{\gamma}\right]^{-1} .$$
(4.1)

In this formulation the low mass slope β , the high mass slope γ , the inflection mass M_1 and the normalisation $\left(\frac{m}{M}\right)_0$ are free parameters that are fit against the observed stellar mass function. Using such a relation, after an appropriate fitting of parameters, each main halo within a simulation can be populated with a central galaxy according to its virial mass M_{vir} , and each subhalo populated according to the peak halo virial mass M_{peak} .

The advantage of this approach is that the observable quantity, the SMF, is automatically reproduced, although scatter must be artificially introduced to mimic observed variance. The disadvantage is that it can only provide information about the galactic population at distinct instants in time, but does not provide a means to connect the galaxies from one halo to another throughout their formation history. In order to study galaxy evolution more completely we require a model that operates within a self consistent framework.

4.1.1 EMERGE

Throughout this work we employ The Empirical ModEl for the foRmation of GalaxiEs EMERGE (Moster et al., 2018)². In the hierarchical view of galaxy formation, each galaxy starts its life at the centre of an isolated halo. As the dark matter haloes grow and cannibalise one another, so too will their occupant galaxies. Similar to semi-analytic models EMERGE populates galaxies into simulated dark matter haloes during post processing. In this model galaxies evolve according to physically *motivated* prescriptions, directly constrained by real observables. Thus, this model provides a statistical link between galaxy and halo properties without the need to directly model baryonic physics. The model also incorporates parameterisations to match observed data beyond the SMF. In this way, EMERGE is able to produce accurate galaxy catalogues exhibiting the range of physical properties observed in large galaxy surveys. Unlike HOD or abundance matching models, this method can self-consistently track galaxies across times steps, providing the opportunity to explore and evaluate their individual growth histories.

In-situ Star formation

The primary avenue for galaxy growth in EMERGE is through *in-situ* star formation. Each galaxy is seeded at the center of a dark matter halo with a SFR directly driven by the growth of the dark matter halo, \dot{M} . On large scales, baryons are assumed to uniformly trace the underlying cosmic dark matter distribution such that each halo contains a fixed baryon fraction $f_{\text{bary}} = \Omega_b / \Omega_m$. From this it follows that the growth rate of each halo \dot{M} ,

 $^{^{2}}$ The model description provided in this section has been adapted from Moster et al. (2018), and O'Leary et al. (2021a)



Figure 4.2: Here we can see how each parameter contributes to the baryon conversion efficiency. Starting from the low mass end we can see the efficiency increase with slope β . We then encounter a peak efficiency. The normalization ϵ_N is the value of the efficiency at the characteristic mass M_1 . The towards the high mass end we have a shallow decreasing efficiency slope as γ .

should be directly proportional to the rate of baryonic growth within the halo, and the SFR in the central galaxy is given by:

$$\dot{m}_*(M,z) = \dot{m}_{\text{bary}}\epsilon(M,z) = f_{\text{bary}}\dot{M}\epsilon(M,z) .$$
(4.2)

Here, $\dot{m}_*(M, z)$ is the star formation rate, $\dot{m}_{\text{bary}}(M, z)$ is the baryonic growth rate which describes how much baryonic material is becoming available, and $\epsilon(M, z)$ is the instantaneous conversion efficiency, which determines how efficiently this material can be converted into stars.

The instantaneous baryon conversion efficiency is impacted by a variety of physical processes, gas cooling, AGN feedback, supernova feedback, etc. (Somerville & Davé, 2015; Naab & Ostriker, 2017) EMERGE seeks to establish the minimally viable parameterisation necessary to replicate observations. In the most basic picture, the instantaneous efficiency is governed only by redshift and halo mass. However, the model remains flexible as additional parameters can be added on an "as-needed" basis. In particular, it was determined that a double power-law parameterisation is sufficient to model the instantaneous baryon conversion efficiency as a function of halo mass at any redshift (Behroozi et al., 2013a; Moster et al., 2018),

$$\epsilon(M, z) = 2 \epsilon_N \left[\left(\frac{M}{M_1} \right)^{-\beta} + \left(\frac{M}{M_1} \right)^{\gamma} \right]^{-1} , \qquad (4.3)$$

where the normalisation ϵ_N , the characteristic mass M_1 , and the low and high-mass slopes β and γ are the free parameters used for the fitting. Figure 4.2 provides a visual schematic for how these parameters contribute to the efficiency at z = 0. Furthermore, the model parameters are linearly dependent on the scale factor:

$$\log_{10} M_1(z) = M_0 + M_z \frac{z}{z+1} , \qquad (4.4)$$

$$\epsilon_N = \epsilon_0 + \epsilon_z \frac{z}{z+1} , \qquad (4.5)$$

$$\beta(z) = \beta_0 + \beta_z \frac{z}{z+1} , \qquad (4.6)$$

$$\gamma(z) = \gamma_0 . \tag{4.7}$$

These parameters are allowed to vary freely within their boundary conditions in order to produce a fit in agreement with observation. Observables are chosen such that model parameters can be isolated and independently constrained, thus avoiding degeneracy. In particular, the characteristic mass $(M_0 \text{ and } M_z)$ is constrained by stellar mass functions (SMFs). The efficiency normalisation parameters (ϵ_0 and ϵ_z) can be constrained by the cosmic star formation rate density (CSFRD). The efficiency slopes (β_0 , β_z and γ_0) are constrained by specific star formation rates (sSFRs).

Rather than computing a singular stellar mass at each time-step, we instead create a stellar population by drawing from a Chabrier (2003) IMF. Stellar populations are tracked

for every simulated galaxy. Subsequently, we can track the age of these populations and compute a loss rate due to stellar death. Stellar death is accounted for through a simple time dependent fraction. For a Chabrier (2003) IMF the fractional loss rate due to stellar death computed with the FSPS package (Conroy & Gunn, 2010) is well fit by

$$f_{loss}(t) = 0.05 \ln \left(1 + \frac{t}{1.4 \text{Myr}} \right)$$
, (4.8)

where t is the cosmological time in Myr. The total stellar mass of each galaxy can thus be computed at anytime by integrating star formation rate while accounting for stellar death:

$$m_*(t) = \int_0^t \mathrm{d}t' [\dot{m}_*(t') \cdot [1 - f_{loss}(t - t')] + \dot{m}_{acc}(t')] = m_{SF}(t) + m_{acc}(t) . \tag{4.9}$$

While the total mass of a galaxy is a function of many different processes the basic components can be easily split. The stellar mass formed in-place m_{SF} , and the mass acquired through accretion of other galaxies m_{acc} . The accreted mass is itself defined by an independent stellar population based on the growth history of the merging satellites. Both of these components provide vital insight into the growth history of all the galaxies we observe. Understanding the relative contributions is thus critical to understanding galaxy formation, and is a central theme of the work presented in §5.

4.1.2 Satellite galaxies

Merging

Aside from *in-situ* star formation, galaxy mergers are the other primary mechanism contributing to galaxy growth. In the context of EMERGE, we specify galaxies of three types; central, satellite and orphan. Central galaxies exist in the center of main haloes. While, satellite galaxies sit at the center of sub-haloes, orbiting within some larger main halo. Orphan galaxies were formed in the same way as satellite galaxies, however, their subhalo has since been stripped below the resolution of the halo finder. As orphans are no longer traceable in the simulation, they require special numerical treatments to address their continued evolution.

When a galaxy first becomes an orphan, a dynamical friction clock is set. We use its last known orbital parameters to compute the dynamical friction time. Dynamical friction can be simply understood as the drag induced on a massive test particle travelling through a field of less massive particles. As this massive particle travels a 'wake' of less massive particles is formed behind it contributing a large gravitational force opposite the direction of travel. This slows the forward motion of the particle, and in the case of orbital motion this will cause orbital decay. Chandrasekhar (1943) provided an analytic prescription for dynamical friction the limiting case that the massive test particle is a point mass. This assumption breaks for extended objects that also experience tidal stripping. Binney & Tremaine (2008) derive a formula to address a satellite orbiting in an isothermal sphere. The treatment for dynamical friction has be a subject of debate, with some simulators choosing to add an additional fudge factor as a free parameter often constrained by clustering observations. We discuss the impact of dynamical friction on merger rates in §5.2.2. For the baseline model implementation we adopt the formulation specified by Boylan-Kolchin et al. (2008) which is tuned to high resolution simulations with the inclusion of baryons:

$$t_{df} = 0.0216H(z)^{-1} \frac{(M_0/M_1)^2}{\ln(1+M_0/M_1)} \exp(1.9\eta) \left(\frac{r_1}{r_{\rm vir}}\right)^2 , \qquad (4.10)$$

where H(z) is the Hubble parameter, $r_{\rm vir}$ is the virial radius of the main halo (M_0) , r_1 is the radial position of the subhalo (M_1) with respect to the center of the main halo, and η is a measure for the orbital circularity of the subhalo. When the dynamical friction time has elapsed the orphan galaxy will be merged with the central system where a portion of the satellite stellar mass will be added to the descendant galaxy as

$$m_{\rm desc} = m_{\rm main} + m_{\rm orphan} \left(1 - f_{\rm esc}\right),\tag{4.11}$$

where m_{desc} is the mass of the descendant galaxy, m_{main} is the mass of the main progenitor galaxy, m_{orphan} is the mass of the progenitor orphan galaxy, and f_{esc} is the fraction of mass that will be distributed to the ICM during the merger. The escape fraction is a free parameter in the model and is largely constrained by the low redshift behaviour on the massive end of SMFs along with the sSFR of massive galaxies.

In addition to determining a merger timescale, orphans also require an approximation to update the positions. In these galaxies position is updated at each simulation timestep until t_{df} has elapsed. Orphans are placed randomly on a sphere with radius $r = r_0 f_{dec}$ where $f_{dec} = \sqrt{1 - \Delta t/t_{df}}$, and Δt is the time elapsed since the subhalo was last resolved (Binney & Tremaine, 1987).

Halo merger trees also allow for sub-subhaloes, therefore a satellite galaxy in a resolved subhalo may also posses its own population orphan galaxies within its orbit. If an orphan is on its way to merge with a satellite galaxy and that satellite itself becomes an orphan before t_{df} has elapsed, then the orphan galaxy will have its dynamical friction clock reset according to the mass of the new central system.

Tidal stripping

Returning to our example of two haloes in the drawn out process of a merger. For simplicity we can assume that the subhalo with mass M_1 is in a circular orbit of radius R_0 , far from the centre of the main halo with mass M_0 . Additionally we can assume that M_1 has already been quenched and is no longer forming stars.³ In this scenario the haloes orbit their common centre of mass with angular speed

$$\Omega = \sqrt{\frac{G(M_0 + M_1)}{R_0^3}}$$
(4.12)

³This section follows closely Binney & Tremaine (2008) and King (1962)

In the restricted three-body problem we can consider the effective potential generated by these two large masses, as well as an effectively massless test particle. When viewed in the equatorial plane we can identify contours of constant potential. The zero-velocity surfaces are centred on each body in the system and remain closed, the surfaces represent the limit for which the test particle, with energy E, cannot cross as it would have zero velocity at that surface. However, There exists a **Lagrange point** between the two bodies where the effective potentials overlaps. This saddle point defines whether a particle can remain in orbit around one body or the other. The limit is called the **tidal radius**, and in the case of our extended masses can be written as

$$r_t = \left[\frac{M_1(r_t)/M_0(R_0)}{2 + \frac{\Omega^2 R_0^2}{GM(R_0)} - \frac{\mathrm{dln}M_0}{\mathrm{dln}R}\Big|_{R_0}}\right]^{1/3} R_0 .$$
(4.13)

Recalling that M_1 is subject to dynamical friction (see §4.1.2), it will slowly spiral toward the centre of M_0 . We can see directly from equation 4.13 that the tidal radius shrinks as the M_1 orbit decays, constantly losing mass throughout the trip. we can imagine that at some point M_1 might lose so much mass that the galaxy at the centre is no longer stable against tidal forces and it too will be stripped. Tidal stripping is a complex processes and the example presented here is a very simplified picture of what happens to a real system. One way to approach this problem would be to compute the tidal radius of the subhalo and create some approximation for the radial extent of the simulated galaxies. However the relation be halo properties and galaxy size remain uncertain (Behroozi et al., 2021), posing a barrier to implementing a physical prescription for tidal disruption. Rather than implementing a more complicated model that gradually strips the satellite galaxy, a simple formulation was found to be sufficient. The stripping model in EMERGE instead strips the galaxy instantaneously when the subhalo has dropped below some fraction of M_{peak}

$$M < f_s \cdot M_{peak} . \tag{4.14}$$

Once stripped, the satellite galaxy has its stellar material distributed to the intracluster medium (ICM) of the host halo. We leave f_s as free parameter in the model, which is predominately constrained by low redshift clustering data.

Orphan-halo mass loss

A prescription for orphan-halo mass loss is important for a few reasons. The first is the dependence of t_{df} on the mass of both systems involved, as shown in eq. 4.10. The other reason is as a means of defining the gravitational potential of the system, which is important for galaxy stripping. To ensure that all galaxies including orphans are subject to stripping, we apply a simplified formula as a stand-in for the physical tidal stripping process experienced by sub-haloes. In this approach orphan halo mass is updated at each time step, declining at the same average rate since peak mass:

$$M_i = M_{i-1} \left(\frac{M_{\text{peak}}}{M_{\text{loss}}}\right)^{-\Delta t/(t_{\text{loss}} - t_{\text{peak}})}$$
(4.15)

Here *i* is the index of the current snapshot, i - 1 is the previous snapshot index and Δt is the amount of time elapsed between these snapshots in Gyr. The halo is assumed to lose mass at the same average rate, with a slope defined by halo peak mass M_{peak} at t_{peak} and subhalo disruption mass M_{loss} at t_{loss} . In the initial EMERGE release, orphans inherited their last resolved halo mass. This static halo mass made orphans impervious to stripping which leads to a resolution dependent model. One consequence of that implementation is that a galaxy would need to be stripped *prior* to entering the orphan phase, necessitating a large value for f_s in order to reproduce the observed clustering values.

The interplay between halo mass, stripping, and clustering make a treatment for orphan halo mass critical, in fitting clustering down to 10 kpc. In the initial release of EMERGE clustering was only constrained down to 100 kpc, for the details of this thesis we extended the data-set to fit down to 10 kpc, which is important for determining merger rates derived through projected galaxy pairs, §6.3.1. Figure 4.3 shows the projected galaxy correlation function in several mass bins under this new model variation and extended data.

Satellite quenching

Galaxy quenching is one other mechanism that affects the growth of galaxies. If a dark matter halo begins to become accreted by a larger halo, its own growth rate will decline. At some point the halo will reach its peak mass M_{peak} after which the halo will not grow, consequently reducing the 'inflow' of gas. After some time the galaxy at the centre of such a halo will deplete the remaining cold gas supply through star formation and become quenched.

To address star formation in these galaxies EMERGE invokes a 'delayed-then-rapid' model for quenching (Wetzel et al., 2013). In this model, after a halo has reached peak mass the central galaxy will continue to form stars at a constant rate equal to the star formation rate at t_{peak} . After a time τ the cold gas supply is assumed depleted and the star formation rate will be set to 0. The quenching timescale can be parameterised as:

$$\tau = t_{\rm dyn} \, \tau_0 \left(\frac{m_*}{10^{10} M_{\odot}} \right)^{-\tau_{\rm s}} \,. \tag{4.16}$$

Here $t_{\rm dyn}$ is the halo's dynamical time and τ_0 is normalisation, which together specifying a minimum quenching time of $t_{\rm dyn} \tau_0$. The normalisation determines the quenching timescale for galaxies with $m_* \geq 10^{10} M_{\odot}$ while the slope $\tau_{\rm s}$ describes the quenching timescale for low mass galaxies. These parameters are largely constrained by the observed fraction of quenched galaxies at several redshifts.



Figure 4.3: Clustering in the updated model down to 10 kpc. Shaded regions illustrate the 1σ uncertainty range determined though jackknife sampling. Red crosses and associated error bars represent the mean observed w_p value for each bin. The observational points represent average quantities constructed from Li et al. (2006); Guo et al. (2011) and Yang et al. (2012).

Additional model options

So far we have described the basic as published version of EMERGE as detailed in Moster et al. (2018), with corrections implemented in O'Leary et al. (2021a). One advantage of empirical models is their flexibility to incorporate additional parameters and data as necessary. The topics discussed in §7.3 discuss the need for a high-z quenching mechanism in order reproduce observed characteristics of low mass galaxies. That model iteration also incorporates a change to the baryon conversion efficiency as suggested by (Moster et al., 2020b). The details of these model updates and expansions are discussed in §7.2.2 and §7.3.1.

4.1.3 Fitting and Model selection

So far we have described an empirical model that reproduces the observed data by construction. That this model reproduces the data is not an indication that this model is the *only* possible avenue. In practice there are many model options that could provide an equally good fit, or better. This requires a discussion of what exactly we mean with goodness of fit, and how this can be used to quantitatively compare model options and implementations. In this case it makes sense to first discuss the basic statistical methods that are used to evaluate a model.

Bayesian statistics

We should start by developing an understanding on the basic arithmetic of statistical analysis as well as some standard terminology⁴. Our goal is to provide a method to evaluate our trust in model, and compare it directly to another. The first ingredient is the sum rule

$$\operatorname{Prob}(X|I) + \operatorname{Prob}(\overline{X}|I) = 1 \tag{4.17}$$

and simply states that the probability of event X given some information I and the probability of not \overline{X} , given the same information, sum to 1. The next concept is the *product* rule and states

$$Prob(X, Y|I) = Prob(X|Y, I) \times Prob(Y|I)$$
(4.18)

the probability that X and Y are true is the product of the probability of X given Y and the probability of Y. These two arithmetic concepts provide the structure for evaluating the relative probability of independent events.

Using these concepts we can construct one of the most powerful tools for statistical analysis *Bayes' theorem* (Bayes, 1763)

$$\operatorname{Prob}(X|Y,I) = \frac{\operatorname{Prob}(Y|X,I) \times \operatorname{Prob}(X|I)}{\operatorname{Prob}(Y|I)}, \qquad (4.19)$$

⁴The description here follows directly from Sivia & Skilling (2006)

which relates the Prob(X|Y, I) to Prob(Y|X, I) and allows us to solve for a quantity of interest using the quantity we are more easily able to calculate. Each component of eq. 4.19 is assigned a formal name

$$posterior = \frac{likelihood \times prior}{evidence} , \qquad (4.20)$$

that we will refer to throughout this section. The *posterior* can be understood as the probability of a hypothesis given some data. Similarly, the *likelihood* describes the probability of the data given the hypothesis. This quantity is often the easiest solve for as it can be directly computed by comparing a data-set constructed from a hypothesis with some observed data. The *prior* encapsulates our initial knowledge on how likely that hypothesis is irrespective of the data. Finally, the *evidence* is the probability of the data, and will play a role in our model selection.

The final topic to discuss is *marginalisation*. The concept of marginalisation comes from the extension of the sum rule (eq. 4.17) beyond a set of binary outcomes to many possible outcomes. Thus we could express the total probability as the integral over all possible outcomes:

$$\operatorname{Prob}(\mathbf{X}|\mathbf{I}) = \int_{-\infty}^{\infty} \operatorname{Prob}(\mathbf{X}, \mathbf{Y}|\mathbf{I}) dY , \qquad (4.21)$$

marginalisation over some distribution can therefore allow us to handle non-contributing parameters as they should subsequently provide no contribution to the final result.

Maximum likelihood

Using the concepts we've discuss so far we can develop the tools needed to evaluate our model parameters in the presence of observed data. As an example, consider the case of a uniform prior; we can see that marginalising over a uniform distribution would result in some constant and eq. 4.19 could reduce to

$$\operatorname{Prob}(\vec{\theta}|D,\mathcal{M}) \propto \operatorname{Prob}(D|\vec{\theta},\mathcal{M}) . \tag{4.22}$$

In this case $\vec{\theta}$ is the vector of our model parameters and D is the vector containing our observed data points and \mathcal{M} is the model⁵. From this we can see that because the posterior is proportional to the likelihood we can determine the probability of some combination of parameters if we can evaluate the likelihood of the data given those parameters!

Furthermore, if we assume that the data is the result of a Gaussian process⁶ we can use repeated application of the product rule (eq. 4.18) to show that

$$\operatorname{Prob}(D|\vec{\theta}, \mathcal{M}) \propto \exp\left(-\frac{\chi^2}{2}\right), \text{ where}$$
 (4.23)

⁵The model \mathcal{M} is *not* the information term I shown previously. In fact I should be carried through but we have omitted it here to reduce clutter.

⁶It is unlikely that all of our data is Gaussian in nature. Nonetheless, making this assumption simplifies the problem making it easier to evaluate model likelihood.

$$\chi^2 = \sum_{k=0}^{N} \frac{(D_k - f(k, \vec{\theta}))}{\sigma_k^2} .$$
(4.24)

Here $f(k, \vec{\theta})$ is our model prediction of the data D_k with measured uncertainty σ_k . From this we can see that the *log-likelihood* can be expressed as

$$\ln \mathcal{L} = \ln(\operatorname{Prob}(D|\vec{\theta}, \mathcal{M})) = \operatorname{const.} - \frac{\chi^2}{2}.$$
(4.25)

From this expression it is clear that we can maximise \mathcal{L} by minimising χ^2 , which is a quantity we can directly compute for a set of model parameters.

For EMERGE we incorporate a suite of observed data to constrain the model: galaxy stellar mass function (SMF), cosmic star formation rate density (CSFRD), specific star formation rates (sSFR), quenched fractions (f_q) and the low redshift projected galaxy 2point correlation function (w_p) . The true observations are contained in the vector $\vec{\omega}$. For each observed data set we construct a mock observation $\mu(\vec{\theta})$ with our set of parameters $\vec{\theta}$. Taking the difference our mocks and the true data we obtain $\vec{\Delta} = \vec{\omega} - \mu(\vec{\theta})$, along with the observation covariance matrix C, a likelihood surface is constructed according to

$$\mathcal{L} = -\frac{\vec{\Delta}^T C^{-1} \vec{\Delta}}{2} = -\frac{\chi^2}{2} \,. \tag{4.26}$$

If the complete covariance matrix is unavailable we calculate it as $C = \text{diag}(\sigma_k, ..., \sigma_N)$ where σ_k is the uncertainty of data point k.

Parameter optimisation

Now that we know how to evaluate any individual set of parameters, we need an efficient way to explore parameter space and locate the most likely set of parameters. By choosing the appropriate algorithm we will be able to iteratively reconstruct the posterior probability density function (pdf) through repeated sampling of the posterior. One of the most common approaches to this problem is through so-called *Markov-chain Monte Carlo sampling* (MCMC). The essentials of the method are to construct a time dependent process (Markov chains) which explores the entire parameter space through some random process (Monte Carlo). The process can be can be summarised as:

- 1. Randomly sample a starting point in parameter space, $\vec{\theta_i}$
- 2. Compute $\mathcal{L}(\vec{\theta_i})$ at that position
- 3. Sample a new candid set of parameters $\vec{\theta'}$ by drawing from some distribution.
- 4. Compute $\mathcal{L}(\vec{\theta}')$ at the new candidate position
- 5. If the candidate position is accepted $\vec{\theta}_{i+1} = \vec{\theta}'$; otherwise $\vec{\theta}_{i+1} = \vec{\theta}_i$

6. Repeat from step 3

One often used method for MCMC is the Metropolis-Hasting's algorithm. The essence of this approach can be boiled down the accept/reject criteria. A random uniform number is drawn $u \in [0, 1]$, if $\mathcal{L}(\vec{\theta'})/\mathcal{L}(\vec{\theta_i}) \geq u$ then the candidate position is accepted, and rejected otherwise. After some number of steps the MCMC chain should make its way toward the minimum of the χ^2 surface. The conditions of step 5 allow for some probability that a worse parameter set can be selected. This serves to prevent the chain from getting stuck in a local minimum. Given enough time the entire parameter will space be explored. The rate at which this can be done depends not only on the topology of the χ^2 surface but also the step size in step 3 and the acceptance criteria set in step 5. To Increase the efficiency of parameter space exploration many chains can be run in parallel, this is called an ensemble. In this work we refer to the leading point of every chain as *walkers*. If we are primarily interested in locating all peaks on the \mathcal{L} -surface there are more efficient algorithms available, though they often share this same core methodology.

The fitting procedure with EMERGE is performed in two stages. First, the broader parameter space is explored with either *parallel tempering* or *hybrid* optimisation methods. These are ensemble methods that increase the rate of parameter exploration by allowing walkers to communicate information about their current state with one another. This way the global maximum likelihood can be more rapidly isolated. After using one of these methods to identify peaks in the posterior, we isolate the maximum peak and perform a followup MCMC run. This followup run is performed with a much tighter allowable parameter range and is intended to better characterise the parameter uncertainty about the most likely peak.

The parallel tempering method builds on standard MCMC by allowing some walkers in the ensemble to explore the parameter space more freely (Sambridge, 2014). In this approach every walker is assigned a temperature. We designate some fraction of the walkers (typically 25 per cent) as cold walkers with T = 1. The remaining walkers have their temperatures spaced log-uniformly between $T = (1, T_{\text{max}}]$. Candidate walkers are selected by drawing from a multivariate Gaussian distribution centred on $\vec{\theta_i}$ with width σ which is dynamically adjusted such that each walker reaches a target acceptance rate, in our case 30 per cent. The acceptance criteria is augmented by multiplying the likelihood ratio by 1/T. Therefore walkers with a higher temperature are more likely to accept step with a 'worse' parameter set, while cold walkers update position according to the standard Metropolis-Hastings accept/reject criteria. At the end of each epoch chains are allowed to randomly exchange locations with another walker of a different temperature. For multimodal parameter spaces this approach provides additional robustness against local minima. Another advantage of this technique is that because cold walkers operate just as in standard MCMC, their chain histories can be used to construct a posterior pdf.

Another technique is the HYBRID method (Elson et al., 2007). This approach combines approaches of simulated annealing (Kirkpatrick et al., 1983) and particle swarm optimisation (PSO) (Kennedy & Eberhart, 1995). In the previous methods we discussed, each walker is treated independently with their step size and acceptance determined only by their individual performance. The HYBRID methods employs PSO which modifies the step size of each walker by providing information on the ensemble performance as well as the current individual walker performance. Just as before candidate walker steps are drawn from a Gaussian distribution, except now the width σ is rescaled by a global performance factor g and an individual performance factor f. The global factor adjusts the step size as

$$g = q^{\alpha} \tag{4.27}$$

$$q = \langle \chi^2 \rangle / \langle \chi_0^2 \rangle,$$

where $\langle \chi^2 \rangle$ is the *current* ensemble average, $\langle \chi_0^2 \rangle$ is the *initial* ensemble average and α is a scaling parameter that we have set to 0.4 for this work. From this we can see that as the ensemble improves the step size will become smaller for all walkers. If a walker is currently performing worse than average the step size will be rescaled linearly as

$$f = 1 - \beta p \qquad (4.28)$$
$$p = \frac{\langle \chi^2 \rangle - \chi_i^2}{\chi_{\max}^2 - \langle \chi^2 \rangle},$$

where χ^2_{max} is the current ensemble maximum, χ^2_i is the walker performance and β is a linear slope that we have set to $\beta = 1$. If a walker is currently performing better than average the step-size will decay as a power law according to

$$f = (p+1)^{-\gamma}$$

$$p = \frac{\langle \chi^2 \rangle - \chi_i^2}{\langle \chi^2 \rangle - \chi_{\min}^2},$$
(4.29)

where χ^2_{\min} is the current ensemble minimum and γ is a power-law slope that we have set to $\gamma = 1$ for this work. Walkers that perform worse than average will be allowed to explore parameter space more broadly while walkers that perform better than average will explore their local space more intensely. Finally, simulated annealing provides a time dependency to walker acceptance. Each walker is assigned a temperature similar to the parallel tempering method, in this case we assign a uniform initial temperature T = 100. Instead of keeping the temperature fixed, simulated annealing allows the walker temperature to drop over time $\propto \ln(1+j)^1$ where j is the epoch number. The accept/reject step is then performed just as in parallel tempering. This method is both robust against local minima and offers rapid convergence.

We perform the final fitting stage using the affine invariant MCMC ensemble method described in Goodman & Weare (2010). This advances the Metropolis-Hasting's algorithm by improving the methods for determining sample step and candidate acceptance. Here the step size is determined by drawing a number Z from the distribution defined by:

$$g(z) \propto \frac{1}{\sqrt{z}} \in [1/a, a] \tag{4.30}$$

where a is a scale parameter that is used to adjust the acceptance rate. For everything here we set a = 4. The candidate point is selected by stretching $\vec{\theta}_i$ in the direction of another randomly selected walker in the ensemble $\vec{\theta}_j$ such that⁷

$$\vec{\theta}' = \vec{\theta}_i + Z(\vec{\theta}_i - \vec{\theta}_j). \tag{4.31}$$

Rather than use ratio of \mathcal{L} directly, instead the candidate point is accepted if

$$Z^{N-1}\mathcal{L}(\vec{\theta'})/\mathcal{L}(\vec{\theta_i}) \ge u$$

Model selection

Now that we understand how to individually evaluate our models, we can compare them with one another to evaluate their relative goodness of fit and balance that with their complexity to find our most likely model. We use techniques from both information theory and Bayesian information theory to evaluate model options side by side.

The first approach we look at is the Akaike Information Criterion (AIC) (Akaike, 1974). The AIC score is attractive because it is simple to compute, requiring only information on the maximum likelihood \mathcal{L}_{max} , and the number of free parameters k. the function can be defined by

$$AIC = -2\ln \mathcal{L}_{max} + 2k . \tag{4.32}$$

In this evaluation the best model is the one that minimises eq. 4.32. We can see that for two models with equivalent \mathcal{L}_{max} the model with fewer free parameters would be preferred. To account for fits against small sample sizes an alternative *corrected* AIC score is available (Sugiura, 1978). This takes the form

AICc = AIC +
$$\frac{2k(k+1)}{N-k-1}$$
, (4.33)

where N is the number of observed data points included in the fitting routine. The AICc score is further penalised when N/k is small but vanishes for $N \gg k$.

Similar to the AICc the Bayes information criterion (BIC) (Schwarz, 1978) incorporates information on the size of data set.

$$BIC = -2\ln \mathcal{L}_{max} + k\ln N \tag{4.34}$$

Again, the preferred model is the one that minimizes the BIC.

The AIC(c) and BIC are attractive for their computational simplicity but they also presume that each data point is independent and the parameters are not degenerate. The Deviance information criterion (DIC) (Spiegelhalter et al., 2002) addresses these issues by incorporating information on the total posterior surface via the Bayesian complexity

$$p_D = \left\langle \chi^2(\theta) \right\rangle - \chi^2_{\min} , \qquad (4.35)$$

⁷This is performed by splitting the ensemble into two subsets of equal number walkers. The secondary walker is selected at random from the other subset.

where $\chi^2 = -2 \ln \mathcal{L}$ as in eq. 4.25. The term p_D can be understood as the an *effective* number of parameters. The DIC can then be defined by

$$DIC = -2\ln \mathcal{L}_{max} + p_D . \tag{4.36}$$

In the case were model parameters are constrained by the data $p_D \rightarrow k$ and the DIC score resembles AIC. Therefore a model is not penalised if additional parameters are added that are not constrained by the data.

The most complete approach involves computing the model evidence directly, which requires knowledge of the complete posterior likelihood. By using eq. 4.19 and eq. 4.21 we can show that model evidence can be determined by marginalising over the available parameter space of a given model which can be written as:

$$\mathcal{Z} \equiv \operatorname{Prob}(D|\mathcal{M}) = \int \operatorname{Prob}(D, \vec{\theta}|\mathcal{M}) d\vec{\theta} = \int \operatorname{Prob}(\vec{\theta}|\mathcal{M}) \operatorname{Prob}(D|\vec{\theta}, \mathcal{M}) d\vec{\theta}$$
(4.37)

Therefore by marginalising over $P(D|\vec{\theta}, \mathcal{M})$, which is the same likelihood we computed with eq. 4.25, then we can determine the probability of the data given some model. To compute \mathcal{Z} we use the method described in Weinberg (2012) and Weinberg et al. (2013).

Best fit model and parameters

Through the methods discussed above Moster et al. (2018) specified a best fit model totalling 11 free parameters. Since the time of publication new data and model options have been added to and EMERGE was refit after these changes were made. The addition of an orphan halo mass loss formulation along with new SMF data naturally resulted in a set of best fit parameters different from previously published results. While the most notable parameter change was f_s , other parameters have moved beyond the 1σ uncertainty ranges quoted in Moster et al. (2018). This movement does not reflect tension in the observed data sets, but is more likely the result of the complex topology of the high-dimensional parameter space. The posterior surfaces contains many local minima separated by many sigma from one another with each exhibiting a reasonable fit to the observations. The best fit parameters following the noted updates are show in table 4.1. These model parameters were used for the analysis discussed in §5 and §6. Additional model options were incorporated for the results discussed in §7. The details of those model changes are discussed in the context of their application to dwarf galaxies.

Parameter	Best-fit	Upper 1σ	lower 1σ
M_0	11.34829	+0.03925	-0.04153
M_z	0.654238	+0.08005	-0.07242
ϵ_0	0.009010	+0.00657	-0.00451
ϵ_z	0.596666	+0.02880	-0.02366
β_0	3.094621	+0.15251	-0.14964
β_z	-2.019841	+0.22206	-0.20921
γ_0	1.107304	+0.05880	-0.05280
f_{esc}	0.562183	+0.02840	-0.03160
f_s	0.004015	+0.00209	-0.00141
$ au_0$	4.461039	+0.42511	-0.40187
$ au_s$	0.346817	+0.04501	-0.04265

Table 4.1: The best fit model parameters used for this work.

Chapter 5

Clashing giants: galaxy growth through mergers

In the hierarchical picture of galaxy formation within the Λ CDM framework, mergers play a critical role in the formation and continued evolution of galaxies. Consequently the galaxy-galaxy merger rate and its dependence on mass, mass ratio, and redshift are of fundamental interest. The frequency of galaxy mergers cannot be observed directly and so we must rely on theoretical models for galaxy formation along with a robust set of observations to ascertain the cosmological galaxy-galaxy merger rate¹.

Many theoretical models build upon the foundation laid by dark matter (DM) only *N*body simulations, with each model applying a different method for populating DM haloes with their constituent galaxies. The underlying halo-halo merger has largely converged among various theoretical models (Fakhouri & Ma, 2008; Fakhouri et al., 2010; Genel et al., 2009, 2010). Despite this agreement in the foundational structure of galaxy evolution, theoretical models for galaxy formation have yet to establish a sufficiently accurate value for the galaxy-galaxy merger rate. There remains as a much as an order of magnitude discrepancy in the predicted values depending on mass, mass-ratio, redshift, and theoretical framework (Bower et al., 2006; Croton et al., 2006; Maller et al., 2006; De Lucia & Blaizot, 2007; Font et al., 2008; Somerville et al., 2008; Khochfar & Silk, 2009; Stewart et al., 2009; Hopkins et al., 2010a,b; Rodriguez-Gomez et al., 2015).

Similarly, observed merger rates have not converged so far, with different rates even derived from the same fields (Mantha et al., 2018; Duncan et al., 2019). Many of the discrepancies can be attributed to varying definitions on the merger rate, including whether galaxy pairs are selected based on their stellar mass or their luminosity (Lotz et al., 2011; Man et al., 2016). Furthermore unreliable redshift measurements introduce considerable uncertainty in the selection of physically associated pairs. Additionally, merging timescales must be separately derived using theoretical models. However, considerable uncertainty remains in these merging timescales and how they might scale with redshift.

¹The contents of this chapter appear as printed in O'Leary et al. (2021a). portions of the introduction have been moved to $\S4.1.1$ and the section on observation timscales has been shifted to $\S6$ to consistency

The primary goal of this chapter is to determine the cosmological galaxy-galaxy merger rate, and its dependence on properties such as the stellar mass of the main galaxy, the stellar mass ratio between both galaxies, and the redshift of the merger. Instead of predicting the merger rate with a model that makes assumptions on the baryonic physics, we derive it empirically, solely based on the evolution of observables such as the stellar mass function and star formation rates, within a Λ CDM cosmology. To this end we employ the empirical galaxy formation model EMERGE² (Moster et al., 2018, 2020b). We further investigate how the stellar mass of galaxies grows over cosmic time, and whether this growth mainly comes from star-formation within the galaxy, major mergers, or minor mergers. In this context, we also study how many major and minor mergers a galaxy typically has over its lifetime. We perform our analysis in the context of our empirical model, but we compare our results to observational evidence and other theoretical work to estimate the robustness of our conclusions³.

This chapter is organised as follows; In $\S5.1$ we provide an overview of the *N*-body simulation we use. We outline our methodologies and fundamental results in $\S5.2$, and discuss how the merger rate scales with stellar mass, mass ratio, redshift, and star formation rate. Here we also we illustrate how our model compares with other theoretical predictions. Finally, in \$5.3 we explore the merging history of present day galaxies. Here, we determine which galaxies are grown through merging, and which type of mergers matter for stellar mass growth.

5.1 Dark Matter Simulations

Our analysis of galaxy-galaxy merger rates relies on producing galaxy merger trees encompassing a large dynamic range, occupying an appropriately large cosmic volume. We employ the empirical model EMERGE to populate dark matter haloes with galaxies based on individual halo growth histories. In this section we discuss the fundamental tools used to ultimately produce galaxy merger trees. Throughout this paper we adopt Planck Λ CDM cosmology (Planck Collaboration, 2016) where $\Omega_m = 0.3070$, $\Omega_{\Lambda} = 0.6930$, $\Omega_b = 0.0485$, where $H_0 = 67.77 \,\mathrm{km \, s^{-1} Mpc^{-1}}$, $n_s = 0.9677$, and $\sigma_8 = 0.8149$.

5.1.1 Obtaining halo merger trees

We utilise a cosmological dark matter only N-body simulation in a periodic box with side lengths of 200 Mpc. The initial conditions for this simulation were generated using MUSIC (Hahn & Abel, 2011) with a power spectrum obtained from CAMB (Lewis et al., 2000). The simulation contains 1024^3 dark matter particles with particle mass $2.92 \times 10^8 M_{\odot}$. The simulation was run from z = 63 to 0 using the Tree-PM code GADGET3 (Springel, 2005). In total 94 snapshots were created evenly spaced in scale factor ($\Delta a = 0.01$). Dark matter haloes are identified in each simulation snapshot using the phase

²The code can be obtained at https://github.com/bmoster/emerge

³Scripts for reproducing our primary results can be obtained at https://github.com/jaoleary

space halo finder, ROCKSTAR (Behroozi et al., 2013a). Halo merger trees are constructed using CONSISTENTTREES (Behroozi et al., 2013c), providing detailed evolution of physical halo properties across time steps. Throughout this paper we use the term 'main halo' to designate haloes which do not reside within some other larger halo, and 'subhalo' to refer to haloes contained within another halo.

5.1.2 Halo-Halo mergers

Prior to evaluating the galaxy-galaxy merger rate we take a look at the halo-halo merger rate. Due to our model's reliance on the individual growth histories of dark matter haloes, it is important to verify that our simulation is assembling haloes in a manner consistent with other theoretical predictions and models (Genel et al., 2009, 2010; Fakhouri et al., 2010). We compute the halo-halo merger rate directly using the trees constructed with CONSISTENTTREES. We define halo mergers at the time when a halo first becomes a subhalo, not when the subhalo becomes disrupted. This is consistent with the definition of halo mergers adopted by Genel et al. (2009, 2010); Fakhouri et al. (2010). The merger rate is then calculated at each redshift as a function of the descendant halo mass M_0 , and the mass ratio $\xi = M_i/M_1$ of the progenitor haloes (for i > 1), where M_1 is the most massive progenitor to M_0 .

Figure 5.1 shows the mean halo-halo merger rate per descendant halo. When taking the halo-halo merger rate per halo, we find rates that adopt the same nearly mass-independent scaling shown in previous works (Genel et al., 2009, 2010; Fakhouri et al., 2010), The top panel shows that the merger rate per Gyr exhibits a strong power-law scaling with redshift. The bottom panel shows that, just as in previous works, we find a scaling $\propto \xi^{-2}$ for a fixed redshift interval. Our results indicate that our underlying N-body simulation is in agreement with other works.

At this point we are free to implement our galaxy formation model. In this process we will see how this simple universal halo merger rate becomes transformed through the complex connection between galaxies and their haloes.

5.2 The Galaxy-Galaxy merger rate

In this section we discuss the galaxy-galaxy merger rates intrinsic to and derived from EMERGE. First we present the *intrinsic* merger rate, that is the rate at which galaxies are merged using the processes outlined in §4.1.2. The intrinsic merger rate provides insight into the actual buildup of stellar material using the internal mechanics of the empirical model. We then present a merger rate derived using mock observations applied to mock galaxy catalogues. This provides a bridge to more completely address any discrepancies between theoretical models and observations.

First we should address some terminology common to both approaches. Each galaxy merger can be classified in terms of stellar mass and stellar mass ratio:

 m_0 : The stellar mass of the *descendant* galaxy at the snapshot following the merger.



Figure 5.1: Halo-halo merger rates. The coloured lines indicate halo-halo merger rates from our simulation for the noted descendant halo masses. For ease of comparison we adopt the mass ratio definition of Fakhouri et al. (2010), where $\xi = M_i/M_1 = 1/\mu$ (see §5.2). *Top panel*: Major ($\xi > 0.25$) Halo-halo merger rates *per Gyr*, scaling with redshift. The bumps for higher mass haloes is due to low number statistics. *Bottom panel*: Halo-halo merger rates scaling with mass ratio ξ at z = 0.1. The black dash-dash and dash-dot lines shows the best fit merger rate for a halo with $M_0 = 10^{12} M_{\odot}$, from Genel et al. (2009) and Fakhouri et al. (2010) respectively.

- m_1 : The stellar mass of the *main* progenitor galaxy defined at the snapshot just prior to the merger.
- m_2 : The stellar mass of the co-progenitor galaxy at the snapshot just prior to the merger.
- μ : The stellar mass ratio taken with respect to the progenitor galaxies, $\mu \equiv m_1/m_2$. In the most general case the main progenitor m_1 is also the most massive progenitor. Due to scatter in the stellar-to-halo mass relation (SHMR) there are some scenarios under which $m_2 \downarrow m_1$. In these cases we invert this relation such that $\mu \geq 1$. These special cases represent fewer than 5 per cent of all mergers with a descendant mass larger than $10^9 M_{\odot}$.

5.2.1 Intrinsic merger rate

Having constructed galaxy merger trees, computing the merger rate is straight forward. In the trees we identify galaxy mergers as any pair of galaxies sharing an identical descendant galaxy. In each case we assume every merger is binary and occurs instantaneously at t_{df} .⁴ We provide two measures for the merger rate: merger rate per comoving volume and merger rate per galaxy. The process is similar in each case, with the difference arising in how the rate is normalised. The first step is to construct bins for time, stellar mass and mass ratio. For each time bin we count the number of mergers which satisfy the mass and mass ratio requirements, and whose merging time (t_{df}) resides within that bin. When computing the merger rate per comoving volume we then divide the number of mergers (N_{merge}) in each bin by the bin widths dt and by the volume of our box, so the merger rate is given by:

$$\Gamma = \frac{N_{\text{merge}}}{\mathrm{d}V\mathrm{d}t} \left[\mathrm{cMpc}^{-3}\mathrm{Gyr}^{-1}\right] \,. \tag{5.1}$$

For the merger rate per galaxy we instead divide the number of mergers in each bin by the bin width, and the number of galaxies (N_{gal}) at the central redshift of the bin. At these redshifts N_{gal} is determined by linearly interpolating counts between the nearest two snapshots. The merger rate per galaxy is then described by:

$$\mathfrak{R} = \frac{N_{\text{merge}}}{N_{\text{gal}} \, \text{d}t} \left[\text{Gyr}^{-1} \right] \tag{5.2}$$

While operating on the same data the two measures for merger rate produce qualitative differences due to the scaling of the merger rate per galaxy with the number density of galaxies. For this reason the merger rate per galaxy is often the preferred measure as it is more robust against cosmic variance. We explore both rate measures to provide a more complete comparison with other works.



Figure 5.2: The galaxy-galaxy merger rates as a function of redshift and mass ratio, $\mu \equiv m_1/m_2$. The top row of panels show the merger rate density $\Gamma(z)$, i.e. the total number of mergers per *comoving* volume, and the lower panels show the merger rate *per galaxy* $\Re(z)$. Solid lines indicate galaxy mergers selected based on the descendant mass of the merging system, m_0 . In this scenario the progenitor galaxies (m_1, m_2) are permitted to have masses outside the noted mass bands. The dashed lines show mergers selected based on the *main* progenitor mass (m_1) of the merging systems. Here only the main progenitor must reside in the specified mass band. The shaded regions indicate Poisson error in the number of mergers.

Scaling with mass and redshift

Figure 5.2 illustrates the merger rate scaling with redshift for three different mass bins. For each mass bin we select mergers based on descendant mass m_0 (solid lines), or main progenitor mass m_1 (dashed lines). Selection based on descendant mass probes galaxy mergers after their mass contributions have been added into the descendant galaxy. This is a common approached taken by other theoretical works exploring galaxy merger rates. This method presumes we have knowledge on how the mass of two galaxies is added together in a merger, something not easily accessible to observations. Selecting mergers based on progenitor properties provides a more observer friendly measure of the merger rate as it can be more directly compared with observationally derived rates (see Section 6.3.1).

In each mass band we show the merger rate for three different mass ratio intervals. The first, and arguably the most important, are called *major mergers* (blue lines), although the precise mass ratio that defines a major merger is not well defined and varies across literature. The key characteristic of major mergers are their transformative properties. Such mergers are very disruptive to both systems, suspected of prompting drastic changes in stellar populations and descendant morphologies. In this paper we take major mergers to be $1 \le \mu \le 4$. The next mass ratio interval are similarly labelled *minor mergers* (green lines). Minor mergers while not as individually disruptive as their larger counterparts, still contribute to the evolutionary process of large galaxies. There is evidence to suggest that such mergers, if occurring at high enough frequency, can produce some of the same morphological changes generated through major mergers (Naab et al., 2009; Oser et al., 2010; Hilz et al., 2012, 2013; Karademir et al., 2019), lead to the thickening of disc galaxies (Abadi et al., 2003; Kazantzidis et al., 2008; Purcell et al., 2009; Moster et al., 2010a, 2013), and even drive the rotation speed of massive early type galaxies (Bois et al., 2011; Moody et al., 2014; Penovre et al., 2017; Schulze et al., 2018). The final category are so called *mini mergers* (orange lines). As their name entails, this category represents the smallest merging events. While not terribly transformative, understanding their frequency is helpful for constructing a complete mass budget and internal radial distribution for the accreted material of a galaxy.

The merger rate per comoving volume (Γ) exhibits a nearly mass independent shape in the number of mergers occurring. For each mass and mass ratio interval we find a sharp increase in the number of mergers at low redshift, with a well defined peak $1 \leq z \leq 2$. Beyond the peak we see a rapid decay in the total number of mergers towards high redshift. When we instead take the merger rate in the context of an evolving galaxy population (\Re_{merge}) we find quite a different trend. In general we find the merger rate increases with redshift at all masses and mass ratios. However, our results do not show a simple power-law scaling with redshift. For major mergers, at low and intermediate mass, we find an excess over a power-law for $z \gtrsim 3$. This break from a power-law redshift scaling is primarily

⁴This binary assumption holds true for > 98 per cent of all mergers with a descendant mass > $10^9 M_{\odot}$. Despite this small number of cases we nonetheless adjust the descendant mass of each merger to ensure the descendant mass only reflects mass contributed by the most massive progenitor and the merging satellite.

evident when selecting mergers based on descendant galaxy stellar mass.

Additionally, We can observe a discrepancy in the merger rate when selecting merger rates based on progenitor mass vs. descendant mass. In the lowest mass bins the difference in these two quantities only becomes manifest at higher redshifts ($z \gtrsim 2$), where selecting based on progenitor mass produces a noticeably lower merger rate. However, for the most massive galaxies the difference is more dramatic. We find that in this mass range both measures produce functionally similar results, with a nearly constant scaling offset, where the descendant mass selected major merger rate is a factor of 1.5 - 2 larger than that produced with progenitor selection. To explain this, we should recall that in our approach the average number of galaxies remains fixed for any time and mass interval, that is N_{gal} of eq. 5.2 remains the same regardless of whether we compute the rate based on descendant mass or progenitor mass. Additionally, for a lower mass threshold, there are in general more galaxies that fit within our mass range at timestep *i* compared to *i*-1. At the highest masses and redshifts this effect is amplified due to the low numbers of galaxies present. This is similar to the argument by Genel et al. (2009) to explain such differences in the context of halo-halo merger rates.

Lastly, we can look across the mass panels to see how the frequency of major mergers changes with redshift. For low-mass galaxies it is clear that minor and mini mergers dominate. When moving to intermediate-mass ranges, galaxies like the Milky Way, we find that for the first half of cosmic time minor and mini mergers occur at greater frequency than major mergers. Near $z \approx 1$ this changes, as major mergers become more frequent than minor mergers and occur at nearly the same frequency as mini mergers by z = 0. Finally, for the most massive galaxies, at $z \leq 3$ major mergers are the most frequent, with mini mergers being the next most common, and minor mergers making up the smallest fraction of mergers. We explore these strong mass ratio dependencies of the merger rate in closer detail in the next section.

Scaling with mass ratio

In this section we address the galaxy-galaxy merger rate scaling with mass ratio. We have already seen in the previous section that the merger rate density and rate per galaxy appear to possess a strong scaling with mass for a fixed mass ratio interval. Most notably, major mergers dominate at high-mass, and mini/minor mergers dominate at low masses. Our goal here is to investigate this inflection with a finer μ binning in order to explain this phenomenon.

Figure 5.3 explores the relationship between descendant stellar mass and merger mass ratio. To provide a more granular view of the mass ratio distribution, without influence from the explicit definition of major(minor) merger, we express these results as the merger rate per galaxy per \log_{10} mass ratio interval $N_{\text{merge}}/\text{dlog}_{10}(\mu)/\text{dt}$. In the top panel we take mergers of a fixed mass ratio at a fixed redshift. We then bin those galaxies according to the stellar mass of their descendant system. This way we are able to see the relative importance of some merger as a function of mass. We selected four target mass ratios with a uniform log space bin of 0.45 dex. In general each mass ratio shows a similar qualitative


Figure 5.3: The galaxy-galaxy merger rate as a function of descendant stellar mass and progenitor mass ratio. Top panel: The merger rate per galaxy as a function of descendant stellar mass, m_0 . The coloured lines specify the merger rate for a specific mass ratio μ . Shaded regions indicate Poisson error in the number of mergers. Bottom panel: The merger rate per galaxy as a function of progenitor mass ratio. Each line includes mergers with a descendant mass noted by the colour. The x-axis shows the distribution of merger mass ratios experienced within each mass band. The vertical dashed black line denotes the threshold for major mergers at $\mu = 4$. Shaded regions surrounding solid lines indicate Poisson error in the number of mergers. The red and blue shaded regions show the best fit halo-halo merger rate for haloes with $11.25 \leq \log_{10}(M/M_{\odot}) < 13$, from Genel et al. (2009) and Fakhouri et al. (2010) respectively.



Figure 5.4: The z = 0 stellar-to-halo mass relation produced by EMERGE. The solid line illustrates the best fit over the mean values of the data. The shaded regions show how a major merger in halo mass ($\mu_M = 4$) translates to galaxy-stellar mass ratio (μ_m) on both the low and high-mass end of the stellar to halo mass relation. Along the best fit curve we can see that a major halo-halo merger on the low-mass end will result in a very minor merger in stellar mass. On the high-mass end a major halo-halo merger will result in a major merger in stellar mass.



Figure 5.5: The *cumulative* galaxy-galaxy merger rate as a function of descendant stellar mass and progenitor mass ratio. The data shown here is the cumulative histogram of the data displayed in Fig. 5.3. The cumulative merger rate can be used to determine the merger rate for a desired mass and mass ratio. Top panel: The cumulative merger rate per galaxy as a function of descendant stellar mass, m_0 . The coloured lines each specify a different mass ratio threshold, and include all mergers below that threshold. Bottom panel: The merger rate per mass ratio interval. Each line includes mergers with a descendant mass noted by the colour. The x-axis therefore shows the distribution of merger mass ratios experienced within each mass band. The black line denotes the threshold for major mergers $\mu = 4$. As an example (square black point) we can see that at z = 0.1 for $\log_{10}(m_0/M_{\odot}) = 11.0$ approximately 4 per cent of galaxies will experience a merger in the next Gyr.

trend. At lower masses the slope stays relatively flat, approaching high masses each curve show an inflection point when more massive galaxies begin to experience more mergers. Looking closely we see that the inflection occurs at lower stellar mass for lower redshifts.

Similarly, The bottom panel of Figure 5.3 shows how merger mass ratios are distributed for a fixed descendant mass and redshift. This plot is analogous to Figure 5.1 for the halohalo merger rate. In this comparison we can immediately note some glaring differences compared to the halo-halo merger rate, most notably the galaxy-galaxy merger rate does not show the same mass independent merger rate with respect to mass ratio. In the case of galaxy-galaxy mergers we break from a simple power law to a more complex relation with a clear redshift, mass ratio and strong mass dependence. For mergers with $\mu \gtrsim 10$ all masses exhibit a similar merger rate, with some minor scaling differences with increasing redshift. Below $\mu \approx 10$ we find a greater dependence on mass. In this regime the most massive galaxies maintain a nearly constant slope for all μ , even experiencing a boost in the major merger rate.. Conversely, for low-mass and intermediate-mass systems we see a dropping merger rate, illustrating suppressed major mergers. Narrowing in on galaxies with $\log_{10}(m_0/M_{\odot}) = 10.5$, we see an interesting evolution. At $z \approx 2$ these galaxies scale with mass ratio similar to the lowest mass galaxies, with a suppressed rate for $\mu \lesssim 10$. As we transition to lower redshift we can see this relationship change, with major mergers becoming more prevalent at lower redshift, and the once decreasing trend bending up to meet the clean scaling seen for massive galaxies. Both of these trends can be explained in the context of an evolving stellar-to-halo mass relation (SHMR).

First, why are major mergers suppressed for low-mass galaxies? The first driver for this effect can be seen directly in the SHMR. Figure 5.4 shows the SHMR present in our model for galaxies at z = 0. Once again recalling the simple relation assumed by the halo-halo merger rate (Figure 5.1), we can trace how a major merger in halo mass would translate to mergers in galaxy stellar mass. In this thought experiment we can make the presumption that any halo-halo merger will eventually result in a galaxy-galaxy merger. Starting along the low-mass slope of the SHMR, we note that if we select the average galaxy masses for a fixed halo mass, a major merger in halo mass would translate to a mini-merger ($\mu \approx 30$) in galaxy stellar mass. Conversely, we can observe the opposite scenario on the high-mass slope, beyond the turnover. Due to the shallow slope in the SHMR for high masses we can see that a major halo-halo merger has a much larger likelihood of also leading to a major merger in galaxy stellar mass. In short, if the slope of the SHMR is unity, we would expect a major halo-halo merger to directly lead to a major galaxy-galaxy merger. Subsequently, where the slope is greater than unity we expect a suppression of major mergers, and where less than unity we expect an enhanced rate of major mergers. The second driving factor is dynamical friction. We see that small satellites orbiting a massive central galaxies have much longer dynamical frictions times (eq. 4.10) and would simply not have had enough time to merge. This effect prevents minor halo mergers from being transformed into major mergers where the stellar mass ratio of the two systems would otherwise be sufficient.

Now, why do we see that intermediate-mass galaxies with $\log_{10}(m_0/M_{\odot}) \approx 10.5$ have suppressed major mergers at high redshift but not at low redshift? Here we move beyond the static low redshift SHMR, and instead focus on how the SHMR evolves. These intermediate-mass galaxies reside close to the turnover on the SHMR. From z = 2 to z = 0 we see the average halo mass for such galaxies increase by ~ 0.08 dex. That is these galaxies tend to live in larger haloes at lower redshift. Additionally, the turnover in the SHMR has a mild shift to lower halo mass by ~ 0.17 dex. These combined effects mean these galaxies tend to sit higher along the turnover where the slope approaches unity. Thus, these galaxies begin to experience more major mergers with decreasing redshift. This can be clearly seen in the much flatter major merger rate scaling with redshift shown in Figure 5.2. The suppression and enhancement of major galaxy mergers as a consequence of the SHMR has been explored in the context of other semi-empirical models (Stewart et al., 2009; Hopkins et al., 2010a). However, it is important to stress the necessity of a model that can accurately reproduce the observed data (e.g. SMFs, cosmic and specific SFRs). Models that show significant deviations from these fundamental observations, or cannot self-consistently track their redshift evolution may arrive at differing conclusions regarding galaxy merger rates or galaxy mass assembly.

Finally, in Figure 5.5 we provide the cumulative galaxy-galaxy merger rates with respect to mass ratio. The information contained in the cumulative merger rates is identical to that of Figure 5.3. Absent a generalised fitting function for our results, the cumulative rates provide a quicker reference for determining the number of mergers occurring at some descendant mass for a given mass ratio interval.

Active vs. Passive galaxies

Finally we address SFR dependencies of the cosmic galaxy-galaxy merger rate. Once again looking back to the baryon conversion efficiency (eq. 4.3) of galaxies we can see a characteristic halo mass (see Table 4.1 and eq. 4.4) at which a galaxy is most efficient at converting gas into stars. One conclusion from this relation is that larger galaxies are inefficient at creating new stars. Thus it is important to understand the merger rate for these specific galaxies to learn how galaxy mergers drive their galaxies' continued mass growth (Khochfar & Silk, 2009). Specifically we would like to know if these galaxies are grown through the merging of other large quenched galaxies, or constructed more slowly through the accretion of smaller star forming satellites. Further, understanding these mergers may help explain how mergers initiate star formation, or power AGN (Hirschmann et al., 2010, 2014, 2017; Weinberger et al., 2017; Choi et al., 2018; Steinborn et al., 2018).

We begin by defining our galaxies in terms of their star formation properties. Broadly this means designating a galaxy as *passive* or *active*, where passive galaxies are quenched and active galaxies are actively star forming. We adopt the quenching criteria of Franx et al. (2008) to make this distinction, where a galaxy is considered quenched if:

$$\Psi < 0.3t_z^{-1} , \tag{5.3}$$

where t_z is the age of the universe at redshift z, and Ψ is the specific star formation rate.

In Figure 5.6 we illustrate how the *major* merger rate scales when selecting mergers based on the star formation properties of their progenitor and descendant systems. We perform this in two mass bands for galaxies with $\log_{10}(m)/M_{\odot} \geq 10$. For each mass



Figure 5.6: The major merger rate ($\mu \leq 4$) as a function redshift for various active/passive progenitor configurations. We show the four possible scenarios for active/passive combinations in progenitor systems. Active-active would specify a merger where both progenitor galaxies are actively star forming, passive-active specifies a configuration where the main progenitor is passive (quenched) and the secondary progenitor is active (star forming), etc. Top Panel: The galaxy merger rate dependency on star forming properties of the progenitor galaxies. The black line shows the total merger rate for all star forming configurations of the descendant system. Bottom Panel: The galaxy merger rates only considering mergers with a quenched descendant galaxy. At high redshift most mergers occur between active progenitors. There is a transition near $z \approx 0.5$ where mergers start to become dominated by passive progenitors.

bin we compute a global merger rate, only considering the star formation of the progenitor galaxies. Additionally, we perform the same analysis only considering mergers with a quenched descendant galaxy. We designate four different scenarios based on the star forming properties of the progenitor galaxies:

Active-active: Both progenitors are active.

Passive-passive: Both progenitors are passive.

Passive-active: The main progenitor is passive and the secondary progenitor is active.

Active-passive: The main progenitor is active and the secondary progenitor is passive.

In Figure 5.6 we compare the total merger rates when considering all galaxies (upper panels) versus only considering mergers with a quenched descendant galaxy (lower panels). For the redshift range shown we find very little difference in the total merger rates. This suggests that most major mergers are occurring in dense environments around an already quenched central galaxy. For the most massive galaxies, by $z \approx 1$ most mergers are occurring between two passive galaxies (red lines), or between a passive central galaxy and an active satellite (yellow lines). Beyond $z \approx 1$ most mergers are occurring between active galaxies (blue lines). When considering only mergers with a quenched descendant (bottom panel) we find a nearly constant merger rate if the major galaxy in the merger is already quenched. There are several different effects that lead to this result. The first is the prevalence of gas-rich (active) galaxies at high redshift making the likelihood of activeactive mergers greater. The second being that if a central galaxy is quenched, it is likely that its descendant galaxy will also be quenched. Due to the lack of passive galaxies at high z these results are more uncertain than the results shown in Figure 5.2. For the lower mass bin, mergers involving any quenched component cannot be constrained to better than a factor $\sim 2-3$ for $z \gtrsim 1.5$. For the higher mass bin we see a similar degree of uncertainty for $z \gtrsim 2.25$

5.2.2 Comparison to other theoretical predictions

Although most theoretical predictions are based on the same Λ CDM framework, the methods used to link DM haloes and galaxy properties has direct consequences on the predicted merger rates. Figure 5.7 displays a side-by-side comparison of galaxy merger rates in three different mass bands produced by a diverse set of models. While our results might initially be surprising, we can see that within the context of other theoretical predictions we are firmly within a previously established range of merger rates.

Theoretical methods for determining the merger rate can be roughly broken down into a few categories:

1. Halo-halo: Assume an average halo-halo merger rate with a dynamical friction delay set at $R_{\rm vir}$. Haloes are populated with galaxies according to a SHMR (e.g. Hopkins et al., 2010a).



Figure 5.7: A comparison of merger rate per galaxy ($\mu \leq 3$) for other theoretical models, adapted from Hopkins et al. (2010b). *Semi-analytic*: Croton et al. (2006); De Lucia & Blaizot (2007); Somerville et al. (2008); Khochfar & Silk (2009); Bower et al. (2006); Font et al. (2008). *Empirical*: Hopkins et al. (2010a); Stewart et al. (2009). *Hydrodynamical simulations*: Rodriguez-Gomez et al. (2015). In general our model produces a rate lower than most other models, and is in agreement with the results of Bower et al. (2006).

- 2. Subhalo disruption: Subhalo disruption rates are convolved with some SHMR with no delay applied (e.g. Stewart et al., 2009; Rodríguez-Puebla et al., 2017).
- 3. Halo merger trees without substructure: Halo merger trees (EPS or N-body) are populated with galaxies according to a model. A dynamical friction delay is applied to satellites at $R_{\rm vir}$ (Somerville et al., 2008; Khochfar & Silk, 2009; Font et al., 2008).
- 4. Halo merger trees with substructure: Sub-haloes are tracked within an N-body simulation, where galaxies are populated according to a model. A dynamical friction delay is applied to satellite galaxies when their N-body subhalo is disrupted (e.g. this work, Bower et al., 2006; Croton et al., 2006; De Lucia & Blaizot, 2007).
- 5. Hydro: The baryonic components of galaxies in hydrodynamical simulations are tracked to determine when they coalesce (e.g. Rodriguez-Gomez et al., 2015).

At low and intermediate masses our model predicts a merger rate as much as an order of magnitude lower than some other predictions. Additionally, due to the more shallow scaling with redshift, our results deviate from other models more strongly at higher redshifts, particularly at intermediate masses. We do, however, find our results to be in good agreement with those of Bower et al. (2006), who employ a semi-analytic model on top of N-body merger trees. At the highest masses models tend to agree more closely in terms of magnitude and redshift scaling of the merger rate.

The categories noted above broadly describe the merging process in each modelling strategy, but do not classify all model options that could impact the merger rate. In particular the calibrated or predicted SMF, as well as the chosen treatments for orphan/satellite stripping could play a factor in the resulting merger rates within each model. Any model



Figure 5.8: A comparison of major merger rates $(1 \le \mu \le 3)$ under various model assumptions. Dashed lines indicate merger rates derived using average halo-halo merger rates convolved with a SMF. Solid lines illustrate results derived directly from EMERGE galaxy merger trees. The solid black line represents the reference results using our standard model shown throughout this paper. Coloured solid lines show our results under different implementations for dynamical friction, satellite evolution, and simulation resolution. Methods using average halo-halo merger rates overpredict galaxy merger rates compared to full tree based methods in all cases. This over-prediction becomes more extreme at lower stellar masses.

implementations that could impact the SHMR or cause premature/prolonged merging would also impact the rates, regardless of the guiding dark matter component of the model. A more complete overview of the models shown in Figure 5.7, as well as the systematic uncertainties in determining merger rates can be found in Hopkins et al. (2010b) and Rodriguez-Gomez et al. (2015). In the next section we will explore how our own model assumptions and treatments for orphans/satellites can impact the merger rate.

Impact of model assumptions

Although the purpose of this work is not to address the systemics that impact the predicted merger rates across the range of available models, we can address known sources of uncertainty, and determine what role our model assumptions play in the resulting merger rates. In this section we will vary core components of EMERGE and analyse the resulting merger rates to determine if any of our model assumptions are suppressing merger rates compared to other empirical approaches (i.e. Stewart et al., 2009; Hopkins et al., 2010a). Throughout this section we take our standard model implementation with the parameters of Table 4.1 as our reference. We only explore the impact these model assumptions have on the *major* merger rate with $1 \leq \mu \leq 3$. We chose this definition of major merger for ease of comparison with Hopkins et al. (2010a). Figure 5.8 illustrates the results of this study alongside Hopkins et al. (2010a). As we vary each model element we do not refit for each model permutation. Consequently, these model variations do not reproduce all observations as accurately as the reference case. Our model permutations cover the following:

- 1. Impact from our satellite stripping implementation.
- Swapping the dynamical friction formula from Boylan-Kolchin et al. (2008) to Binney & Tremaine (1987).
- 3. No orphan galaxies. Satellites galaxies are merged onto their host when their subhaloes are lost in the simulation.
- 4. All satellites as orphans. Orphan galaxies are initiated when one halo enters the virial radius of another for the first time.
- 5. No satellites. Galaxies are merged at the same time as their host haloes.
- 6. Increased resolution. Standard EMERGE options applied to a simulation with a much higher particle/halo resolution.
- 7. Average halo-halo merger rates with rank ordered abundance matching, using the SHMR derived directly from our reference model.
- 8. Average halo-halo merger rates and abundance matching with scatter allowed in the SHMR.

In §4.1.2 we described our methodology for stripping satellite galaxies, and our newly implemented approximation for halo mass-loss in orphan galaxies, which presents the possibility that our implementation may strip orphan satellites too strongly, suppressing the merger rate. We can provide a simple check for this scenario by setting the stripping parameter $f_s = 0$. This ensures that all satellites with a short enough dynamical friction will merge and contribute to the computed merger rate. In Figure 5.8 this case is illustrated by the blue line labelled ' $f_s = 0$ '. We find a ~ 33 per cent boost in the major merger rate at intermediate masses, but at low and high masses we see no discernible change in the merger rate. Though not displayed in Figure 5.8 we also verified that our merger rates are nearly unchanged within the $f_s \pm 1\sigma$ parameter range noted in Table 4.1. Also, setting $f_s = 0.1$ did not reduce the merger by more than a factor ~ 2 compared to the reference, though we note that both the $f_s = 0$ and $f_s = 0.1$ cases do not reproduce the local clustering data well. We can conclude from these tests that our major merger rates are not strongly impacted by our current implementation for halo mass loss or stripping.

The models shown in Figure 5.7 adopt a wide range of dynamical friction formulations. The use of Boylan-Kolchin et al. (2008) dynamical friction has by now become the standard for controlling satellite decay, and builds upon earlier work by tuning the merging timescales using a suite of idealised N-body mergers that track haloes and their baryonic components from infall to final coalescence. While this formulation should provide a more physical description for the satellite decay process, we can nonetheless explore merger rate with the more classical dynamical friction formula provided by Binney & Tremaine (1987). We find that with respect to major mergers the choice of dynamical friction makes little difference for massive galaxies. At intermediate and low mass we find lower merger rates compared to our reference case. This effect is most pronounced at high redshift and intermediate masses where we see a ~ 56 per cent difference from the reference.

The need for orphans in modelling has been, by this point, thoroughly discussed (e.g. Campbell et al., 2018; Behroozi et al., 2019). Where models continue to differ is in the precise treatment of orphan galaxies, with a core difference being *when* orphans are placed into the simulation. It has been argued that models relying on N-body trees that track substructure experience lower merger rates due to overly effective tidal stripping of subhaloes with the absence of baryons (Hopkins et al., 2010b). Furthermore, it has been argued that applying a dynamical friction recipe at the time of subhalo disruption can introduce additional uncertainty as these formulations are calibrated from the initial halohalo merger. Some tests have shown that without sufficiently resolved subhaloes this can artificially increase merger timescales by as much as a factor 8 (Hopkins et al., 2010b). Now, we probe whether our treatment for orphans is suppressing the merger rate compared to other methods.

In Figure 5.8 the green line labelled 'No delay from subhalo destruction' indicates the scenario where we 'turn off' orphans and simply merge satellite galaxies with their host when the subhalo is lost in the simulation. This test can best be compared with methods that determine the merger rate by convolving a SHMR with the average subhalo destruction rate (e.g. Stewart et al., 2009; Rodríguez-Puebla et al., 2017). It has been shown that the majority of haloes are disrupted in the inner halo (Wetzel & White, 2010), where the remaining time before final merger is relatively short. Thus by merging galaxies at the time of subhalo disruption we are placing an upper limit on the merger rate compared to our reference case. Indeed we do see that this scenario produces a increase in major merger rates for each mass bin. Overall this contributes as much as a ~ 58 per cent increase compared to the reference model. From this we can see that by implementing dynamical friction at subhalo loss we are not creating a substantially longer lived population of orphan satellites.

In §5.1.2 we showed that the halo-halo merger rates from our halo merger trees are consistent with the fits shown in Genel et al. (2009); Fakhouri et al. (2010). As an additional check on this model variation we compare the subhalo destruction rate measured from our input halo merger trees against the fitted relation provided in Behroozi et al. (2013c). Subhalo disruption is defined when a subhalo can no longer be tracked in the simulation. The mass ratio is taken between the host halo mass and the subhalo *peak* mass. Having shown that removing orphans does not strongly impact our merger rates, one could see how altering the subhalo destruction rate. Generally, we find that our subhalo destruction rates show a similar mass ratio scaling as Behroozi et al. (2013c). We can also note that our rates show a weaker scaling with primary halo mass, and tend to be lower when compared to Behroozi et al. (2013c). Such a result is consistent with the lower galaxy merger rates we exhibit compared to some other models relying on similar methods. However, we have not explored in detail the source of this discrepancy. It is not the goal

of this work to probe how strongly this metric is affected by cosmology, definition for peak mass, box size or resolution.

Next, we alter the model such that orphan galaxies are initiated at the virial radius at the time of halo-halo merger. By doing so we can evaluate whether our implementation for dynamical friction is inducing prolonged merging timescales compared to the use case for which the dynamical friction formula was calibrated. To perform this test we remove all subhaloes from the input merger trees. The resulting merger rates from this model are indicated by the yellow line labelled 'Delay applied at $R_{\rm vir}$ ' in Figure 5.8. In this case we actually observe a slightly lower merger rate at all masses compared to the reference. We find a ~ 40 per cent lower merger rate at intermediate masses by z = 6.

As a final check on our orphan treatment we run our standard model with a more highly resolved N-body simulation. For this test we utilise TNG100-Dark, the dark-matter only run of the Illustris TNG100 simulation (Springel et al., 2018), with merger trees constructed using ROCKSTAR and CONSISTENTTREES. This simulation uses a Planck cosmology in a periodic volume with side lengths of 110.7 Mpc. This simulation contains 1820^3 dark matter particles with particle mass $8.86 \times 10^6 M_{\odot}$. With this test we can determine whether our major merger rates are significantly affected by resolution. The results from this model are labelled as 'TNG100-Dark' in Figure 5.8. We find that our results are robust to a substantially increased mass resolution. While this volume does produce a marginal increase in merger rates at high redshift, we find that the broader trends are fundamentally equivalent, and any difference is within the range of sample variance between the dark matter simulations.

So far the model variations we have introduced have not resulted in extraordinary changes to our standard implementation. To place an upper limit on the merger rates we can expect from this model we explore an extreme case where galaxies are merged at the same time as their parent haloes merge. For this test we once again use our modified halo merger trees where substructure has been removed. This modification results in extreme changes for some model predictions that completely disagree with observational constraints. In particular the SMFs under-predict the abundances of galaxies with $\log_{10}(m/M_{\odot}) \leq 11.3$, and small-scale clustering all but vanishes. As we are probing the merger rate *per galaxy* the lower abundances can serve as an additional boost to the merger in the affected mass ranges. The end result is that we get higher merger rates at all masses, with the increase in rates becoming most pronounced for low mass galaxies. At the most extreme this approach results in a factor ~ 5 higher merger rate than the reference case.

Finally, we determine merger rates using the method described in Hopkins et al. (2010a) but with updated halo-halo merger rates and SHMR to better match EMERGE. In this approach we derive galaxy-galaxy merger rates by convolving average halo-halo merger rates with the SHMR produced by our reference model SMF via abundance matching. Hopkins et al. (2010a) showed that this method is robust to changes in dynamical friction, inclusion of substructure, and choice of mass function. In the last decade the quality of observational estimates of galaxy properties has improved substantially; in particular we now have reliable measurements of galaxy properties beyond $z \approx 2$, a noted limitation of the abundance matching model they employed for their core model (Conroy & Wechsler,



Figure 5.9: The subhalo destruction rate from our input halo merger trees (solid lines) compared to the best fit formula presented in Behroozi et al. (2013c) (dashed lines). Each panel adopts a uniform binning in scale factor were da = 0.1. The mass ratio is taken with respect to the host halo mass and the subhalo *peak* halo mass, $\mu_{\text{peak}} \equiv M_{\text{host}}/M_{\text{sub,peak}}$

2009). In this final test we investigate if the galaxy merger rate can be constrained using only the average halo-halo merger rate, and observed galaxy mass functions, or if we require the self consistent growth history contained in complete merger trees.

Beyond the updated SMF we also make a few other changes to the model described in Hopkins et al. (2010a). For average halo-halo merger rates we adopt the model presented in Fakhouri et al. (2010) along with their best fit parameters. This formulation opts for a more simplified redshift evolution compared to Fakhouri & Ma (2008). The basic procedure for this abundance matching method can be summarised as follows:

- 1. Create mock halo catalogues by sampling the HMF at any redshift for which we want to know the merger rate. We sample our *total* N-body HMF directly, enforcing the same halo mass resolution limit as our reference model.
- 2. Determine the number of mergers $N(M, z, \mu_H)$ for each mock halo. We only consider halo mergers with $\mu_H \leq 100$, with redshift bins centred on each 'snapshot'.
- 3. Compute the dynamical friction time for each halo-halo merger according to Boylan-Kolchin et al. (2008).
- 4. Evolve the main halo mass forward to the time of t_{df} . As described in Conroy & Wechsler (2009) we use the halo mass growth formula described in Wechsler et al. (2002)

$$M_{\rm vir}(a) = M_0 \exp\left[-2a_c \left(\frac{a_0}{a} - 1\right)\right] , \qquad (5.4)$$

where the average formation scale factor a_c is parameterised as

$$a_c(M_{\rm vir}) = \frac{4.1}{c(M_{\rm vir})(1+z)}$$
 (5.5)

Here $c(M_{\rm vir})$ is the halo mass-concentration relation at z = 0. We use the updated form to Bullock et al. (2001), as presented by Macciò et al. (2008) which takes the form

$$c(M_{\rm vir}, z) = K \left[\frac{\Delta_{\rm vir}(z_c)}{\Delta_{\rm vir}(z)} \frac{\rho_{\rm u}(z_c)}{\rho_{\rm u}(z)} \right]^{1/3} , \qquad (5.6)$$

where Δ_{vir} is the overdensity of the halo relative to the mean density of the Universe ρ_{u} . The parameter K = 3.8 is the halo concentration at the collapse redshift z_c and is fit to numerical simulations. We allow for $\Delta \log_{10}(c_{\text{vir}}) = 0.14$ when computing halo growth.

- 5. We populate haloes with galaxies at t_{df} using a simple rank ordered abundance matching, using the evolved halo mass for the central galaxy and the infall halo mass for the satellite galaxy.
- 6. Finally the merger rate per galaxy is computed using the same strategy described in 5.2.1.

The results of this test are displayed with dashed lines labelled 'FMBK10 + AM', we performed this analysis both with (magenta) and without (blue) scatter in the SMHR. In either case we find that this method over-predicts intrinsic major merger rates at all masses. The difference is most extreme at low masses where there is an order of magnitude over-prediction in the major merger rate. Given this discrepancy we conclude that average halo merger rates are ill-suited for deriving galaxy merger rates.

These differences in predicted merger rate among various models makes make clear that the methodology chosen to link halo and galaxy properties has tangible impact on the assembly pathway of galaxies. In the next section we will utilise mock observations of our simulated galaxy catalogues to gain a better understanding of galaxy assembly in our framework.

5.3 Which galaxies grow through mergers?

In this last section we will move on from addressing the galaxy-galaxy merger rate to exploring the role that mergers play for the growth of galaxies. We aim to answer two key questions:

- Where does a galaxy's stellar mass come from?
- Are all types of mergers equally important?

We approach these questions in the context of the main branch evolution of the z = 0 galaxy population. Here we explore the merging history of individual galaxies.

5.3.1 Stellar mass fraction accreted through different merger types

First we investigate the assembly of galaxies and whether the stellar mass has grown mainly through star formation (in-situ) or through mergers (ex-situ). Previous work has shown that the accreted stellar mass fraction, $f_{\rm acc}$, ranges between less than 2 per cent for low-mass galaxies to more than 50 per cent for massive galaxies (Lackner et al., 2012; Cooper et al., 2013; Lee & Yi, 2013; Rodriguez-Gomez et al., 2016).

Figure 5.10 illustrates the fraction of accreted material as a function of z = 0 stellar mass, divided into different merger mass ratios. The top panel illustrates the accreted mass fraction with respect to total stellar mass, the bottom panels shows the accreted fraction delivered by merger type with respect to the total *accreted* stellar mass. In the previous sections we showed that for the global merger rate more massive galaxies are biased to experience major mergers, while low-mass galaxies are biased to experience mostly mini mergers. Massive galaxies will pass through both of these regimes during their lifetime, so we can reasonably ask ourselves which mergers ultimately built the galaxies we see? Were galaxies quickly assembled through successive minor mergers, or are the most massive galaxies assembled (late) through major mergers?



Figure 5.10: Contribution of accreted mass by merger ratio as a function of z = 0 stellar mass $m_{z=0}$. Top panel: The ex-situ mass fraction $(m_{ex}/m_{z=0})$ with respect to the z = 0 stellar mass. The solid black line represents the average accreted mass fraction across all z = 0 galaxies with $\log_{10}(m/M_{\odot}) \ge 9.0$ within the simulation volume. Lines of different colour and line style illustrate the average accreted mass fraction broken down by merger mass ratio μ . Lines of the same colour sum to the total accreted mass fraction, solid black line. Bottom panel: The fraction of the z = 0 accreted stellar mass, $m_{ex,tot}$ broken down by merger mass ratio. Similar to the top panel, line type and colour show the contributions by merger mass ratio with like colours summing to the total average. For instance, the solid blue line illustrates the fraction of all accreted mass deposited through mergers with $1 \le \mu \le 2$. At the lowest mass this means that only ~ 20 per cent of accreted material is deposited by mergers of this type. Conversely, the dashed blue line shows the fraction history these lines sum to $m_{ex}/m_{ex,tot} = 1$, the solid black line.

Looking at the top panel of Figure 5.10 we first find that the largest galaxies are constructed primarily through accreted material. On average we expect upwards of $f_{\rm acc} \approx$ 80 per cent at the massive end, and as little as $f_{\rm acc} \approx 1.5$ per cent at the low-mass end for z = 0 galaxies. There exists a strong mass dependence in the accreted mass fraction between $\log_{10}(m/M_{\odot}) = 10.25$ and $\log_{10}(m/M_{\odot}) = 11.25$. In this regime we see a corresponding inversion in the relative contribution of major and minor mergers to the final system. Looking at the solid red line we can see clearly that the most massive galaxies are indeed assembled by successive major merging events. On the massive end, we see that these major mergers contribute as much as 90 per cent to the total accreted mass budget of a galaxy (bottom panel). Additionally, find that these mergers begin to dominate the accreted mass budget at around $\log_{10}(m/M_{\odot}) = 10.3$. If we look at the classical merger mass ratio definitions, we see that up until $\log_{10}(m/M_{\odot}) = 10$, major ($\mu < 4$), minor $(4 \le \mu < 10)$, and mini $(10 \le \mu)$ mergers contribute roughly equal quantities to the total accreted mass budget. Beyond this point we see the relative contributions diverge. This contrasts strongly with recent results that indicate major mergers contribute a roughly flat 50 per cent of the accreted mass fraction at all mass scales, and minor/mini mergers show a roughly constant 20 per cent contribution (Rodriguez-Gomez et al., 2016).

Understanding the source of accreted stellar material can have direct consequences on the internal kinematics, and stellar mass distributions of a galaxy. Some models indicate that major mergers deposit stellar material at the center of the descendant galaxy (Deason et al., 2013; Pillepich et al., 2015), while minor mergers tend to deposit material at larger radii, growing the stellar halo (Hilz et al., 2012, 2013; Karademir et al., 2019). Subsequently observations of stellar populations could be used to determine the merging history of galaxies (Merritt et al., 2016; Amorisco, 2017, 2019; Bernardi et al., 2019; Ferré-Mateu et al., 2019; Hendel et al., 2019).

We can more directly probe which mergers contribute the most to the stellar mass growth of a galaxy by evaluating the *mass-weighted* mass ratio defined as:

$$\mu_{\rm mw} \equiv \frac{\sum_{i}^{N_{\rm m}} m_{i,2}}{\sum_{i}^{N_{\rm m}} \frac{m_{i,2}}{m_{i,1}} m_{i,2}}$$
(5.7)

In this approach, each merger has its mass-ratio weighted by the amount of stellar mass contributed to the final system. This way we see what types of mergers were on average most important for the growth of galaxies at a given mass scale. In Figure 5.11, we show the median $\mu_{\rm mw}$ for all z = 0 galaxies in our simulation. In agreement with the results displayed in Figure 5.10, we see that low-mass galaxies on average experience mergers with $\mu_{\rm mw} \approx 40$, once again illustrating that major mergers are not important for the growth of low-mass galaxies. Previous works (Naab et al., 2009; Oser et al., 2012; Hilz et al., 2012, 2013) have suggested that successive minor mergers can be an effective pathway to form large galaxies. Figure 5.11 shows clearly that there is only a narrow transition region where these minor mergers play a significant role in stellar mass assembly. For massive systems we once again see that most of the stellar mass is delivered through major mergers. We can see that on average the mergers that bring the most mass into the system are very



Figure 5.11: The mass-weighted mass-ratio as a function of z = 0 stellar mass $m_{z=0}$. The solid line illustrates the median mass weighted mass ratio of mergers along the galaxy main branch for each galaxy in the specified mass range that experienced at least one merger in its lifetime. The shaded region shows the 68^{th} percentile surrounding the median.



Figure 5.12: Probability of experiencing N mergers in a z = 0 galaxy's history. Each line shows the number of mergers occurring along the main branch per z = 0 galaxy. Each line sums to 1, representing the merging history of all z = 0 galaxies. For galaxies with $\log_{10}(m_{z=0}/M_{\odot}) < 11$ more than 90 per cent of galaxies do not experience any major mergers $\mu < 4$. For galaxies with $\log_{10}(m_{z=0}/M_{\odot}) \geq 11$ only ~ 10 per cent of galaxies do not experience any major mergers.

major mergers with $\mu \approx 2$.

5.3.2 How frequent are different kinds of mergers?

So far this work has focused on the average merger rate, or the average merging history across an entire population of galaxies in a cosmological volume. In this section we investigate the individual merging history of galaxies to see how many mergers a galaxy has experienced in its (main-branch) lifetime.

Figure 5.12 shows the number of mergers experienced along the main branch per galaxy. This accounts for the complete main branch merging history of each z = 0 galaxy in the simulation and indicates the probability distribution of a galaxy having a given number of mergers with a certain mass ratio. Similar to our merger rate analysis we separate the galaxies into three stellar mass bins. In the lowest mass bin (left panel) we see that more than 90 per cent of galaxies experience no major mergers along their main branch, consistent with the expectations set in the previous section, though more than 85 per cent of galaxies in this mass bin will experience *at least* one merging event in their lifetime.

At intermediate masses (central panel), we start to observe the larger frequency of major mergers. However, we find that under the loosest definition of a major merger, more than 80 per cent of galaxies will experience no major mergers in their lifetime. This suggests that a galaxy like the Milky Way has a low likelihood of having ever been impacted by a major merger. Probing in the more narrow range $10.6 \leq \log_{10}(m/M_{\odot}) < 10.8$ we still find that ~ 70 per cent of galaxies will have no major merger in their lifetime. Conversely, we find that the majority (97 per cent) will have experienced *some* merger in their life, no matter how small. This is in agreement with recent observations (D'Souza & Bell, 2018; Helmi et al., 2018).

This trend changes for the highest mass bin where mergers play a substantial role in stellar mass growth. In this range we can see that 90 per cent of galaxies will experience at least one major merger in their lifetime, with one galaxy experiencing as many as 11 major mergers. This particular galaxy is the largest galaxy in our box with a stellar mass of $\log_{10}(m_{z=0}/M_{\odot}) = 12.17$, and had already grown to $\log_{10}(m/M_{\odot}) \approx 11$ before encountering its first major merger (near $z \approx 2.5$). However, on average the most massive galaxies experience ~ 1.8 major mergers in their lifetime.

While mergers play very different roles for the evolution of low-mass and massive galaxies, we find that most galaxies are subject to a merger in their lifetime. For z = 0 galaxies with $\log_{10}(m/M_{\odot}) \ge 9.0$ we find that ~ 92 per cent experience a merger with any mass ratio, while ~ 8 per cent of galaxies experience no merging event along their main branch. However, the complete history of a galaxy is more complicated. Here we have only probed galaxy mergers along the main branch for the most evolved galaxies. This is *not* a complete accounting of the number of mergers that occurred within a galaxy's complete evolutionary tree.

5.4 Discussions & Conclusions

In this work we have presented our analysis of the galaxy-galaxy merger rate within the context of the empirical model for galaxy formation EMERGE. This model connects galaxy growth directly to the halo growth in N-body simulations using simple relations constrained by a suite of observables (see Moster et al., 2018, for more details). We investigated a range of properties associated with galaxy merger rates, including: scaling with stellar mass and mass ratio, the relationship between the merger rates and observed galaxy pairs, the merging history of large galaxies, and the role of merging in galaxy quenching. We also presented a brief comparison of our results to other theoretical models.

We find a galaxy merger rate density Γ that increases with redshift until $z \approx 1.5$, followed by a sharp decline in the rate towards higher redshift. This general trend holds for the three mass bins we explored between $9 \leq \log_{10}(m^*) < 12$. For the merger rate per galaxy \Re , unlike previous works, we do not exhibit universal power law scaling with redshift. We find that merger rates show an excess over a power-law scaling for $z \gtrsim 3$. This effect is most apparent when determining merger rates based on descendant galaxy mass. Generally mergers occur at a higher frequency with increasing galaxy mass. When exploring how merger rates scale with mass ratio, we find that the largest galaxies are biased to experience major mergers and subsequently show an enhanced major merger rate compared to lower mass galaxies. This effect can be seen even for $\mu \leq 10$. This effect is a departure of the nearly mass independent scaling in the numerically derived halo-halo merger rate. We conclude that the self similarity shown in halo-halo mergers is broken through the complex connection between galaxies and their haloes. In this view a major halo-halo merger occurring along the high mass slope of the SHMR are likely to result in an eventual major merger in stellar mass as well. Conversely, a major halo-halo merger along the low mass slope of the SHMR is generally more likely to eventually produce a very minor merger in galaxy stellar mass. The influence of the SHMR on major merger rates has been explored in the past (Stewart et al., 2009; Hopkins et al., 2010a), due to advance in observations and modelling techniques we are able to show that such a phenomenon causes a break from a power-law redshift scaling, particularly at intermediate masses.

Our model predicts merger rates that are consistent with other theoretical models. However, within the range of previous works, our results tend to sit lower than the average. The merger rates produced from our model are in closest agreement with Bower et al. (2006). Additionally, our mass dependent bias towards major mergers is an effect absent in some other models (Rodriguez-Gomez et al., 2015). A more complete discussion of the intricacies and differences in these models can be found in Hopkins et al. (2010b).

Additionally, we also explored a diverse set of model variations to determine the merger rate sensitivity to a range of model options. We show that our newly implemented model for galaxy stripping does not result in overly aggressive satellite loss, suppressing the merger rate. Further, we set reasonable upper bounds on the expected merger rate in the context of our model by systematically varying the treatment of orphan galaxies. In particular we illustrate that application of dynamical friction delay at subhalo destruction does not result in a prolonged merger timescales compared to application at $R_{\rm vir}$ for the simulation volume we tested. We also verified that major satellites are not spending spending too much time in the orphan phase by merging galaxies directly at the time of subhalo destruction. we also showed that our model is robust to large increase in particle resolution by running our model on top of the Illustris TNG100-Dark merger trees. Ultimately none of the reasonable model changes we tested resulted in a substantial change to major merger rate normalisation or redshift evolution. We can conclude that our results are reliable to within a factor ~ 2 for the current model and parameter set.

When testing model variations we also explored galaxy merger rates derived from average halo-halo merger rates. This test reproduced the methods described by (Hopkins et al., 2010a), with updates to better reflect our improved fit to the relevant observables. We found that these halo averaged merger rates over predict galaxy merger rates by as much as an order of magnitude compared with rates derived from complete galaxy merger trees.

We also explored the merging history of the z = 0 galaxy population to determine what role mergers play in the buildup of stellar mass. Our model shows that galaxies with $\log_{10}(m/M_{\odot}) \lesssim 11$ grow almost entirely through *in-situ* star formation, with accreted material accounting for ≤ 10 per cent of the total stellar mass. Furthermore, this accreted material is overwhelmingly deposited in minor or mini mergers, with ≥ 90 per cent of accreted stellar mass attributed to mergers with $\mu \geq 4$. For more massive galaxies $\log_{10}(m/M_{\odot}) \gtrsim 11$, galaxy-galaxy mergers play a critical role in the buildup of mass. In these galaxies, accreted material accounts for as much as 80 per cent of the total stellar mass for some galaxies. In these cases the stellar mass is largely deposited through major mergers, where as much as 90 per cent of the total accreted mass is delivered through mergers with $\mu \leq 4$. Our results indicate that EMERGE can accurately predict the galaxy merger rate out to high redshift. We are able to not only compute the cosmic merger rate of galaxies but explore the individual merging history of each galaxy in our simulation volume. Lastly, at the time of the analysis our model did not include any information about the gas properties or orbital configurations of merging systems. These additional details may be necessary to form a complete understanding of the impact of galaxy merging on the star formation properties, and radial distribution of mass in observed galaxies.

Chapter 6

The view from afar: galaxy pair merging probabilities and observation timescales

6.1 Introduction

Galaxy mergers play a crucial role in the buildup of stellar material under the current hierarchical view of galaxy formation. The galaxy merger process is responsible for not only stellar mass growth, but is also invoked to explain many observed phenomena, such as AGN (Choi et al., 2018; Steinborn et al., 2018; Gao et al., 2020; Marian et al., 2020; Sharma et al., 2021), stellar streams, disturbed morphologies (Conselice et al., 2003; Lotz et al., 2008, 2010; Wen & Zheng, 2016; Martin et al., 2018; Bluck et al., 2019; Mantha et al., 2019; Yoon & Lim, 2020), and quenching (Khalatyan et al., 2008; Jesseit et al., 2009; Bois et al., 2011; Moody et al., 2014; Naab et al., 2014). Due to the large timescales involved in galaxy evolution, large volume simulations are necessary to understand vital aspects of galaxy formation such as the galaxy merger rate. In order to validate these theoretical predictions we require theoretical methods to determine these key quantities in an observational context.

When extracting the merger rate of galaxies from observations, the so called observation timescale T_{obs} is of central importance. This parameter specifies how long a galaxy pair will remain observable under some specified selection criteria. Despite the importance of this value, theoretical models have not yet converged on how this quantity should scale with redshift, stellar mass, mass ratio, projected separation, or redshift proximity (Kitzbichler & White, 2008; Lotz et al., 2008, 2011; Jiang et al., 2014; Snyder et al., 2017; O'Leary et al., 2021a). Efforts to constrain this parameter are hampered by lack of consensus regarding which pair selection criteria should be used, as well as observable limitations which force various groups to adopt different stellar mass, and mass ratio constraints for their analysis (Lin et al., 2008; Bundy et al., 2009; Man et al., 2016; Ventou et al., 2017; Mantha et al., 2018; Duncan et al., 2019; Ventou et al., 2019).

The intent of this work is to explore observation timescales to determine more precisely what drives this quantity. To that end we employ the empirical model EMERGE, as we can readily translate pair fractions into rates as both can easily be extracted. These can be used to place tighter constraints on $T_{\rm obs}$, which can be used to determine rates from observed pairs. Through this work we will provide utilities for observers to translate observed pair fractions into galaxy merger rates. In this effort we will answer three key questions concerning observed merger rates.

- What is the probability that two galaxies observed as a close pair will merge by z = 0 and what is that dependency on radial projected separation Δr , line of sight velocity difference Δv and z?
- On what timescale will an observed pair merge, and for how long is that pair observable given the pair selection criteria?
- What determines pair observation timescales?

This chapter is organised as follows: In §6.3.2 we evaluate the pair fraction evolution from our model and discuss methods to directly relate this quantity to the intrinsic merger rates discussed in §5¹. In §6.4 we discuss how we construct the mock catalogues which can more accurately track the evolving galaxy pair fraction. Our pair selection criteria and approach to fitting functions are described in §6.5, with fitting formulas for pair merging probability and merging timescales explained in §6.5.1 and §6.5.2 respectively. In §6.6 we show that the results of §6.5 can be used to reconstruct the merger rates shown in O'Leary et al. (2021a) when applied to mock catalogues. Finally, in §6.7 we discuss caveats of our analysis and summarise our key conclusions.

6.2 The simulation and relevant emerge details

Our analysis relies on producing galaxy merger trees encompassing a large dynamic range, occupying an appropriately large cosmic volume. We employ the empirical model EMERGE to populate dark matter haloes with galaxies based on individual halo growth histories. The analysis of this chapter relies on the same N-body simulation described in §5.1

6.2.1 The relation between merging and clustering

As discussed in §4.1.2, merging in this model occurs between galaxies residing at the centre of a resolved N-body halo and so-called orphan galaxies. Orphan galaxies are those systems whose host halo has fallen below the resolution limit of the halo finder due to real mass stripping in the simulation. Rather than remove/merge these systems from the simulation

¹This section was originally published in O'Leary et al. (2021a) but has been grouped with work of O'Leary et al. (2021b) for clarity. The remainder of this chapter is taken directly from O'Leary et al. (2021b)

when their halo is lost, we continue to track these galaxies within the empirical model using approximate formulae to update their halo mass and position within their host halo system. The orbits of these orphan galaxies will continue to decay and we merge them with their host system according to some dynamical friction formula. In O'Leary et al. (2021a) we showed that our merger rates are not strongly driven by our choice of dynamical friction formulation, or our treatment of orphans. This formulation does however play a role in how we update the position of orphan galaxies in emerge, and is thus relevant to the discussion of pair fractions, and their merging timescales.

The focus of this work involves galaxies in close pairs. The galaxy pair fraction is related to, but distinct from, the projected galaxy correlation function (w_p) , which is one of the observations used to constrain EMERGE. In order to compute clustering we need to know the position of each galaxy in the simulation volume, including orphans. For galaxies in resolved haloes we use the position of the N-body halo directly to compute clustering. In the case of orphans, positions are approximated by placing the orphan galaxy randomly on a sphere of radius

$$r = r_0 \sqrt{1 - \Delta t / t_{\rm df}} \tag{6.1}$$

centred on the main halo. Here r_0 is the radial position of the orphan when its halo was last resolved and Δt is the time elapsed since the subhalo was last resolved (Binney & Tremaine, 1987).

6.3 Observed merger rates

So far we have established the intrinsic galaxy assembly process within the context of our model. The next step is to take this knowledge of galaxy assembly and translate that into something we might observe. Observationally, the galaxy-galaxy merger rate is difficult to ascertain. Additionally, the dynamic process of merging takes place on the scale of hundreds of Myrs to Gyrs.

Obvious physical tracers of a recent merger such as disturbed morphologies present one option for deducing the galaxy merger rate. Methods invoking quantitative morphology such as $G - M_{20}$ or asymmetry are not equally sensitive to all merger mass ratios. Furthermore, these morphological methods are sensitive to total mass, gas properties, orbital parameters, merger stage, and viewing angles (Abraham et al., 2003; Conselice et al., 2003; Lotz et al., 2008, 2011; Scarlata et al., 2007). These additional difficulties present a greater barrier to identifying mergers and determining a cosmological merger rate (Kampczyk et al., 2007; Scarlata et al., 2007; López-Sanjuan et al., 2009; Shi et al., 2009; Kartaltepe et al., 2010; Abruzzo et al., 2018; Nevin et al., 2019). One common observational method for deriving the galaxy merger rate is through the analysis of galaxies in close pairs. The foundation of this approach is simple, as galaxies found in close proximity are expected to merge within some finite predictable time scale.

Within theoretical models we have the possibility of investigating the complete growth history of galaxies in a cosmological volume, and by performing mock observations on our simulated catalogue we are able to provide guidance on how physical observables can be translated into a true merger rate. The standard galaxy lists produced with EMERGE provide an ideal sandbox for comparing observed merger rates with theoretical predictions.

In this section we will be studying two particular quantities: the evolution of the galaxy pair fraction, and the observation timescale of close pairs. The close pair fraction is a measurement of how galaxies cluster, not unlike the projected correlation function w_p . While these observables are related they are not directly interchangeable. In particular as we will soon see the pair fraction is used as a proxy for the merger rate and is subject to additional selection criteria to maximise the likelihood that observed pairs are physically associated and expected to merge. Furthermore, EMERGE is only fit to the stellar mass projected correlation function at z = 0, wheras the pair fraction must be measured to high redshift. In §4.1.2 we noted model improvements that needed to be made in order to fit small scale clustering. This highlights the fact that this particular observable is sensitive to the model. Additionally, in §5.2.2 we explored several variations of our model and how those changes impact the merger rate. In most cases the resulting SMF under these variations was different from the reference, but largely within observed errorbars. With regards to clustering this is not the case, we find while these changes have only a small influence on galaxy abundances and merger rates, they show far greater influence on galaxy clustering. As we do not refit the model for each of these variations we cannot say whether all models produce the same pair fraction evolution provided an equally good fit to the z = 0 projected correlation function.

An advantage of empirical modelling is in the 'observables first' approach. Because these models match observation by design, they are the ideal test-bed for relating the coevolution of large scale scale observables with other extrinsic galaxy properties such as their *ex-situ* stellar mass growth. Numerical simulations are a necessary tool in bridging this gap, so in this next section we will see how our model compares with current observations, and whether our results are in agreement with the expectations set by other theoretical models.

6.3.1 Close galaxy pairs

A typical pair count requires two quantities: a projected galaxy separation radius $R_{\rm proj}$, and some additional redshift proximity criterion. While it is in principle possible to use a 3D deprojected radius to determine physical proximity, this method is prone to error for galaxies with large uncertainties in redshift. When reliable spectroscopic redshift data are available and relative proper motions of companion galaxies can be determined, a common criteria is to use a maximum line of sight velocity difference Δv to establish physically associated pairs. Pairs with a small enough relative velocity are assumed to be gravitationally bound and will eventually merge. We use the complete information available in our catalogues to perform such an observation. To maintain the most transferable results we selected pairs by stellar masses, $R_{\rm proj}$ and Δv according to values commonly used by observers (Lotz et al., 2011; Man et al., 2016; Mantha et al., 2018; Duncan et al., 2019). At each simulation snapshot we compute the fraction of galaxies hosting a major ($\mu \leq 4$) companion, $f_p = N_{\rm pairs}/N_{\rm gal}$. We do not construct light-cone catalogs, as such our analysis



Figure 6.1: Redshift evolution of the pair fraction. We present the major pair fraction $(\mu \leq 4)$ for our simulation alongside observed pair fractions (Man et al., 2016; Mundy et al., 2017; Ventou et al., 2017; Mantha et al., 2018; Duncan et al., 2019). Solid lines illustrate the pair fraction from our model under four different combinations of radial projected separation $(R_{\rm proj})$ and mass threshold. The shaded regions indicate Poisson error in the number of pairs. These combinations where chosen to provide the best comparison with observations displayed. Line and marker colours indicate which results share the same selection criteria. Our methods best match observations that incorporate spectroscopic redshift data, these data points are noted by diamond markers. The black line is employs the pair selection criteria of Ventou et al. (2017) with $R_{\rm proj} = 3 - 25$ kpc h⁻¹ $\approx 5 - 37$ kpc, only this work provides spectroscopic data out to z = 5. The majority of results indicate an increasing pair fraction at low redshifts, and either a flat or decreasing pair fraction at high redshift.

does not incorporate field variance, nor do we impose volume restrictions at low redshift to approximate sample incompleteness from a narrow field. Our analysis only considers major galaxy pairs, where the most massive galaxy in each pair must reside above a specified mass threshold (Table 6.2 indicates the mass thresholds used in this work). All measurements adopt a fixed $\Delta v = 500 \,\mathrm{km \, s^{-1}}$ for redshift proximity, consistent with previous works (Patton et al., 2000; Lin et al., 2004, 2008; de Ravel et al., 2009; Lotz et al., 2011; López-Sanjuan et al., 2012; Mantha et al., 2018; Mundy et al., 2017).

In Figure 6.1 we display the pair fraction of our simulation alongside recent observations. In this figure we show our pair fractions using five different criteria for mass threshold and projected separation. In all cases, regardless of the radial separation chosen, our results express similar features. In each case we see an increase in the pair fraction with redshift, with a peak at $z \approx 2.5$, followed by a shallow decrease in pair fraction toward even higher redshifts. Previously published observed pair fractions have shown a power law increase with redshift (Kartaltepe et al., 2007; Lin et al., 2008; Bundy et al., 2009; Conselice et al., 2009; de Ravel et al., 2009; López-Sanjuan et al., 2009, 2013; Shi et al., 2009; Xu et al., 2012), while more recent observations indicate a flattening or even decreasing pair fraction at higher redshifts (Man et al., 2016; Mundy et al., 2017; Ventou et al., 2017, 2019; Mantha et al., 2018). In this respect our results more closely align with more recent works. However, the precise functional form remains a point of contention. We find that our pair fractions are most appropriately fit with a modified power-law exponential function (Carlberg, 1990; Conselice et al., 2008):

$$f_p(z) = \alpha (1+z)^m \exp[\beta (1+z)] .$$
(6.2)

When comparing our results to observations or other models it is important to note some of the inconsistencies that might prevent a more accurate comparison. We chose our pair selection criteria to provide the most direct comparison possible with observations. Though we find qualitatively similar results between our selected apertures, the differences produced are immediately noticeable. The pair fraction is sensitive to the selection criteria applied and in the case of observations, sensitive to the methods used to account for sample completeness. Additionally, in this work we only compare with fractions derived using stellar mass and stellar mass ratio of pairs. Previous works have shown that pair fractions determined using flux ratio, or baryon mass ratio pairs, produce results very different results than stellar mass selected pairs (Lotz et al., 2011; Man et al., 2016). Furthermore, observations often lack robust spectroscopic redshift data, instead relying on photometric redshifts. Under the best circumstances scatter in photometric redshift estimates is $\delta z/(1+z) \approx 0.01$ (Molino et al., 2014; Duncan et al., 2019). This level of precision is insufficient to determine relative velocity differences down to $\Delta v = 500 \,\mathrm{km \, s^{-1}}$. Instead of using relative velocities between galaxies, photometric redshift differences along with their associated uncertainties are utilised. One approach is to use $\Delta z^2 \leq \sigma_1^2 + \sigma_2^2$, where the σ_1 and σ_2 are the photometric redshift uncertainties for the major and minor galaxy in each pair, respectively (Bundy et al., 2009; Mantha et al., 2018). Otherwise, probabilistic methods can also be employed to determine physically associated pairs (López-Sanjuan et al., 2015; Mundy et al., 2017; Duncan et al., 2019). We can see the impact of many

of these differences in pair counting methodology if we compare the pair fractions derived from the same field data (Mantha et al. 2018, blue circles; Duncan et al. 2019, red squares). While both of these analyses are based on the same underlying image data, they come to very different conclusions regarding both the normalisation and functional form of the pair fraction.

Considering these difficulties in measuring the pair fraction, the most comparable set of observations for our results are those from Ventou et al. (2017) based on MUSE (Bacon et al., 2010) data, who have spectroscopic data out to high redshift. To make a direct comparison with their work we adopt their pair selection criteria where $R_{\rm proj} = 3 - 25$ kpc h⁻¹, and $\log_{10}(m/M_{\odot}) \geq 9.5$ (black line). Although we are in reasonable agreement with Ventou et al. (2017) we note that the MUSE fields are very small. Consequently their uncertainty due to cosmic variance is large, ranging from $\sigma_v = 0.15$ at $z \approx 0.6$ up to $\sigma_v = 0.52$ by $z \approx 5$ (Moster et al., 2011). The small field size also results in a limited pair sample; in the redshift range where our results disagree the most $1.5 \leq z \leq 3$, only 9 pairs were observed in a sample of 152 galaxies. These large uncertainties in the observed data hinder a more detailed study to describe the differences with respect to our analysis. Similarly, we are able to make a direct comparison to the low redshift SDSS (York et al., 2000) data point presented in Mantha et al. (2018), as well as the low redshift GAMA data point presented in Mundy et al. (2017). In these instances we are once again in close agreement where spectroscopic redshifts are available.

In Figure 6.1 we displayed select results most comparable to some recent observational and theoretical predictions. A short summary of the of the pair selection criteria for the data of Figure 6.1 can be found in Table 6.1. Table 6.2 provides the best-fit to our simulated pair fractions for an additional set of stellar mass thresholds and $R_{\rm proj}$. In each of these fits we assume the functional form of eq. 6.2.

6.3.2 The merger rate from close pairs

Determining a merger rate from pair fractions is *conceptually* straight forward. Observed pairs are assumed to result in a merger on some finite time scale. Therefore, one (mathematically) simple approach to convert a measured pair fraction to a rate is to simply divide the pair fraction by an average observation timescale T_{obs} . Which specifies the amount of time that a pair *could* be identified by the established pair selection criteria. The actual timescale needed for a given galaxy pair to merge can depend on properties other than stellar mass, and projected radial separation. To account for the possibility that not all pairs will merge in the expected timescale or that some pairs are a result of chance projection, an additional correction factor C_{merge} is often introduced to specify which fraction of the observed pairs will actually end up merging. These quantities can be combined, resulting in the merger rate per galaxy formulation:

$$\mathfrak{R} = \frac{C_{\text{merge}} \times f_{\text{p}}}{\langle T_{\text{obs}} \rangle} [\text{Gyr}^{-1}] .$$
(6.3)



Figure 6.2: A comparison of observed merger rates (red lines) alongside the merger rate produced through the mock observations shown in Fig. 6.1. Our results assume $\log_{10}(m/M_{\odot}) \ge 10.3, 1 \le \mu \le 4$ and $R_{\rm proj} = 5 - 50$ kpc. Magenta lines show merger rates derived from our pair fractions assuming previously published results for $T_{\rm obs}$. In particular $T_{\rm obs} = 0.65$ Lotz et al. (2011), $T_{\rm obs} \propto H(z)^{-1/3}$ Jiang et al. (2014), and $T_{\rm obs} \propto (1+z)^{-2}$ Snyder et al. (2017). Black lines illustrate our predicted results from pair fractions under different best fit forms for $T_{\rm obs}$. The solid blue curve is the underlying intrinsic merger rate produced using the methods described in Section 5.2, and the shaded blue region depicts poisson error in the merger count. Deriving merger rates using our pair fractions and published values for $T_{\rm obs} \propto (1+z)^m$ provide a better reproduction of merger rates, but under-predict the intrinsic vales for $z \gtrsim 2.5$. Our model intrinsic merger rate is most accurately reproduced assuming a $T_{\rm obs}$ with a linear redshift scaling, solid black line.

This formulation is contingent upon having a pair selection criteria that does adequately identify physically associated galaxy pairs in the early stages of a merger. Additionally, it assumes that adopting an average observation timescale is a suitable method for converting a sample of galaxy pairs into a rate. Under this formulation the observation timescale is a crucial quantity in translating pair fractions to merger rates. Work seeking to characterise this quantity remains in tension regarding the functional form. A common approach is to take $T_{\rm obs}$ as a constant. Suggested values for a range of stellar masses and $R_{\rm proj}$ have been proposed by Lotz et al. (2011). Conversely, recent work has suggested formulations for a redshift dependent observation timescale. For instance, Snyder et al. (2017) have proposed $T_{\rm obs} \propto (1 + z)^{-2}$, while Jiang et al. (2014) suggest $T_{\rm obs} \propto H(z)^{-1/3}$.

In Figure 6.2 we compare rates derived from two recent observational results (Mantha et al., 2018; Duncan et al., 2019, red lines) to our intrinsic merger rates (i.e. the true merger rate measured in our simulation) and our mock pair fraction derived rates. For clarity we show our results only for $\log_{10}(m/M_{\odot}) \ge 10.3$, $1 \le \mu < 4$, and $R_{\rm proj} = 5 - 50$ kpc. This aligns with the selection criteria of the displayed observations. The pair selection criteria and $T_{\rm obs}$ scaling for the observed data can be found in Table 6.1. Here, low redshift results agree within a factor of ~ 2 . However, predictions deviate heavily towards higher redshift. By z = 3 there is as much as an order of magnitude difference between predicted and observed major merger rates. Generally, these recently published observations overpredict the merger rate compared to our intrinsic values for $z \gtrsim 0.5$. While these methods draw from the same fields, they come to very different conclusions regarding pair fraction evolution. Additionally, different observation timescales are adopted as the default choice to produce each result. If we apply the same observation timescales used in these works (magenta lines) we are unable to reproduce our intrinsic values, from our pair fractions. To better understand these deviations we instead fit $T_{\rm obs}$ using these proposed formulations. From here we can determine if any can provide a meaningful mapping of the pair fraction into the intrinsic merger rate based on our results.

Directly comparing the redshift evolution for merger rates and the pair fraction we can see that a constant value for $T_{\rm obs}$ is insufficient. The increasing merger rate and decreasing pair fraction at high z require that the observation timescale decrease with increasing redshift. If we impose $T_{\rm obs} \propto H(z)^{-1/3}$ (Jiang et al., 2014) we find that we are able to reproduce the intrinsic merger rate until $z \sim 2$ beyond which the predicted merger rate flattens, under-predicting the intrinsic rate. Similarly, the best fit power-law scaling $T_{\rm obs} \propto (1 + z)^{\alpha}$ reproduces the low redshift merger rate scaling, but again flattens and under-predicts the merger rate at high z. We find that the most simple scaling that can recover the intrinsic merger rate to high z is linear, $T_{\rm obs} = w(1 + z) + b$. Additionally, we find that such a scaling provides a better fit for high stellar mass galaxies than for low stellar masses. In the case of lower masses the linear fit begins to deviate for z > 4. However, it's clear that this formulation could fail at any mass if the best fit values result in a negative $T_{\rm obs}$ at the desired redshift.

Absent a generalised fitting formula, we find that for any mass threshold and mass ratio our intrinsic merger rates and pair fractions can be well fit by a power-law exponential as eq. 6.2. Table 6.2 shows our best fit intrinsic merger rates, pair fractions, and observations

Table 6.1: Summary of selection criteria for observed pair fractions. Observations that use a have 'CDF' as their redshift proximity indicate a cumulative probability that two galaxies are a pair. Where observations are included in Figure 6.2 we list the redshift depended observation timescale used to translate observed pair fractions into merger rates.

Publication	$\log_{10}(m/M_{\odot})$	$R_{\rm proj} \; [{\rm kpc}]$	Redshift proximity	$T_{\rm obs} [{\rm Gyr}]$
Man+(2016)	≥ 10.8	$\sim 14 - 43$	$\Delta z_{\rm photo} < 0.2(1+z_1)$	-
Mundy+(2017)	≥ 10.0	5 - 30	$CDF(z_1, z_2)$	-
Ventou+ (2017)	≥ 9.5	$\sim 5 - 37$	$\Delta v \le 500 \text{ km/s}$	-
Mantha+ (2018)	≥ 10.3	5 - 50	$\Delta z_{\rm photo}^2 \le \sigma_1^2 + \sigma_2^2$	0.65
Duncan+(2019)	≥ 10.3	5 - 30	$CDF(z_1, z_2)$	$2.4(1+z)^{-2}$

timescales for a few common stellar mass thresholds.

Our findings conflict with those recent works suggesting a strong redshift evolution for observation timescales. In the case of Snyder et al. (2017) the proposed scaling where $T_{\rm obs} \propto (1+z)^{-2}$ provides a mapping from a flat pair fraction to an underlying merger rate that scales as a power-law with increasing redshift. However, as noted in Snyder et al. (2017), the measured pair fractions from their work rely on a mass ratio calculated using galaxy properties at the same redshift for which the mock observation was performed, while the intrinsic merger rate as measured by Rodriguez-Gomez et al. (2015) takes the mass ratio definitions makes a direct translation between the intrinsic merger rate and the measured pair fraction troublesome. Subsequently, their proposed scaling for the observation timescale is not necessarily reflecting a physical mechanism driving such a formulation.

When finding the best fit value for $T_{\rm obs}$ we assume $C_{\rm merge} = 1$. An accurate determination of $C_{\rm merge}$ is beyond the scope of this work, thus the best fit observation time scale represents an upper limit to the true underlying value. Furthermore, our analysis does not perform a complete light cone analysis, nor do we attempt to reproduce any observational uncertainties in our redshifts or stellar masses. All fits are performed assuming Poisson error in the number of pairs or number of mergers. We leave a more detailed description and analysis of $T_{\rm obs}$ and $C_{\rm merge}$ to future works.

ese	t a	ith	
Τh	nou	y w	
3.2.	wit]	onl	
g. (SWO	gers	
of (Щ	nerg	
orm	ers.	or 1	
al f	serv	ma	
tion	/ ob	ıme	
func	d b	assu	
the	use	ues	
ng t	ons	val	
lowi	ecti	The	
fol :	proj	SS.	
ates	ial	ma	
ger 1	rad	ated	
nerg	and	ıdicâ	; -
sic r	SUC	le ir	erge
trins	ectic	or th	$C_{ m m e}$
l in	sele	te fc	ing
anc	lass	r ra	ums
ions	d n	erge	t as:
ract	use	C me	st fi
uir fi	only	insid	be.
r pê	nmc	intr	are
is fo	COI	the	+
leter	d to	to	$\widehat{(z)}$
ran	pone	e fit	$v(1 \cdot$
t pa	rres]	s ar	i =
st-fi	CO	$T_{\rm ob}$	$T_{\rm obs}$
Be	ulue	oj OI	ted
6.2:	t va	$R_{\rm prc}$	list
ble	st fi	ted	∧ 4.
$\mathbf{I}_{\mathbf{a}}$	be	no	Ħ

$\log_{10}(m/M_{\odot})$	$R_{\rm moi}$ [kpc]	$\log_{10}(\alpha)$	m	β	m	9
	5-30	-1.489 ± 0.006	1.827 ± 0.037	-0.542 ± 0.017	-0.185 ± 0.006	1.528 ± 0.024
1 0 /	14 - 43	-1.325 ± 0.006	1.824 ± 0.036	-0.565 ± 0.017	-0.262 ± 0.008	2.139 ± 0.036
≥ 9.7	5 - 50	-1.154 ± 0.005	1.800 ± 0.031	-0.552 ± 0.015	-0.498 ± 0.026	3.446 ± 0.092
	Ι	-1.785 ± 0.018	1.277 ± 0.127	-0.100 ± 0.055	Ι	Ι
	5 - 30	-1.333 ± 0.006	1.972 ± 0.034	-0.619 ± 0.016	-0.177 ± 0.007	1.205 ± 0.023
	14 - 43	-1.194 ± 0.007	1.955 ± 0.041	-0.643 ± 0.020	-0.241 ± 0.009	1.586 ± 0.030
∠ 10.3	5 - 50	-1.022 ± 0.005	1.933 ± 0.033	-0.630 ± 0.016	-0.360 ± 0.013	2.385 ± 0.045
	Ι	-1.530 ± 0.020	1.519 ± 0.134	-0.160 ± 0.059	Ι	Ι
	5 - 30	-1.143 ± 0.014	2.166 ± 0.082	-0.561 ± 0.040	-0.126 ± 0.007	0.865 ± 0.019
/ 11.0	14 - 43	-1.000 ± 0.015	2.270 ± 0.084	-0.657 ± 0.042	-0.171 ± 0.007	1.097 ± 0.024
Z 11.0	5 - 50	-0.832 ± 0.009	2.172 ± 0.055	-0.607 ± 0.027	-0.257 ± 0.012	1.666 ± 0.037
		-1.189 ± 0.028	1.774 ± 0.186	-0.133 ± 0.085	I	I

6.4 From simulation to observation

Simulations are a tool needed to interpret observed data. The galaxy-galaxy merger rate is a particularly difficult quantity to derive from the relatively static view of the the universe we see in galaxy surveys. As we cannot view the complete merging of two galaxies in real time a proxy is required as stand in. Obvious physical tracers of a recent merger such as disturbed morphologies present one option for deducing the galaxy merger rate. Methods invoking quantitative morphology such as $G - M_{20}$ or asymmetry are not equally sensitive to all merger mass ratios. Furthermore, these morphological methods are sensitive to total mass, gas properties, orbital parameters, merger stage, and viewing angles (Abraham et al., 2003; Conselice et al., 2003; Lotz et al., 2008, 2011; Scarlata et al., 2007). These additional difficulties present a greater barrier to identifying mergers and determining a cosmological merger rate (Kampczyk et al., 2007; Scarlata et al., 2007; López-Sanjuan et al., 2009; Shi et al., 2009; Kartaltepe et al., 2010; Abruzzo et al., 2018; Nevin et al., 2019).

One common observational method for deriving the galaxy merger rate is through the analysis of galaxies in close pairs. The foundation of this approach is simple, as galaxies found in close proximity are expected to merge within some finite predictable time scale. Three key questions must be answered to utilise this method: How should we select galaxy pairs, do those two pairs merge, and on what timescale does an observed pair merge?

6.4.1 Light-cone construction

Throughout this work we will make use of mock light-cone galaxy catalogues. Working with light-cones offers a few advantages to working with snapshot catalogues directly. The first reason is that these catalogues provide a more natural comparison between simulation and observation. In a snapshot catalogue all galaxies exist at the same cosmological stage of evolution. In contrast, real observations have to contend with galaxies samples that span often large redshift ranges encompassing galaxies at various stages of evolution. Our catalogues are constructed in the same way, such that a galaxy is placed at its cosmologically relevant stage of evolution according to the comoving distance from the observer. Another advantage is that these catalogues inherently adopt the same constraints seen in observation due to limited viewing angles. While this does limit the sample size we can use for analysis, it provides a more appropriate environment to test how the models we develop and the conclusions drawn are impacted by these real limitations.

Our cone geometry is set using the method described by Kitzbichler & White (2008). In this method light-cone geometry is defined by two integers m and n. The line of sight vector u_3 , is defined by a line drawn from the origin though the point $(L_{\text{box}}/m, L_{\text{box}}/n, L_{\text{box}})$, where L_{box} is the comoving side length of our simulation volume (200 Mpc). The second vector u_1 is defined to be orthogonal to u_3 and the coordinate axis corresponding to the smaller value of m and n. The final vector u_2 is defined to be orthogonal to u_3 and u_1 , where all three taken together form a right handed coordinate system. The observation area of he light-cone is then covered by $(m^2n)^{-1} \times (mn^2)^{-1} \text{rad}^2$, which is centred along u_3 with edges aligned along u_1 and u_2 . When galaxies are placed into light cones we draw from snapshots according to

$$\frac{D_i + D_{i-1}}{2} \le D_{\text{gal}} < \frac{D_i + D_{i+1}}{2} \tag{6.4}$$

where D_i is the cosmological distance to simulation snapshot with index *i*, and D_{gal} is the comoving distance to the galaxy within the light-cone. The 'cosmological' redshift of each galaxy is then set by that comoving distance from the observer. Additionally, we apply redshiftspace distortions to each catalogue galaxy to obtain its observed redshift z_{obs} .

We construct a series of light-cones intended to reproduce the observation area of the five CANDLES (Koekemoer et al., 2011) survey fields. These fields serve as the test bed for applying our fitting formula. In addition to these mock catalogues we construct an additional 'full width' light-cone catalogue. This data-set has galaxies placed according to eq. 6.4 just as in the standard light-cone catalogues. However the line of sight vector u_3 is aligned with the coordinate z-axis, and no view restricting cone geometry is applied. This results in a catalogue with galaxies occupying a rectangular volume and creates a 'smoothly' evolving galaxy catalogue without the restrictions due to limited viewing angle, which is helpful when fitting data at low redshift.

6.4.2 Identifying close pairs

In this work pairs are identified and defined in terms of the following:

- m_1 : The stellar mass of the main (more massive) galaxy in each pair.
- μ : The stellar mass ratio taken with respect to the two galaxies forming the pair, $\mu \equiv m_1/m_2$. Here m_2 is the *minor* (less massive) galaxy in the pair. The mass ratio is defined such that $\mu \geq 1$.
- z: Pair redshift measured at the observed redshift of the main galaxy.
- Δr : The *projected* radial separation between the two galaxies in pkpc.
- Δv : The line-of-sight velocity difference between the two galaxies, as measured by the difference in their respective z_{obs}

In this work we provide fits for a variety of mass and mass ratio combinations. For simplicity of conveying key concepts and conclusions we use a single reference case when displaying results. Which refers to pairs with $\log_{10}(m_1) \ge 10.3$, $\Delta v \le 500 \text{ km s}^{-1}$ and $1 \le \mu < 4$. This selection criteria is used for ease of comparison with O'Leary et al. (2021a). In Figure 6.3 we compare the pair fractions determined from the data sets described in §6.4.1 with the pair fractions of O'Leary et al. (2021a), which computed the pairs on a snapshot by snapshot basis. With this we are able to see that our underlying results agree, and can get a good idea over the amount of uncertainty in the pair fraction under more realistic observable constraints.



Figure 6.3: The pair fraction evolution in our simulated galaxy catalogues for pairs with $\log_{10}(m_1/M_{\odot}) \geq 10.3$, $5 \leq \Delta r < 50$ kpc projected separation, $\Delta v \leq 500$ km s⁻¹ and $1 \leq \mu < 4$. The solid black line indicates the pair fraction computed at each simulation snapshot as in O'Leary et al. (2021a). Coloured lines show the pair fraction evolution with our mock light-cones. Poission error in the number count of pairs is reflected in the error bars for the mock light-cones, and the grey shaded region for the results of O'Leary et al. (2021a). Here we do not place any redshift restraints on the light cone catalogues that would more closely resemble the observables limits of the noted surveys.


Figure 6.4: The distribution of pairs for $\log_{10}(m_1) \ge 10.3$ and $1 \le \mu < 4$. Each panel includes pairs where the observed redshift of the primary galaxy falls within ± 0.5 of the noted central redshift. The percentages shown correspond indicate pairs within that redshift bin, not percentages to the entire catalogue of pairs.

6.5 Merging probabilities and timescales

Once we have a handle on the pair fraction, we can determine the galaxy merger rates: for a given mass range, observable aperture and redshift proximity criteria we compute the pair fraction and divide that value by the average time that pairs under that selection criteria will be observable. The galaxy rate can then be expressed as:

$$\Re = C_{\text{merge}} \frac{f_p}{\langle T_{\text{obs}} \rangle} , \qquad (6.5)$$

where $\langle T_{obs} \rangle$ is the average observation timescale. The term C_{merge} is an optional correction factor to account for pairs that do not merge before z = 0 or at all. In general these two quantities are dependent on the pair selection criteria imposed by the observer.

The goal of this work is to characterise these values, and provide meaningful formulations that reduce the need to establish fitting functions for each specific observation. In the theoretical framework there are a few ways we can do this. One approach is to use idealised galaxy merger simulations, making mock pair observations and tracking these merging pairs to final coalescence (Lotz et al., 2011). Alternatively, one could employ some self-consistent model for galaxy formation in large volumes, and track pairs identified in mock observations (Snyder et al., 2017). Previous efforts have focused on setting a fixed pair selection criteria, and finding some fitting function for the observation timescale under that criteria. The issue with that approach is that those fitting functions cannot be easily applied to pairs identified under some other criteria.

In the next two sections we provide fitting functions for both C_{merge} and T_{merge} . We will describe the process we use to determine these formulae that can describe these values for a range of Δr , Δv , and redshift commonly used in the literature. In both cases parameter space is explored and best fit values determined using the affine invariant ensemble sampler described in Goodman & Weare (2010) as implemented in Emcee (Foreman-Mackey et al., 2013). For each mass range explored we fit to three mass ratio intervals 1 : 4 (major), 4 : 10 (minor), and 1 : 10 (all). Throughout, we only included pairs with $\Delta v \leq 500 \text{ km s}^{-1}$ and $\Delta r \leq 100 \text{ kpc}$.



Figure 6.5: The per cent of pairs that merge by z = 0 for pairs with $\log_{10}(m_1) \ge 10.3$ with $1 \le \mu < 4$. Each panel includes pairs where the observed redshift of the primary galaxy falls within ± 0.5 of the noted central redshift.

As an aside, the nature of this problem makes it attractive to deal with from a machine learning perspective. If our central interest is determining the observation timescale of any pair regardless of the mechanisms driving those timescales it would appear on its surface to be an ideal problem for machine learning algorithms. We tested this approach using a random forest regressor to predict merging probabilities and timescales. In practice we found there was not enough information in the *observable* pair features to meaningfully predict the desired values on an individual basis, due to information loss in projected quantities. Other works have had greater success on this front using more advanced networks (Pfister et al., 2020).

6.5.1 Merging Probability

First we need to address the merging probability. Traditionally chance pairs, and pairs that did not have enough time to merge with in the average timescale were captured through a correction factor C_{merge} . This correction factor typically takes some value between 0.4 and 1.0. For instance Lotz et al. (2011) adopt a constant 0.6 for all scenarios they tested. Other works have chosen to marignalize over this parameter by including it directly into the T_{obs} formulation. This the approach taken in Snyder et al. (2017) as well as in O'Leary et al. (2021a).

Several recent works have adopted a probabilistic approach (Duncan et al., 2019; Ventou et al., 2019). In these works each pair is assigned a weight, based on some observed properties, that the pair is physically associated and will merge on in the expected timescale. Ventou et al. (2019) derived their weighting function through pairs selected in the Illustris simulation, in their results they determined the merger probability could be well described by an exponential function in both Δr and Δv with little redshift dependence. Looking at Figure 6.5 we can immediately see a redshift dependent formulation is required. Just probing this coarse redshift bins shown in Figure 6.5 we find the probability appears to saturate near $z \approx 3.5$. However it should not be surprising that there is a redshift dependence to the merger probability. This value is defined as the probability that two galaxies will merge by z = 0 which directly implies that a pair at z = 0 would have no

chance to merge. Similarly, if two pairs at high z are in fact physically associated, there is simply more time available where they *could* merge before present day. Redshift dependencies not withstanding, we found that the fitting formula of Ventou et al. (2019) can only reasonably fit our data for $z \leq 1$.

After testing several candidate fitting functions we found that the merger probability can be well described by a logistic function in velocity, with redshift and radial dependencies, eq. 6.6.

$$W(\Delta r, \Delta v, z) = \frac{\exp(b\Delta r)}{1 + \exp[c_0(\Delta v - a)]}$$

$$a = a_0(1+z)^{a_z} + a_r\Delta r$$

$$b = b_0 + (1+z)b_z$$
(6.6)

In this formulation the maximum of the curve is set by the term $\exp(b\Delta r)$ where the exponential slope *b* has a linear dependence on redshift. The logistic midpoint *a* was found to have dependencies in both redshift and Δr . For this parameter a linear relation to Δr and a power-law relation to redshift produced the best results. We did not find that the logistic growth rate c_0 , varies strongly with either Δr redshift. Eq. 6.6 reproduces the data best for smaller Δr and Δv .

Table 6.3 Shows the best fit parameters from eq. 6.6 for a range of primary stellar mass thresholds, and mass ratios. For each mass range we fit to major $(1 \le \mu < 4)$, minor $(4 \le \mu < 10)$ and major+minor $(1 \le \mu < 10)$ pair mass ratios.

6.5.2 Merging Timescales

The observation timescales shown in O'Leary et al. (2021a) are derived by mapping the pair fraction directly onto the merger rate without making any consideration over which pairs actually ended up merging. There we found that $T_{\rm obs} \propto (1+z)^{-1}$ provided a reasonable translation from pair fractions to intrinsic galaxy merger rates. This is a weaker scaling than the $T_{\rm obs} \propto (1+z)^{-2}$ proposed by Snyder et al. (2017), and the $T_{\rm obs} \propto H(z)^{-1/3}$ scaling suggested by Jiang et al. (2014).

In this work, we identify individual pairs and track them until the point they finally merge. Here we should highlight that the value that we are measuring is the *merging* timescale not the *observation* timescale. The observation timescale tracks how long a pair would remain in the aperture set by the observer. The observation timescale should in principle also incorporate information on how long a pair remains in the mass and mass ratio criteria that has been set. Because $T_{\rm obs}$ is dependent on the particular observation, we have elected to fit to the merging timescales as it can be more readily applied to a broader range of selection criteria.

9.
9
eq.
probability
merging
for
parameters
fit
Best
6.3:
Table

Mass ratio μ and a_0	$a_0 \left[\frac{1}{41} \right]$	$\tan s^{-1}$	a_z	$\frac{a_r \left[\text{ km s}^{-1} \text{ kpc}^{-1} \right]}{\frac{1}{2} \frac{1}{1} \frac{1}$	$b_0 \left[{ m kpc}^{-1} ight]$	$b_{z} [\mathrm{kpc}^{-1}]$	$c_0 [\mathrm{skm^{-1}}]$
$\begin{array}{ccc} 1 \leq \mu < 4 & 41.1^{-2} \\ 4 \leq \mu < 10 & 72.5^{+2} \\ -3 \end{array}$	41.1_{-2} 72.5_{-3}	0.3 0.3 0.3	$1.20_{-0.35}^{-0.35}$ $0.95_{-0.35}^{+0.30}$	-1.10 -1.08 -1.43 $+1.13$ -1.43 -0.98	$-0.0277_{-0.0096}$ $-0.0248_{-0.0124}$	$0.0033^{+0.0016}_{-0.0026}$	$0.01083_{-0.0016}^{0.0016}$ $0.0108_{-0.0025}^{0.0019}$
$1 \le \mu < 10 \qquad 57.6^{+23.6}_{-31.6}$	$57.6^{+23.6}_{-31.6}$		$1.06\substack{+0.31\\-0.35}$	$-1.28_{-0.97}^{+1.15}$	$-0.0265_{-0.0108}^{+0.0086}$	$0.0040\substack{+0.0022\\-0.0018}$	$0.0098_{-0.0020}^{+0.0016}$
$1 \le \mu < 4$ 234^{+49}_{-48}	234^{+49}_{-48}		$0.629\substack{+0.167\\-0.210}$	$-1.84^{+1.25}_{-1.00}$	$-0.0141\substack{+0.0056\\-0.0065}$	$0.0021\substack{+0.0014\\-0.0012}$	$0.0084_{-0.0029}^{+0.0021}$
$4 \le \mu < 10 \qquad 197^{+47}_{-48}$	197^{+47}_{-48}		$0.665_{-0.221}^{+0.176}$	$-2.00^{+1.18}_{-0.95}$	$-0.0145_{-0.0079}^{+0.0068}$	$0.0019\substack{+0.0017\\-0.0015}$	$0.0093^{+0.0022}_{-0.0029}$
$1 \le \mu < 10 \qquad 215^{+44}_{-44}$	215^{+44}_{-44}	- 1	$0.644_{-0.192}^{+0.153}$	$-1.89^{+1.13}_{-0.92}$	$-0.0147^{+0.0057}_{-0.0067}$	$0.0021\substack{+0.0013\\-0.0012}$	$0.0089\substack{+0.0020\\-0.0027}$
$1 \le \mu < 4 \qquad 400^{+120}_{-217}$	400^{+120}_{-217}	1	$1.17_{-0.41}^{+0.43}$	$-1.58^{+2.20}_{-2.08}$	$-0.0141^{+0.0041}_{-0.0042}$	$0.0027_{-0.0011}^{+0.0011}$	$0.0030_{-0.0026}^{+0.0014}$
$4 \le \mu < 10 \qquad 152^{+43}_{-60}$	152^{+43}_{-60}		$1.33\substack{+0.36\\-0.31}$	$-0.54_{-1.48}^{+2.20}$	$-0.0264^{+0.0062}_{-0.0067}$	$0.0045_{-0.0015}^{+0.0016}$	$0.0050\substack{+0.0017\\-0.0026}$
$1 \le \mu < 10 \qquad 339^{+110}_{-211}$	$339\substack{+110\\-211}$		$1.12\substack{+0.40\\-0.41}$	$-1.33^{+2.30}_{-2.02}$	$-0.0170\substack{+0.0045\\-0.0046}$	$0.0031\substack{+0.0011\\-0.0011}$	$0.0026\substack{+0.0012\\-0.0020}$
$1 \le \mu < 4$ 238^{+66}_{-71}	238^{+66}_{-71}		$0.467\substack{+0.200\\-0.248}$	$-2.38^{+1.47}_{-1.35}$	$-0.0165_{-0.0073}^{+0.0067}$	$0.0029\substack{+0.0015\\-0.0013}$	$0.0057\substack{+0.0012\\-0.0017}$
$4 \le \mu < 10 \qquad 163^{+45}_{-47}$	163^{+45}_{-47}		$0.716\substack{+0.196\\-0.245}$	$-2.17\substack{+1.20\-1.04}$	$-0.0165^{+0.0081}_{-0.0092}$	$0.0021\substack{+0.0019\\-0.0017}$	$0.0086\substack{+0.0017\\-0.0024}$
$1 \le \mu < 10$ 200^{+55}_{-59}	200^{+55}_{-59}		$0.584\substack{+0.200\\-0.246}$	$-2.18\substack{+1.43\\-1.23}$	$-0.0170\substack{+0.0074\\-0.0080}$	$0.0026\substack{+0.0016\\-0.0015}$	$0.0068\substack{+0.0014\\-0.0020}$
$1 \le \mu < 4 \qquad 262^{+64}_{-68}$	262^{+64}_{-68}		$0.494\substack{+0.186\\-0.230}$	$-2.35^{+1.46}_{-1.30}$	$-0.0151_{-0.0069}^{+0.0063}$	$0.0025_{-0.0013}^{+0.0014}$	$0.0061\substack{+0.0013\\-0.0020}$
$4 \leq \mu < 10$ 176^{+45}_{-47}	176^{+45}_{-47}		$0.718\substack{+0.184\\-0.228}$	$-2.14^{+1.25}_{-1.03}$	$-0.0159^{+0.0074}_{-0.0085}$	$0.0021\substack{+0.0017\\-0.0015}$	$0.0085^{+0.0018}_{-0.0024}$
$1 \le \mu < 10$ 218^{+55}_{-58}	218^{+55}_{-58}		$0.604\substack{+0.186\\-0.232}$	$-2.27^{+1.40}_{-1.21}$	$-0.0153^{+0.0067}_{-0.0074}$	$0.0023\substack{+0.0015\\-0.0014}$	$0.0068\substack{+0.0015\\-0.0021}$
$1 \le \mu < 4$ 282^{+65}_{-69}	282_{-69}^{+65}		$0.592_{-0.237}^{+0.195}$	$-2.39^{+1.51}_{-1.41}$	$-0.0131^{+0.0056}_{-0.0061}$	$0.0021\substack{+0.0013\\-0.0012}$	$0.0060_{-0.0014}^{+0.0014}$
$4 \le \mu < 10$ 188^{+45}_{-47}	188^{+45}_{-47}		$0.752_{-0.229}^{+0.184}$	$-2.08^{+1.32}_{-1.12}$	$-0.0152^{+0.0068}_{-0.0077}$	$0.0021\substack{+0.0016\\-0.0014}$	$0.0082\substack{+0.0018\\-0.0025}$
$1 \le \mu < 10$ 249^{+55}_{-57}	249^{+55}_{-57}		$0.617\substack{+0.175\\-0.212}$	$-2.41^{+1.38}_{-1.23}$	$-0.0140^{+0.0058}_{-0.0065}$	$0.0021\substack{+0.0013\\-0.0012}$	$0.0067\substack{+0.0015\\-0.0023}$
$1 \le \mu < 4 \qquad 289^{+65}_{-71}$	289^{+65}_{-71}		$0.71\substack{+0.22\\-0.27}$	$-2.10^{+1.65}_{-1.57}$	$-0.0126^{+0.0050}_{-0.0054}$	$0.0020\substack{+0.0012\\-0.0011}$	$0.0061\substack{+0.0017\\-0.0027}$
$4 \le \mu < 10$ 190^{+50}_{-47}	190^{+50}_{-47}		$-2.17^{+1.29}_{-1.44}$	$0.84\substack{+0.24\\-0.20}$	$-0.0144_{-0.0064}^{+0.0072}$	$0.0020\substack{+0.0013\\-0.0015}$	$0.0074\substack{+0.0025\\-0.0017}$
$1 \le \mu < 10 \qquad 260^{+58}_{-62}$	260^{+58}_{-62}		$0.70\substack{+0.19\\-0.23}$	$-2.52^{+1.43}_{-1.40}$	$-0.0126^{+0.0053}_{-0.0057}$	$0.0019\substack{+0.0012\\-0.0011}$	$0.0061\substack{+0.0014\\-0.0023}$
$1 \le \mu < 4$ 295^{+67}_{-76}	295^{+67}_{-76}		$0.813^{+0.241}_{-0.315}$	$-1.97^{+1.75}_{-1.67}$	$-0.0124_{-0.0049}^{+0.0044}$	$0.0021^{+0.0011}_{-0.0010}$	$0.0057\substack{+0.0017\\-0.0028}$
$4 \le \mu < 10 \qquad 175^{+49}_{-54}$	175^{+49}_{-54}		$0.966_{-0.309}^{+0.253}$	$-1.52^{+1.70}_{-1.49}$	$-0.0164\substack{+0.0065\\-0.0074}$	$0.0024\substack{+0.0016\\-0.0014}$	$0.0071\substack{+0.0018\\-0.0027}$
$1 \le \mu < 10 \qquad 255^{+62}_{-67}$	255^{+62}_{-67}		$0.816\substack{+0.222\\-0.279}$	$-2.04^{+1.74}_{-1.65}$	$-0.0133^{+0.0051}_{-0.0055}$	$0.0020\substack{+0.0012\\-0.0011}$	$0.0057\substack{+0.0015\\-0.0024}$
$1 \le \mu < 4$ 321^{+80}_{-111}	321^{+80}_{-111}		$1.04\substack{+0.34\\-0.42}$	$-1.82^{+1.98}_{-1.97}$	$-0.0128^{+0.0040}_{-0.0045}$	$0.0023^{+0.0011}_{-0.0010}$	$0.0044_{-0.0027}^{+0.0016}$
$4 \le \mu < 10 \qquad 152^{+44}_{-56}$	152^{+44}_{-56}		$1.21\substack{+0.33\\-0.35}$	$-0.65^{+2.01}_{-1.46}$	$-0.0209^{+0.0061}_{-0.0068}$	$0.0033^{+0.0015}_{-0.0014}$	$0.0061\substack{+0.0018\\-0.0027}$
$1 \leq \mu < 10$ 267^{+72}_{-93}	267^{+72}_{-93}		$1.03\substack{+0.32\\-0.39}$	$-1.53^{+2.08}_{-1.90}$	$-0.0149\substack{+0.0044\\-0.0049}$	$0.0024\substack{+0.0011\\-0.0010}$	$0.0042\substack{+0.0015\\-0.0024}$
$1 \le \mu < 4 \qquad 400^{+120}_{-220}$	400^{+120}_{-220}		$1.16\substack{+0.43\\-0.41}$	$-1.51_{-2.05}^{+2.22}$	$-0.0141_{-0.0042}^{+0.0040}$	$0.0028_{-0.0011}^{+0.0011}$	$0.0030^{+0.0014}_{-0.0027}$
$4 \le \mu < 10$ 153^{+43}_{-62}	153_{-62}^{+43}		$1.33_{-0.32}^{+0.36}$	$-0.54_{-1.49}^{+2.17}$	$-0.0264_{-0.0068}^{+0.0062}$	$0.0045_{-0.0015}^{+0.0015}$	$0.0050_{-0.0027}^{+0.0017}$
$1 \le \mu < 10$ 339^{+112}_{-200}	339^{+112}_{-200}		$1.12_{-0.40}^{+0.40}$	$-1.29^{+2.34}_{-1.98}$	$-0.0169^{+0.0044}_{-0.0045}$	$0.0031^{+0.0011}_{-0.0011}$	$0.0026_{-0.0020}^{+0.0012}$

94

$\log_{10}(m_1/{ m M}_\odot)$	Mass ratio μ	$a_0 [Gyr]$	a_{z} [Gyr]	$b_0 \left[\text{ Gyr kpc}^{-1} \right]$	b_{z} [Gyr kpc ⁻¹]	$c_0 \left[\mathrm{Gyr s km^{-1}} \right]$	$c_{z} \left[\text{ Gyr s km}^{-1} \right]$
	$1 \leq \mu < 4$	$0.533_{-0.420}^{+0.427}$	$0.0312^{+0.0423}_{-0.0420}$	$0.0426\substack{+0.0252\\-0.0252}$	$-0.0050^{+0.0064}_{-0.0064}$	$0.0118_{-0.0089}^{+0.0091}$	$-0.0026_{-0.0022}^{+0.0021}$
9.0 - 10.0	$4 \leq \mu < 10$	$0.498_{-0.398}^{+0.401}$	$-0.0597^{+0.0808}_{-0.0797}$	$0.0343_{-0.0090}^{+0.0089}$	$-0.0011_{-0.0015}^{+0.0015}$	$0.0077_{-0.0058}^{+0.0059}$	$-0.0015_{-0.0014}^{+0.0014}$
	$1 \leq \mu < 10$	$0.483_{-0.375}^{+0.372}$	$-0.0289\substack{+0.0385\\-0.0390}$	$0.0429\substack{+0.0202\\-0.0204}$	$-0.0040^{+0.0051}_{-0.0051}$	$0.0081\substack{+0.0066\\-0.0065}$	$-0.0016\substack{+0.0016\\-0.0016}$
	$1 \leq \mu < 4$	$-0.284_{-0.330}^{+0.330}$	$0.224\substack{+0.105\\-0.105}$	$0.0315_{-0.0079}^{+0.0079}$	$-0.0036^{+0.0023}_{-0.0023}$	$0.0068^{+0.0015}_{-0.0015}$	$-0.0019^{+0.0004}_{-0.0004}$
10.0 - 11.0	$4 \leq \mu < 10$	$0.083\substack{+0.105\\-0.106}$	$0.024\substack{+0.031\\-0.031}$	$0.0334_{-0.0093}^{+0.0093}$	$-0.0028^{+0.0029}_{-0.0029}$	$0.0069^{+0.0017}_{-0.0017}$	$-0.0015^{+0.0005}_{-0.0005}$
	$1 \leq \mu < 10$	$-0.164_{-0.214}^{+0.211}$	$0.146\substack{+0.078\\-0.078}$	$0.0334\substack{+0.0085\\-0.0085}$	$-0.0036^{+0.0026}_{-0.0026}$	$0.0069_{-0.0016}^{+0.0016}$	$-0.0017_{-0.0004}^{+0.0005}$
	$1 \leq \mu < 4$	$-0.265_{-0.277}^{+0.277}$	$0.179\substack{+0.084\\-0.083}$	$0.0311_{-0.0052}^{+0.0053}$	$-0.0058^{+0.0016}_{-0.0016}$	$0.0038_{-0.0009}^{+0.0009}$	$-0.0011^{+0.0003}_{-0.0003}$
11.0 - 12.0	$4 \leq \mu < 10$	$0.450^{+0.178}_{-0.176}$	$-0.041_{-0.052}^{+0.050}$	$0.0452_{-0.0078}^{+0.0084}$	$-0.0093^{+0.0025}_{-0.0026}$	-0.0005 ± 0.0005	$0.0003 \substack{+0.0001\\-0.0002}$
	$1 \leq \mu < 10$	$0.133_{-0.094}^{+0.094}$	$0.0003^{+0.0005}_{-0.0005}$	$0.0272\substack{+0.0051\\-0.0054}$	$-0.0032^{+0.0014}_{-0.0015}$	$0.0030^{+0.0008}_{-0.0008}$	$-0.0007 \substack{+0.0002\\-0.0002}$
	$1 \leq \mu < 4$	$-0.137^{+0.181}_{-0.181}$	$0.192\substack{+0.081\\-0.082}$	$0.0315\substack{+0.0091\\-0.0092}$	$-0.0028^{+0.0028}_{-0.0028}$	$0.0053_{-0.0014}^{+0.0014}$	$-0.0014_{-0.0004}^{+0.0004}$
≥ 9.5	$4 \leq \mu < 10$	$0.135\substack{+0.161\\-0.161}$	$0.027\substack{+0.035\\-0.035}$	$0.0323^{+0.0078}_{-0.0078}$	$-0.0018\substack{+0.0022\\-0.0022}$	$0.0064\substack{+0.0021\\-0.0021}$	$-0.0014\substack{+0.0006\\-0.0006}$
	$1 \leq \mu < 10$	$0.012\substack{+0.016\\-0.016}$	$0.101_{-0.066}^{+0.065}$	$0.0327^{+0.0092}_{-0.0093}$	$-0.0025_{-0.0028}^{+0.0028}$	$0.0055_{-0.0016}^{+0.0016}$	$-0.0013_{-0.0005}^{+0.0005}$
	$1 \leq \mu < 4$	$-0.058_{-0.077}^{+0.077}$	$0.149\substack{+0.063\\-0.063}$	$0.0276\substack{+0.0064\\-0.0063}$	$-0.0017\substack{+0.0019\\-0.0019}$	$0.0051\substack{+0.0011\\-0.0012}$	$-0.0014\substack{+0.0004\\-0.0003}$
≥ 9.7	$4 \leq \mu < 10$	$0.094^{+0.121}_{-0.123}$	$0.034_{-0.043}^{+0.042}$	$0.0322_{-0.0078}^{+0.0078}$	$-0.0019_{-0.0023}^{+0.0023}$	$0.0067\substack{+0.0018\\-0.0018}$	$-0.0015_{-0.0005}^{+0.0005}$
	$1 \leq \mu < 10$	$0.015\substack{+0.020\\-0.020}$	$0.098\substack{+0.060\\-0.060}$	$0.0313^{+0.0083}_{-0.0085}$	$-0.0024^{+0.0025}_{-0.0025}$	$0.0053\substack{+0.0014\\-0.0014}$	$-0.0013^{+0.0004}_{-0.0004}$
	$1 \leq \mu < 4$	$-0.187\substack{+0.233\\-0.238}$	$0.187\substack{+0.082\\-0.082}$	$0.0291\substack{+0.0071\\-0.0070}$	$-0.0030^{+0.0021}_{-0.0021}$	$0.0055^{+0.0012}_{-0.0012}$	$-0.0015^{+0.0004}_{-0.0004}$
≥ 10.0	$4 \leq \mu < 10$	$0.102\substack{+0.124\\-0.124}$	$0.025\substack{+0.031\\-0.031}$	$0.0336\substack{+0.0091\\-0.0090}$	$-0.0029^{+0.0027}_{-0.0027}$	$0.0061\substack{+0.0016\\-0.0016}$	$-0.0013^{+0.0004}_{-0.0004}$
	$1 \leq \mu < 10$	$-0.044\substack{+0.059\\-0.059}$	$0.109\substack{+0.053\\-0.054}$	$0.0309\substack{+0.0079\\-0.0079}$	$-0.0029^{+0.0024}_{-0.0024}$	$0.0055^{+0.0013}_{-0.0013}$	$-0.0014\substack{+0.0004\\-0.0004}$
	$1 \leq \mu < 4$	$-0.226^{+0.267}_{-0.265}$	$0.179\substack{+0.081\\-0.080}$	$0.0288^{+0.0061}_{-0.0061}$	$-0.0036^{+0.0018}_{-0.0018}$	$0.0058_{-0.0011}^{+0.0011}$	$-0.0016^{+0.0003}_{-0.0003}$
≥ 10.3	$4 \leq \mu < 10$	$0.068\substack{+0.089\\-0.090}$	$0.033\substack{+0.038\\-0.038}$	$0.0334^{+0.0087}_{-0.0086}$	$-0.0034^{+0.0025}_{-0.0025}$	$0.0063\substack{+0.0015\\-0.0015}$	$-0.0014\substack{+0.0004\\-0.0004}$
	$1 \leq \mu < 10$	$0.008\substack{+0.010\\-0.010}$	$0.077\substack{+0.043\\-0.043}$	$0.0289\substack{+0.0067\\-0.0067}$	$-0.0026^{+0.0020}_{-0.0021}$	$0.0052\substack{+0.0011\\-0.0011}$	$-0.0013^{+0.0003}_{-0.0003}$
	$1 \leq \mu < 4$	$-0.103^{+0.133}_{-0.134}$	$0.113\substack{+0.050\\-0.050}$	$0.0273\substack{+0.0055\\-0.0054}$	$-0.0031^{+0.0017}_{-0.0017}$	$0.0050_{-0.0009}^{+0.0009}$	$-0.0013^{+0.0003}_{-0.0003}$
≥ 10.5	$4 \leq \mu < 10$	$0.043\substack{+0.058\\-0.058}$	$0.047\substack{+0.043\\-0.043}$	$0.0363_{-0.0090}^{+0.0090}$	$-0.0048^{+0.0026}_{-0.0026}$	$0.0062\substack{+0.0014\\-0.0014}$	$-0.0014\substack{+0.0004\\-0.0004}$
	$1 \leq \mu < 10$	$-0.042^{+0.056}_{-0.057}$	$0.087\substack{+0.040\\-0.040}$	$0.0307\substack{+0.0064\\-0.0064}$	$-0.0037^{+0.0019}_{-0.0019}$	$0.0051\substack{+0.0011\\-0.0011}$	$-0.0012^{+0.0003}_{-0.0003}$
	$1 \le \mu < 4$	$-0.097\substack{+0.125\\-0.126}$	$0.096\substack{+0.041\\-0.042}$	$0.0280\substack{+0.0051\\-0.0050}$	$-0.0039^{+0.0015}_{-0.0015}$	$0.0045\substack{+0.008\\-0.008}$	$-0.0012^{+0.0002}_{-0.0002}$
≥ 10.8	$4 \leq \mu < 10$	$0.239\substack{+0.158\\-0.156}$	$0.023\substack{+0.030\\-0.030}$	$0.0330\substack{+0.0106\\-0.0106}$	$-0.0049^{+0.0032}_{-0.0033}$	$0.0045^{+0.0016}_{-0.0016}$	$-0.0011\substack{+0.0005\\-0.0005}$
	$1 \leq \mu < 10$	$0.085_{-0.099}^{+0.099}$	$0.020\substack{+0.025\\-0.025}$	$0.0295^{+0.0059}_{-0.0060}$	$-0.0036^{+0.0017}_{-0.0017}$	$0.0042\substack{+0.0009\\-0.0009}$	$-0.0009^{+0.0003}_{-0.0003}$
	$1 \leq \mu < 4$	$-0.267\substack{+0.275\\-0.277}$	$0.180\substack{+0.083\\-0.084}$	$0.0312^{+0.0052}_{-0.0051}$	$-0.0059^{+0.0015}_{-0.0016}$	$0.0038^{+0.0009}_{-0.0009}$	$-0.0011^{+0.0003}_{-0.0003}$
≥ 11.0	$4 \leq \mu < 10$	$0.449\substack{+0.169\\-0.170}$	$-0.040\substack{+0.049\\-0.050}$	$0.0453^{+0.0083}_{-0.0075}$	$-0.0088^{+0.0026}_{-0.0020}$	$-0.0005^{+0.0005}_{-0.0005}$	$0.0003\substack{+0.0001\\-0.0002}$
	$1 \leq \mu < 10$	$0.133\substack{+0.093\\-0.095}$	$0.0004\substack{+0.0006\\-0.0006}$	$0.0276\substack{+0.0054\\-0.0051}$	$-0.0032^{+0.0014}_{-0.0015}$	$0.0031\substack{+0.0007\-0.0008}$	$-0.0007\substack{+0.0002\\-0.0002}$

Table 6.4: Best fit parameters for merging timescale eq. 6.7.



Figure 6.6: The mean merging timescales for pairs with $\log_{10}(m_1) \ge 10.3$ with $1 \le \mu < 4$. Each panel includes pairs where the observed redshift of the primary galaxy falls within ± 0.5 of the noted central redshift.

Figure 6.6 shows the mean merging timescale as a function of Δr and Δv in several redshift bins for our reference case. For $z \gtrsim 0.5$ the data indicates merging timescales that decrease with redshift, as expected from previous results. We also find that the dependence on Δv has a stronger redshift scaling that the dependence on Δr . We found that a flat plane was sufficient to reproduce the data across a wide range of redshifts:

$$T_{\text{merge}}(z, \Delta r, \Delta v) = a + (b\Delta r) + (c\Delta v)$$

$$a = a_0 + (1+z)a_z$$

$$b = b_0 + (1+z)b_z$$

$$c = c_0 + (1+z)c_z$$
(6.7)

In this formulation allowing each parameter to scale linearly with redshift provided the best reproduction of merger rates.

Table 6.4 Shows the best fit parameters from eq. 6.7 for a range of the same mass ranges shown in Table 6.3. Generally eq. 6.7 provides an accurate description of the data for $1 \leq z \leq 3.5$ for most mass ranges and mass ratios tested. Within this redshift range merging timescales increase with increasing Δv as seen in the data. However, outside this range the scaling with respect to Δv undergoes a sign inversion due to the linear redshift scaling of the parameter c. This inversion results in a considerable under prediction of the merging timescales for small radii and $\Delta v \geq 200$ km s⁻¹. However, since most of the major pairs sit below this Δv and at a more moderate Δr , we find this limited fit to the data does not strongly impact our reproduced major merger rates (see §6.6).

We tested a variety of fitting functions and parameter redshift scaling to fit the pair merging timescales. Notably, allowing for $c \propto (1+z)^{c_z}$ alleviates the severity of underpredicted merging timescales for high Δv pairs. However this improved fit at high-*z* comes at the expense of a significantly worse reproduction of merger rates for $z \leq 3$. The Δv slope scaling with redshift is the dominate parameter determining the goodness of fit for this parameterisation. The scaling with Δr is relatively stable with redshift and a suitable merger rate reproduction is possible if *b* is kept static, though in that scenario we find a stronger over prediction of the merging timescales towards low-*z*, resulting in a more pronounced under prediction of the merger rate. Just as in Jiang et al. (2014)



Figure 6.7: Galaxy merger rates reconstructed from mock pairs, the fitting functions of eq. 6.6 and eq. 6.7. The leftmost panel shows the reconstructed rates under various apertures (coloured lines) against the intrinsic merger rates derived from trees in O'Leary et al. (2021a) (black lines). The right panels show how well our fitting functions (solid lines) reproduce the observation timescales (blue lines) and merging probabilities (red lines) seen in the data for the noted apertures (dashed lines). We also include the *effective* observation timescales (blue dash-dot lines) defined as $T_{\rm obs}^{\rm eff} \equiv T_{\rm obs}(z)/C_{\rm merge}(z)$.

we express merging timescales that scale linearly in Δr , however in their work they elect for a physically motivated redshift scaling with $b \propto H(z)^{-1/3}$. When measuring average merging timescales for fixed Δr_{max} we find that our results are statistically consistent with that scaling. In practice we found this formulation did not offer an improved fit to the data, or reproduction of merger rates. If we allow a more free scaling where $b \propto H(z)^{\alpha}$ there is a slight improvement over $H(z)^{-1/3}$. Although, allowing that free scaling in H(z)makes the physical interpretation of that term ambiguous, so we opt for the more simple linear redshift scaling, which produces the the better reproduction of merger rates.

In the next section we will combine these fits for merging probabilities and timescales and recover the merger rate of galaxies under a range of pair selection criteria.

6.6 Reconstructing merger rates

Now that we have a handle on the merger probability, and the merging timescales for individual pairs we can apply these to our mock observations and recover the intrinsic merger rates. If we blindly use T_{merge} in-place of T_{obs} the resulting merger rates we produce will be lower than those predicted in O'Leary et al. (2021a). First we need to approximate the observation timescale by estimating how long each of our pairs reside in the observable aperture. Taking inspiration from the orphan position formula (eq. 6.1) we



Figure 6.8: The ratio of predicted to intrinsic merger rates as a function of radius, with various assumptions for area correction. In both panels, lines labelled "fixed r_{inner} " indicate the rate ratio assuming fixed inner aperture $r_{\text{inner}} = 0$ kpc and a variable outer aperture r_{outer} . Similarly, lines labelled "fixed r_{outer} " indicate the rate ratio assuming fixed outer aperture $r_{\text{outer}} = 100$ kpc and a variable inner aperture r_{inner} . Line colour indicates the redshift at which the ratio is taken. Values near 1 indicate better agreement with intrinsic values. Top panel: Solid and dashed lines show the rate ratio with no area corrections applied. The dotted line shows the rate ratio with the standard inner area correction of eq. 6.9. Bottom panel: Dash-dot lines show prediction accuracy with the additionally correction factor shown in eq. 6.11 which addresses incompleteness for limited r_{outer} . The long-dash lines show prediction accuracy using a newly implemented inner area correction eq. 6.12.

get:

$$T_{\rm obs} = T_{\rm merge} \left[1 - \left(\frac{r_{\rm inner}}{\Delta r} \right)^2 \right] ,$$
 (6.8)

where r_{inner} is the inner radius of the observable aperture. Additionally, we find it helpful to include an area correction to account for pairs that sit below the chosen aperture. We use the same correction shown in Ventou et al. (2019):

$$C_1 = \frac{r_{\rm inner}^2}{r_{\rm outer}^2 - r_{\rm inner}^2} \tag{6.9}$$

Incorporating our weighting scheme (eq. 6.6), individualised observation timescales (eq. 6.8) and the area correction (eq. 6.9), the merger rate formula from eq. 6.5 can be rewritten as;

$$\Re_{\text{merge}} = C_1 f_p \frac{\sum_{i}^{N_p} W(\Delta r, \Delta v, z)}{\sum_{i}^{N_p} T_{\text{obs}}(\Delta r, \Delta v, z)}$$
(6.10)

Where f_p is the pair fraction and N_p is the number of pairs in the sample.

Figure 6.7 compares the galaxy merger rate derived from eq. 6.10 (coloured lines) with the merger rates shown in O'Leary et al. (2021a) (black line). In general the results from this are in excellent agreement with intrinsic rates. We find that we are consistently able to reproduce major merger rates to at least z = 4 for a range of projected separations. We find a notable exception in the often used $\Delta r = 5$: 30 kpc aperture (Mundy et al., 2017; Duncan et al., 2019), which under predicts the intrinsic merger rate at nearly every redshift. If we view the corresponding (upper right) panel of Figure 6.7 we can see that the fitting functions reproduce the observation timescales and merging probabilities of the data just as well as any of the other apertures we tested. Taken in the context of Figure 6.4 and absent any other effects that might influence the observation timescales for this selection criteria we can deduce that this aperture is ill-suited for reconstructing the underlying merger rate as the small outer radius excludes a significant fraction of pairs undergoing a merger. The under prediction exhibited by this particular selection criteria is present for all major merger rates that we tested.

In Figure 6.8 we show the accuracy of our predictions, exhibited as the ratio of predicted to intrinsic merger rate, as a function of observable aperture at three redshifts. For lines labelled "fixed r_{inner} " we set a constant $r_{\text{inner}} = 0$ kpc while increasing r_{outer} out to 100 kpc. Lines labelled "fixed r_{outer} " hold a fixed out radius at $r_{\text{outer}} = 100$ kpc with a variable inner radius r_{inner} . In the top panel, solid and dashed lines show the accuracy of our model absent any area corrections. For "fixed r_{inner} ", solid lines, we can see clearly if $r_{\text{outer}} \leq 50$ kpc our predictions would undercut the intrinsic merger rate. This coincides directly with the under prediction shown in the $\Delta r = 5$: 30 kpc range of Figure 6.7, and illustrates the necessity for an additional corrective factor. To address the pair incompleteness due to limited outer aperture we introduce the following correction:

$$C_2^{\text{logistic}} = \left[\frac{L}{1 + \exp(-kr_{\text{outer}})} - 1\right]^{-1} , \qquad (6.11)$$

where L and k are free parameters. This functions once again takes the form of a (half) logistic curve and is fit to the solid lines in the top panel of Figure 6.8. When fitting we assume z = 1.0, $\log(m/M_{\odot}) \ge 10.3$, and $1 \le \mu < 4$. Under these conditions, and a simple non-linear least squares fit, we find $L = 2.13 \pm 0.02$ and $k = 0.050 \pm 0.002 \text{ kpc}^{-1}$. The bottom panel illustrates the impact of this new correction, dash-dotted lines. Here we can see that this correction substantially improves the accuracy of predicted results, particularly for $r_{\text{outer}} \lesssim 50 \text{ kpc}$. We found these best fit parameters are suitable to correct major merger rates for $\log(m/M_{\odot}) \ge 9.5$ and $\log(m/M_{\odot}) \ge 11.0$, indicating minimal mass dependency.

Additionally, The top panel of Figure 6.8 shows that that eq. 6.9 provides a reasonable correction for non-optimal r_{outer} (dotted lines), but predictably deviates as $r_{\text{inner}} \rightarrow r_{\text{outer}}$. The corresponding uncorrected result, dashed lines, suggest an exponential correction maybe be more appropriate. Thus we introduce:

$$C_1^{\exp} = \alpha \exp\left(\frac{r_{\text{inner}}}{\beta}\right) ,$$
 (6.12)

where α and β are free parameters. Using the same fitting criteria noted above we find $\alpha = 0.889 \pm 0.003$ and $\beta = 106 \pm 1$ kpc. In the lower panel of Figure 6.8 (long-dash lines) we can see this updated correction factor improves the accuracy of predicted rates even when r_{inner} is similar to r_{outer} . Just as before we found this correction factor broadly applicable to major merger rates without the need to refit for higher or lower mass cuts. However, in practice this updated formulation does not provide significantly improved results compared to eq. 6.9 when applied to commonly used apertures.

Further for the lowest mass bin we fit, $9.0 \leq \log_{10}(m_1/M_{\odot}) < 10.0$, we found our fitting functions fail to accurately reproduce the underlying major merger rate at any redshift. Here we generally find a consistent over prediction in merger rates by a factor of ~ 3 to ~ 5 . We nonetheless include fits to this mass range for completeness.

When considering minor mergers we find that rates can be well reproduced out to $z \approx 3.5$ for all mass ranges tested. However, due to the poorer performance of the fitting function at high z and Δv , we find a significant under prediction of the merging timescales. This translates into an over prediction of the merger rate by a factor ~ 1.5 near $z \approx 3.5$ to a factor ~ 5 near $z \approx 5$. This mismatch is most pronounced where $\log_{10}(m_1/M_{\odot}) \geq 10.0$.

Figure 6.9 illustrates the derived merger rates from the mock light-cone catalogues described in §6.4.1. Here we can see that our formulae show excellent reproduction of underlying merger rates for a range of cone geometry, where each sample contains a unique set of pairs.

6.7 Discussion

6.7.1 Fitting at low and high redshift

Although we are able to confidently reproduce merger rates using our fitting functions there are some regions where care should be taken in how results can be interpreted. Notably,



Figure 6.9: The merger rates determined from mock light-cones. Rates are computed for $\log_{10}(m_1/M_{\odot}) \geq 10.3$, $5 \leq \Delta r < 50$ kpc projected separation, $\Delta v \leq 500$ km s⁻¹ and $1 \leq \mu < 4$. The solid black line indicates intrinsic merger rates as shown in O'Leary et al. (2021a). Coloured lines show merger rates determined from mock light-cones using eq 6.10. Here we do not place any redshift restraints on the light-cone catalogues that would more closely resemble the observables limits of the noted surveys.

the data indicates that for $z \leq 1$ the merger probability tends to zero. This trend makes sense recalling that under our definition of merger probability, that a pair observed at zero should have zero per cent chance to merge by zero. Looking back to the merging timescales, here we also note that the merging timescale turns over and begins to sharply decrease for $z \leq 1$. This turn over does not indicate a physical process that suppresses the merging timescales for low z. This data is constructed only from pairs that *did* merge by z = 0 so must necessarily decrease to reflect the decreasing time remaining for a pair to merge.

With this data we are left to decide whether we should fit to this low redshift regime given the bias towards pairs that merged quickly. In selecting fitting functions, we also tested variations that reproduced the trend to zero for both merging probabilities and timescales. Those fitting functions that worked for that low redshift region typically struggled more to fit the data, and reproduce merger rates at intermediate and high redshift. For this reason we opted for fitting functions that over shoot the data at low-z as they showed better performance when reconstructing merger rates.

6.7.2 Completeness

In §6.6 we illustrated how improper selection of observable aperture can result in large under predictions in merger rates due to missing pairs. Although corrections already exist to address incompleteness due to large r_{inner} , there are currently no widely used treatments to counteract incompleteness due to small r_{outer} . In this section we provided updated correction functions that successfully improve the accuracy of predicted results for a large range of observable area, and mass cuts.

The primary goal of this work is to place better constraints on T_{obs} and C_{merge} , not address other observable restrictions that might impact pair counts or their translation to rates. Therefore, for these completeness corrections we only tested their application to major mergers $1 \le \mu < 4$ for select mass ranges. Although we do not see evidence for any strong mass scaling in the function parameters, future work should explore the mass and mass ratio dependencies in greater detail. Additionally, when fitting these functions we elected to fit only at a single redshift, z = 1.0. The results shown in Figure 6.8 suggest either of these functions might have a weak redshift dependence. Future work with these functions would benefit from a fitting routine that considers the entire redshift range of analysis, or even an additional parameterisation for redshift dependency.

6.7.3 Sources of uncertainty

The results shown here stand as an extension to the results of O'Leary et al. (2021a), what we show here does not explore the complete dependencies of these formulations on our model assumptions. Thus, there remain several sources of uncertainty that could impact these results that we do not quantify in this work.

We have seen that merging timescales, and merging probabilities are strongly dependent on the relative line of sight velocities between each galaxy in the pair. At present EMERGE has no prescription for updating velocities of orphan galaxies. Currently, orphans simply inherit the velocity of their last resolved halo. Particularly at low masses, where orphans make up a large portion of the galaxy stellar mass function, this may alter the distribution of velocities in observed pairs, as well as the assumed fitted parameters to eq. 6.6 and eq. 6.7. For the example case we show in this work, major pairs consisting of at least 1 orphan galaxy constitute as much as ~ 30 per cent of all pairs at $z \approx 0$ falling to around ~ 5 per cent by $z \approx 4$. For the lowest mass bin shown in Table 6.3 and Table 6.4 the orphan pair fraction increases to around ~ 60 per cent by $z \approx 0$. This may be one of the sources for the poor reproduction of major merger rates in this mass range.

Additionally, merging in EMERGE is entirely defined by the dynamical friction formula chosen. In practice there are other aspects of the physical system that should be considered. In this model a satellite galaxy can be placed arbitrarily close to its host system but will only be merged at t_{df} . In practice these satellites may be affected by the radial extent of the host system resulting in a galaxy merging sooner than t_{df} . Such mechanisms could be a driver in the lower galaxy merger rates exhibited by our model compared with others. If merging timescales are artificially long, this would similarly extend the assumed observation timescales of satellite galaxies under our current formulation. Further, including mechanisms that reduce merging timescales may also impact the distribution of observed pairs, our chosen fitting functions may not be suitable under such model variations.

These results would benefit from a more complete study on how model assumptions drive pair fractions and merging timescales. Our fitting functions assume a narrow description of the observation timescale. In this work we assume the observation timescale is driven entirely by the time a galaxy pair fulfils the Δr criteria chosen, which linked directly to the average merging timescale of such a pair. This notably neglects the star formation properties in each pair. For $z \gtrsim 2$ high star formation rates result in a strong mass evolution in the galaxy population. This should reduce the amount of time any given pair can satisfying the mass and mass ratio criteria. In this work we largely reference pair selections set by a lower stellar mass cut, which mitigates the impact on observation timescales by pairs moving out of the noted mass bin. However it places no constraints on the impact from pairs that evolve outside of the mass ratio criteria. Under the broad assumption that these mechanisms would only serve to reduce the observation timescale, our results indicate that the observation timescale is instead largely driven by the time spent in the chosen aperture, as our fitting functions generally do not *under* predict the merger rate.

6.8 Conclusions

In this work we compare simulated pair fractions from EMERGE with observations. We use the simulated data two determine the observation timescale of merging pairs. Here we employ two methods for determining the evolution of observation timescales: first by directing mapping pair fractions to simulated merger rates, then by measuring the merging probability and timescales of simulated pairs directly. We show that our model produces galaxy pair fractions consistent with observations out to high redshift. Despite general agreement in the redshift scaling of the pair fraction, there remains considerable tension between observation and theoretical predictions. Discrepancies in methodologies make a direct comparison between models and even between observations difficult. Subsequently, predictions have not converged to a single functional form for pair fraction evolution. Our model can most reasonably be fit with a power-law exponential form, consistent with the observations from Jiang et al. (2014); Man et al. (2016); Mundy et al. (2017); Ventou et al. (2017); Mantha et al. (2018).Our results best match those of Ventou et al. (2017), who employ a redshift proximity criterion most similar to ours owing to their use of spectroscopic redshift information. Following the pair selection criteria of Ventou et al. (2017) we find a pair fraction that ranges between 2 per cent and 7 per cent.

Differences are further compounded when translating observed (simulated) pair fractions into galaxy-galaxy mergers rates due to the necessity of a well defined observation timescale T_{obs} . When using published values for T_{obs} (Lotz et al., 2011), we find a merger rate that over predicts our model intrinsic results by more than a factor of 2. Further, we find that utilising an observation timescale that scales $\propto (1 + z)^{-2}$ (Snyder et al., 2017) over-predicts our predicted merger rates by nearly an order of magnitude. Converting our simulated pair fractions to merger rates is most consistent with a linearly evolving observation time scale $T_{obs} = w(1 + z) + b$. The results presented here are a first pass at confining the observation timescales through our model. However, we do not expect future work within this model to produce the strong scaling seen by Snyder et al. (2017). Additionally, more complete inspection of the pair fraction sensitivity to observables and model variations is necessary. Other recent works have shown the pair fraction sensitivity to the SHMR (Grylls et al., 2020), understanding how changes in these statistical relations impacts observed pair fractions is vital to understanding galaxy clustering and merging timescales at high redshift.

In second part of this work we show that with our empirical model EMERGE model we are able to construct fitting functions for galaxy merging probabilities and timescales independently that together can accurately recover the intrinsic merger rate across range of selection criteria. To that end we provide best fit model parameters for a wide range of commonly used stellar mass intervals, and mass ratios.

For a given aperture, the data indicates $T_{\rm obs}$ is approximately linear between $z \sim 1$ and $z \sim 4$, but $T_{\rm obs}^{\rm eff}$ can range between a linear scaling and power-law scaling depending on the aperture chosen. However, we reinforce a $T_{\rm obs}$ that does not evolve as strongly as Snyder et al. (2017) or Jiang et al. (2014). We can conclude from our results that pairs undergoing a major merger have a $T_{\rm obs}$ that is primarily driven by dynamical processes to at least ~ 3.5 .

We further show that it is not necessary to fit the observation timescales directly, but instead fit to merging timescales and derive the selection criteria dependent observation timescales using a formula that approximates radial decay in halos. Additionally, we can show with these methods that not all pair selection criteria are equally suited to determining merger rates. If the outer radius of the observable aperture is too small, a non-negligible fraction of merging systems can be missed, resulting in an under prediction of the galaxy merger rate. To combat this outer area incompleteness we suggest a new correction factor that should be applied to observations with sub-optimal apertures.

Finally, in the last section we discussed areas where these results could be improved with further study. At high- $z \ (\geq 4)$ box size limitations prevent us from placing tighter constraints on merger timescale evolution of massive objects due to low number counts. At the lowest masses our fitting functions are inadequate to reliably reproduce the underlying galaxy merger rate, at these masses and at high redshift $T_{\rm obs}$ may be dominated by high start formation rates, which reduce the time a pair spends in the mass and mass ratio bin.

Chapter 7

Here be dragons: exploring the stellar-to-halo mass relation in the dwarf regime

7.1 Introduction

Galaxy formation in the ACDM framework predicts that dwarf galaxies are expected to be the most abundant galaxies in the universe, however their low luminosities make them particularly difficult to observe in practice. Meanwhile their sensitivity to feedback processes makes them difficult to model. Additionally, their shallow gravitational potential wells make them not only sensitive to internal feedback processes but also to assumptions on cosmology. Together this makes dwarf galaxies one of the best test-beds for the our understanding of both cosmology and the fundamentals of galaxy formation.

Recent advancements in observational techniques have improved both the quantity and quality of dwarf galaxy observations. In particular these observations have probed to lower magnitudes offering data completeness to lower masses than previously available. Further followup measurements have advanced accuracy in measuring the distance (Putman et al., 2021), mass (Woo et al., 2008) and star formation histories (Weisz et al., 2014) of these systems. These advancements present the possibility to better compare observations with high resolution theoretical models and open their use as direct constraining data numerical models.

Most recent theoretical models have focused on utilising high resolution hydrodynamical zoom-in simulations to explore dwarf galaxies (Sawala et al., 2015, 2016b; Garrison-Kimmel et al., 2019; Fattahi et al., 2020; Applebaum et al., 2021; Munshi et al., 2021; Wang et al., 2021). These approaches simulate large cosmological volumes then re-simulate the Milky Way like haloes at higher resolution. The advantage of zoom simulations in studies of dwarves is that it provides a large scale cosmological context while providing the resolution necessary to probe low mass satellites in systems similar to the Milk Way, where we have the most observational data. The draw back to this approach is that the SHMR can

only be addressed within the context of Milky Way like haloes and generally do not explore how assumptions made impact the global SHMR. Furthermore, despite the improved resolution in hydrodynamical zooms, these models are still restricted by uncertain subgrid implementations that could impact the resulting dwarf population.

Empirical models offer several distinct advantages over hydrodynamical and semianalytic models. Because these models operate by relating galaxies to the host halo in post processing, computational power can instead be devoted to increasing mass resolution without the need to sacrifice statics by simulating small volumes. And unlike either hydrodynamical simulations or semi-analytic models, empirical models make fewer assumptions on the relevant subgrid processes that impact galaxy formation. In this way the relationship between galaxies and haloes can be explored without being polluted by the personal priors imposed by subgrid models. The big caveat to empirical models is their need for observational data to constrain the model.

The lack of observational data has made empirical approaches to this problem difficult. There have been several recent attempts to quantify the SHMR using empirical techniques. Nadler et al. (2019) employed abundance matching on zoom-in simulations tuned to hydro-dynamical simulations to make predictions on the abundance of low mass satellites down to $\log_{10}(m_*/M_{\odot})2$. Other recent works (Wang et al., 2021) have taken an exploratory approach by extrapolating the UNIVERSEMACHINE model (Behroozi et al., 2019) into the ultra-faint regime. This is a useful technique to determine where the model must be improved to reproduce the observed characteristics of observed dwarves. We expand on these approaches by using existing observations to directly constrain our own empirical model EMERGE .

The goal of this work is to utilise real observables to constrain an empirical model that self consistently relates galaxy properties to dark matter halo properties at dwarf scales. Our aim is to better understand how low mass galaxies populate their haloes, and by doing so gain a better understanding of their star formation histories. This chapter is organised as follows. First, in §7.2.1 we will introduce the *N*-body simulations we use. We discuss updates to the EMERGE model implementation in §7.2.2. Additionally, in §7.2.3 we discuss the observational data used to constrain our model. In §7.3.1 we introduce the model variations we explore in order to reproduce the observed properties of dwarf galaxies in the local Universe. §7.4 compares our model implementations with one another and we discuss the resulting stellar to halo mass relation and star formation histories from our preferred model. Finally, in §7.5 we discuss how our model assumptions might impact our results, in this context we discuss opportunities opportunities for future work and model improvements.

7.2 Methods, Observations, Simulations etc.

7.2.1 N-body simulations

We utilise a cosmological dark matter only N-body simulation in a periodic box with side lengths of 60 Mpc. This simulation adopts Λ CDM cosmology consistent with (Planck Collaboration, 2016) results where $\Omega_m = 0.3080$, $\Omega_{\Lambda} = 0.6920$, $\Omega_b = 0.0484$, where $H_0 = 67.81 \,\mathrm{km \, s^{-1} Mpc^{-1}}$, $n_s = 0.9677$, and $\sigma_8 = 0.8149$. The initial conditions for this simulation were generated using MUSIC (Hahn & Abel, 2011) with a power spectrum obtained from CAMB (Lewis et al., 2000). The simulation contains 2048³ dark matter particles with particle mass $9.88 \times 10^5 \mathrm{M}_{\odot}$. The simulation was run from z = 124 to 0 using the Tree-PM code OPENGADGET3 (Springel, 2005). In total 147 snapshots were created. Dark matter haloes are identified in each simulation snapshot using the phase space halo finder, ROCKSTAR (Behroozi et al., 2013a). Halo merger trees are constructed using CONSISTENTTREES (Behroozi et al., 2013c), providing detailed evolution of physical halo properties across time steps.¹

7.2.2 Updates to EMERGE

For the work shown in this chapter we rely on recent improvements to the baryon conversion efficiency model in EMERGE. Other recent works exploring the galaxy-halo connection with neural networks (Moster et al., 2020a) have shown that a linear-max scaling for ϵ_N can provide an improved fit to observed stellar mass functions, especially at intermediate redshift. We have incorporated this proposed change into the version of EMERGE used in this work.

$$\epsilon_N = \begin{cases} \epsilon_0 + \epsilon_z \frac{z}{z+1} & \text{if } \epsilon_N \le \epsilon_{\text{peak}} \\ \epsilon_{\text{peak}} & \text{otherwise} \end{cases}$$
(7.1)

Here ϵ_{peak} is a maximum allowed value for ϵ_N . This new parameter is free in the model and has been fit using the methods and data described in §4.1.1. The remainder of the model operates as described in §4.1.1.

7.2.3 Observational data

The empirical model is directly constrained by observational data. In addition to the data described in Moster et al. (2018) and O'Leary et al. (2021a) we extend the galaxy stellar mass function (SMF) data down to $\log_{10}(m_*/M_{\odot}) = 5$ through the inclusion of Local Group dwarf galaxies. In this work all galaxies within 2 Mpc are defined as belonging the the Local Group (Putman et al., 2021). We assign a galaxy as a satellite if it is positioned within 300 kpc of either the Milky Way or Andromeda. We construct the SMF using the positively identified dwarf galaxies listed in online database of McConnachie (2012). Where

¹Additional information on these simulations, including: configuration files, build info, and parameter files can be found at https://github.com/jaoleary.

Galaxy Name	$\log_{10}(m^*/M_{\odot})$	D_{\odot} [kpc]	$D_{\rm MW}$ [kpc]	$D_{\rm M31}[$ kpc $]$	τ_{90} [Gyr]	Reference
Sagittariusd-Sph	7.526	27	19	787	$10.3^{+0.33}_{-1.82}$	Weisz et al. (2014)
LMC	9.380	51	50	807	-	-
SMC	8.867	64	61	807	-	-
UrsaMinor	5.748	76	78	754	$4.63^{+3.27}_{-1.60}$	Weisz et al. (2014)
Draco	5.681	82	82	748	$3.55^{+2.50}_{-1.52}$	Weisz et al. (2014)
Sculptor	6.459	86	86	761	$3.09^{+3.53}_{-1.29}$	Weisz et al. (2014)
Sextans(1)	5.709	95	98	841	-	-
Carina	5.920	106	108	838	$11.46^{+0.07}_{-1.49}$	Weisz et al. (2014)
Crater 2	5.408	118	116	886	-	-
Antlia 2	5.748	132	133	889	-	-
Fornax	7.483	139	141	768	$11.46^{+0.20}_{-0.27}$	Weisz et al. (2014)
Canes Venatici(1)	5.635	211	211	856	$5.38^{+2.01}_{-1.13}$	Weisz et al. (2014)
Leo 2	6.030	233	236	897	$7.29^{+0.63}_{-0.75}$	Weisz et al. (2014)
Leo 1	6.848	254	258	918	$12.02^{+0.06}_{-0.20}$	Weisz et al. (2014)
Leo T	5.350	409	414	982	$12.12_{-0.06}^{+0.12}$	Weisz et al. (2014)
Phoenix	6.091	415	415	864	$10.56^{+0.63}_{-1.03}$	Gallart et al. (2015)
NGC6822	8.204	459	452	894	- 1.05	-
Andromeda XVI	5.736	476	480	319	$7.88^{+0.56}_{-0.49}$	Skillman et al. (2017)
Andromeda XXIV	5.173	600	605	204	-0.45	-
NGC185	8.037	617	621	184	-	-
Andromeda XV	5.885	625	630	175	$4.24^{+0.87}_{-2.12}$	Skillman et al. (2017)
Andromeda II	7.163	652	656	181	$7.39^{+0.60}_{-0.51}$	Skillman et al. (2017)
Andromeda XXVIII	5 526	661	661	365	$6.13^{+0.28}$	Skillman et al. (2017)
Andromeda X	5 149	670	674	130	-	
NGC147	7 997	676	680	139	_	-
Andromeda XXX	5 318	682	686	144	_	-
Andromeda XVII	5 547	728	732	66	_	-
Andromeda XXIX	5 459	731	734	187	_	-
Andromeda I	6.876	745	749	55	$6.20^{+0.67}$	Skillman et al. (2017)
Andromeda III	6 204	748	752	73	$4.93^{+0.67}$	Skillman et al. (2017)
IC1613	8 204	755	758	518	$4.35_{-1.47}$ 11 46 ^{+0.63}	Callart et al. (2017)
Cotue	6.651	755	756	678	$_{4 2^{\pm 0.63}}^{11.40_{-1.03}}$	Callart et al. (2015)
Andromoda XXXI	6.817	750	760	262	4.0-1.03	Gallart et al. (2010)
Andromeda VII	7 408	769	765	202	-	-
Andromeda IV	5 280	766	705	217	-	-
LCS2	6 186	760	772	39	7 00+0.63	Collect et al. (2015)
LG55 Andromodo VVIII	0.100	709	774	200	1.02-1.03	Gallart et al. (2013)
Andromeda XXVIII	0.240	709	774	249	-	-
Andromeda V	0.200	779	779	348 100	-	-
Andromeda v	0.902	115	700	109	-	-
Andromeda AAAII	7.037	770	780	140	-	-
Andromeda VI	6.723	783	785	268	-	-
Andromeda AIV	5.584	794	798	161	-	-
IC10	8.139	794	798	252	-	
Leo A	6.982	798	803	1197	$12.55_{-1.03}$	Gallart et al. (2015)
M32	8.709	805	809	27	-	-
Andromeda XXV	6.037	813	817	90	-	-
Andromeda XIX	5.723	820	824	115	-	-
NGC205	8.723	824	828	46	-	-
Andromeda XXI	6.049	828	831	135	-	-
Andromeda XXVII	5.283	828	832	77	-	-
Tucana	5.952	887	883	1352	$4.11_{-1.03}^{+0.05}$	Gallart et al. (2015)
Pegasus-dIrr	7.024	920	921	474	-	-
WLM	7.838	933	933	835	-	-
Sagittarius-dIrr	6.748	1067	1059	1354	-	-
Aquarius	6.408	1072	1066	1170	$11.1^{+0.02}_{-0.04}$	Cole et al. (2014)
Andromeda XVIII	5.903	1213	1217	457	-	-
Antlia B	6.003	1294	1296	1963	-	-
NGC3109	8.085	1300	1301	1984	-	-
Antlia	6.318	1349	1350	2036	-	-
UGC4879	7.123	1361	1367	1394	-	-
Sextans B	7.920	1426	1429	1940	-	-
Sextans A	7.848	1432	1435	2024	-	-
Leo P	5.795	1622	1625	2048	-	-
KKR25	6.505	1923	1922	1869	-	-
ESO410-G005	6.748	1923	1922	1861	-	-
IC5152	8.350	1950	1945	2209	-	-

Table 7.1: Catalogue of Local Group dwarf galaxies and their properties used in this work. The reference column indicates the source for the τ_{90} measurements.

available we use galaxy stellar masses from Woo et al. (2008) and a 1.6 mass-to-light ratio otherwise (Bell & de Jong, 2001; Martin et al., 2008). We assume 0.8 dex uncertainty in $\log_{10}(m_*/M_{\odot})$ for each system (Woo et al., 2008).

The dwarf mass function is constructed using 7 bins evenly spaced in $\log_{10}(m_*/M_{\odot})$. We then create 10000 random realisations of the locally observed dwarf population by sampling within the mass uncertainty range, each instance is then sampled again assuming a Poisson error. From these random realisations we compute the average and 1 σ interval mass function. Finally, we renormalise the dwarf mass function such that the average value matches the observed average SMF at $\log_{10}(m_*/M_{\odot}) = 7$. This requires a vertical adjustment of -0.83 dex to the locally observed SMF.

Before moving on we should take a moment to consider the implications of relying on Local Group observations to characterise the dwarf galaxies in large volumes. The first clear draw back to this approach is due to the limited sample size. When probing down to $\log_{10}(m_*/M_{\odot}) = 5$ observations are only complete out to ~ 2 Mpc, this practical limitation restricts our sample of galaxies to 64 confirmed objects. This low number of systems results in large error bars on the resulting SMF, which may reduce our ability to make a statistical distinction on the goodness of fit in our proposed model implementations. Nonetheless, we can still make arguments for or against each model variation based on the resulting systems and whether their properties make sense with our current understanding of galaxy growth at low masses. Additionally, an argument could be made that locally observed trends are not representative of the universe at large; however this data is cutting edge and probes the limit of our observational capabilities (see §7.5).

7.3 High redshift quenching

In the discussion of dwarf galaxies there are two particular observational facts we need to contend with. The first is the abundance of dwarf galaxies. Following the paradigm we discussed in §2 we would expect that low mass haloes should appear in the greatest abundances, and that haloes with $M_{\rm h} \gtrsim 10^5 {\rm M}_{\odot}$ might contain a galaxy. If every low mass halo hosted a bright galaxy we would expect to observe many thousands of dwarf galaxies in the local volume. This prediction conflicts with the relatively low number of dwarf galaxies that have actually been observed and catalogued (McConnachie, 2012). The discrepancy between number of observed dwarves, and the number predicted by ACDM is what is known as the missing satellite problem (e.g. Klypin et al., 1999; Moore et al., 1999).

The second complication is in the star formation history of these galaxies. Investigation of dwarves in the Local Group reveals that these galaxies possess low star formation rates and tend to be much older. At the lowest masses observations indicate that for most stars formed by z = 6 (Weisz et al., 2014, 2015). This leaves us with the question: why are dwarf galaxies so old?

There are several possible solutions to these problems. One possible solution that can address both is that star formation efficiency is suppressed in low mass haloes at late times due to an additional feedback mechanism that we have not discussed, reionization. In §1.1.3 we discussed the concept of recombination, when the universe cooled sufficiently to form neutral hydrogen. The first stars and galaxies therefore formed in a largely neutral intergalactic medium (IGM). Observations of Lyman- α absorption features provide one of the most vital probes of neutral hydrogen in the IGM. The Lyman- α series refers the range of absorption/emission lines from/to the ground level of neutral hydrogen and can be characterised by:

$$\lambda_{\alpha} = \frac{hc}{13.6\text{eV}} \left(1 - \frac{1}{n^2}\right)^{-1} \tag{7.2}$$

where n is a natural number ≥ 2 corresponding to electron energy levels. From this we can see that the Lyman- α series resides in the ultra-violet spectrum. Combined with the large absorption cross section, Lyman- α lines can be used to probe for neutral hydrogen in the IGM. It stands to reason that if there is an emitter of UV photons at high-z and if there is neutral hydrogen along the line of sight to that emitter, then there should be prominent absorption features in the observed spectra corresponding to the redshift of the neutral hydrogen cloud. These absorption features become manifest in the spectra of high-z quasars and galaxies in the form of the Lyman- α forest (Lynds, 1971). That these lines exist at all is an indicator that the IGM is largely ionized. Leveraging the strong absorption of UV photons by neutral hydrogen Gunn & Peterson (1965) predicted that if the Universe was in-fact filled with neutral hydrogen then the flux from a UV source would be highly suppressed below λ_{α} . Becker et al. (2001) validated this prediction with the discovery of a quasar at z = 6.28 which featured a crushed Lyman- α forest (known as the Gunn-Peterson Trough) indicating a substantial fraction of neutral hydrogen in the IGM for $z \gtrsim 6$.

With the understanding that reionization is capable of quenching star formation through both the heating of otherwise cold gas and through the reduction in cooling rates by reducing the neutral fraction. Early work investigating the impact of reionization on the abundance of dwarves showed that including this feedback process in models for galaxy formation could reconcile the discrepancy of dwarf abundances between Λ CDM predictions and observation (Thoul & Weinberg, 1996; Bullock et al., 2000; Somerville, 2002). These models generally predicted that reionization inhibits star formation in haloes with a maximum circular velocity $V_{\text{peak}} \leq 20 \text{ km s}^{-1}$ ($M_{\rm h} \leq 10^9 M_{\odot}$).

In this work we build on these techniques by incorporating a model for high-z quenching into EMERGE. This model option will complement the already existing mechanisms that impact galaxies in low mass haloes such as environmental quenching and tidal disruption (see §4.1.1).

7.3.1 Model variations

In this section we introduce three empirically motivated models for a high-z quenching process that impacts low mass haloes. Throughout this work we compare our results to the 'reference' model which is our unaltered EMERGE model with the parameters shown in

table 7.2. The reference case we show uses the same 60 Mpc box as described in §7.2.1. However, the parameters used were fit on the 200 Mpc box described in §5.1. The Local Group dwarf data was not incorporated into those fits.

In addition to the standard EMERGE model we introduce three model variations that suppress star formation in low mass-high redshift haloes in a manner consistent with expectations from reionization. the shared characteristic of these methods is that they specify a minimum halo mass $M_{\rm h}^{\rm min}$ required to form stars at some scale factor a = 1/(1+z). When a halo does not meet that threshold its star formation will be set to zero and will remain zero for the remainder of that galaxy's lifetime. If a halo first in appears in the simulation (leaf halo) below the specified threshold, no galaxy will be seeded and the halo will remain dark (see §7.5.4).

The first, most simple, model we test treats reionization as an instantaneous process eq. 7.3. Here star formation is shut off in haloes of insufficient mass M_q by the reionization scale factor a_q . We linearly interpolate halo masses between snapshots to avoid imprinting a preferred reionization scale due to time-step discreteness.

$$M_{\rm h}^{\rm min}(a) = \begin{cases} M_{\rm q} & \text{if } a > a_{\rm q} \\ 0 & \text{otherwise} \end{cases}$$
(7.3)

Both of these parameters are free in the model. This model is referred to as *instantaneous* for the remainder of this chapter.

Our next model is a linear-max construction eq. 7.4. This describes a process where the minimum halo mass needed to form stars increases linearly with increasing scale factor up to some maximum scale factor, after which the minimum mass remains constant.

$$M_{\rm h}^{\rm min}(a) = \begin{cases} R_{\rm q}(a) + M_{\rm q} & \text{if } a \le a_{\rm q} \\ R_{\rm q}(a_{\rm q}) + M_{\rm q} & \text{otherwise} \end{cases}$$
(7.4)

Here R_q indicates the rate at which the mass threshold increases with scale factor and a_q indicates the scale factor where the threshold reaches its maximum. This model is referred to as *lin-max* hereafter.

Finally, we test a logistic model which allows for continuously increasing reionization threshold up to some max eq. 7.5.

$$M_{\rm h}^{\rm min}(a) = \frac{M_{\rm q}}{1 + \exp\left[-R_{\rm q}(a - a_{\rm q})\right]},\tag{7.5}$$

here M_q is the maximum threshold mass, R_q is the transition strength, a_q is the midpoint scale factor where the rate reaches its maximum. We implement eq. 7.5 in log-space such that the minimum mass threshold is $10^0 M_{\odot}$. Our simulation is of insufficient resolution to probe galaxy formation on those scales so this floor is somewhat artificial. If the floor value is changed we could anticipate possible changes in R_q , a_q or both. While this formulation contributes the same number of free parameters as lin-max, the smooth continuous transition delivers a more physically interpretable view of the reionization process, while also offering greater flexibility to incorporate additional parameters if needed. This model will be referred as *logistic* for the remainder this work.

7.3.2 Fitting

When performing an N-body simulation there is typically a trade off between the size of the simulation volume and the particle resolution. Our simulation parameters were chosen to maximise the number of Milky Way like systems while simultaneously providing the resolution necessary to capture the stellar mass to halo mass relation and the associated scatter down to $\log_{10}(m_*/M_{\odot}) = 5$. This decision, while necessary for this study, limits which model parameters can explored in the fitting process largely due to the limited number of more massive systems. Subsequently we only allow the low mass slope of the baryon conversion efficiency β_0 , its redshift evolution β_z and the stellar mass dependent quenching timescale τ_s as free parameters in addition to those introduced by our new model variations. All remaining free parameters of the model are fixed to values indicated in table 7.2. For these model fits we utilised binned observational data as opposed to the raw observations shown in §A. By fitting to binned data we minimise chance of our MCMC (see §4.1.3) getting stuck in a local minima due to conflicting observations. Finally, we apply a prior to the binned data. This prior applies increased weight to data at low-z and near strong inflection points in order to preserve observed trends where the data is most robustly measured².

7.4 Results

7.4.1 Fits and model selection

Here we discuss how each model performs in reproducing the range of the observed data, as well as evaluate some evidence based model selection criteria. Table 7.2 shows the best fit parameters for each model. Empty fields indicate that the parameter is either free or not relevant to the model, otherwise the parameters for the reference model have been used. The most notable change to the model parameters is the low mass baryon conversion efficiency slope β_0 . With the extended SMF all models tend to prefer a steeper conversion slope indicating less effective star formation compared to the reference. All three of our model variations converge to $M_q \approx 9.3 - 9.4$. With the quenching mass saturating between $z \approx 3.4 - 3.75$. The lin-max and logistic model options, which allow for a time dependent M_h^{\min} , do not agree on the rate of increase. These models nonetheless show nearly identical reproduction of the observational data used to constrain EMERGE as well as in the resulting SHMR and star formation histories of dwarves (see§7.4.2 and §7.4.4 respectively). This likely means there is insufficient data to constrain the time evolution of M_h^{\min} and that it is not required to explore galaxies with $\log_{10}(m_*/M_{\odot}) > 5$.

Figure 7.2 shows how each of our models reproduces the z = 0 mass function (coloured lines), compared with the reference (orange line). Each of our model variations successfully reproduces relevant observables, and qualitatively no distinction can be visually identified

²The data used for fitting along with associated weights can be obtained at https://bitbucket.org/ bmoster/emerge.



Figure 7.1: The evolution of $M_{\rm h}^{\rm min}(a)$ under our three model variations. haloes with mass below $M_{\rm h}^{\rm min}(a)$ (coloured lines) have star formation instantaneously quenched. The grey region blocks our sub-resolution halo masses.

Parameter	Reference	Instantaneous	Lin-max	Logistic
M_0	$11.3228^{+0.0311}_{-0.0183}$	-	-	-
M_z	$1.4451_{-0.0706}^{+0.0572}$	-	-	-
ϵ_0	$0.0204^{+0.0096}_{-0.0080}$	-	-	-
ϵ_z	$1.6972^{+0.0097}_{-0.0153}$	-	-	-
ϵ_{peak}	$0.3039^{+0.0173}_{-0.0139}$	-	-	-
β_0	$2.9833^{+0.2030}_{-0.2156}$	$2.3701_{-0.0221}^{+0.0342}$	$1.8849^{+0.0355}_{-0.0162}$	$2.2240^{+0.0872}_{-0.0469}$
β_z	$-2.6788^{+0.2629}_{-0.2225}$	$-1.6983^{+0.0305}_{-0.0456}$	$-1.1766^{+0.0211}_{-0.0451}$	$-1.4993^{+0.0431}_{-0.1069}$
γ_0	$1.2491^{+0.0142}_{-0.0140}$	-	-	-
f_{esc}	$0.5866^{+0.0199}_{-0.0172}$	-	-	-
f_s	$0.0044^{+0.0009}_{-0.0006}$	-	-	-
$ au_0$	$0.8118^{+0.1028}_{-0.1821}$	-	-	-
$ au_s$	$0.5242_{-0.0452}^{+0.0574}$	$0.3954_{-0.0066}^{+0.0103}$	$0.4150^{+0.0041}_{-0.0038}$	$0.3954_{-0.0085}^{+0.0193}$
$a_{\mathbf{q}}$	-	$0.2317^{+0.0256}_{-0.0177}$	$0.2380^{+0.0292}_{-0.0290}$	$0.1938^{+0.0302}_{-0.0279}$
$M_{\mathbf{q}}$	-	$9.4027^{+0.0890}_{-0.0715}$	$9.3400^{+0.0538}_{-0.1870}$	$9.3271_{-0.0779}^{+0.1089}$
R_{q}	-	-	$1.1440^{+0.0563}_{-0.0549}$	$2.5626^{+0.3810}_{-0.5155}$

Table 7.2: The best fit model parameters for each model variant used in this work. empty fields indicate the parameter was not left free when fitting that model or was not available in that model. Noted confidence intervals correspond to the $\pm 1\sigma$ range.

Table 7.3: Model statistics

Model	$\chi^2_{ m min}$	$\chi^2_{ m mean}$	$N_{\rm p}$	$p_{\rm D}$	AIC	BIC	DIC	$-2\ln(Z)$
Instantaneous	822.38	826.59	5	4.20	832.58	851.160	830.79	858.36
Lin-max	781.83	788.06	6	6.23	794.10	816.37	794.29	828.48
Logistic	828.51	833.13	6	4.62	840.78	863.05	837.76	860.11

with respect to the SMF. The lower star formation efficiency in the new models does however result in a lower density in the $7 \leq \log_{10}(m_*/M_{\odot}) < 8$ range, but all remain consistent with observed data.

A more quantitative comparison can be performed using an information criterion. This way we can weight the goodness of fit for each model against the added complexity of additional free parameters. Higher order models can fit the data better, but at some point there will not be enough information to further constrain additional parameters and increasing the order will not provide a better fit to the data. In these schemes models are penalised as additional parameters are added, providing a pathway to selecting the most simple model that can reproduce the data. Table 7.3 list the statistical characteristics of each model. For each model we compute statics using 80 walkers over 335 epochs totalling 26800 permutations per model³.

³This is a relatively low number of data points that may not provide sufficient coverage of the parameter space. This is a result of the computational limitations discussed in §7.5.4



Figure 7.2: A comparison of the global galaxy stellar mass function for our model variations against the observed average stellar mass function. Solid lines illustrate the resulting z = 0 stellar mass function using best fit parameters for each model variant. The grey line and shaded region indicate the observed SMF average from a suite of observational estimates along with the 68 per cent confidence interval. The observed data to the left of the vertical black line illustrates the region where we have extended the observed SMF using the Local Group dwarves listed in 7.1

From a statistical point of view the model lin-max provides the best reproduction of observed data. Beyond these statistical measures we find very little quantitative differences between each model. For the remainder of this work we will make all of our comparisons with respect to the logistic model. While the logistic model does not provide any additional predictive power with respect to the lin-max, the logistic model provides the greatest flexibility for incorporating new observables along with the possibility to set a halo mass floor for galaxy formation. Rather than selecting a model which will need to be fundamentally overhauled given new data, it may be preferable to select the model which readily accepts additional parameters to increase its complexity when needed. Ultimately, we found that our conclusions are unaffected by which high-z quenching model variation we chose.

7.4.2 The stellar-to-halo mass relation

The primary purpose of this work is not to explore the physics of reionization, but to better understand how observed dwarves came to be, and how these low mass objects fit into the Λ CDM paradigm. In particular we are interested in the relationship between dwarf galaxies and their dark matter haloes and the associated scatter in that relationship.

Here we discuss how the SHMR from our preferred high-z the quenching model compares with the reference case. Figure 7.3 shows as side by side comparison of these two models in terms of halo peak mass $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}})$ (blue contours) and present day halo mass $\log_{10}(M_{\rm h}^{z=0}/{\rm M_{\odot}})$ (orange contours). As mentioned in §4.1 galaxy stellar mass at any epoch is closely related to the peak halo mass of that galaxy, and the SHMR can be well fit by a double power-law. We can see from the black solid lines in Figure 7.3 that the average SHMR can be well approximated by a power law relation down to at least $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}}) \approx 10$. The results of these model are consistent with an extrapolation of the power-law relation shown in Moster et al. (2018) (solid red lines).

How this relationship evolves for $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}}) \lesssim 10$ is difficult to ascertain. The reference model produces a relation that extends as a power-law to lower masses. In this case the relationship is primarily determined by the low mass slope of the baryon conversion efficiency. Here, altering the low mass conversion slope can change the slope but the SHMR remains a power-law, by construction. Introducing high-z quenching breaks this relationship by preventing some haloes from forming galaxies in a way that is proportional to the halo growth. This can be clearly seen in the blue contours of the high-z quenching model (right panel) of Figure 7.3. Here we can see a rapid cut off in the SHMR as we approach $M_{\rm q}$. Comparing the blue contours of the reference (left panel) and high-z quenching models we can see that although the models share a similar average SHMR, they diverge strongly in the scatter of $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M}_{\odot})$ at fixed $\log_{10}(m_*/{\rm M}_{\odot})$. We can conclude that the suppression of the SMF at low masses is primarily due to the elimination of late forming low mass ratio systems, $m_*/M_{\rm h}^{\rm peak}$, due to a high redshift quenching process.

Previous work has shown that the scatter in the SHMR takes a log-normal distribution at a fixed halo mass (Cooray, 2006), similarly we find that our results exhibit this same trend down to $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}}) \approx 10$. Due to the lack of observable constraints for $\log_{10}(m_*/{\rm M_{\odot}}) < 5$ we are unable to verify that this trend extends to lower halo masses.



Figure 7.3: The stellar-to-halo mass relation under the reference model compared with the logistic reionization model. The solid lines indicate the best fit SHMR assuming a double-power law relation. The black line shows the average relation for our logistic model and the red line shows the relation from the (Moster et al., 2018). The blue and orange contours show the iso-proportion contours of the SHMR for our logistic and reference models respectively. The left panel represents the SHMR in terms of halo peak mass $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}})$ while the right panel shows the same data expressed in terms of the present day halo mass $\log_{10}(M_{\rm h}^{z=0}/{\rm M_{\odot}})$.



Figure 7.4: The distribution of halo masses at fixed stellar mass. Solid lines correspond to the logistic model and dashed lines are for the reference model. Colors indicate the bass bin for each distribution. The left panel shows distributions in peak halo mass $\log_{10}(M_{\rm h}^{\rm peak}/M_{\odot})$ and the right panel shows distributions in present day halo mass $\log_{10}(M_{\rm h}^{z=0}/M_{\odot})$.

Instead we can compare the scatter in $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M}_{\odot})$ at fixed stellar mass to better evaluate out how high-z quenching impacts scatter compared the the reference treatment.

Figure 7.4 illustrates the scatter in halo mass at fixed stellar mass for $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}})$ (left panel) and $\log_{10}(M_{\rm h}^{z=0}/{\rm M_{\odot}})$ (right panel). We can see that down to $\log_{10}(m_*/{\rm M_{\odot}}) \approx 7$ the distribution in $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}})$ remains approximately Gaussian with slight tail extension toward lower masses. For $\log_{10}(m_*/{\rm M_{\odot}}) \lesssim 7$ we find that the standard implementation (dashed lines) produces a highly asymmetric distribution with a pronounced tail toward low halo masses. Conversely, introducing high-z quenching (solid lines) not only reduces the range in scatter but results in something *more* Gaussian. Looking at the same distributions with respect to present day halo mass (right panel) we can see that in the reference model low-mass galaxies occupy a much larger range of halo mass than in the high-z quenching model, this can also clearly be seen in the orange contours of figure 7.3.

Table 7.4 compares the typical halo mass for a fixed stellar mass interval for each model, measured at the distribution peak along with 68 per cent interval. In general we find that for the stellar mass ranges evaluated, dwarves in the high-z quenching model tend to reside in more massive haloes on average compared to the reference model. The difference becomes more pronounced at lower stellar mass, peaking with 0.15 dex difference for the lowest mass range. Additionally, when evaluating the SHMR using present day halo masses we should recall that our formulation for orphan mass loss §4.1.2 may play a roll in artificially broadening the distribution, if that formulation strips too aggressively. We find that in our reference model orphans contribute between ~ 13 and ~ 30 per cent of the dwarf population $f_{\rm orph}$, while the high-z quenching model tends to produce a lower orphan fraction ranging between ~ 10 and ~ 20 per cent.

Some recent works (Nadler et al., 2019, 2020; Wang et al., 2021) suggest that the singular power-law relation can be extrapolate into the ultra-faint dwarf regime for satellites down to $\log_{10}(m_*/M_{\odot}) \approx 2$. However these models have only been constrained down to $\log_{10}(m_*/M_{\odot}) \approx 8$ and are unable to self consistently reproduce the observed star formation history of these systems (see §7.4.4). We find the introduction of high-z quenching substantially alters the SHMR at low masses. In particular we find these models produce an over-abundance of galaxies in the $\log_{10}(m_*/M_{\odot}) \in [3, 5)$ largely residing in haloes with $M_{\rm h}^{\rm peak} \leq M_{\rm q}$. This is an obvious consequence of our model which quenches star formation in these haloes but does not disrupt their constituent galaxy, effectively preventing their growth to higher masses along the standard SHMR. Without observables to validate such trends the SHMR should not be extrapolated beyond its range of constraint.

7.4.3 Satellite populations

Although our models have been fit assuming a global mass function, the majority of observed low mass systems are satellite galaxies of either the Milky Way or Andromeda. In this section we provide a comparison been the locally observed satellite mass functions and radial distributions compared with simulated systems of similar mass. Our simulated systems are all isolated galaxies (type 0) and selected based on their stellar mass only. When searching for simulated Milky Way and Andromeda analogues we use the stellar and halo

Table 7.4: Typical halo peak $(M_{\rm h}^{\rm peak})$ and current $(M_{\rm h}^{z=0})$ mass for a range of stellar mass intervals in each model. Masses are measured at the distribution peak with 68 per cent interval measured at iso-density levels. $N_{\rm gal}$ shows the number of galaxies in each mass band and $f_{\rm orph}$ indicates the fraction of those galaxies that are orphans. All mass values are expressed in $\log_{10}(m/M_{\odot})$ units

<i>~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~</i>		F	Reference			Hig	gh-z quench	
m_*	$N_{\rm gal}$	$f_{\rm orph}$	$M_{\rm h}^{\rm peak}$	$M_{\rm h}^{z=0}$	$N_{\rm gal}$	$f_{\rm orph}$	$M_{\rm h}^{\rm peak}$	$M_{\rm h}^{z=0}$
$\in [5.0, 5.5)$	133120	0.31	$9.44_{-0.82}^{+0.32}$	$9.55_{-2.07}^{+0.28}$	9717	0.20	$9.60^{+0.23}_{-0.14}$	$9.80^{+0.21}_{-1.31}$
$\in [5.5, 6.0)$	65644	0.26	$9.77^{+0.28}_{-0.79}$	$9.85_{-2.04}^{+0.28}$	11959	0.14	$9.97^{+0.13}_{-0.24}$	$10.02^{+0.21}_{-0.98}$
$\in [6.0, 6.5)$	32887	0.22	$10.09^{+0.24}_{-0.74}$	$10.13^{+0.29}_{-1.91}$	10445	0.12	$10.21_{-0.26}^{+0.12}$	$10.23^{+0.22}_{-0.87}$
$\in [6.5, 7.0)$	16474	0.21	$10.39^{+0.19}_{-0.66}$	$10.40^{+0.30}_{-1.76}$	7815	0.12	$10.42^{+0.13}_{-0.25}$	$10.43_{-0.82}^{+0.23}$
$\in [7.0, 7.5)$	8453	0.19	$10.63^{+0.18}_{-0.54}$	$10.62^{+0.31}_{-1.59}$	5516	0.12	$10.61^{+0.13}_{-0.23}$	$10.61\substack{+0.25\\-0.78}$
$\in [7.5, 8.0)$	4524	0.19	$10.85_{-0.42}^{+0.16}$	$10.84_{-1.39}^{+0.34}$	3586	0.13	$10.81^{+0.13}_{-0.21}$	$10.81_{-0.71}^{+0.27}$
$\in [8.0, 8.5)$	2417	0.17	$11.05_{-0.27}^{+0.15}$	$11.03^{+0.36}_{-0.95}$	2360	0.11	$11.00^{+0.13}_{-0.17}$	$10.99^{+0.28}_{-0.60}$
$\in [8.5, 9.0)$	1497	0.13	$11.23_{-0.18}^{+0.14}$	$11.22_{-0.68}^{+0.35}$	1586	0.10	$11.21_{-0.16}^{+0.10}$	$11.18^{+0.27}_{-0.45}$

Table 7.5: Observational estimates of Milky Way and Andromeda stellar and halo masses. Milky Way stellar and halo mass estimates are from Bland-Hawthorn & Gerhard (2016). Andromeda stellar and halo mass estimates are from Sick et al. (2015) and Diaz et al. (2014) respectively.

Name	m_*	$M_{\rm h}^{\rm obs}$	$M_{\rm h}^{\rm sim}$
Milky Way	$10.70^{+0.08}_{-0.10}$	$12.11_{-0.11}^{+0.09}$	$12.43_{-0.31}^{+0.27}$
Andromeda	$11.01\substack{+0.09\\-0.08}$	$12.23_{-0.08}^{+0.07}$	$12.86_{-0.29}^{+0.27}$

mass ranges noted in Table 7.5.

Figure 7.5 shows the cumulative satellite mass function for the Milky Way (left panel) and Andromeda (right panel) under three different host selection criteria. In the first case we select simulated hosts based on stellar mass only (solid lines). Here we find that the reference model significantly over predicts the number of satellites for both Milky Way and Andromeda analogues. Enabling high-z quenching substantially improves agreement with observation, exhibiting only a mild over prediction abundances for satellites $\log_{10}(m_*/M_{\odot}) \lesssim 7.0$ around Andromeda analogues. In the second case we select simulated hosts based on halo mass (dashed lines). With respect to Milky Way the reference model is consistent with observation for $\log_{10}(m_*/M_{\odot}) \gtrsim 7.0$ with a significant over prediction of satellite abundances at lower masses. Although halo mass selection does improve overall normalisation compared with observations there is a much steeper slope which drives the over abundances at the lowest masses. Meanwhile, the high-z quenching model reproduces the general trend of the observed SMF but tends to undercut at all masses compared to the observed data. In the case of comparison to Andromeda we find that selecting hosts based on halo mass under predicts the abundances of satellites in both the reference case and high-z quenching case. Finally, we check the case where we select hosts based on both



Figure 7.5: Cumulative satellite stellar mass function for satellites with $5 \leq \log_{10}(m_*/M_{\odot}) < 9$ around Milky Way (left panel) and Andromeda (right panel) like hosts. The grey line and shaded region shows the observed stellar mass functions with Poisson error about the number count in each mass bin. Line colour indicates whether the mass function was generated from the reference model (orange) or the high-z quenching model (blue). Solid lines show the mass function when selecting hosts based on stellar mass, dashed lines by halo mass and dotted lines by both stellar and halo mass.



Figure 7.6: Normalised cumulative radial distribution of satellites with $5 \leq \log_{10}(m_*/M_{\odot}) < 9$ within 300 kpc of Milky Way (left panel) and Andromeda (right panel) like hosts. The grey line and shaded region shows the observed radial distribution of satellites with Poisson error about the number count in each mass bin. Line colour indicates whether the sample was generated from the reference model (orange) or the logistic model (blue). Solid lines show the distribution when selecting hosts based on stellar mass, dashed lines by halo mass and dotted lines by both stellar and halo mass. Each line is normalised by the total number of satellites identified in the specified mass and radial ranges.

stellar mass and halo mass. For the Milky Way we find that selecting based on both mass measures produces a mass function very similar to the halo mass only case. For Andromeda like systems we are unable to locate any simulated systems that meet both the stellar mass *and* halo mass estimates from observation. In general, our simulated analogues (for both the Milky Way and Andromeda) reside in more massive haloes on average. Table 7.5 shows the average simulated halo masses where we can see there is no overlap in the observed and simulated halo ranges for Andromeda analogues. The implications of this mismatch are discussed in §7.5.2

Figure 7.6 illustrates the cumulative satellite distribution around Milky Way (left panel) and Andromeda (right panel) like hosts. Here we find that when selecting hosts based on stellar mass both the reference and high-z quenching model produce a satellite distribution in agreement with observation. For Milky Way analogues even selecting based on halo mass continues to produce a satellite distribution in line with those observed. For Andromeda analogues we find that selecting based on halo mass tends to produce a slightly more centrally concentrated satellite distribution. This could of course be driven by the large mismatch in the simulated vs observed halo masses for these analogues.



Figure 7.7: Star formation history in galaxies over time under the reference model (orange panels) and high-z quenching model (blue panels) in three mass intervals. Upper panels show the average star formation rate as a function of redshift. Lower panels show the average cumulative mass growth for the same sample of galaxies. Grey lines illustrate the observed star formation history and 68 per cent confidence interval for local dwarves as computed by Weisz et al. (2014)

7.4.4 Star formation history

In this section we evaluate the star formation history (SFH) of low mass systems in this model. We trace back the main branches of the z = 0 galaxy population and compute the average SFH, and stellar mass build up in discrete mass bins.

The top row panel of Figure 7.7 compares the star formation history of the reference EMERGE model (orange lines) with that of our high-z quenching model (blue lines). Both models share some qualitative similarities. In both model variants we see that less massive systems experience peak SFR at higher redshift. For equivalent mass ranges the reference model experiences peak SFR earlier than the high-z quenching model. For the two most massive bins the reference case also exhibits a higher peak SFR, for the lowest mass range the high-z quenching model has a higher peak SFR. Further, we can see that in the case of high-z quenching, peak SFR tends to occur very near to the redshift where M_q^{peak} is reached.

The bottom panels on Figure 7.7 illustrates the cumulative star formation history as a fraction of the z = 0 stellar mass for the same sample of simulated galaxies as in the top panel. Although lowest mass systems in the reference model experience their peak SFR at higher redshift than the high-z quenching model, we find that low mass systems in the



Figure 7.8: The 90 per cent formation timescale τ_{90} [Gyr], red contours indicated galaxies where the formation timescale occurred before halo peak mass. Blue contours show systems where formation occurred after peak halo mass. Black lines show the average formation timescale for all systems. Grey points show the formation timescale for individual observed dwarf satellites (see table 7.1)

high-z quenching model are build up more rapidly than in the standard model. However, we find that both models are largely consistent with the observed average SFH determined by Weisz et al. (2014) (grey lines and shaded region).

Another, perhaps more revealing measure of galaxy growth is the formation time. As a galaxy does not form in an instant we instead attribute the formation time as the 90 per cent formation timescale τ_{90} . This value specifies the cosmological time when a galaxy reached 90 per cent of its present day mass. Figure 7.8 compares the formation time of Milky Way and Andromeda satellites (grey points), with the satellite populations around our simulated analogues. In the first panel we can see that the reference case in general produces a mean formation time that is too long compared with observation, while also failing to reproduce the diversity of formation times seen in the data. The coloured lines indicated the iso-proportion contours for the formation time of galaxies in our simulations. Particularly in the case of the reference model we can see that EMERGE produces a strongly bi-modal distribution of formation times at low masses. As previously stated, galaxy mass can be directly linked with its associated peak halo mass M_h^{peak} . In this figure we separate the galaxy sample to distinguish whether a galaxy formed before or after M_h^{peak} . The blue contours indicate systems that formed after the halo reached M_h^{peak} and the red contours show galaxies that formed prior to M_h^{peak} . The reference case shows that a majority of low


Figure 7.9: Quenched fractions as a function of mass at z = 0 for our reference (orange line) and high-z quenching model (blue line).

mass systems are formed at late times after halo peak mass indicating that these objects experience significant stellar growth, coasting on their remaining cold gas reservoir, even while their host halo is no longer gaining mass. When high-z quenching is enabled we can see that although the majority of galaxies still reach τ_{90} after peak mass, we can see that a significant portion of these systems form at much earlier times with a mean formation time in much better agreement with observed data. Furthermore, we can see that enabling high-z quenching does a much better job reproducing the diversity of formation times seen in local satellites. That these low mass galaxies primarily form after M_h^{peak} but at early times, indicates that the majority of these systems have their star formation shut off during their coasting phase as opposed to during active buildup.

The consequences of these early formation times and premature quenching should therefore be observable in the present day star forming properties of our simulated galaxies. Figure 7.9 shows the z = 0 quenched fractions as a function of mass between our two model options. Although quenched fractions are not constrained for $\log_{10}(m_*/M_{\odot}) < 8.5$ there are still observed trends at low masses that we should consider. Unsurprisingly, both models exhibit nearly identical quenched fractions in the range used to constrain the model. For dwarf galaxies we see that the reference model possess a higher quenched fraction than the high-z quenching model for $6 \leq \log_{10}(m_*/M_{\odot}) \leq 8.0$. Toward lower masses the reference model starts to saturate at near 50 per cent quenched. Conversely, the high-z quenching model shows a rapidly increasing quenched fraction toward low masses, nearing 100 per cent quenched by $\log_{10}(m_*/M_{\odot}) = 5$.

7.5 Discussion and Conclusions

7.5.1 Halo mass threshold for galaxy formation

Our model suggests that there should be very few galaxies $\log_{10}(m_*/M_{\odot}) \gtrsim 5$ with $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M}_{\odot}) \lesssim 9$ which is in agreement with other theoretical models that show high-z quenching becomes important for haloes $\log_{10}(M_{\rm h}^{\rm peak}/M_{\odot}) \lesssim 9 - 10$ (Thoul & Weinberg, 1996; Bullock et al., 2000; Somerville, 2002; Kuhlen et al., 2013; Sawala et al., 2015; Nadler et al., 2019). Our model differs from other empirical models that include high-z quenching in that we do not enforce that all haloes below some threshold should remain dark. The quenching mechanism we implement suppresses star formation as a function of halo mass and time, effectively penalising star formation in late forming haloes. Recent hydrodynamical models indicate that this distinction maybe be necessary in order to reproduce the properties of the ultra-faint dwarf (UFD) galaxies with $\log_{10}(m_*/M_{\odot}) \lesssim 5$ (Garrison-Kimmel et al., 2019; Munshi et al., 2021; Applebaum et al., 2021). Additionally, Using high resolution hydrodynamical simulations Schauer et al. (2020) suggest that a minimum halo mass between $10^6 \sim 10^7 M_{\odot}$ is required to begin star formation, depending on free streaming velocities and the strength of the Lyman-Werner Background. With respect to our work, that would indicate that whatever UFDs exist were likely strongly impacted by a high-z quenching process and reside in haloes with a halo peak mass spanning only ~ 2 orders of magnitude.

7.5.2 Is the Local Group representative?

The work presented so far has assumes that the Local Group is somewhat average with a local mass function that is representative of larger volumes. However we have so far not provided any quantitative analysis of how 'normal' the Local Group is. Figure 7.10 shows how the distribution in local SMF slope relates to the vertical offset (in dex) from the global average for stellar mass selected Milky Way analogues. The local slope is defined by the power-law index measured in a 2 Mpc sphere around each analogue. Solid points indicate Milky Way like systems that *also* host one Andromeda like companion within 1 Mpc. The locally measured slope and offset derived from the data in table 7.1 is indicated by the red point. If our reference model (orange contours/points) is correct it would indicated that our locally observed SMF slope is a substantial outlier compared to simulated Milky Way analogues. However, the locally observed offset from the global SMF is average with respect to simulated systems. Additionally, we can see the that simulated systems that

host an Andromeda companion are more likely to reside closer to the global slope than those without. What is possibly more interesting is that even after fitting our model to Local Group data (blue contours/points) we find that the local SMF still does not posses an average local SMF slope compared with simulated systems, but is still within the simulated range. In this case the observed slope is steeper than the majority of simulated analogues. These deviations could indicate that assuming a global slope from local data is not adequate to constrain dwarf systems. To complicate the matter, halo masses for our simulated analogues tend to be higher than estimates for both Milky Way or Andromeda. This mismatch may could indicate that the Local Group is indeed less dense than comparable counterparts and be subject to an accretion history that is not representative for haloes with comparable mass.

Finally, we have assumed that adopting a spherical volume is sufficient to describe the Local Group. Putman et al. (2021) instead propose a prolate Local Group surface to more accurately define group membership. It is unclear how our simplification might impact our goodness of fit or alter the demographics of Local Group analogues. This would be an interesting area for future work. Additionally, we have assumed that currently confirmed dwarf observations provide sample completeness down to $\log_{10}(m_*/M_{\odot}) \sim 5$ (Tollerud et al., 2008). Other recent works suggest that the current sample of galaxies is incomplete due to selection bias and the local observed mass function ought to be higher (Loveday et al., 2015; Jethwa et al., 2018; Newton et al., 2018; Nadler et al., 2019, 2020; Drlica-Wagner et al., 2020). In testing we found that allowing for a steeper mass function at low masses can negate the need for a high-z quenching model and tend to prefer an earlier quenching with a lower $M_{\rm q}$.

7.5.3 Other model options

The foundation for this work is the introduction of a basic model for high-z quenching that has been implemented in order to match the characteristics of locally observed dwarf systems. This high-z quenching model only shows that we can reproduce locally observed populations by imposing restrictions on star formation at high redshift in a manner consistent with expectations from reionizaton. Limitations in our own implementation might restrict our ability to use this model to place direct constraints on the epoch of reionization. As implemented this model does not and cannot accurately capture the richness of evolution imposed by a reionization process. In particular, all models we tested suppress galaxy formation via an instantaneous quenching mechanism although there is evidence to suggest star formation may be periodic or even continue long after reionization (Geha et al., 2012; Skillman et al., 2017). The need for stars to form beyond reionization may explain why our models tend to favour such a late quenching scale $a_{\rm q}$.

A key tenet of the empirical approach is to select the most simple model required to reproduce the observed data only increasing the complexity when the data demands it. For the purposes of the analysis in this work we have instead adopted the slightly more complicated model in favour of its flexibility to constrain with the inclusion of future observables. There are however several sensible model options that could be explored



Figure 7.10: The relation between the locally measured density offset from the global averaged and the locally observed power law slope of the mass function for Milky Way like systems. Coloured contours show the distribution of these two properties based on the 2 Mpc spherical volume around simulated systems similar to the Milky Way, based only on stellar mass selection. Solid points indicate Milky Way analogues that also host a massive Andromeda like companion within 1 Mpc. The red marker shows the locally measured slope and offset based on the galaxies shown in Table 7.1 and illustrated by the observed mass function in Figure 7.2.

if more information on the abundances and star formation histories of dwarves becomes available.

Augmented baryon conversion efficiency

Another reasonable approach to this problem would be to modify the standard baryon conversion efficiency with an additional term to further suppress the SFR in low mass haloes.

$$\epsilon(M) = 2\epsilon_{\rm N} \left[\left(\frac{M}{M_{\rm q}} \right)^{-\alpha} + \left(\frac{M}{M_{\rm 1}} \right)^{-\beta} + \left(\frac{M}{M_{\rm 1}} \right)^{\gamma} \right]^{-1}$$
(7.6)

While this approach is sensible and fits well within the current paradigm of the model, in testing we found that SMF data alone was insufficient to appropriately constrain the larger number of free parameters of this variant. In particular the redshift evolution of the terms α and M_q is difficult to ascertain without more detailed measurements on the star formation history of dwarves. While the other models we test can quench and prevent star formation in haloes entirely, this method does not allow for empty haloes and does not introduce any explicit quenching mechanism.

Quenching timescales

Instead of shutting down star formation instantaneously we could instead set a timer similar eq. 4.16, allowing a more prolonged period for star formation after 'reionization'. This could likely result in lower values for $M_{\rm q}$ and $a_{\rm q}$.

Disk destruction

As noted in §5 one of the lacking features of this model is the absence of proximity merging. Figure 7.6 does not indicate that we have excessive satellites at small radii, at least not in the mass ranges we explore. However, Garrison-Kimmel et al. (2017) and Sawala et al. (2017) indicate that the presence of disk potential is significant for depletion of satellites in the inner halo. In our model the suppression of the SMF is accomplished between the baryon conversion efficiency, the high-z quenching and through tidal stripping. In reality this is likely a more complex process that includes contributions from mergers or disruption due to the extended physical size of a host system. Here we have not explored the relative contributions of each of our model options on the abundance on dwarves, but including additional mechanisms may alter our best fit model parameters. Although the parameters needed for these various mechanisms may be degenerate and locating observations that can independently constrain each parameter are problematic.

7.5.4 Further analysis

Beyond the exploration of the SHMR there are some additional studies that could be performed to better understand the lives and ultimate fate of galaxies in low mass haloes. The bulk of this work has focused on the global statistics of dwarf galaxies. A deeper study of local satellite dwarves may provide additional validation as these objects are most readily observables. In particular, a better understanding of when the Milky Way and Andromeda acquired their satellites may help relate the distribution and star formation histories of these systems. Understanding these correlations could help verify our predictions or constrain the model.

One direct prediction from our model is the existence of dark haloes. Other works explore this topic (Sawala et al., 2015, 2016a; Fitts et al., 2018) and establish an occupation fraction around ≤ 10 per cent for $\log_{10}(M_{\rm h}^{\rm peak}/M_{\odot}) \leq 9$. The observational confirmation of dark haloes would provide vital information for breaking parameter degeneracies, more accurately determining the impact of reionization physics on the SMHR and the broader relationship between galaxies and haloes. Further work in this area with EMERGE should include studies of dark haloes and their properties.

Furthermore, our models for high-z quenching are rather ad hoc and may not accurately represent how this process should evolve with time. An alternative approach would be to utilise machine learning to a inform a more appropriate parameterisation. GALAXYNET (Moster et al., 2020a) has already shown success for developing more accurate empirical prescriptions (see §7.2.2). Future improvements may make it possible to do the same with the high-z quenching model. The Primary developments needed would be for the neural network to process complete merger trees instead of isolated snapshots, this is necessary as the dwarf population is constrained by a low redshift observable that is driven by a high redshift process.

Finally, for this work our ability to fit the model with our chosen simulation was severely restricted by computational limitations of EMERGE . At present EMERGE , and all other N-body based galaxy formation codes, parallelise by distributing the input halo merger trees across cpu cores such that each individual tree is in shared memory and accessed by only a single task. For simulations in moderate volumes of moderate resolution this is a completely adequate approach. When scaling these codes to more highly resolved simulations there are often a handful of very massive trees that dominate the computation time leaving all other tasks idle while the massive trees are processed. This places a hard limitation on the rate that parameter space can be explored and by extension our ability quantitatively compare models variations. Future studies of dwarf galaxy formation in EMERGE would benefit massively from refactoring the code to perform a more advance branch based parallelism that can allow for a more even distribution of the workload across many cores.

7.5.5 Conclusions

Over the course of this work we have developed an empirical model for galaxy formation with the goal of placing constraints on the stellar mass-halo mass relationship for galaxies down to $\log_{10}(m_*/M_{\odot}) = 5$. We achieve this goal by introducing an empirical quenching mechanism that approximates the effects of a reionization process. This model is constrained by our extension of the global galaxy stellar mass function using the mass function observed in the Local Group volume out to 2 Mpc.

While our model has been fit assuming a global mass function, we show that our simulated galaxy sample reproduces both the number density and radial distributions of satellites for the Milky Way and Andromeda. Additionally, our simulated galaxies are able to reproduce properties of observed systems beyond stellar mass. In particular we find that for the z = 0 sample, nearly 100 per cent of galaxies are quenched at $\log_{10}(m_*/M_{\odot}) \sim 5$. This is 50 per cent higher than a model that does not include high-z quenching. Further, we show that introducing high-z quenching produces a population of galaxies that better matches the star formation histories and formation timescales of observes dwarf satellites.

We show that the SHMR can be extended as a power-law down to at least $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}}) \approx 10$. Models including high-z quenching indicated substantially reduced scatter in peak halo mass at a fixed stellar mass and show that the elimination of high scatter galaxies galaxies in the tail of the distribution are the primary mechanism that flattens the SMF at low masses. Further inspection of the SHMR shows that there should be almost no galaxies down to $\log_{10}(m_*/{\rm M_{\odot}}) = 5$ with a peak halo mass lower than $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}}) \lesssim 9.0$. This is consistent with our model implementations which predict quenching in haloes with $\log_{10}(M_{\rm h}^{\rm peak}/{\rm M_{\odot}}) \lesssim 9.3$ by z = 4 (see table 7.2)

Incorporating a model which suppresses star formation in low mass-high redshift haloes produces a population of galaxies that is consistent with observables in both the total number density and in the bulk star formation history of galaxies. Our results indicate that as few as two additional free parameters are necessary to reproduce the number density of dwarf galaxies. However, we propose a more complex 3 parameter logistic model as the framework for model variations due to its adaptability.

Chapter 8 Summary and Outlook

Over the course of this thesis we have addressed several open problems in galaxy formation physics through the use of the empirical model for galaxy formation EMERGE. We discussed the benefits and inner workings of this model in §4.1.1. Besides their ability to fit observed data by design, one of the main advantages of the empirical approach is the ability to efficiently compare many model options to qualitatively determine which implementation is able to best represent the data. In §4.1.3 we showcased how statistical analysis can be used to determine the most likely option. Tools we put to use when developing model options to incorporate new data.

Chapter 5 showed the results of O'Leary et al. (2021a) where we investigated galaxygalaxy merger rates that results from constructing a self consistent model for galaxy formation that accurately reproduces observed data. Here we showed that EMERGE predicts merger rates that increase with redshift, consistent with other theoretical models. While models converge on merger rates at high masses (within a factor ~ 2) our model exhibits an order of magnitude lower merger rates at low masses. We show that our lower merger rates are not a result of limited N-body resolution or model assumptions on the treatment of orphan galaxies. We conclude that empirical methods that make merger rate predictions based on halo-halo merger rates convolved with the SHMR result in over-predicted merger rates. The precise reason for this over prediction remains uncertain and indicates the current understanding on the relation between halo-halo mergers and galaxy assembly is incomplete. Additionally, we show that most massive galaxies are constructed predominantly through major mergers while low mass galaxies tend to experience mini mergers, if they experience a merger at all. Subsequently we show that the classical minor mergers are relatively infrequent and unimportant for the build-up of mass for almost all mass ranges.

In Chapter 6 we explored pair fractions as an analogue to measure galaxy merger rates from observation. Our model produces a pair fraction that evolves as a power-law exponential function with a rapid rise in the observed pair fraction until $z \approx 2$ then decreases toward higher redshift. We find that a simple mapping from pair fractions to the merger rates can be best accomplished with an observation timescale that decreases linearly with increasing redshift $T_{\text{obs}} \propto (1 + z)$. This observation timescales stands in contrast to other recent works which propose an observation timescale that decreases as $T_{\rm obs} \propto (1+z)^{-2}$ (Snyder et al., 2017) or $T_{\rm obs} \propto H(z)^{-1/3}$ (Jiang et al., 2014). Additionally, we further break down our pair fraction analysis to address merging probabilities and merging timescales on a pair by pair basis. We propose functional forms for each of these values that are able to successfully reproduce the intrinsic merger rate found from merger trees directly. Lastly, our results show that the commonly used 30 kpc outer aperture for observed pair selection does not deliver an adequate sample for reproducing merger rates. To combat this we propose an outer area incompleteness correction which we show is able to successfully improve the quality of results when a poorly selected outer aperture is employed.

Finally, Chapter 7 addresses the relationship between galaxies and halos in the dwarf regime $\log_{10}(m_*/M_{\odot}) \leq 7$. For this research we extend the observed stellar mass function down to $\log_{10}(m_*/M_{\odot}) = 5$ by leveraging observations of dwarf galaxies in the Local Group. We use this data to constrain a new model option in EMERGE that quenches star formation in low mass halos at high redshift. Here we investigate three model options; while all model variants can reproduce the observed SMF we ultimately prefer a more complex model for its robustness to include future data, as well as its ability to reproduce observed trends that were not explicitly included in the fitting routine. The results of this section indicate that the SHMR continues as a power-law relation down to at-least $\log_{10}(M_{\rm h}^{\rm peak}/M_{\odot}) = 10$. However we stress that the inclusion of high-*z* primarily regulates the number density of dwarves by eliminating high scatter galaxies (i.e. high mass to light ratio), ultimately breaking the trends in scatter to the SHMR shown in Moster et al. (2018).

The research discussed in this thesis has improved our understanding of galaxy formation, particularly in areas relating galaxies to dark matter halos. That said, there are several areas in which our work could be improved and where further research could be performed. First and foremost is the disconnect between galaxy-galaxy merger rates and halo-halo merger rates. Despite the fact that many theoretical frameworks can meet baseline observables the order of magnitude disagreement in predicted merger rates creates a gap in our understanding of galaxy formation. We may require additional methods of measuring galaxy merger rates from observed data in order to break this stalemate. In future work it might be possible to determine the merger rates of galaxies by tracing the merger rate of super-massive black holes. Having a model that accurately links black holes to their host systems is *crucial* for making these predictions. Making such definite links between galaxies and their central black holes may enable GW observations to relieve some tension in predictions of galaxy merger rates, especially for low and intermediate mass galaxies. Developing an empirical model to address the black hole galaxy connection could prove a vital tool in these investigations.

Similarly, our proposed fitting functions for $T_{\rm obs}$ begins to fail at low masses. Here a deeper investigation on what drives the observation timescales at low mass would be worthwhile. Understanding why these formulas fail could help in the development of a more universal fitting function for merging probability and timescales and more cleaning accepts primary mass and mass ratio as an argument.

Last, our investigation of the SHMR at dwarf masses offers several promising areas for development. In particular it would be important to include star formation history data directly into the fitting routine. This would enable our model to better constrain the ramp up in mass on the high-z quenching model. Weak lensing observations in search of dark halos could also help determine what process is responsible for the depletion of dwarf galaxies within Λ CDM. One of the largest barriers faced when fitting these models was computational expense. Currently, most tree based semi-analytic and empirical models operate by grouping individual merger trees into shared memory. In the case of very massive boxes or very high resolution simulations this method breaks down as an few massive merger trees can dominate the computation time, leaving all other tasks idle in wait. To address these problems a new parallelisation scheme should be developed for the next generation of cosmological simulations.

Appendix A Observational data

The data tables shown in this section have been adapted from Moster et al. (2018)

Publication	Redshift	Area	IMF	A/P
Wetzel et al. (2012)	0.0 - 0.1	7,97	С	sSFR
Drory et al. (2009)	0.2 - 1.0	1.73	\mathbf{C}	NUV-R-J
Ilbert et al. (2013)	0.2 - 4.0	1.52	\mathbf{C}	$NUV-r^+-J$
Moustakas et al. (2013)	0.2 - 1.0	9	\mathbf{C}	sSFR
Lin et al. (2014)	0.2 - 0.8	70	\mathbf{S}	sSFR
Muzzin et al. (2013)	0.2 - 4.0	1.62	Κ	U-V-J

Table A.1: Observed quenched fractions

Table A.2: Observed clustering

Publication	Survey	Stellar mass	IMF
Li et al. (2006)	SDSS/DR2	K03	С
Guo et al. (2011)	SDSS/DR7	K03	\mathbf{C}
Yang et al. (2012)	SDSS/DR7	B03	С

Table A.3: Observed stellar mass functions

Publication	Redshift	Area	IMF
Baldry et al. (2012)	0.0 - 0.1	143	С
Li & White (2009)	0.0 - 0.2	$6,\!437$	\mathbf{C}
Bernardi et al. (2013)	0.0 - 0.2	4,681	\mathbf{C}
Pérez-González et al. (2008)	0.0 - 4.0	0.184	\mathbf{S}
Ilbert et al. (2010)	0.0 - 2.0	2	\mathbf{C}
Pozzetti et al. (2010)	0.1 - 1.0	1.4	\mathbf{C}
Ilbert et al. (2013)	0.2 - 4.0	1.52	\mathbf{C}
Moustakas et al. (2013)	0.2 - 1.0	9	\mathbf{C}
Muzzin et al. (2013)	0.2 - 4.0	1.62	Κ
Kajisawa et al. (2009)	0.5 - 3.5	0.0364	\mathbf{S}
Santini et al. (2012)	0.6 - 4.5	0.0092	\mathbf{S}
Mortlock et al. (2011)	1.0 - 3.5	0.0121	\mathbf{S}
Marchesini et al. (2009)	1.3 - 4.0	0.1620	Κ
Caputi et al. (2011)	3.0 - 5.0	0.6	\mathbf{S}
Grazian et al. (2015)	3.5 - 7.5	0.1025	\mathbf{S}
Lee et al. (2012)	3.7 - 5.0	0.0889	\mathbf{C}
González et al. (2011)	4.0 - 7.0	0.0150	\mathbf{S}
Duncan et al. (2014)	4.0 - 7.0	0.0167	\mathbf{C}
Song et al. (2016)	4.0 - 8.0	0.0778	\mathbf{S}

Publication	Redshift	Area	IMF	λ
Robotham & Driver (2011)	0.0 - 0.1	833.13	S	UV
Salim et al. (2007)	0.0 - 0.2	741	\mathbf{C}	UV
Gunawardhana et al. (2015)	0.0 - 0.4	144	\mathbf{C}	$H\alpha$
Ly et al. (2011a)	0.8	0.82	\mathbf{S}	$H\alpha$
Zheng et al. (2007)	0.2 - 1.0	0.458	\mathbf{C}	UV/IR
Rujopakarn et al. (2010)	0.0 - 1.2	≤ 9	\mathbf{S}	FIR
Smolčić et al. (2009)	0.1 - 1.3	2	\mathbf{S}	$1.4~\mathrm{GHz}$
Shim et al. (2009)	0.7 - 1.9	0.029	\mathbf{S}	$H\alpha$
Tadaki et al. (2011)	0.0 - 0.2	0.016	\mathbf{S}	$H\alpha$
Sobral et al. (2013)	2.2	≤ 1.68	\mathbf{S}	$H\alpha$
Magnelli et al. (2011)	1.3 - 2.3	0.079	\mathbf{S}	IR
Hayes et al. (2010)	2.2	0.016	\mathbf{S}	$H\alpha$
Karim et al. (2011)	0.2 - 3.0	1.72	\mathbf{C}	$1.4~\mathrm{GHz}$
Ly et al. $(2011b)$	1 - 3	0.242	\mathbf{S}	UV
Kajisawa et al. (2010)	0.5 - 3.5	0.029	\mathbf{S}	UV/IR
Reddy & Steidel (2009)	1.9 - 3.4	0.906	Κ	UV
Burgarella et al. (2013)	0 - 4	≤ 0.6	\mathbf{S}	UV/IR
Cucciati et al. (2012)	0 - 4.5	0.611	\mathbf{S}	UV
Dunne et al. (2009)	0 - 5	0.8	\mathbf{S}	$1.4~\mathrm{GHz}$
Le Borgne et al. (2009)	1 - 5	0.07	\mathbf{S}	IR/mm
van der Burg et al. (2010)	3 - 5	4	\mathbf{S}	UV
Bourne et al. (2017)	0.5 - 6	0.064	\mathbf{C}	UV/IR
Duncan et al. (2014)	4 - 7	≤ 0.017	\mathbf{C}	UV
Oesch et al. (2013)	3.8 - 11	0.045	\mathbf{S}	UV
McLure et al. (2013)	6 - 10	≤ 0.05	S	UV

Table A.4: Observed cosmic star formation rate densities

Table A.5: Observed specific star formation rates

Publication	Redshift	Area	IMF	λ
Salim et al. (2007)	0.0 - 0.2	741	С	UV
Zheng et al. (2007)	0.2 - 1.0	0.458	\mathbf{C}	UV/IR
Twite et al. (2012)	1.0	1.4	\mathbf{C}	$H\alpha$
Noeske et al. (2007)	0.2 - 1.1	0.5	Κ	UV/IR
Tadaki et al. (2011)	2.2	0.016	\mathbf{S}	$H\alpha$
Whitaker et al. (2012)	0.0 - 2.5	0.4	\mathbf{C}	UV/IR
Daddi et al. (2007)	1.4 - 2.5	0.06	\mathbf{S}	UV-1.4 GHz
Salmi et al. (2012)	0.9 - 1.3	0.06	\mathbf{C}	UV
Karim et al. (2011)	0.2 - 3.0	1.72	\mathbf{C}	$1.4~\mathrm{GHz}$
Kajisawa et al. (2010)	0.5 - 3.5	0.029	\mathbf{S}	UV/IR
Reddy et al. (2012)	1.4 - 3.7	0.44	\mathbf{S}	$\rm UV/IR$
Feulner et al. (2005)	0.4 - 5.0	0.014	\mathbf{S}	UV/IR
Lee et al. (2012)	3.3 - 4.3	5.3	\mathbf{C}	UV/IR
González et al. (2011)	4 - 6	0.015	\mathbf{S}	UV/IR
Schaerer & de Barros (2010)	6 - 8	2	\mathbf{S}	UV
Labbé et al. (2013)	8	0.04	\mathbf{S}	$\rm UV/IR$
McLure et al. (2011)	6 - 8.7	0.013	\mathbf{C}	UV
Duncan et al. (2014)	4 - 7	0.017	С	UV

Appendix B Displacing particles on a grid

Setting up initial conditions is a critical and necessary step to performing an N-body simulation. He we go through the complete process of setting up initial particle positions and velocities. Here we assume a generic Λ CDM cosmological box with periodic boundary conditions. After specifying the desired cosmology, box-size and particle count particles are place in a uniform grid pattern¹ throughout the simulation volume.

First we must understand how cosmological perturbations evolve with time². We are able to explore the evolution or perturbations in the linear regime with the fluid equations:

$$\frac{\mathrm{D}\rho}{\mathrm{D}t} = -\rho \nabla_{\mathbf{r}} \cdot \mathbf{u} \quad \text{(continuity)}, \tag{B.1}$$

$$\frac{\mathrm{D}\mathbf{u}}{\mathrm{D}t} = -\frac{\nabla_{\mathbf{r}}P}{\rho} - \nabla_{\mathbf{r}}\phi \quad \text{(Euler)}, \tag{B.2}$$

$$\nabla_{\mathbf{r}}^2 \phi = 4\pi G \rho \quad (\text{Poisson}), \tag{B.3}$$

where **u** is velocity, ρ is the density, **r** is the proper coordinate, ϕ is the potential and

$$\frac{\mathrm{D}}{\mathrm{D}t} \equiv \frac{\partial}{\partial_t} + \mathbf{u} \cdot \nabla_{\mathbf{r}} \tag{B.4}$$

is the convective derivative. We can then introduce the perturbations as $\rho = \rho_0 + \delta \rho$ and $\mathbf{u} = \mathbf{u}_0 + \delta \mathbf{u}$, as well as the dimensionless over-density field

$$\delta(\mathbf{x},t) \equiv \frac{\bar{\rho}(t) - \rho(\mathbf{x},t)}{\bar{\rho}(t)} . \tag{B.5}$$

We can then transfer into a co-moving coordinate system, that is coordinates which remain fixed with respect to an expanding universe:

$$\mathbf{r} = a\mathbf{x},\tag{B.6}$$

¹More sophisticated initial position methods do exist in order to avoid introducing geometric effects from an initial grid like layout.

²This derivation follows the notes of Dr. Benjamin P. Moster as compiled for lectures in *Numeri*cal Cosmology and Galaxy Formation, taught during summer semester 2016 at the Ludwig-Maximilians-Universität München.

$$\mathbf{u} = \dot{a}\mathbf{x} + \mathbf{x}, \quad \mathbf{v} \equiv a\dot{\mathbf{x}} . \tag{B.7}$$

After applying these to our fluid equations we can collect first order terms to show that the dynamical equations become

$$\dot{\delta} = -\frac{1}{a} \nabla_{\mathbf{x}} \cdot \mathbf{v}$$
 (continuity), (B.8)

$$\dot{\mathbf{v}} + 2\frac{\dot{a}}{a}\mathbf{v} = -\frac{\nabla_{\mathbf{x}}\delta\phi}{a} - \frac{\nabla_{\mathbf{x}}\delta P}{\bar{\rho}a}$$
 (Euler), (B.9)

$$\nabla_{\mathbf{x}}^2 \delta \phi = 4\pi G \bar{\rho} a^2 \delta \quad \text{(Poisson)}. \tag{B.10}$$

Using the sound speed $c_s^2 \equiv \partial P / \partial \rho$, and considering a plane wave disturbance

$$\delta(\mathbf{x}) = \int \frac{\mathrm{d}^3 \mathbf{k}}{(2\pi)^3} \delta \tilde{\mathbf{k}} e^{-i\mathbf{k}\mathbf{x}} , \qquad (B.11)$$

we can show the perturbations grow according to

$$\ddot{\delta} + 2\frac{\dot{a}}{a}\dot{\delta} - \frac{c_s}{a^2}\nabla^2\delta = 4\pi G\bar{\rho}\delta.$$
(B.12)

If all of the density fluctuations are small at very early times, we can assume that only growing modes are significant at later times. The density field can that be written in the simplified form

$$\delta(\mathbf{x}, a) = D(a)\delta_0(\mathbf{x}) \tag{B.13}$$

where D(a) is the linear growth function.

This is known as the **Zel'dovich Approximation**, and this allows us to determine the density field at any time if we have information on the initial field $\delta_0(\mathbf{x})$. We can use equation B.13, to solve B.12,

$$\ddot{D} + 2\frac{\dot{a}}{a}\dot{D} - \frac{c_s^2}{a}\nabla^2 D = 4\pi G\bar{\rho}D \tag{B.14}$$

Plugging this into equation B.10,

$$\nabla^2 \phi = 4\pi G \bar{\rho} a^2 \delta = 4\pi G \frac{\bar{\rho}}{a^3} a^2 \delta_0 = \frac{D}{a} \nabla^2 \phi_0 \tag{B.15}$$

thus,

$$\phi(a) = \frac{D(a)}{a}\phi_0 . \tag{B.16}$$

In the case where $P \ll \rho$ equation B.9 can be re-written as,

$$\dot{\mathbf{v}} + \frac{\dot{a}}{a}\mathbf{v} = -\frac{\nabla_{\mathbf{x}}\phi}{a} \tag{B.17}$$

which can be integrated to obtain

$$\mathbf{v} = -\frac{\nabla_{\mathbf{x}}\phi_0}{a} \int \frac{D}{a} \mathrm{d}t \;. \tag{B.18}$$

Under the same assumption we can see that $c_s^2 = 0$ and equation B.14 becomes

$$\ddot{D} + 2\frac{\dot{a}}{a}\dot{D} = 4\pi G\bar{\rho}D , \qquad (B.19)$$

which can similarly be integrated to produce

$$a^2 \dot{D} = 4\pi G \bar{\rho}_0 \int \frac{D}{a} \mathrm{d}t \;. \tag{B.20}$$

We are subsequently able to combine equations B.17 and B.14 to yield

$$\dot{\mathbf{x}} = \frac{\mathbf{v}}{a} = -\frac{\nabla_{\mathbf{x}}\phi_0 \dot{D}}{4\pi G\bar{\rho}_0} \ . \tag{B.21}$$

Finally we can integrate this to acquire the displacements to particle coordinates, known as the *Zel'dovich Approximation*:

$$\Psi = \mathbf{x} - \mathbf{x}_0 = -\frac{a\nabla_{\mathbf{x}}\phi}{4\phi G\bar{\rho}_0} \,. \tag{B.22}$$

We obtain the velocity displacements by first transforming our position displacements into Fourier space,

$$\tilde{\Psi}_{\mathbf{k}} = -\frac{ai\mathbf{k}\phi_{\mathbf{k}}}{4\phi G\bar{\rho}_0} , \qquad (B.23)$$

which when applied to equation B.10 produces

$$\tilde{\Psi}_{\mathbf{k}} = \frac{i\mathbf{k}}{\mathbf{k}^2}\tilde{\delta}_{\mathbf{k}} . \tag{B.24}$$

We can relate this expression to the velocity displacements as

$$\tilde{\mathbf{v}}_{\mathbf{k}} = a \frac{D}{D} \tilde{\boldsymbol{\Psi}}_{\mathbf{k}} . \tag{B.25}$$

By definition, we know that the power spectrum is related to the Fourier space over-density as

$$P(\mathbf{k}) = \langle |\delta_{\mathbf{k}}|^2 \rangle , \qquad (B.26)$$

which leads to

$$\tilde{\delta}_{\mathbf{k}} = \sqrt{P(\mathbf{k})} \cdot Re^{i\phi} = \sqrt{P(\mathbf{k})}(R_1 + iR_2) .$$
(B.27)

Here, in Fourier space, we can sample a random Gaussian distribution for R_1 and R_2 . We can then transform these displacements back into real-space and apply he corrections to our initial grid of particles.

Appendix C

high-z quenching model fits

C.1 Parameter correlations



Figure C.1: Parameter correlations for the 'instantaneous' high-z quenching model. In this model we find a positive correlation between a_q and M_q



Figure C.2: Parameter correlations for the 'lin-max' high-z quenching model. In this model we find a positive correlation between a_q and M_q , negative correlation between a_q and R_q , and negative correlation between M_q and R_q .



Figure C.3: Parameter correlations for the 'logistic' high-z quenching model. In this model we find a positive correlation between $a_{\rm q}$ and $M_{\rm q}$, negative correlation between $a_{\rm q}$ and $R_{\rm q}$, and no correlation between $M_{\rm q}$ and $R_{\rm q}$.

C.2 Comparison to observed data



Figure C.4: Comparison of the reference EMERGE configuration, high-z quenching models and the observed SMF data from table A.3.



Figure C.5: Comparison of the reference EMERGE configuration, high-z quenching models and the observed CSFRD data from table A.4.



Figure C.6: Comparison of the reference EMERGE configuration, high-z quenching models and the observed SSFR data from table A.5.



Figure C.7: Comparison of the reference EMERGE configuration, high-z quenching models and the observed quenched fraction data from table A.1.

Bibliography

- Abadi M. G., Navarro J. F., Steinmetz M., Eke V. R., 2003, ApJ, 591, 499
- Abraham R. G., van den Bergh S., Nair P., 2003, ApJ, 588, 218
- Abruzzo M. W., Narayanan D., Davé R., Thompson R., 2018, arXiv e-prints, p. arXiv:1803.02374
- Akaike H., 1974, IEEE Transactions on Automatic Control, 19, 716
- Amorisco N. C., 2017, MNRAS, 464, 2882
- Amorisco N. C., 2019, MNRAS, 482, 2978
- Applebaum E., Brooks A. M., Christensen C. R., Munshi F., Quinn T. R., Shen S., Tremmel M., 2021, ApJ, 906, 96
- Arnett D., 1996, Supernovae and Nucleosynthesis: An Investigation of the History of Matter from the Big Bang to the Present. Princeton University Press
- Bacon R., et al., 2010, in McLean I. S., Ramsay S. K., Takami H., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III. p. 773508, doi:10.1117/12.856027
- Baldry I. K., et al., 2012, MNRAS, 421, 621
- Barnes J., Hut P., 1986, Nature, 324, 446
- Bayes T., 1763, Philosophical Transactions Roy. Soc., 53, 370
- Becker R. H., et al., 2001, AJ, 122, 2850
- Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013a, ApJ, 762, 109
- Behroozi P. S., Wechsler R. H., Conroy C., 2013b, ApJ, 762, L31
- Behroozi P. S., Wechsler R. H., Wu H.-Y., Busha M. T., Klypin A. A., Primack J. R., 2013c, ApJ, 763, 18
- Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, MNRAS, 488, 3143

- Behroozi P., Hearin A., Moster B. P., 2021, arXiv e-prints, p. arXiv:2101.05280
- Bell E. F., de Jong R. S., 2001, ApJ, 550, 212
- Benson A. J., 2012, New Astron., 17, 175
- Benson A. J., Cole S., Frenk C. S., Baugh C. M., Lacey C. G., 2000, MNRAS, 311, 793
- Bernardi M., Meert A., Sheth R. K., Vikram V., Huertas-Company M., Mei S., Shankar F., 2013, MNRAS, 436, 697
- Bernardi M., Domínguez Sánchez H., Brownstein J. R., Drory N., Sheth R. K., 2019, arXiv e-prints, p. arXiv:1904.11996
- Binney J., Tremaine S., 1987, Galactic dynamics. Princeton University Press
- Binney J., Tremaine S., 2008, Galactic Dynamics: Second Edition. Princeton University Press
- Bland-Hawthorn J., Gerhard O., 2016, ARA&A, 54, 529
- Bluck A. F. L., et al., 2019, MNRAS, 485, 666
- Bois M., et al., 2011, MNRAS, 416, 1654
- Bourne N., et al., 2017, MNRAS, 467, 1360
- Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, MNRAS, 370, 645
- Boylan-Kolchin M., Ma C.-P., Quataert E., 2008, MNRAS, 383, 93
- Brigham E. O., Morrow R. E., 1967, IEEE Spectrum, 4, 63
- Bryan G. L., Norman M. L., 1998, ApJ, 495, 80
- Bullock J. S., Kravtsov A. V., Weinberg D. H., 2000, ApJ, 539, 517
- Bullock J. S., Kolatt T. S., Sigad Y., Somerville R. S., Kravtsov A. V., Klypin A. A., Primack J. R., Dekel A., 2001, MNRAS, 321, 559
- Bundy K., Fukugita M., Ellis R. S., Targett T. A., Belli S., Kodama T., 2009, ApJ, 697, 1369

Burgarella D., et al., 2013, A&A, 554, A70

Campbell D., van den Bosch F. C., Padmanabhan N., Mao Y.-Y., Zentner A. R., Lange J. U., Jiang F., Villarreal A., 2018, MNRAS, 477, 359

- Caputi K. I., Cirasuolo M., Dunlop J. S., McLure R. J., Farrah D., Almaini O., 2011, MNRAS, 413, 162
- Carlberg R. G., 1990, ApJ, 359, L1
- Chabrier G., 2003, PASP, 115, 763
- Chandrasekhar S., 1943, Reviews of Modern Physics, 15, 1
- Choi E., Somerville R. S., Ostriker J. P., Naab T., Hirschmann M., 2018, ApJ, 866, 91
- Cole A. A., Weisz D. R., Dolphin A. E., Skillman E. D., McConnachie A. W., Brooks A. M., Leaman R., 2014, ApJ, 795, 54
- Conroy C., Gunn J. E., 2010, ApJ, 712, 833
- Conroy C., Wechsler R. H., 2009, ApJ, 696, 620
- Conselice C. J., Bershady M. A., Dickinson M., Papovich C., 2003, AJ, 126, 1183
- Conselice C. J., Rajgor S., Myers R., 2008, MNRAS, 386, 909
- Conselice C. J., Yang C., Bluck A. F. L., 2009, MNRAS, 394, 1956
- Cooper A. P., D'Souza R., Kauffmann G., Wang J., Boylan-Kolchin M., Guo Q., Frenk C. S., White S. D. M., 2013, MNRAS, 434, 3348
- Cooray A., 2006, MNRAS, 365, 842
- Croton D. J., et al., 2006, MNRAS, 365, 11
- Cucciati O., et al., 2012, A&A, 539, A31
- D'Souza R., Bell E. F., 2018, Nature Astronomy, 2, 737
- Daddi E., et al., 2007, ApJ, 670, 156
- Davis S. W., Laor A., 2011, ApJ, 728, 98
- De Lucia G., Blaizot J., 2007, MNRAS, 375, 2
- Deason A. J., Belokurov V., Evans N. W., Johnston K. V., 2013, ApJ, 763, 113
- Diaz J. D., Koposov S. E., Irwin M., Belokurov V., Evans N. W., 2014, MNRAS, 443, 1688
- Dicke R. H., Peebles P. J. E., Roll P. G., Wilkinson D. T., 1965, ApJ, 142, 414
- Drlica-Wagner A., et al., 2020, ApJ, 893, 47
- Drory N., et al., 2009, ApJ, 707, 1595

- Dubois Y., et al., 2014, MNRAS, 444, 1453
- Duncan K., et al., 2014, MNRAS, 444, 2960
- Duncan K., et al., 2019, ApJ, 876, 110
- Dunne L., et al., 2009, MNRAS, 394, 3
- Eggen O. J., Lynden-Bell D., Sandage A. R., 1962, ApJ, 136, 748
- Einasto J., Kaasik A., Saar E., 1974, Nature, 250, 309
- Einstein A., 1916, Annalen der Physik, 354, 769
- Elson E. C., Bassett B. A., van der Heyden K., Vilakazi Z. Z., 2007, A&A, 464, 1167
- Fakhouri O., Ma C.-P., 2008, MNRAS, 386, 577
- Fakhouri O., Ma C.-P., Boylan-Kolchin M., 2010, MNRAS, 406, 2267
- Fattahi A., Navarro J. F., Frenk C. S., 2020, MNRAS, 493, 2596
- Ferrarese L., Merritt D., 2000, ApJ, 539, L9
- Ferré-Mateu A., Forbes D. A., McDermid R. M., Romanowsky A. J., Brodie J. P., 2019, ApJ, 878, 129
- Feulner G., Gabasch A., Salvato M., Drory N., Hopp U., Bender R., 2005, ApJ, 633, L9
- Fitts A., et al., 2018, MNRAS, 479, 319
- Fixsen D. J., 2009, ApJ, 707, 916
- Font A. S., et al., 2008, MNRAS, 389, 1619
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
- Franx M., van Dokkum P. G., Förster Schreiber N. M., Wuyts S., Labbé I., Toft S., 2008, ApJ, 688, 770
- Friedmann A., 1922, Zeitschrift fur Physik, 10, 377
- Friedmann A., 1924, Zeitschrift fur Physik, 21, 326
- Gallart C., et al., 2015, ApJ, 811, L18
- Gamow G., 1946, Physical Review, 70, 572
- Gao F., et al., 2020, A&A, 637, A94
- Garrison-Kimmel S., et al., 2017, MNRAS, 471, 1709

BIBLIOGRAPHY

- Garrison-Kimmel S., et al., 2019, MNRAS, 487, 1380
- Gebhardt K., et al., 2000, ApJ, 539, L13
- Geha M., Blanton M. R., Yan R., Tinker J. L., 2012, ApJ, 757, 85
- Genel S., Genzel R., Bouché N., Naab T., Sternberg A., 2009, ApJ, 701, 2002
- Genel S., Bouché N., Naab T., Sternberg A., Genzel R., 2010, ApJ, 719, 229
- Georgy C., et al., 2013, A&A, 558, A103
- González V., Labbé I., Bouwens R. J., Illingworth G., Franx M., Kriek M., 2011, ApJ, 735, L34
- Goodman J., Weare J., 2010, Communications in Applied Mathematics and Computational Science, 5, 65
- Gott III J. R., Thuan T. X., 1976, ApJ, 204, 649
- Grazian A., et al., 2015, A&A, 575, A96
- Grylls P. J., Shankar F., Conselice C., 2020, arXiv e-prints, p. arXiv:2001.06017
- Gunawardhana M. L. P., et al., 2015, MNRAS, 447, 875
- Gunn J. E., Gott III J. R., 1972, ApJ, 176, 1
- Gunn J. E., Peterson B. A., 1965, ApJ, 142, 1633
- Guo Q., et al., 2011, MNRAS, 413, 101
- Hahn O., Abel T., 2011, MNRAS, 415, 2101
- Haid S., Walch S., Seifried D., Wünsch R., Dinnbier F., Naab T., 2018, MNRAS, 478, 4799
- Harrison C. M., 2017, Nature Astronomy, 1, 0165
- Hayes M., Schaerer D., Ostlin G., 2010, A&A, 509, L5
- Helmi A., Babusiaux C., Koppelman H. H., Massari D., Veljanoski J., Brown A. G. A., 2018, Nature, 563, 85
- Hendel D., Johnston K. V., Patra R. K., Sen B., 2019, MNRAS, 486, 3604
- Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R., Guo Q., Lemson G., Springel V., Overzier R., 2015, MNRAS, 451, 2663
- Hilz M., Naab T., Ostriker J. P., Thomas J., Burkert A., Jesseit R., 2012, MNRAS, 425, 3119

- Hilz M., Naab T., Ostriker J. P., 2013, MNRAS, 429, 2924
- Hirschmann M., Khochfar S., Burkert A., Naab T., Genel S., Somerville R. S., 2010, MNRAS, 407, 1016
- Hirschmann M., Dolag K., Saro A., Bachmann L., Borgani S., Burkert A., 2014, MNRAS, 442, 2304
- Hirschmann M., Charlot S., Feltre A., Naab T., Choi E., Ostriker J. P., Somerville R. S., 2017, MNRAS, 472, 2468
- Hopkins P. F., et al., 2010a, ApJ, 715, 202
- Hopkins P. F., et al., 2010b, ApJ, 724, 915
- Hopkins P. F., et al., 2018, MNRAS, 480, 800
- Howlett C., Lewis A., Hall A., Challinor A., 2012, J. Cosmology Astropart. Phys., 4, 027
- Hu W., Fukugita M., Zaldarriaga M., Tegmark M., 2001, ApJ, 549, 669
- Hubble E. P., 1925, Popular Astronomy, 33, 252
- Hubble E., 1929, Proceedings of the National Academy of Science, 15, 168
- Ibáñez-Mejía J. C., Mac Low M.-M., Klessen R. S., 2021, arXiv e-prints, p. arXiv:2108.04967
- Ilbert O., et al., 2010, ApJ, 709, 644
- Ilbert O., et al., 2013, A&A, 556, A55
- Janka H.-T., 2012, Annual Review of Nuclear and Particle Science, 62, 407
- Jeans J. H., 1902, Philosophical Transactions of the Royal Society of London Series A, 199, 1
- Jesseit R., Cappellari M., Naab T., Emsellem E., Burkert A., 2009, MNRAS, 397, 1202
- Jethwa P., Erkal D., Belokurov V., 2018, MNRAS, 473, 2060
- Jiang C. Y., Jing Y. P., Han J., 2014, ApJ, 790, 7
- Kajisawa M., et al., 2009, ApJ, 702, 1393
- Kajisawa M., Ichikawa T., Yamada T., Uchimoto Y. K., Yoshikawa T., Akiyama M., Onodera M., 2010, ApJ, 723, 129
- Kampczyk P., et al., 2007, ApJS, 172, 329

- Karachentsev I. D., Kudrya Y. N., Karachentseva V. E., Mitronova S. N., 2006, Astrophysics, 49, 450
- Karademir G. S., Remus R.-S., Burkert A., Dolag K., Hoffmann T. L., Moster B. P., Steinwandel U. P., Zhang J., 2019, MNRAS, 487, 318
- Karim A., et al., 2011, ApJ, 730, 61
- Kartaltepe J. S., et al., 2007, ApJS, 172, 320
- Kartaltepe J. S., et al., 2010, ApJ, 721, 98
- Kazantzidis S., Bullock J. S., Zentner A. R., Kravtsov A. V., Moustakas L. A., 2008, ApJ, 688, 254
- Kennedy J., Eberhart R., 1995, in Proceedings of ICNN'95 International Conference on Neural Networks. pp 1942–1948 vol.4, doi:10.1109/ICNN.1995.488968
- Khalatyan A., Cattaneo A., Schramm M., Gottlöber S., Steinmetz M., Wisotzki L., 2008, MNRAS, 387, 13
- Khochfar S., Silk J., 2009, MNRAS, 397, 506
- King I., 1962, AJ, 67, 471
- King A., Pounds K., 2015, ARA&A, 53, 115
- Kirkpatrick S., Gelatt C. D., Vecchi M. P., 1983, Science, 220, 671
- Kitzbichler M. G., White S. D. M., 2008, MNRAS, 391, 1489
- Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, ApJ, 522, 82
- Knebe A., et al., 2011, MNRAS, 415, 2293
- Koekemoer A. M., et al., 2011, ApJS, 197, 36
- Komatsu E., et al., 2011, ApJS, 192, 18
- Kudritzki R.-P., Puls J., 2000, ARA&A, 38, 613
- Kuhlen M., Madau P., Krumholz M. R., 2013, ApJ, 776, 34
- Labbé I., et al., 2013, ApJ, 777, L19
- Lackner C. N., Cen R., Ostriker J. P., Joung M. R., 2012, MNRAS, 425, 641
- Larson R. B., 1974, MNRAS, 169, 229
- Le Borgne D., Elbaz D., Ocvirk P., Pichon C., 2009, A&A, 504, 727

- Lee J., Yi S. K., 2013, ApJ, 766, 38
- Lee K.-S., et al., 2012, ApJ, 752, 66
- Lemaître G., 1927, Annales de la Société Scientifique de Bruxelles, 47, 49
- Lewis A., 2014, CAMB Notes, http://cosmologist.info/notes/CAMB.pdf
- Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473
- Li C., White S. D. M., 2009, MNRAS, 398, 2177
- Li C., Kauffmann G., Jing Y. P., White S. D. M., Börner G., Cheng F. Z., 2006, MNRAS, 368, 21
- Lin L., et al., 2004, ApJ, 617, L9
- Lin L., et al., 2008, ApJ, 681, 232
- Lin L., et al., 2014, ApJ, 782, 33
- López-Sanjuan C., Balcells M., Pérez-González P. G., Barro G., García-Dabó C. E., Gallego J., Zamorano J., 2009, A&A, 501, 505
- López-Sanjuan C., et al., 2012, A&A, 548, A7
- López-Sanjuan C., et al., 2013, A&A, 553, A78
- López-Sanjuan C., et al., 2015, A&A, 576, A53
- Lotz J. M., Jonsson P., Cox T. J., Primack J. R., 2008, MNRAS, 391, 1137
- Lotz J. M., Jonsson P., Cox T. J., Primack J. R., 2010, MNRAS, 404, 575
- Lotz J. M., Jonsson P., Cox T. J., Croton D., Primack J. R., Somerville R. S., Stewart K., 2011, ApJ, 742, 103
- Loveday J., et al., 2015, MNRAS, 451, 1540
- Ly C., Lee J. C., Dale D. A., Momcheva I., Salim S., Staudaher S., Moore C. A., Finn R., 2011a, ApJ, 726, 109
- Ly C., Malkan M. A., Hayashi M., Motohara K., Kashikawa N., Shimasaku K., Nagao T., Grady C., 2011b, ApJ, 735, 91
- Lynds R., 1971, ApJ, 164, L73
- Macciò A. V., Dutton A. A., van den Bosch F. C., 2008, MNRAS, 391, 1940

- Magnelli B., Elbaz D., Chary R. R., Dickinson M., Le Borgne D., Frayer D. T., Willmer C. N. A., 2011, A&A, 528, A35
- Maller A. H., Katz N., Kereš D., Davé R., Weinberg D. H., 2006, ApJ, 647, 763
- Man A. W. S., Zirm A. W., Toft S., 2016, ApJ, 830, 89
- Mantha K. B., et al., 2018, MNRAS, 475, 1549
- Mantha K. B., et al., 2019, MNRAS, 486, 2643
- Maoz D., 2016, Astrophysics in a Nutshell. Princeton University Press
- Marchesini D., van Dokkum P. G., Förster Schreiber N. M., Franx M., Labbé I., Wuyts S., 2009, ApJ, 701, 1765
- Marian V., et al., 2020, ApJ, 904, 79
- Martin N. F., de Jong J. T. A., Rix H.-W., 2008, ApJ, 684, 1075
- Martin G., Kaviraj S., Devriendt J. E. G., Dubois Y., Pichon C., 2018, MNRAS, 480, 2266
- McConnachie A. W., 2012, AJ, 144, 4
- McLure R. J., et al., 2011, MNRAS, 418, 2074
- McLure R. J., et al., 2013, MNRAS, 432, 2696
- Merritt A., van Dokkum P., Abraham R., Zhang J., 2016, ApJ, 830, 62
- Mo H., van den Bosch F. C., White S., 2010, Galaxy Formation and Evolution. Cambridge University Press
- Molino A., et al., 2014, MNRAS, 441, 2891
- Moody C. E., Romanowsky A. J., Cox T. J., Novak G. S., Primack J. R., 2014, MNRAS, 444, 1475
- Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, ApJ, 524, L19
- Mortlock A., Conselice C. J., Bluck A. F. L., Bauer A. E., Grützbauch R., Buitrago F., Ownsworth J., 2011, MNRAS, 413, 2845
- Moster B. P., Macciò A. V., Somerville R. S., Johansson P. H., Naab T., 2010a, MNRAS, 403, 1009
- Moster B. P., Somerville R. S., Maulbetsch C., van den Bosch F. C., Macciò A. V., Naab T., Oser L., 2010b, ApJ, 710, 903

- Moster B. P., Somerville R. S., Newman J. A., Rix H.-W., 2011, ApJ, 731, 113
- Moster B. P., Naab T., White S. D. M., 2013, MNRAS, 428, 3121
- Moster B. P., Naab T., White S. D. M., 2018, MNRAS, 477, 1822
- Moster B. P., Naab T., Lindström M., O'Leary J. A., 2020a, arXiv e-prints, p. arXiv:2005.12276
- Moster B. P., Naab T., White S. D. M., 2020b, MNRAS, 499, 4748
- Moustakas J., et al., 2013, ApJ, 767, 50
- Mundy C. J., Conselice C. J., Duncan K. J., Almaini O., Häußler B., Hartley W. G., 2017, MNRAS, 470, 3507
- Munshi F., Brooks A., Applebaum E., Christensen C., Sligh J. P., Quinn T., 2021, arXiv e-prints, p. arXiv:2101.05822
- Muzzin A., et al., 2013, ApJ, 777, 18
- Naab T., Ostriker J. P., 2017, ARA&A, 55, 59
- Naab T., Johansson P. H., Ostriker J. P., 2009, ApJ, 699, L178
- Naab T., et al., 2014, MNRAS, 444, 3357
- Nadler E. O., Mao Y.-Y., Green G. M., Wechsler R. H., 2019, ApJ, 873, 34
- Nadler E. O., et al., 2020, ApJ, 893, 48
- Nevin R., Blecha L., Comerford J., Greene J., 2019, ApJ, 872, 76
- Newton O., Cautun M., Jenkins A., Frenk C. S., Helly J. C., 2018, MNRAS, 479, 2853
- Noeske K. G., et al., 2007, ApJ, 660, L43
- O'Leary J. A., Moster B. P., Naab T., Somerville R. S., 2021a, MNRAS, 501, 3215
- O'Leary J. A., Moster B. P., Krämer E., 2021b, MNRAS, 503, 5646
- Oesch P. A., et al., 2013, ApJ, 773, 75
- Oser L., Ostriker J. P., Naab T., Johansson P. H., Burkert A., 2010, ApJ, 725, 2312
- Oser L., Naab T., Ostriker J. P., Johansson P. H., 2012, ApJ, 744, 63
- Ostriker J. P., Peebles P. J. E., Yahil A., 1974, ApJ, 193, L1
- Patton D. R., Carlberg R. G., Marzke R. O., Pritchet C. J., da Costa L. N., Pellegrini P. S., 2000, ApJ, 536, 153
- Peacock J. A., Smith R. E., 2000, MNRAS, 318, 1144
- Peebles P. J. E., 1982, ApJ, 263, L1
- Penoyre Z., Moster B. P., Sijacki D., Genel S., 2017, MNRAS, 468, 3883
- Penzias A. A., Wilson R. W., 1965, ApJ, 142, 419
- Pérez-González P. G., et al., 2008, ApJ, 675, 234
- Pfister H., Dotti M., Laigle C., Dubois Y., Volonteri M., 2020, MNRAS, 493, 922
- Pillepich A., Madau P., Mayer L., 2015, ApJ, 799, 184
- Pillepich A., et al., 2018, MNRAS, 473, 4077
- Planck Collaboration 2016, A&A, 594, A13
- Planck Collaboration et al., 2020a, A&A, 641, A4
- Planck Collaboration et al., 2020b, A&A, 641, A6
- Pozzetti L., et al., 2010, A&A, 523, A13
- Press W. H., Schechter P., 1974, ApJ, 187, 425
- Puls J., Vink J. S., Najarro F., 2008, A&ARv, 16, 209
- Purcell C. W., Kazantzidis S., Bullock J. S., 2009, ApJ, 694, L98
- Putman M. E., Zheng Y., Price-Whelan A. M., Grcevich J., Johnson A. C., Tollerud E., Peek J. E. G., 2021, ApJ, 913, 53
- Reddy N. A., Steidel C. C., 2009, ApJ, 692, 778
- Reddy N. A., Pettini M., Steidel C. C., Shapley A. E., Erb D. K., Law D. R., 2012, ApJ, 754, 25
- Riess A. G., et al., 1998, AJ, 116, 1009
- Robertson H. P., 1935, ApJ, 82, 284
- Robertson H. P., 1936a, ApJ, 83, 187
- Robertson H. P., 1936b, ApJ, 83, 257
- Robotham A. S. G., Driver S. P., 2011, MNRAS, 413, 2570

- Rodriguez-Gomez V., et al., 2015, MNRAS, 449, 49
- Rodriguez-Gomez V., et al., 2016, MNRAS, 458, 2371
- Rodríguez-Puebla A., Primack J. R., Avila-Reese V., Faber S. M., 2017, MNRAS, 470, 651
- Rujopakarn W., et al., 2010, ApJ, 718, 1171
- Salim S., et al., 2007, ApJS, 173, 267
- Salmi F., Daddi E., Elbaz D., Sargent M. T., Dickinson M., Renzini A., Bethermin M., Le Borgne D., 2012, ApJ, 754, L14
- Sambridge M., 2014, Geophysical Journal International, 196, 357
- Santini P., et al., 2012, A&A, 538, A33
- Sawala T., et al., 2015, MNRAS, 448, 2941
- Sawala T., et al., 2016a, MNRAS, 456, 85
- Sawala T., et al., 2016b, MNRAS, 457, 1931
- Sawala T., Pihajoki P., Johansson P. H., Frenk C. S., Navarro J. F., Oman K. A., White S. D. M., 2017, MNRAS, 467, 4383
- Scarlata C., et al., 2007, ApJS, 172, 406
- Schaerer D., de Barros S., 2010, A&A, 515, A73
- Schauer A. T. P., Glover S. C. O., Klessen R. S., Clark P., 2020, arXiv e-prints, p. arXiv:2008.05663
- Schawinski K., Koss M., Berney S., Sartori L. F., 2015, MNRAS, 451, 2517
- Schaye J., et al., 2015, MNRAS, 446, 521
- Schechter P., 1976, ApJ, 203, 297
- Schulze F., Remus R.-S., Dolag K., Burkert A., Emsellem E., van de Ven G., 2018, MNRAS, 480, 4636
- Schwarz G., 1978, Annals of Statistics, 6, 461
- Scoccimarro R., Sheth R. K., Hui L., Jain B., 2001, ApJ, 546, 20
- Sharma R. S., et al., 2021, arXiv e-prints, p. arXiv:2101.01729
- Shi Y., Rieke G., Lotz J., Perez-Gonzalez P. G., 2009, ApJ, 697, 1764

- Shim H., Colbert J., Teplitz H., Henry A., Malkan M., McCarthy P., Yan L., 2009, ApJ, 696, 785
- Sick J., Courteau S., Cuillandre J.-C., Dalcanton J., de Jong R., McDonald M., Simard D., Tully R. B., 2015, in Cappellari M., Courteau S., eds, IAU Symposium Vol. 311, Galaxy Masses as Constraints of Formation Models. pp 82–85 (arXiv:1410.0017), doi:10.1017/S1743921315003440
- Sivia D. S., Skilling J., 2006, Data Analysis A Bayesian Tutorial, 2nd edn. Oxford Science Publications, Oxford University Press
- Skillman E. D., et al., 2017, ApJ, 837, 102
- Smolčić V., et al., 2009, ApJ, 690, 610
- Smoot G. F., et al., 1992, ApJ, 396, L1
- Snyder G. F., Lotz J. M., Rodriguez-Gomez V., Guimarães R. d. S., Torrey P., Hernquist L., 2017, MNRAS, 468, 207
- Sobral D., Smail I., Best P. N., Geach J. E., Matsuda Y., Stott J. P., Cirasuolo M., Kurk J., 2013, MNRAS, 428, 1128
- Soltan A., 1982, MNRAS, 200, 115
- Somerville R. S., 2002, ApJ, 572, L23
- Somerville R. S., Davé R., 2015, ARA&A, 53, 51
- Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, MNRAS, 391, 481
- Song M., et al., 2016, ApJ, 825, 5
- Spiegelhalter D. J., Best N. G., Carlin B. P., Van Der Linde A., 2002, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 64, 583
- Springel V., 2005, MNRAS, 364, 1105
- Springel V., Yoshida N., White S. D. M., 2001, New Astron., 6, 79
- Springel V., et al., 2018, MNRAS, 475, 676
- Steinborn L. K., Hirschmann M., Dolag K., Shankar F., Juneau S., Krumpe M., Remus R.-S., Teklu A. F., 2018, MNRAS, 481, 341
- Steinwandel U. P., Moster B. P., Naab T., Hu C.-Y., Walch S., 2020, MNRAS, 495, 1035
- Stewart K. R., Bullock J. S., Barton E. J., Wechsler R. H., 2009, ApJ, 702, 1005

- Sugiura N., 1978, Communications in Statistics Theory and Methods, 7, 13
- Tadaki K.-I., Kodama T., Koyama Y., Hayashi M., Tanaka I., Tokoku C., 2011, PASJ, 63, 437
- Thoul A. A., Weinberg D. H., 1996, ApJ, 465, 608
- Tollerud E. J., Bullock J. S., Strigari L. E., Willman B., 2008, ApJ, 688, 277
- Toomre A., Toomre J., 1972, ApJ, 178, 623
- Twite J. W., Conselice C. J., Buitrago F., Noeske K., Weiner B. J., Acosta-Pulido J. A., Bauer A. E., 2012, MNRAS, 420, 1061
- Ventou E., et al., 2017, A&A, 608, A9
- Ventou E., et al., 2019, arXiv e-prints, p. arXiv:1909.03706
- Vogelsberger M., et al., 2014, MNRAS, 444, 1518
- Walker A. G., 1937, Proceedings of the London Mathematical Society, 42, 90
- Wang Y., Nadler E. O., Mao Y.-Y., Adhikari S., Wechsler R. H., Behroozi P., 2021, arXiv e-prints, p. arXiv:2102.11876
- Wechsler R. H., Bullock J. S., Primack J. R., Kravtsov A. V., Dekel A., 2002, ApJ, 568, 52
- Weinberg M. D., 2012, Bayesian Analysis, 7, 737
- Weinberg M. D., Yoon I., Katz N., 2013, arXiv e-prints, p. arXiv:1301.3156
- Weinberger R., et al., 2017, MNRAS, 465, 3291
- Weisz D. R., Dolphin A. E., Skillman E. D., Holtzman J., Gilbert K. M., Dalcanton J. J., Williams B. F., 2014, ApJ, 789, 147
- Weisz D. R., Dolphin A. E., Skillman E. D., Holtzman J., Gilbert K. M., Dalcanton J. J., Williams B. F., 2015, ApJ, 804, 136
- Wen Z. Z., Zheng X. Z., 2016, ApJ, 832, 90
- Wetzel A. R., White M., 2010, MNRAS, 403, 1072
- Wetzel A. R., Tinker J. L., Conroy C., 2012, MNRAS, 424, 232
- Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2013, MNRAS, 432, 336
- Whitaker K. E., van Dokkum P. G., Brammer G., Franx M., 2012, ApJ, 754, L29

- White S. D. M., Rees M. J., 1978, MNRAS, 183, 341
- Woo J., Courteau S., Dekel A., 2008, MNRAS, 390, 1453
- Xu C. K., Zhao Y., Scoville N., Capak P., Drory N., Gao Y., 2012, ApJ, 747, 85
- Yang X., Mo H. J., van den Bosch F. C., 2003, MNRAS, 339, 1057
- Yang X., Mo H. J., van den Bosch F. C., Zhang Y., Han J., 2012, ApJ, 752, 41
- Yoon Y., Lim G., 2020, ApJ, 905, 154
- York D. G., et al., 2000, AJ, 120, 1579
- Zel'Dovich Y. B., 1970, A&A, 500, 13
- Zheng X. Z., Bell E. F., Papovich C., Wolf C., Meisenheimer K., Rix H.-W., Rieke G. H., Somerville R., 2007, ApJ, 661, L41
- Zwicky F., 1933, Helvetica Physica Acta, 6, 110
- de Ravel L., et al., 2009, A&A, 498, 379
- van der Burg R. F. J., Hildebrandt H., Erben T., 2010, A&A, 523, A74

Acknowledgements

First and foremost, I would like to thank my girlfriend Laura Terberl for giving me the courage to undertake this journey, and for supporting me through every moment. I would also like to thank my advisers Ben Moster and Klaus Dolag for their guidance and support. Additionally, I thank all of those who struggled along with me in their own work so that we could support one another through success. Especially, I would like to thank Ulrich Steinwandel, Ludwig Böss, Julien Wolf, Kristina Monsch, Felix Schulze, Rhea-Silvia Remus, Stefan Heigl, Aura Obreja, and Apostolos Zormpas for their friendship and (almost) always enlightening discussions. Finally, I would like to thank the coauthors of my published works Thorsten Naab, Rachel Somerville and Eva Krämer.